

Lange, Kai-Robin et al.

**Working Paper**

## Towards extracting collective economic narratives from texts

Ruhr Economic Papers, No. 963

**Provided in Cooperation with:**

RWI – Leibniz-Institut für Wirtschaftsforschung, Essen

*Suggested Citation:* Lange, Kai-Robin et al. (2022) : Towards extracting collective economic narratives from texts, Ruhr Economic Papers, No. 963, ISBN 978-3-96973-127-7, RWI - Leibniz-Institut für Wirtschaftsforschung, Essen, <https://doi.org/10.4419/96973127>

This Version is available at:

<https://hdl.handle.net/10419/264399>

**Standard-Nutzungsbedingungen:**

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

**Terms of use:**

*Documents in EconStor may be saved and copied for your personal and scholarly purposes.*

*You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.*

*If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.*



# RUHR

## ECONOMIC PAPERS

Kai-Robin Lange  
Matthias Reccius  
Tobias Schmidt  
Henrik Müller  
Michael Roos  
Carsten Jentsch

### **Towards Extracting Collective Economic Narratives from Texts**

## Imprint

---

### Ruhr Economic Papers

Published by

RWI – Leibniz-Institut für Wirtschaftsforschung

Hohenzollernstr. 1-3, 45128 Essen, Germany

Ruhr-Universität Bochum (RUB), Department of Economics

Universitätsstr. 150, 44801 Bochum, Germany

Technische Universität Dortmund, Department of Economic and Social Sciences

Vogelpothsweg 87, 44227 Dortmund, Germany

Universität Duisburg-Essen, Department of Economics

Universitätsstr. 12, 45117 Essen, Germany

### Editors

Prof. Dr. Thomas K. Bauer

RUB, Department of Economics, Empirical Economics

Phone: +49 (0) 234/3 22 83 41, e-mail: [thomas.bauer@rub.de](mailto:thomas.bauer@rub.de)

Prof. Dr. Ludger Linnemann

Technische Universität Dortmund, Department of Business and Economics

Economics – Applied Economics

Phone: +49 (0) 231/7 55-3102, e-mail: [Ludger.Linnemann@tu-dortmund.de](mailto:Ludger.Linnemann@tu-dortmund.de)

Prof. Dr. Volker Clausen

University of Duisburg-Essen, Department of Economics

International Economics

Phone: +49 (0) 201/1 83-3655, e-mail: [vclausen@vwl.uni-due.de](mailto:vclausen@vwl.uni-due.de)

Prof. Dr. Ronald Bachmann, Prof. Dr. Manuel Frondel, Prof. Dr. Torsten Schmidt,

Prof. Dr. Ansgar Wübker

RWI, Phone: +49 (0) 201/81 49-213, e-mail: [presse@rwi-essen.de](mailto:presse@rwi-essen.de)

### Editorial Office

Sabine Weiler

RWI, Phone: +49 (0) 201/81 49-213, e-mail: [sabine.weiler@rwi-essen.de](mailto:sabine.weiler@rwi-essen.de)

### Ruhr Economic Papers #963

Responsible Editor: Ludger Linnemann

All rights reserved. Essen, Germany, 2022

ISSN 1864-4872 (online) – ISBN 978-3-96973-127-7

The working papers published in the series constitute work in progress circulated to stimulate discussion and critical comments. Views expressed represent exclusively the authors' own opinions and do not necessarily reflect those of the editors.

---

## **Ruhr Economic Papers #963**

Kai-Robin Lange, Matthias Reccius, Tobias Schmidt, Henrik Müller, Michael Roos,  
and Carsten Jentsch

# **Towards Extracting Collective Economic Narratives from Texts**

## **Bibliografische Informationen der Deutschen Nationalbibliothek**

---

The Deutsche Nationalbibliothek lists this publication in the Deutsche Nationalbibliografie;  
detailed bibliographic data are available on the Internet at <http://dnb.dnb.de>

RWI is funded by the Federal Government and the federal state of North Rhine-Westphalia.

<http://dx.doi.org/10.4419/96973127>

ISSN 1864-4872 (online)

ISBN 978-3-96973-127-7

---

Kai-Robin Lange, Matthias Reccius, Tobias Schmidt, Henrik Müller, Michael Roos,  
and Carsten Jentsch<sup>1</sup>

