



City Research Online

City, University of London Institutional Repository

Citation: Fang, L., Lanzolla, G. & Tsanakas, A. (2024). Shared exposures or management fashions? Drivers of cross-industry convergence of textual risk disclosures. .

This is the preprint version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/32319/>

Link to published version:

Copyright: City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

Reuse: Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Shared exposures or management fashions? Drivers of cross-industry convergence of textual risk disclosures

Lei Fang¹, Gianvito Lanzolla², and Andreas Tsanakas²

¹School of Mathematical Sciences, Queen Mary, University of London

²Bayes Business School, City, University of London

This version: 19 February 2024

Abstract

We contribute to the study of systemic risk, by introducing quantitative measures for the convergence in the attention of decision-makers across industries, and for its contextual antecedents. Our analysis is based on textual data (10-K reports from Securities and Exchange Commission of insurers and banks, 2006-2018), giving a snapshot of decision-makers priorities and contexts. First, we theoretically connect systemic risk with the context of decision making through the lens of the Attention-Based View of the firm. Second, through strategic management theory, we establish a framework for antecedents of convergence in attention, and therefore, systemic risk. These antecedents include common trends in the macro environment, threat of substitution, and management fashions. Third, we combine the theoretical framework with machine learning tools to create quantitative measures of convergence in attention and its antecedents. Finally, based on regression and sensitivity analyses, we identify the dominant antecedent of inter-industry systemic risk. We show that shared risk management fashions largely drive convergence in attention. This finding underlines the need for responses to systemic risk engaging a wide range of stakeholders through inclusive governance and raises the spectre of a type of systemic risk being induced by shared legitimacy concerns.

Keywords: Risk analysis, Systemic risk, Sensitivity analysis, Text analysis, Management attention

1 Introduction

While it is well understood that risks within a system can be interlinked, their specifically systemic nature arising from such linkages has recently attracted increasing attention. *Systemic risks* are characterised by complexity, uncertainty, ambiguity, and wide ripple-effects (Aven and Renn, 2020; Renn et al., 2022), with examples ranging from the 2007-08 Financial Crisis (Earle, 2009; Ivashina and Scharfstein, 2010) and the (still unfolding) COVID-19 pandemic (Botzen et al., 2022), to the current and accelerating climate crisis, which interconnects risks in complex ways on a global (and trans-generational) scale (indicatively: Aglietta and Espagne', 2016; Chaudhry et al., 2023; Bachner et al., 2023).

Within that context, important challenges arise for the risk management and governance responses to systemic risk. First, how can systemic risk be measured? Quantitative measures of systemic risk measures mainly focus on econometric approaches, e.g., by reflecting observed statistical associations between traded instruments (see indicatively: [Billio et al., 2012](#); [Brownlees and Engle, 2017](#); [Acharya et al., 2017, 2012](#); [Huang et al., 2009](#)). The associated methodological concerns have become of continued interest in the operational research literature; see the more recent contributions of [Ye et al. \(2024\)](#) and [Calabrese and Osmetti \(2019\)](#). Furthermore, a fundamental element of systemic risks is the structural connectedness of affected institutions, which goes beyond statistical association. Addressing this point is a burgeoning literature that combines econometric study of systemic risk with network analysis ([Bonaccolto et al., 2022](#); [Simaan et al., 2020](#); [Tang et al., 2022](#)). A second challenge revolves around the identification of signals that can serve in the construction of early warning systems ([Renn et al., 2022](#)). Here, one would have to go beyond the study of financial data and systems and take into account the wider impact and relationships between different types of risks (e.g., [Crona et al., 2021](#)). Furthermore, since the response to systemic risks – and consequently, their amplification or attenuation – is contingent on their perception by stakeholders, it is imperative to understand how “signals get filtered and modulated, transmitted, and interpreted” ([Schweizer et al., 2022](#)).

In this paper, we aim to take a different (in some sense a ‘sideways’) perspective on these challenges and aim to contribute a distinct set of quantitative tools helpful for the understanding of systemic risk and its antecedents. In the design of those tools, we mobilize the Attention-Based View (ABV) of the firm ([Ocasio, 1997](#)), thus associating a literature strand of management research with problems studied within the operational research literature. The ABV offers a conceptual framework for examining the focused attention of stakeholders, which is itself an antecedent of the risk perceptions that they form. As risk does not exist independently of our minds and culture ([Slovic, 1992](#)), risk perception itself influences human behavior and can be a driver of (systemic) risk ([Renn et al., 2022](#); [Schweizer et al., 2022](#)). Furthermore, risk perception operates through information systems and channels of communication, based on which interpretations are formed, leading to different behavioral responses, ripple effects, and finally social impacts (e.g., [Pidgeon et al., 2003](#); [Schweizer et al., 2022](#)). While these processes are multifaceted and complex, it is empirically more tractable to study their result: the limited set of issues that decision-makers focus their attention on ([Ocasio, 1997](#)). Furthermore, as decision-makers’ attention is contextually situated, we use insights from the strategy literature to in order to identify salient features of the decision context ([Porter, 1985, 1986](#); [Drucker, 1995](#); [Cornelius et al., 2005](#); [Grant, 2021](#)).

We apply these ideas in the context of operational research and risk analysis, in the specific context of systemic risk arising across two adjacent but distinct industries: Insurance and Banking. Specifically, we study the cross-boundary convergence (or divergence) in the attention of risk managers in those industries. The insurance and banking industries have different business models, but at the same time, are subject to common pressures from the economic and broader environment, while their adjacency means that they are often able to enter each other’s markets. For example, the near-collapse of the insurance giant AIG in 2008 demonstrated interconnectedness across industry boundaries, with insurers investing in mortgage-backed

securities, which generated shared vulnerabilities and, ultimately, extreme downside risks (Billio et al., 2012; Bushman et al., 2017). Furthermore, the cultural commonalities of the financial services industries make issues such as groupthink and herding behaviour an important concern (e.g., Trueman, 1994; Clement and Tse, 2005), with such co-behaviors causing systemic risk (Piškorec et al., 2014; Haldane and May, 2011).

We use text data from risk disclosures as empirical manifestations of decision-makers' attention. Specifically, we use companies' 10-K reports, spanning the 2006-2018 period. Such data have been widely used as a data source in management (e.g., Guo et al., 2017; Dutt and Joseph, 2019), accounting (e.g., Cazier et al., 2021), risk management (e.g., Bao and Datta, 2014), bankruptcy prediction (Mai et al., 2019), and – importantly – systemic risk measurement (Bushman et al., 2017). In contrast with previous literature, we focus on convergence in attention *across industries*, which contrasts with a growing literature using financial reports to extract information about individual firms or risks arising *within* a single industry (e.g., Bushman et al., 2017; Mai et al., 2019; Gupta et al., 2021).

We analyse these textual data quantitatively using Machine Learning techniques. Our approach has four steps. First, we vectorize the risk disclosure documents using Doc2vec embeddings (Mikolov et al., 2013a,b), and then use silhouette values (Rousseeuw, 1987; Bao and Datta, 2014) to quantify the semantic similarity between reports from individual companies and the industries that are external to them – we call this semantic similarity *convergence of attention*. Second, we provide quantitative measures of the context in which decision-makers' attention is situated, through topic analysis (Blei et al., 2003) of the Business section of 10-K reports. We focus on three contextual antecedents: *common trends* in the macro environment, managers' attention to an adjacent industry because of the threat of *substitution*, and management *fashions* or fads. Third, we associate convergence in attention with its contextual antecedents, via Random Forest (Breiman, 2001) regression analysis. Fourth, we use sensitivity analysis to gain insights on model response to changes in the inputs. Sensitivity analysis (for a review see Borgonovo and Plischke, 2016) has been widely applied in fields such as epidemiological modeling (Lu and Borgonovo, 2023) and financial investment studies (Borgonovo et al., 2010). We carry out extensive sensitivity analysis of the convergence of attention measure to its modelled antecedents, using Accumulated Local Effects (Apley and Zhu, 2020), feature importance (Breiman, 2001; Ishwaran, 2015), Shapley values (Lundberg and Lee, 2017a,b), and marginal attribution by conditioning on quantiles (Merz et al., 2022).

The empirical analysis leads to two key findings. First, the regressions show that our measure of convergence is increasing in all three antecedents (common trends, substitution and fashions), thus validating our theoretical framing. Second, the variable importance and sensitivity analyses consistently find that management fashions have a dominant impact on convergence in attention. This finding underlines the need for responses to systemic risk engaging a wide range of stakeholders by inclusive governance approaches (Renn et al., 2022; Schweizer et al., 2022). Furthermore, it opens serious questions on the qualitative features of systemic risk induced by risk disclosures being potentially driven by legitimacy rather than substantive concerns.

The rest of the paper is structured as follows. In Section 2, we develop our theoretical framework based on the Attention-Based View of the firm and identify three key antecedents of convergence in attention across

industries. In Section 3, we report our text data sources and construct the text-based measure of convergence across industries. Variables for the proposed antecedents of convergence are constructed in Section 4. In Section 5, convergence in management attention is modelled through its measured antecedents. We offer concluding remarks in Section 6.

2 Attention-Based View and antecedents of convergence

2.1 Management attention and 10-K data

The Attention-Based View of the firm (ABV) (Ocasio, 1997; Ocasio and Joseph, 2005) is an established theoretical perspective in management research on firm strategic behavior. The ABV theorizes firm strategic behavior as an outcome of the focusing and channelling of decision-makers' attention. Here, attention is understood as "the noticing, encoding, interpreting, and focusing of time and effort by organizational decision makers on both problems and solutions" (Ocasio, 1997). The theoretical tenet that the focus of decision-makers' attention leads to firm strategic behaviour is supported by empirical evidence in variety of contexts and for several strategic behaviours including: responses to institutional change (Ocasio and Radoynovska, 2016); multinational strategy (Bouquet and Birkinshaw, 2008); technology strategy (Eggers and Kaplan, 2009); strategic adaptation (Joseph and Ocasio, 2012), and corporate governance (Tuggle et al., 2010).

Furthermore, the focus of managers' attention is situated in the firm's context, which includes the environmental stimuli for decision-making, the interactions among participants in the context, and the embodiment of issues and answers in cultural symbols, artefacts, and narratives (Ocasio, 1997). It follows that, if decision makers are subject to a context that shares similarities, they are also likely to focus their attention towards similar issues, hence their firms are likely to exhibit similar strategic behaviors. Therefore, the ABV provides a theoretical framework to understand and anticipate strategic co-behaviours across firms and industries, through the lens of the convergence, or divergence, of the constituent elements of the contexts in which management attention is situated.

A first challenge, then, is the identification of salient descriptors of the context in which risk managers' attention is situated, which are also available and comparable for individual firms and across time. We suggest that a firm's 10-K report – the document that all listed companies have to submit yearly to the US Securities and Exchange Commission (SEC) – can play such a role. In fact, 10-Ks have been extensively used as data sources for describing managers' and firms' behaviours (e.g., Dutt and Joseph, 2019; Bushman et al., 2017; Bao and Datta, 2014; Guo et al., 2017; Cazier et al., 2021; Mai et al., 2019). The 10-K reports are structured in several sections, describing a company's business, the risks it faces, and the operating and financial results for the fiscal year. To capture the context of risk managers' situated attention, we focus on the 10-K's Risk Factors section, which includes information about the most significant risks that apply to the company or to its assets. The data have limitations, stemming from their formal regulatory purpose. Still, we choose 10-K filings for our research because: (a) every US listed company has it in each year, (b) the format is consistent,

and (c) there is a section specific on Risk Factors, which is pertinent for our paper.

Building on these premises, we argue that the extent to which companies' 10-K Risk Factor sections converge with each other is a proxy for convergence in the focus of risk managers' attention, hence a potential driver of more similar firm strategic behaviours in response to the risk environment. In that sense, convergence among the content of companies' 10-K Risk Factor reports is an indicator of systemic risk. In Section 3, we offer a methodology to operationalize these theoretical insights, by introducing a text-based cross-industry metric of convergence.

2.2 Antecedents of convergence in risk managers' attention

The extent of convergence among companies' 10-K Risk Factors reports may signal an increase or decrease in systemic risk, but does not in itself specify the drivers of such a change. Hence, the convergence of attention, without reference to its drivers, does not give clear a signal to risk professionals and policymakers as to what might be suitable prevention or mitigation responses. To address this issue, we note that risk managers are exposed to considerations around a company's overall strategy. Then, the broader management literature can help us identify some specific dimensions of the context in which risk managers' attention is situated.

First, the strategic management literature has highlighted that, in developing their strategy, firms should look at the trends in their macro environment (Grant, 2021; Drucker, 1995) and describe scenarios that a firm is likely to face (Cornelius et al., 2005). These *common trends* are not limited to specific industries, and apply to the whole economy, e.g., increased attention to customer services, operation of firms, products, asset management, gains and loss of investment, etc. In the context of risk management, we argue that firm's business narrative around common trends is an important dimension that risk managers consider.

Second, Porter (1985, 1986) and the literature spurred from his seminal books on competitive strategy and competitive advantage have identified some factors – often referred as Porter's Five forces – that affect an industry's profitability; namely, suppliers' bargaining power; clients' bargaining power; intensity of competition; threat of new entrants; and threat of substitution. *Substitution* occurs when companies of one industry have to compete with those in other industries producing substitute products or services (for example, insurers selling financial derivatives that are traditionally sold by banks). While narratives around, e.g., competition and buyers/clients' bargaining power are related to dynamics within a firm's legacy industry, narratives around substitutes are more likely to be related to cross-industry dynamics and hence, pertinent for our analysis.

Finally, the management literature has also revealed that strategic choices often find their antecedents in management *fashions* (Abrahamson, 1996) and institutional legitimacy considerations (Suchman, 1995). Management fashions bear relevance for studying the antecedents of systemic risk: for example, it has been argued that the adoption of (risk) management systems can simply follow fashion rather than address real need (Milos et al., 2008), which reduces the usefulness of such systems and may even render them sources of risk (Power, 2009).

Hence we identify common trends, threat of substitution and management fashions as salient aspects

of risk managers' context of attention. In Section 4, we present in detail the construction of quantitative measures of those antecedents based on the 10-K reports.

2.3 Implications for risk management

Our theoretical framework has direct implications for risk professionals and policy makers. If different antecedents of attention are dominant as drivers of convergence, different types of systemic threats will prevail, which in turn would necessitate alternative risk management responses. We explore this below, with a summary of our discussion given in Table 1.

Prevailing antecedent	Implications	Risk management response
<u>Common Trends</u>	Cross-industry and economy-wide risk, by common risk exposures across industries.	Risk transfer / diversification; precautionary measures; regulation on risk limits.
<u>Substitution</u>	Cross-industry risk. Firms in adjacent industries facing common vulnerabilities.	Regulation on firewalls to limit the blurring of industry boundaries.
<u>Fashion</u>	Idiosyncratic and economy-wide risk. Idiosyncratic risks nor reflected in prevailing fashions are ignored. Risk disclosure documents become less useful for anticipation of potential threats.	Inclusive governance; diversity within in boards and risk management functions.

Table 1: Antecedents of convergence, implications and risk management response

If the first antecedent – attention on common trends – prevails, there will be both cross-industry and economy-wide risk. Decision-makers' attention to common issues in the environment, e.g., stock market performance, supply chain operation, customer service, will be the main cause of convergence. Systemic risk caused by this antecedent is difficult to fully control (i.e. reduce a risk event's probability of happening) because it is related to system-wide challenges; nonetheless, it can be to an extent mitigated (Hopkin, 2018). At individual firm level, risk can be diversified (e.g. the risk supply chain failure is managed by diversifying suppliers (Gornall and Strebulaev, 2018)) or partially transferred (e.g., by reinsurance). Regulators can help individual firms control the risk by setting risk limits, and guiding them to prepare for emerging system-wide risks.

If the second antecedent – attention on the adjacent industry and threat of substitution – is the main driver of the convergence, there will be cross-industry risk. If this is the case, managers in insurers and banks are preoccupied by firms' selling products or adopting strategies traditionally associated with the other industry. As a consequence, risk contagion arises, whereby shocks in one industry spill over to the other. For instance, in the wake of the 2007-08 financial crisis insurers' involvement in banking business, e.g., selling structured credit products, was seen as a key cause of systemic risk (Cummins and Weiss, 2009). To control

and mitigate such systemic risk, regulators build firewalls to limit the blurring of industry boundaries.

If the third antecedent – risk managers’ exposure to similar (risk) management fashions – prevails, then both idiosyncratic and economy-wide risks arise. Managers’ attention may be led to ‘fashionable’ topics, e.g., implementing enterprise risk management (ERM) frameworks without sufficiently reflecting on non-quantifiable uncertainties (Power, 2009). Attention and thus risk perception can become biased (Schweizer et al., 2022), possibly away from crucial risks (idiosyncratic or otherwise) towards less material ones. Furthermore, groupthink (Janis, 2008) may also lead to similar actions adopted when facing the same event, which further exacerbates systemic risk. A compounding issue is that then risk disclosure itself, by reflecting the biases of risk perception, becomes less informative in reflecting and anticipating the risks that firms are exposed to. Such problems invite a risk governance response. Previous studies show that corporate governance, especially the role of risk professionals, has a big impact on the performance of firms during crises (Hunjra et al., 2020; Aebi et al., 2012) and the need for inclusive governance (Schweizer et al., 2022; Renn et al., 2022) is well established given the complexities and ambiguities of systemic risk management.

3 Measuring inter-industry convergence

3.1 Data

Our dataset contains 10-K submissions to the SEC of all the 214 banks and 94 insurers listed on the New York Stock Exchange from 2006 to 2018. These are the industries across which we aim to measure convergence in attention. We focus on the two most substantial sections of those reports: the Business and Risk Factors sections. We collected all the 10-K reports manually from Filings Expert and classified the companies as Insurers or Banks based on the North American Industry Classification System (NAICS) codes. In addition, we use text data from 223 pharmaceuticals’ 10-K reports over the same period – data from pharmaceuticals will be used in Section 4 for constructing a measure of management fashion.

Thus the text data cover $T = 13$ years and $I = 531$ companies. Define the sets $\mathcal{T} = \{1, \dots, T\}$ and $\mathcal{I} = \{1, \dots, I\} = \mathcal{I}_{IN} \cup \mathcal{I}_{BK} \cup \mathcal{I}_{PH}$, where $i \in \mathcal{I}_{IN}$ if the i -th company is an insurer, $i \in \mathcal{I}_{BK}$ if it is a bank, and $i \in \mathcal{I}_{PH}$ if it is a pharmaceutical. Thus each 10-K report in the data corresponds to a pair (i, t) , $i \in \mathcal{I}$, $t \in \mathcal{T}$. Let J_{it} be equal to one if there is a document for firm i in year t and zero otherwise. The numbers of documents for in year t for each industry are then denoted as

$$n_{A,t} = \sum_{j \in \mathcal{I}_A} J_{j,t}, \quad A \in \{IN, BK, PH\}.$$

The total number of firms in each industry and year can be found in Table 2.

Before any text analysis is carried out, some standard pre-processing steps are applied: (a) Stop words are those words most likely to appear in all documents in the corpus (e.g., “a,” “the,” “of”, etc.), and they carry little semantic meaning (Bengfort et al., 2018, p.65). Stop words are removed, which improves the performance of algorithms, as there are fewer and more meaningful tokens left (Bengfort et al., 2018, p.72-74). (b) Bigrams,

Table 2: Total number of firms for each industry and year.

	Insurers	Banks	Pharmaceuticals
2006	53	173	94
2007	55	171	99
2008	54	177	105
2009	54	182	113
2010	76	191	116
2011	82	194	118
2012	61	193	110
2013	84	202	103
2014	90	204	102
2015	89	200	94
2016	84	189	100
2017	60	171	92
2018	74	131	91

possible contiguous subsequences of two words, are constructed. For example, ‘real estate’ is understood as a one term when the two words appear together (Bengfort et al., 2018, p.132-145). (c) All the words and bigrams are lemmatized, that is, we remove inflectional endings and return the base or dictionary form of a word (Stanford NLP Group, 2009).

3.2 Construction of the convergence measure

Here we construct a measure of convergence, which reflects the similarity of an individual bank’s (resp. insurer’s) Risk Factors section in their 10-K report, to the insurance (resp. banking) industry as a whole. This measure is meant to represent managers’ cross-industry convergence of attention. Constructing such a measure entails two distinct steps. First, each individual document (Risk Factor sections of 10-K reports), needs to be converted to a vector representation. Second, a measure of similarity is employed to calculate inter-industry convergence, based on those vectors.

To convert each document into a vector, we use word embeddings constructed by Doc2vec, which is an unsupervised algorithm that learns fixed-length feature representations from documents of varying length (Le and Mikolov, 2014). Word embeddings are well suited for mapping documents’ semantic content onto a vector space, by explicitly considering the context in which individual words occur. Such context is not captured in standard approaches that are based purely on word frequencies e.g. frequency, one-hot, and TFIDF encoding (Kusner et al., 2015; Bengfort et al., 2018, p.65-66).

To derive word embeddings from our data, we use Gensim’s Doc2vec class in Python. The Doc2vec algorithm represents each document by a dense vector which is trained to predict words in the document, using a neural network. We choose Doc2vec over alternatives such as Word2Vec embeddings (Mikolov et al.,

2013b) (which are word- rather than document- specific), because Doc2vec has been found to outperform, e.g., simple averaging of Word2Vec vectors, in terms of the error rates of an information retrieval task (Le and Mikolov, 2014).

Given the representation of each document by a vector, we measure the convergence in management attention across organizations in distinct industries using silhouette values (Rousseeuw, 1987). The silhouette value is a measure, used in unsupervised learning, of how similar an object is to its own cluster (cohesion) compared to other clusters (separation). It is calculated using the mean intra-cluster (here: intra-industry) distance and the mean nearest-cluster (here: inter-industry) distance for each observation. A low value indicates that the object is poorly matched to its own cluster and well matched to neighboring clusters, indicating, in our context, high convergence.

We define the vector representing each Risk Factors document for insurers and banks by D_{it} , $i \in \mathcal{I}_{IN} \cup \mathcal{I}_{BK}$, $t \in \mathcal{T}$. To calculate convergence for insurers, for each document with $i \in \mathcal{I}_{IN}$ and $J_{it} = 1$, we define the quantities:

$$a_{it} = \frac{1}{n_{IN,t} - 1} \sum_{j \in \{\mathcal{I}_{IN} \setminus i\}} J_{jt} \cdot d(D_{it}, D_{jt}),$$

$$b_{it} = \frac{1}{n_{BK,t}} \sum_{j \in \mathcal{I}_{BK}} J_{jt} \cdot d(D_{it}, D_{jt}),$$

where d represents the Euclidean distance between document vectors. Here, a_{it} represents the average dissimilarity of an insurer’s Risk Factors section to other insurers’ reports (across years) while b_{it} reflects the average dissimilarity to banks’ reports. For $i \in \mathcal{I}_{BK}$ we define the quantities a_{it}, b_{it} analogously. Then, our measure of inter-industry convergence is given by the *negative silhouette value* defined as

$$Convergence_{it} = -\frac{b_{it} - a_{it}}{\max\{a_{it}, b_{it}\}}, \quad i \in \mathcal{I}_{IN} \cup \mathcal{I}_{BK}. \quad (1)$$

This measure reflects the extent to which a document is close to documents from a different industry, compared to the baseline of the document’s closeness to other documents from its own industry.

In Figure 1 we show boxplots of the *Convergence* variable by year; the horizontal line is set at the overall median. It can be seen that the level of convergence varies widely across firms, with many outliers present. A slight positive trend towards higher convergence in attention may be observed, though this is clearly dominated by intra-year variability. Explaining this variability in convergence in terms of its theorized antecedents is the aim of the next two sections.

4 Measures for the antecedents of convergence

The calculation of the negative silhouette value (1), based on vector representations of firms’ Risk Factor sections, enabled us to construct a measure of convergence in management attention, as reflected in textual risk disclosures. Now we turn our attention to constructing firm/year-specific measures for the antecedents of attention. We first perform a topic modelling analysis, in order to identify general themes from the analysed

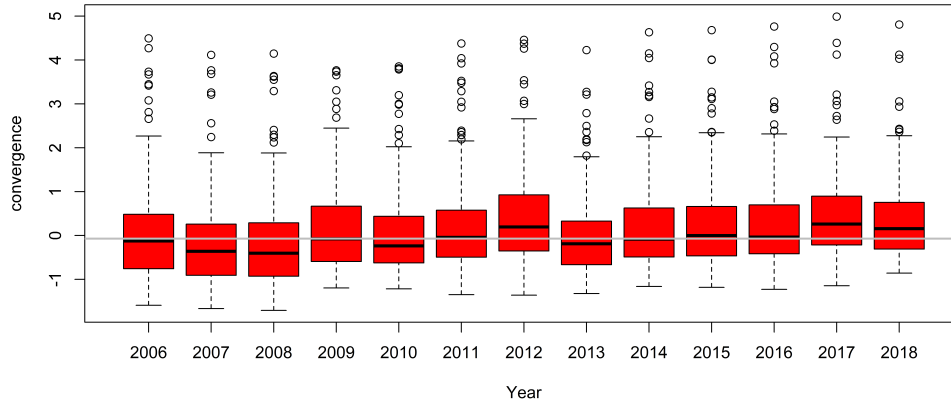


Figure 1: Boxplot of the *Convergence* variable by year; the horizontal line is at the overall median.

text documents. This helps us to derive interpretable descriptors of the composition of individual documents, which in turn will be used to construct the required measures.

4.1 Topic modelling of 10-K Business sections

A topic model is a type of statistical model for discovering a set of topics that are shared by a collection of documents. A well-established topic model the Latent Dirichlet Allocation (LDA) of [Blei et al. \(2003\)](#). Examples of application of LDA in risk analysis and management science include [Bao and Datta \(2014\)](#); [Bellstam et al. \(2021\)](#); [Taeuscher et al. \(2021\)](#). LDA generates summaries of topics in terms of a discrete probability distribution over words for each topic, and further infers, for each document, the distribution of its word content over topics.

LDA is a generative model, with the following underlying assumptions. A fixed number of topics, K exist. Each document is assumed to be randomly constructed, by choosing its composition by topic according to a Multinomial distribution with Dirichlet prior. Furthermore, for each topic, a Multinomial distribution over the vocabulary of size N is assumed, again with Dirichlet prior, giving the probability of any word in the vocabulary belonging to this topic. LDA estimates the following quantities:

- a) For each document (i, t) , $i \in \mathcal{I}$, $t \in \mathcal{T}$, $k = 1, \dots, K$, p_{itk} gives the proportion of the topic k in that document.
- b) For each word $n = 1, \dots, N$ in the vocabulary, q_{nk} gives the relative frequency by which the word will appear as part of of topic k .

We run the LDA model on all the Business section from 10-K reports of all the three industries in all years. We use Business section rather than Risk Factors section (which was used to measure overall convergence) to (a) reflect the broader strategic focus of the firms, giving the context within risk managers'

Table 3: Average percentage distribution of topics, for documents within each industry.

		Insurance	Banking
Topic 1	Insurance: liabilities	50.83	0.34
Topic 2	Banking: loans and deposits	3.02	33.33
Topic 3	Insurance: life policies	23.50	1.27
Topic 4	Banking: regulation and capital	4.39	53.59
Topic 5	Products, operation and customers	10.07	4.54
Topic 6	Investment	7.91	6.82

4.2 Antecedent 1: Common trends

The first contextual antecedent of firm strategy relates to the challenges and opportunities in the industry environment. To quantify this antecedent we introduce the variable $CommonTrend_{it}$. From the LDA analysis of Section 4.1, Figure 2, we identify two topics that reflect aspects of the environment that cut across industries, specifically Topics 5 ('Products, operations, and customers') and 6 ('Investment'). Hence we define:

$$CommonTrend_{it} = p_{it5} + p_{it6}, \quad i \in \mathcal{I}_{IN} \cup \mathcal{I}_{BK}.$$

4.3 Antecedent 2: Threat of substitution

The second contextual antecedent of firm strategy relates to the potential for substitution and new industry entry. We measure the relevance of this antecedent for specific firms, via the variable $Substitution_{it}$, defined as follows. From the LDA analysis of Section 4.1, we have seen that there are insurance-specific and bank-specific topics. The variable $Substitution_{it}$ reflects the extent to which a firm (e.g. an insurer) places attention on topics specific to the industry it does not belong to (e.g. banking), measured by the percentage of the firm's 10-K Business section spent on those topics. Hence we have that:

$$Substitution_{it} = \begin{cases} p_{it2} + p_{it4}, & i \in \mathcal{I}_{IN}, \\ p_{it1} + p_{it3}, & i \in \mathcal{I}_{BK}. \end{cases}$$

4.4 Antecedent 3: Management fashions

Finally, we identified as the third contextual antecedent of firm strategy, related to industry convergence, the level of attention to management fashions. We quantify such attention, via the variable $Fashion_{it}$.

To construct this variable, there are two steps. In the first step, we measure the negative silhouette value between the Risk Factors sections of 10-K reports, but now with clusters formed by financial firms (insurers and banks combined) and pharmaceuticals. Specifically, we let

$$SVFP_{it} = -\frac{\tilde{b}_{it} - \tilde{a}_{it}}{\max\{\tilde{a}_{it}, \tilde{b}_{it}\}}, \quad i \in \mathcal{I}_{IN} \cup \mathcal{I}_{BK},$$

where,

$$\begin{aligned}\tilde{a}_{it} &= \frac{1}{n_{IN,t} + n_{BK,t} - 1} \sum_{j \in \{\mathcal{I}_{IN} \cup \mathcal{I}_{BK} \setminus i\}} J_{jt} \cdot d(D_{it}, D_{jt}), \\ \tilde{b}_{it} &= \frac{1}{n_{PH,t}} \sum_{j \in \mathcal{I}_{PH}} J_{jt} \cdot d(D_{it}, D_{jt}).\end{aligned}$$

The logic of this a construction is as follows. The variable $SVFP_{it}$ represents the convergence of attention of risk managers in individual financial firms (insurers or banks) to the pharmaceutical industry. We consider the pharmaceutical industry to be highly distinct from the financial services industry, with, for example, no opportunities for financial firms to enter the pharma industry (and vice versa). Hence, $SVFP_{it}$ reflects the extent to which attention is placed on more general narrative about risk, which cut across the boundaries of industries that are not closely related. In particular, we claim that there is not a high level of connectivity between pharmaceutical and financial services companies. Then, if the convergence between, say, an individual bank and the pharmaceuticals industry explains the convergence between that bank and the insurance industry (the phenomenon we are interested in), then this is likely the case because the bank managers' attention is captured by (e.g., risk management) fashions that cut across sectors.

It may be argued that $SVFP_{it}$ and $CommonTrend_{it}$ intertwined, as $SVFP_{it}$ may also capture broader trends, besides fashions. To address this issue, in a second step, we regress $SVFP_{it}$ against $CommonTrend_{it}$ using Ordinary Least Squares regression (OLS), and measure $Fashion_{it}$ as the residual of this regression.

4.5 Control variables

In Section 5 we will present non-linear regressions of $Convergence_{it}$ on the metrics we just introduced to represent the theoretically derived antecedents of convergence. For this regression we introduce a number of statistical controls. All control variables are defined for $i \in \mathcal{I}_{IN} \cup \mathcal{I}_{BK}$, $t \in \mathcal{T}$.

Year fixed effects. The time $t \in \mathcal{T}$ corresponding to each 10-K report is used as a categorical variable to control for trends that are not reflected in the convergence antecedents.

Industry fixed effects. The industry a firm belongs to is used as a categorical variable, to control for firms in different industries potentially responding differently to convergence drivers.

Inter-industry stock return correlation. We measure Kendall's rank correlation between an individual insurer's (resp. bank's) daily stock return and the return on the S&P 500 Bank Index (resp. Insurance Index). This variable, which we call $StockCorr_{it}$, is meant to control for 10-K reports responding to external shocks that affect both industries and are already reflected in stock movements. For each (i, t) , the value of $StockCorr_{it}$ is calculated on stock market data over the one-year time period before the disclosure date of the corresponding 10-K report. We use Kendall's rank correlation because it is robust to extreme observations (Lindskog et al., 2003). We preferred it over tail risk measures (e.g., Bushman et al., 2017) for reasons of estimation stability, as we are not specifically focused on extreme levels of stock co-movement.

Data on the stock price of each company come from the Quandl database. Data for the industry indexes are downloaded from S&P Global Market Intelligence. In regressions, $StockCorr_{it}$ is standardized by industry.

Boilerplate language. It may be plausible that the 10-K reports of a firm is similar to reports in a different industries, because of its use of non-informative language in risk disclosures, which is heavy with cliches, that is, standardized expressions carrying little meaning. To control for this effect, we quantify language non-informativeness, by slightly adapting the ‘boilerplate’ language metric of [Lang and Stice-Lawrence \(2015\)](#). The resulting variable, $Boiler_{it}$, is calculated by the following process based on [Lang and Stice-Lawrence \(2015\)](#).

- We count all tetragrams contained in each document in my sample, where a tetragram is an ordered group of four words within a single sentence. We aggregate these counts by year of issue.
- We identify tetragrams that occur in at least 30% of the documents or on average at least 5 times per document across all the three industries in a year. The identified tetragrams are considered to reflect non-informative (‘boilerplate’) language.
- Then the percentage of the common boilerplate tetragrams out of all the tetragrams in each document is calculated, and set equal to $Boiler_{it}$.

In regressions $Boiler_{it}$ is standardized by industry.

The relationship between *Convergence* and all covariates (excluding industry fixed effects) is depicted in the scatter plot matrix of Figure 3. Time trends can be glanced in the first column; there are a noticeable positive trends in the *Boiler* and *Convergence* variables. Bivariate relationships between *Convergence* and all covariates are seen in the last row. It is clear that *Convergence* increases in the measures of its three postulated antecedents: *Substitution*, *CommonTrend* and *Fashion*.

5 Explaining convergence of risk disclosures in terms of contextual antecedents

In this section we empirically investigate the extent to which the contextual antecedents of firm strategy drive convergence of attention. Empirically, we carry out regression models of the form $\mathbb{E}[Convergence_{it} | \mathbf{X}_{it}] = g(\mathbf{X}_{it})$, where g is a (non-linear) regression function and \mathbf{X}_{it} includes the covariates and controls developed in Section 4.

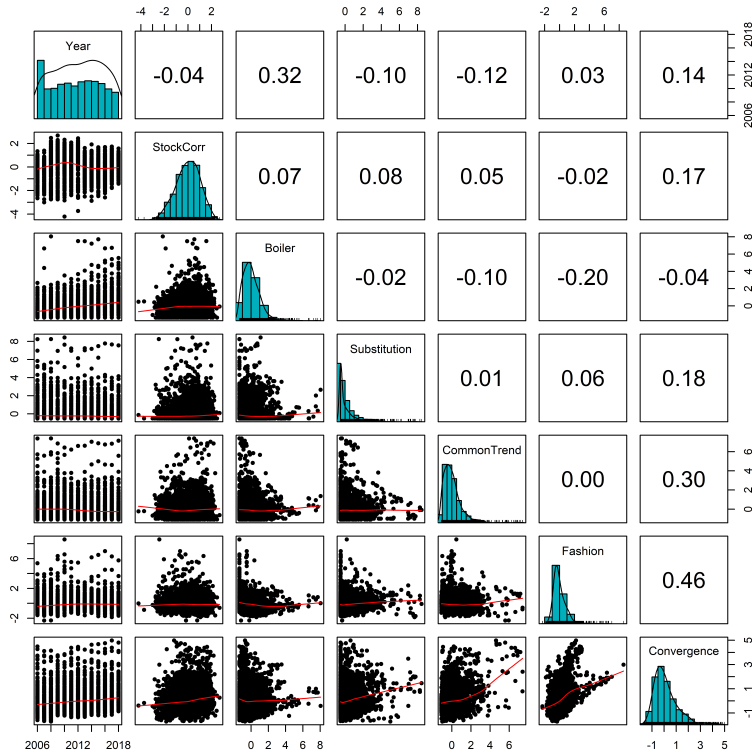


Figure 3: Scatter plot matrix of *Covergence*, control variables and antecedents.

5.1 Model selection

First, we select the best regression model from a set of approaches. Here, we focus on three types of regression models: a regularised linear model (Elastic Net Regression) and two non-linear models (Random Forests and Gradient Boosting Trees). All regressions are carried out in Python, using the scikit-learn package (Pedregosa et al., 2011). To evaluate the models, we use cross-validation with 20 folds. As we have far more banks than insurers in our sample, we make sure all folds are stratified and balanced, to have the same proportion of insurers and banks in each fold. We run each model for different hyperparameter settings and report the performance results using Mean Squared Errors (MSE), averaged over all folds. These prediction errors are reported in Table 4, for different hyperparameter choices, in an approach similar to Agarwal et al. (2019) (all non-reported hyperparameter values are set at their default values).

In the following analysis, we use the results from the Random Forest regression model with 100 trees due to its lowest reported MSE. The superior predictive performance on non-linear compared to linear models indicates the presence of important non-linear effects and variable interactions. To disentangle these relationships, we carry out a detailed discussion of variable importance and sensitivity analysis in Section 5.2.

Table 4: Results of the performance of different regression models with different parameters

Regression model	Parameter	Param value	Average MSE
ElasticNet	Regularization parameter	0.0001	0.5410
		0.001	0.5410
		0.01	0.5408
		0.1	0.5679
		0.5	0.7204
RandomForest	Number of trees in the forest	10	0.4345
		50	0.3687
		100	0.3550
		1000	0.3604
		10000	0.3622
GradientBoosting	Number of boosting stages to perform	10	0.6117
		50	0.4395
		100	0.4203
		1000	0.3847
		10000	0.4186

5.2 Variable importance and sensitivity

5.2.1 Accumulated Local Effects

We visualize the results of the chosen regression model, via Accumulated Local Effects (ALE) plots and variable importance measures. ALEs, proposed by [Apley and Zhu \(2020\)](#), describe how features influence the prediction of a machine learning model on average. ALE plots shows how the model predictions change in a small window of the feature around a certain grid value for data instances in that window. ALEs can be efficiently calculated and are not distorted by correlation between features ([Molnar, 2020](#)), which affects alternative measures such as Partial Dependence Plots.

The ALE plots for the chosen regression model are displayed in [Figure 4](#). The impact of extreme values (bottom and top 5% of feature values) of each variable is not shown (except for *Year*), as these are very noisy. We observe the following:

- For all three the covariates, *CommonTrend*, *Substitution* and *Fashion*, we see a clear positive impact on the response variable *Convergence*, which confirms our theoretical argument of [Section 2](#).
- From observing the scale of the respective ALEs, we note that *Fashion* is the most impactful variable.
- We also see that the control *StockCorr* is positively associated with the response, which reflects our assumption that co-movements of stock prices would have a positive impact on attention convergence, which is a factor we control for.
- The ALE for *Year* also shows an increasing pattern, reflecting an underlying positive trend in convergence in general.
- Finally, there is no clear pattern for *BoilerPlate*, but the scale of the ALE is much smaller, hence we consider this variable inconsequential. Hence, lack of informativeness in the language used does not

appear to be a driver of text similarity between 10-K reports.

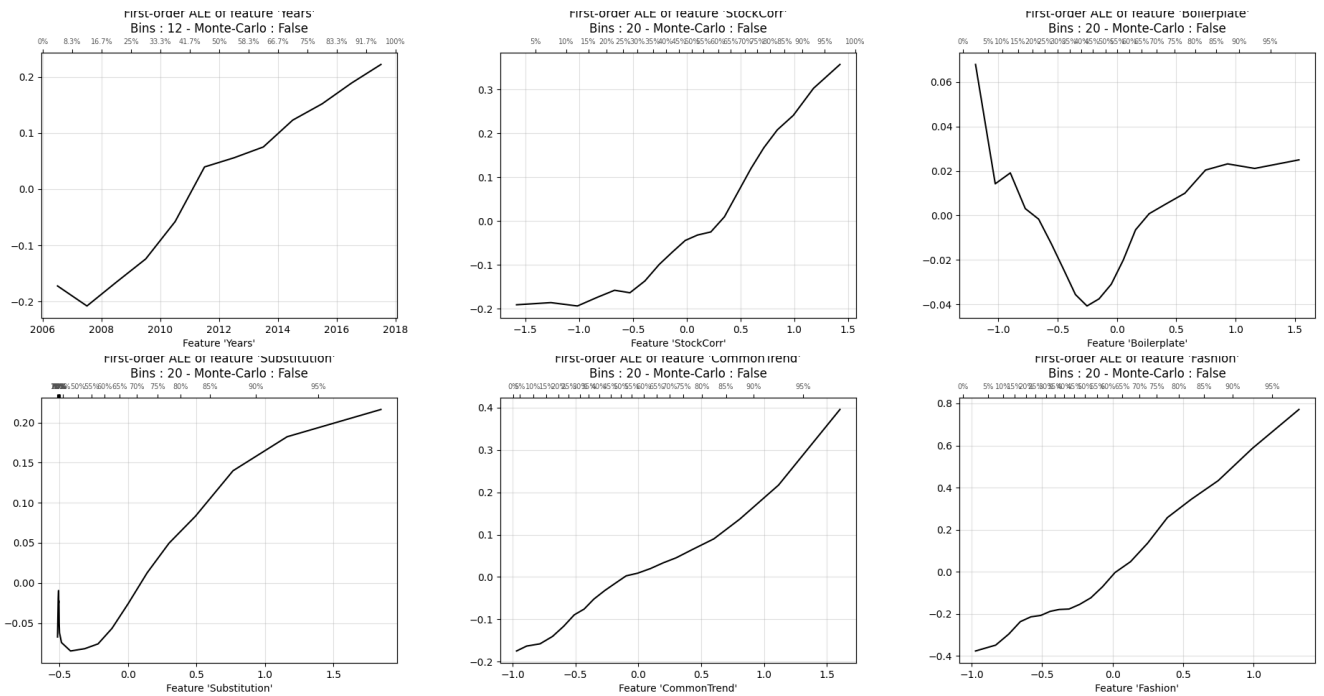


Figure 4: ALE plots for the Random Forest regression model.

5.2.2 Feature importance

We quantify the relative importance of the different covariates in our regression model. For that purpose we use two standard variable importance measures in Machine Learning: the impurity-based feature importance and the variance permutation importance (Breiman, 2001). For each of the two methods, the higher the value of the measure is, the more important the feature. Impurity-based feature importance is specific to tree-based methods and is computed as the (normalized) total reduction of the error criterion (e.g., Gini or Mean Squared Error) brought by that feature. A split (i.e., the separation of data into different nodes in a particular tree of the Random Forest) with a large decrease of impurity is considered important, and therefore, variables used for splitting at important splits are also considered important. The impurity is usually measured by the Gini impurity (Ishwaran, 2015). Permutation importance is defined to be the difference between the baseline prediction error of the fitted model and the error arising after randomly permuting a given feature column (Breiman, 2001). Impurity-based feature importance can be misleading for high cardinality features and permutation importance has been shown as an appropriate alternative to solve this problem (Breiman, 2001).

The feature importance and permutation importance of all the drivers of the Random Forest regression model can be found in Figure 5. According to both importance measures, the most important feature *Fashion*. This is consistent with the picture provided by ALEs and establishes *Fashion* as the most important driver of convergence.

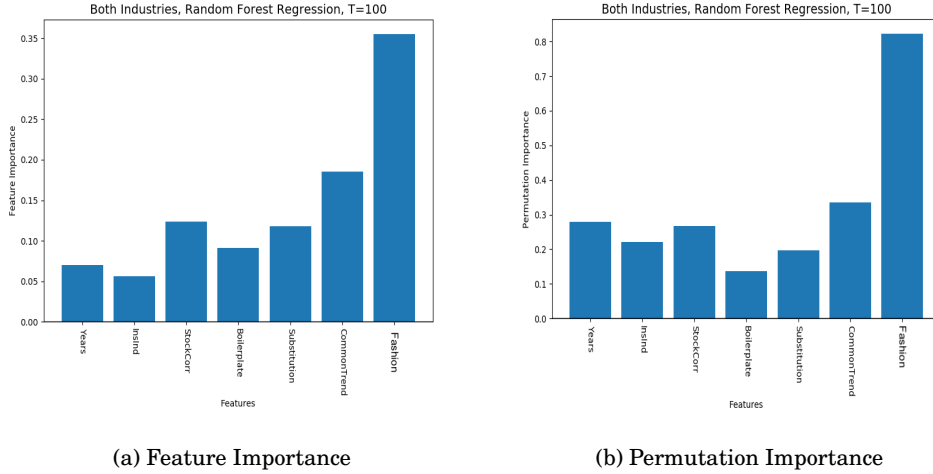


Figure 5: Variable importance of each antecedent: (a) impurity-based feature importance; (b) permutation importance.

5.2.3 Shapley values and quantile contributions

Finally, we evaluate the contribution of individual antecedents to predictions at different – in particular, high – levels of convergence. We follow the Marginal Attribution by Conditioning on Quantiles (MACQ) framework of Merz et al. (2022), drawing from Tsanakas and Millosovich (2016). For this, we require an additive decomposition of individual predictions, such that for the i -th observation in the sample, with features \mathbf{x}_i and predictions $\hat{g}(\mathbf{x}_i)$, we have

$$\hat{g}(\mathbf{x}_i) = \sum_j \phi_j(\mathbf{x}_i), \quad (2)$$

where $\phi_j(\mathbf{x}_i)$ represents the contribution of the j -th covariate to the prediction of the i -th observation. In Merz et al. (2022) such a decomposition was derived by a quadratic approximation to smooth prediction functions of deep learning models. Here, given the non-smoothness of the fitted Random Forest’s prediction function, we instead carry out decomposition in equation (2) via Shapley values.

The Shapley value for a feature-observation combination is derived from the marginal contribution of the feature to individual predictions. Such marginal contributions are averaged across ‘coalitions’ of features to which the one under focus is added. Thus, differently to feature importance measures, Shapley values quantify the contribution of each covariate at the *individual* observation level. Shapley values were originally defined in the context of co-operative game theory (Shapley, 1953) and have become popular for interpreting the predictions of machine learning algorithms following the SHAP framework of Lundberg and Lee (2017a,b) – for a technical overview see Aas et al. (2021).

We plot in Figure 6 (top) for each of the features *CommonTrend*, *Substitution*, and *Fashion*, the attributions $\phi_j(\mathbf{x}_i)$ on the vertical axis, against the quantile level u_i of the prediction $\hat{g}(\mathbf{x}_i)$, that is, $u_i = \hat{F}(\hat{g}(\mathbf{x}_i))$, where \hat{F} is the empirical distribution of predictions $\hat{g}(\mathbf{X})$. The red curves show non-parametric estimates of the function $u \mapsto \mathbb{E}[\phi_j(\mathbf{X}) \mid \hat{g}(\mathbf{X}_i) = \hat{F}^{-1}(u)]$, with the expectation taken over the empirical measure.

Finally the color of the plotted points represents the value of the feature examined. If we focus specifically on the right of each plot in Figure 6, e.g. $u \in [0.8, 1)$, we consider those firm-years for which the highest convergence is observed. We can see that for such companies, which are most relevant for considerations of systemic risk, the contribution of *Fashion* is by far the most important one.

In Figure 6 (bottom) we plot the attributions $\phi_j(\mathbf{x}_i)$ against the feature values $x_{i,j}$ – the colours now represent the response quantile level u_i .¹ we observe that the relationships between covariates $x_{i,j}$ and attributions $\phi_j(\mathbf{x}_i)$ are clearly increasing, with the points at the top right typically displaying high values of the response variable, convergence. The steeper increase in the plot for *Fashion* once again confirms the dominant effect of this variable.

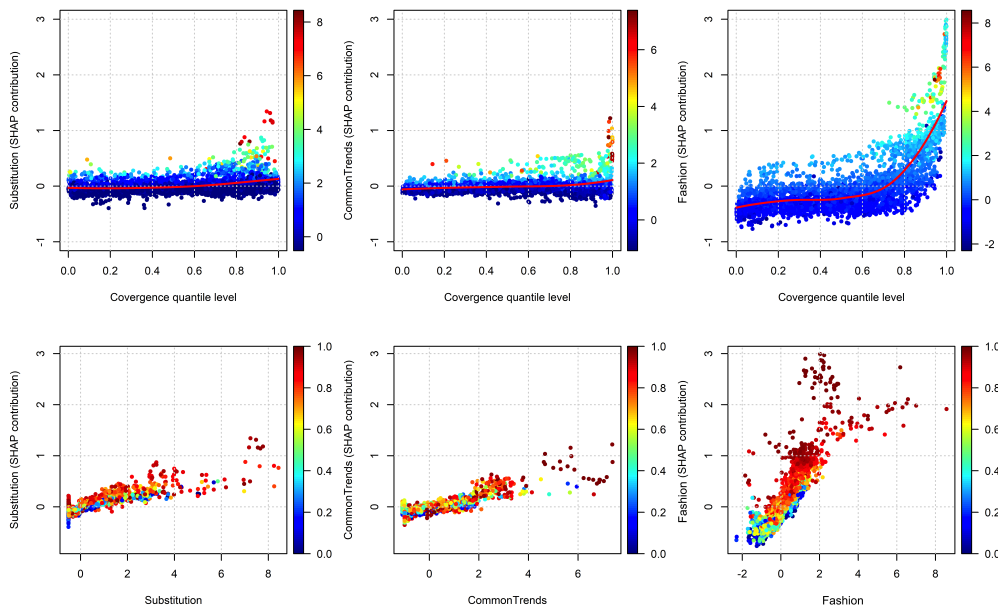


Figure 6: Results of sensitivity analysis based on Shapley values. Top: feature contributions to predicted value, against quantile level of predictions; color represents the feature value. Bottom: feature contributions to predicted value, against feature values; the colors represent the quantile level of predictions.

5.3 Implications of empirical results

Our analysis showed that *Convergence* is increasing in all three variables representing the antecedents of attention that we identified. This lends empirical support to the theoretical framework developed in Section 2. Furthermore, we have found that the most important antecedent of convergence in the attention of risk managers *Fashion* – this is consistently implied by all three importance and sensitivity measures that we used. In summary, if the measured convergence in attention between insurers and banks is largely explained

¹Note that for a linear model of the form $\hat{g}(\mathbf{x}_i) = \sum_j \beta_j x_{i,j}$, Shapley value attributions take the form $\phi_j(\mathbf{x}_i) = \beta_j x_{i,j}$, thus reproducing the modelled linear effects.

by the convergence of the same firms' 10-K reports to those of pharmaceutical companies, then we may conclude that the main driver of convergence in attention has less to do with shared stresses and risk exposures and more with the tendency of firms to speak in a language that spans unrelated industries, potentially driven by institutional legitimacy concerns. Hence, the empirical analysis raises the danger of groupthink as most pressing, driving firms' attention and thus resources to similar issues, with substantial risks (idiosyncratic or otherwise) ignored. As a corollary, literal reading of risk management disclosures becomes less useful in detecting the risks specific to individual firms.

6 Concluding remarks

In this paper we aimed to contribute to the study of systemic risk, by weaving together a qualitative conceptualization with quantitative measures of cross-industry convergence in attention and its contextual antecedents. By developing and analysing those measures we aspire to provide methodological tools that reflect systemic risk in a more holistic way than standard (e.g., financial) measures. Furthermore, the use of textual data speaks to the need to consider the information ecosystem and biases of decision makers (Schweizer et al., 2022)

The specific contribution of this paper is fourfold. First, we theoretically connect systemic risk with the context of decision making through the lens of the Attention-Based View of the firm. Second, building on strategic management theory, we establish a framework for antecedents of convergence in attention, and therefore, systemic risk. Third, we operationalize the theoretical framework using text data from risk disclosures and tools from Machine Learning, to create measures of convergence in attention and its antecedents. Finally, based on the theoretical framework and the regression and sensitivity analyses, we identify the dominant antecedent (and hence the particular "flavour") of cross-industry systemic risk.

Among the three antecedents considered, management fashion has the highest explanatory power. This raises the spectre of systemic risk induced by groupthink (Janis, 2008), leading both to herding behaviour and the disregard of idiosyncratic threats. The amelioration of such effects is contingent on improved corporate governance (O'Connor, 2002; Howard, 2010) and, in particular, inclusive and participatory approaches (Renn et al., 2022; Schweizer et al., 2022).

Nonetheless, the importance of management fashion in explaining the convergence in attention poses something of a puzzle. Namely, if risk disclosures are driven by fashions, that is, by a concern for institutional legitimacy rather than the actual risks that firms face, where does this leave our textual analysis of 10-K reports? Does it mean that these formal risk disclosures are not useful for understanding systemic risk? We believe that analysis of risk disclosures remains pertinent since the dominance of fashions as an explanatory factor is itself a source of systemic risk, within a wider conceptualisation. This argument is also salient to other risk contexts. For example, climate risk reporting frameworks are increasingly mandated (Task Force on Climate-related Financial Disclosures, 2022), while at the same time, scholars have criticised the extent to which such disclosures tend to become ceremonial practices (Di Marco et al., 2023). Such practices may not only distract from effecting transformative change, but also open up companies like insurers to the new

systemic risk of ‘climate litigation’ (Doering et al., 2023).

As with all quantitative empirical analyses of complex social phenomena, important caveats apply to our work: not least the extent to which our datasets and mathematical constructs appropriately reflect the theoretical concepts we discuss. Given the interdisciplinarity of the field of systemic risk (Renn et al., 2022), qualitative approaches are necessary to complement and validate (and indeed criticize) our analysis. Nonetheless, we believe that our methods and results provide some clues on the framing of questions that can only be addressed by qualitative means and case-specific contexts.

References

- Aas, K., Jullum, M., and Løland, A. (2021). Explaining individual predictions when features are dependent: More accurate approximations to shapley values. *Artificial Intelligence*, 298:103502.
- Abrahamson, E. (1996). Management fashion. *Academy of Management Review*, 21(1):254–285.
- Acharya, V., Engle, R., and Richardson, M. (2012). Capital shortfall: A new approach to ranking and regulating systemic risks. *American Economic Review*, 102(3):59–64.
- Acharya, V. V., Pedersen, L. H., Philippon, T., and Richardson, M. (2017). Measuring systemic risk. *The Review of Financial Studies*, 30(1):2–47.
- Aebi, V., Sabato, G., and Schmid, M. (2012). Risk management, corporate governance, and bank performance in the financial crisis. *Journal of Banking & Finance*, 36(12):3213–3226.
- Agarwal, A., Gupta, A., Kumar, A., and Tamilselvam, S. G. (2019). Learning risk culture of banks using news analytics. *European Journal of Operational Research*, 277(2):770–783.
- Aglietta, M. and Espagne, E. (2016). Climate and finance systemic risks, more than an analogy? the climate fragility hypothesis. *CEPII Working Paper*.
- Apley, D. W. and Zhu, J. (2020). Visualizing the effects of predictor variables in black box supervised learning models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 82(4):1059–1086.
- Aven, T. and Renn, O. (2020). Some foundational issues related to risk governance and different types of risks. *Journal of Risk Research*, 23(9):1121–1134.
- Bachner, G., Knittel, N., Poledna, S., Hochrainer-Stigler, S., and Reiter, K. (2023). Revealing indirect risks in complex socioeconomic systems: A highly detailed multi-model analysis of flood events in Austria. *Risk Analysis*.
- Bao, Y. and Datta, A. (2014). Simultaneously discovering and quantifying risk types from textual risk disclosures. *Management Science*, 60(6):1371–1391.
- Bellstam, G., Bhagat, S., and Cookson, J. A. (2021). A text-based analysis of corporate innovation. *Management Science*, 67(7):4004–4031.
- Bengfort, B., Bilbro, R., and Ojeda, T. (2018). *Applied text analysis with python: Enabling language-aware data products with machine learning*. O’Reilly.
- Billio, M., Getmansky, M., Lo, A. W., and Pelizzon, L. (2012). Econometric measures of connectedness and systemic risk in the finance and insurance sectors. *Journal of Financial Economics*, 104(3):535–559.
- Blei, D. M., Ng, A. Y., and Jordan, M. I. (2003). Latent dirichlet allocation. *The Journal of Machine Learning Research*, 3:993–1022.
- Bonaccolto, G., Caporin, M., and Maillet, B. B. (2022). Dynamic large financial networks via conditional expected shortfalls. *European Journal of Operational Research*, 298(1):322–336.
- Borgonovo, E., Gatti, S., and Peccati, L. (2010). What drives value creation in investment projects? an application of sensitivity analysis to project finance transactions. *European Journal of Operational Research*, 205(1):227–236.

- Borgonovo, E. and Plischke, E. (2016). Sensitivity analysis: A review of recent advances. *European Journal of Operational Research*, 248(3):869–887.
- Botzen, W. W., Duijndam, S. J., Robinson, P. J., and van Beukering, P. (2022). Behavioral biases and heuristics in perceptions of covid-19 risks and prevention decisions. *Risk Analysis*, 42(12):2671–2690.
- Bouquet, C. and Birkinshaw, J. (2008). Weight versus voice: How foreign subsidiaries gain attention from corporate headquarters. *Academy of Management Journal*, 51(3):577–601.
- Breiman, L. (2001). Random forests. *Machine Learning*, 45(1):5–32.
- Brownlees, C. and Engle, R. F. (2017). Srisk: A conditional capital shortfall measure of systemic risk. *The Review of Financial Studies*, 30(1):48–79.
- Bushman, R. M., Chen, J. V., and Williams, C. D. (2017). Informativeness and timeliness of 10-K text similarity for predicting tail-risk comovement. *Available at SSRN 2983315*.
- Calabrese, R. and Osmetti, S. A. (2019). A new approach to measure systemic risk: A bivariate copula model for dependent censored data. *European Journal of Operational Research*, 279(3):1053–1064.
- Cazier, R. A., McMullin, J. L., and Treu, J. S. (2021). Are lengthy and boilerplate risk factor disclosures inadequate? an examination of judicial and regulatory assessments of risk factor language. *The Accounting Review*, 96(4):131–155.
- Chaudhry, S. M., Chen, X. H., Ahmed, R., and Nasir, M. A. (2023). Risk modelling of ESG (environmental, social, and governance), healthcare, and financial sectors. *Risk Analysis*.
- Clement, M. B. and Tse, S. Y. (2005). Financial analyst characteristics and herding behavior in forecasting. *The Journal of Finance*, 60(1):307–341.
- Cornelius, P., Van de Putte, A., and Romani, M. (2005). Three decades of scenario planning in shell. *California Management Review*, 48(1):92–109.
- Crona, B., Folke, C., and Galaz, V. (2021). The anthropocene reality of financial risk. *One Earth*, 4(5):618–628.
- Cummins, J. D. and Weiss, M. A. (2009). Convergence of insurance and financial markets: Hybrid and securitized risk-transfer solutions. *Journal of Risk and Insurance*, 76(3):493–545.
- Di Marco, R., Dong, T., Malatincová, R., Reuter, M., and Strömsten, T. (2023). Symbol or substance? scrutinizing the ‘risk transparency premise’ in marketized sustainable finance: The case of TCFD reporting. *Business Strategy and the Environment*, 32(6):3027–3052.
- Doering, S., Radcliffe, M., and Lockman, M. (2023). Climate litigation risk: Is there shelter from the storm? *WTW*.
- Drucker, P. (1995). *Managing in a Time of Great Change*. Talley Books.
- Dutt, N. and Joseph, J. (2019). Regulatory uncertainty, corporate structure, and strategic agendas: Evidence from the us renewable electricity industry. *Academy of Management Journal*, 62(3):800–827.
- Earle, T. C. (2009). Trust, confidence, and the 2008 global financial crisis. *Risk Analysis: An International Journal*, 29(6):785–792.
- Eggers, J. P. and Kaplan, S. (2009). Cognition and renewal: Comparing ceo and organizational effects on incumbent adaptation to technical change. *Organization Science*, 20(2):461–477.
- Gornall, W. and Strebulaev, I. A. (2018). Financing as a supply chain: The capital structure of banks and borrowers. *Journal of Financial Economics*, 129(3):510–530.
- Grant, R. M. (2021). *Contemporary strategy analysis*. John Wiley & Sons.
- Guo, W., Yu, T., and Gimeno, J. (2017). Language and competition: Communication vagueness, interpretation difficulties, and market entry. *Academy of Management Journal*, 60(6):2073–2098.
- Gupta, A., Wang, R., and Lu, Y. (2021). Addressing systemic risk using contingent convertible debt—a network analysis. *European Journal of Operational Research*, 290(1):263–277.
- Haldane, A. G. and May, R. M. (2011). Systemic risk in banking ecosystems. *Nature*, 469(7330):351–355.

- Hopkin, P. (2018). *Fundamentals of risk management: understanding, evaluating and implementing effective risk management*. Kogan Page Publishers.
- Howard, A. (2010). Groupthink and corporate governance reform: Changing the formal and informal decisionmaking processes of corporate boards. *Southern California Interdisciplinary Law Journal*, 20:425.
- Huang, X., Zhou, H., and Zhu, H. (2009). A framework for assessing the systemic risk of major financial institutions. *Journal of Banking & Finance*, 33(11):2036–2049.
- Hunjra, A. I., Hanif, M., Mehmood, R., and Nguyen, L. V. (2020). Diversification, corporate governance, regulation and bank risk-taking. *Journal of Financial Reporting and Accounting*.
- Ishwaran, H. (2015). The effect of splitting on random forests. *Machine Learning*, 99(1):75–118.
- Ivashina, V. and Scharfstein, D. (2010). Bank lending during the financial crisis of 2008. *Journal of Financial Economics*, 97(3):319–338.
- Janis, I. L. (2008). Groupthink. *IEEE Engineering Management Review*, 36(1):36.
- Joseph, J. and Ocasio, W. (2012). Architecture, attention, and adaptation in the multibusiness firm: General electric from 1951 to 2001. *Strategic Management Journal*, 33(6):633–660.
- Kusner, M., Sun, Y., Kolkin, N., and Weinberger, K. (2015). From word embeddings to document distances. In *International Conference on Machine Learning*, pages 957–966. The Proceedings of Machine Learning Research.
- Lang, M. and Stice-Lawrence, L. (2015). Textual analysis and international financial reporting: Large sample evidence. *Journal of Accounting and Economics*, 60(2-3):110–135.
- Le, Q. and Mikolov, T. (2014). Distributed representations of sentences and documents. In *International Conference on Machine Learning*, pages 1188–1196. PMLR.
- Lindskog, F., McNeil, A., and Schmock, U. (2003). Kendall’s tau for elliptical distributions. In *Credit Risk*. Physica-Verlag.
- Lu, X. and Borgonovo, E. (2023). Global sensitivity analysis in epidemiological modeling. *European journal of operational research*, 304(1):9–24.
- Lundberg, S. M. and Lee, S.-I. (2017a). Consistent feature attribution for tree ensembles. *arXiv preprint arXiv:1706.06060*.
- Lundberg, S. M. and Lee, S.-I. (2017b). A unified approach to interpreting model predictions. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 4768–4777.
- Mai, F., Tian, S., Lee, C., and Ma, L. (2019). Deep learning models for bankruptcy prediction using textual disclosures. *European journal of operational research*, 274(2):743–758.
- Merz, M., Richman, R., Tsanakas, A., and Wüthrich, M. V. (2022). Interpreting deep learning models with marginal attribution by conditioning on quantiles. *Data Mining and Knowledge Discovery*, 36(4):1335–1370.
- Mikolov, T., Chen, K., Corrado, G., and Dean, J. (2013a). Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., and Dean, J. (2013b). Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems*, pages 3111–3119.
- Milos, D., Sevic, Z., and Tekavčić, M. (2008). A review of the rationales for corporate risk management: Fashion or the need? *International Journal of Economic Sciences and Applied Research*, 1(1).
- Molnar, C. (2020). *Interpretable machine learning*. Leanpubbook.
- Ocasio, W. (1997). Towards an attention-based view of the firm. *Strategic Management Journal*, 18(S1):187–206.
- Ocasio, W. and Joseph, J. (2005). An attention-based theory of strategy formulation: Linking micro-and macroperspectives in strategy processes. In *Strategy Process*. Emerald Group Publishing Limited.
- Ocasio, W. and Radoynovska, N. (2016). Strategy and commitments to institutional logics: Organizational heterogeneity in business models and governance. *Strategic Organization*, 14(4):287–309.
- O’Connor, M. A. (2002). The enron board: The perils of groupthink. *University of Cincinnati Law Review*, 71:1233.

- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.
- Pidgeon, N., Kasperson, R. E., and Slovic, P. (2003). *The social amplification of risk*. Cambridge University Press.
- Piškorc, M., Antulov-Fantulin, N., Novak, P. K., Mozetič, I., Grčar, M., Vodenska, I., and Šmuc, T. (2014). News cohesiveness: an indicator of systemic risk in financial markets. *arXiv preprint arXiv:1402.3483*.
- Porter, M. E. (1985). Technology and competitive advantage. *Journal of Business Strategy*.
- Porter, M. E. P. M. (1986). *Competition in global industries*. Harvard Business Press.
- Power, M. (2009). The risk management of nothing. *Accounting, Organizations and Society*, 34(6-7):849–855.
- Renn, O., Laubichler, M., Lucas, K., Kröger, W., Schanze, J., Scholz, R. W., and Schweizer, P.-J. (2022). Systemic risks from different perspectives. *Risk Analysis*, 42(9):1902–1920.
- Rousseeuw, P. J. (1987). Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics*, 20:53–65.
- Schweizer, P.-J., Goble, R., and Renn, O. (2022). Social perception of systemic risks. *Risk Analysis*, 42(7):1455–1471.
- Shapley, L. (1953). A value for n -person games. In *Contributions to the Theory of Games*, pages 307–317. Princeton: Princeton University Press.
- Simaan, M., Gupta, A., and Kar, K. (2020). Filtering for risk assessment of interbank network. *European Journal of Operational Research*, 280(1):279–294.
- Slovic, P. (1992). Perception of risk: Reflections on the psychometric paradigm. In Krinsky, S. and Golding, D., editors, *Social Theories of Risk*. Praeger.
- Stanford NLP Group (2009). Stemming and lemmatization. <https://nlp.stanford.edu/IR-book/html/htmledition/stemming-and-lemmatization-1.html>. Accessed: 2021-09-05.
- Suchman, M. C. (1995). Managing legitimacy: Strategic and institutional approaches. *Academy of Management Review*, 20(3):571–610.
- Tauscher, K., Bouncken, R., and Pesch, R. (2021). Gaining legitimacy by being different: Optimal distinctiveness in crowdfunding platforms. *Academy of Management Journal*, 64(1):149–179.
- Tang, Q., Tong, Z., and Xun, L. (2022). Insurance risk analysis of financial networks vulnerable to a shock. *European journal of operational research*, 301(2):756–771.
- Task Force on Climate-related Financial Disclosures (2022). 2022 Status Report. <https://www.fsb-tcfd.org/publications/>. Accessed: 2023-10-06.
- Trueman, B. (1994). Analyst forecasts and herding behavior. *The Review of Financial Studies*, 7(1):97–124.
- Tsanakas, A. and Millosovich, P. (2016). Sensitivity analysis using risk measures. *Risk Analysis*, 36(1):30–48.
- Tuggle, C. S., Sirmon, D. G., Reutzel, C. R., and Bierman, L. (2010). Commanding board of director attention: investigating how organizational performance and ceo duality affect board members’ attention to monitoring. *Strategic Management Journal*, 31(9):946–968.
- Ye, W., Zhou, Y., Chen, P., and Wu, B. (2024). A simulation-based method for estimating systemic risk measures. *European Journal of Operational Research*, 313(1):312–324.