# Towards Extracting Collective Economic Narratives from Texts

## Abstract

*Identifying narratives in texts is a challenging task, as not only narrative elements such as the actors and events have to be identified but their semantic relation has to be explained as well. Despite this complexity, an effective technique to extract narratives from texts can have a great impact on how we view political and economical developments. By analyzing narratives, one can get a better understanding of how such narratives spread across the media landscape and change our world views as a result. In this paper, we take a closer look into a recently proposed definition of a collective economic narrative that is characterized by containing a cause-effect relation which is used to explain a situation for a given world view. For the extraction of such collective economic narratives, we propose a novel pipeline that improves the RELATIO-method for statement detection. By filtering the corpus for causal articles and connecting statements by detecting causality between them, our augmented RELATIO approach adapts well to identify more complex narratives following our definition. Our approach also improves the consistency of the RELATIO-method by augmenting it with additional pre- and post-processing steps that enhance the statement detection by the means of Coreference Resolution and automatically filters out unwanted noise in the form of uninterpretable statements. We illustrate the performance of this new pipeline in detecting collective economic narratives by analyzing a Financial Times data set that we filtered for economic and inflation-related terms as well as causal indicators.*

JEL-Codes: C18, C55, C87, E70

Keywords: Econometrics; narrative; text mining; coreference resolution; named entity recognition; causal linking

August 2022

---

<sup>1</sup> Kai-Robin Lange, TU Dortmund; Matthias Reccius, RUB; Tobias Schmidt, TU Dortmund; Henrik Müller, TU Dortmund; Michael Roos, RUB; Carsten Jentsch, TU Dortmund. – This study is part of a project of the Dortmund Center for data-based Media Analysis (DoCMA) at TU Dortmund University. The work was supported by the Mercator Research Center Ruhr (MERCUR) with project number Pe-2019-0044. In addition, the authors gratefully acknowledge the computing time provided on the Linux HPC cluster at TU Dortmund University (LiDO3), partially funded in the course of the Large-Scale Equipment Initiative by the German Research Foundation (DFG) as project 271512359. – All correspondence to: Kai-Robin Lange, TU Dortmund University, 44221 Dortmund, Germany, e-mail: [kalange@statistik.tu-dortmund.de](mailto:kalange@statistik.tu-dortmund.de)

# 1 Introduction

*What has caused the rising inflation in early 2022? What exactly has lead to Russia's invasion of Ukraine in February 2022? Are the measures of the of the US government combating Covid 19 appropriate?* Questions like these are so complex that individual reasoning will not reflect the full set of information that is relevant to them. Instead, opinion formation will be guided by narratives that abstract from the multitude of facts available. Narratives are cognitive instruments (Bruner, 1991) through which humans organize their experiences, explain the past and predict the future (Bénabou et al., 2018). In complex and uncertain situations, it is convincing narratives rather than facts and figures that enable agents to take action (Tuckett and Nikolic, 2017). As a result, the study of economic narratives is becoming an important component of empirical macroeconomics.

Fortunately from an empirical perspective, narratives are not only psychological. They also manifest as linguistic structures and are often used in textual and other media (Meer, 2022). However, despite the recent advances in quantitative language modeling, no empirical model yet exists that is able to automatically identify narratives in a corpus of textual data. One reason why such one-step analyses using word embeddings (Mikolov et al., 2013) or topic models (Blei et al., 2003) are not up the task is that narratives are quite particular constructs. Even though state-of-the-art language models can deliver impressive results in related and seemingly similar tasks like for instance text summarization, a narrative is more specific than a summary. Thus, any attempt to identify narratives inevitably necessitates a prior definition of the term.

In this paper, we leverage the definition of collective economic narratives (CENs) put forth by Roos and Reccius (2021) to extract inflation narratives from newspaper articles. According to this definition, a CEN is a sense-making story about the economy that is shared within a social group and suggests action. Based on this definition, we identify text passages that establish a causal link between economic events and actors and thus allow the reader to make sense of a concrete economic issue. In a general sense, the aim of this study is to start closing the gap between the theoretical understanding of narratives as provided by the cognitive psychology literature and the definition in Roos and Reccius (2021) and their empirical identification.

Often research on narrative extraction focuses on determining narratives through named entity recognition and event detection (Benner et al., 2022). Ultimately, these named entities and events can be collected and presented by displaying

their connection (e.g. the verbs that connect a named entity with an event). One such example is the RELATIO-method by Ash et al. (2021), which clusters named entities with similar meanings. Words such as "health insurance", "health human service" and "health care system" can be collectively represented by the word "healthcare". As a result of this dimensionality reduction, RELATIO is less prone to tagging synonyms of the same named entity as different entities. Each extracted text passage corresponds to the structure "AGENT VERB PATIENT" (for instance "worker lose jobs"), so that all statements contained in the corpus can later be displayed as a network with the agents and patients as nodes and the verbs as edges (see Figure 2). The edges can then be weighted based on how often that exact narrative occurs. RELATIO works well for identifying economic events, phrases and statements which is why we use a modified version of it as part of our pipeline.

However, the structure that is recognized by RELATIO is missing the sense-making property that is central to the concept of CENs. RELATIO will commonly detect statements like "worker lose jobs" or "god bless america", that are either common phrases or simply facts. Such statements are not necessarily narratives in the sense of being driven by the narrator's world view or interacting with the recipient's belief system. Thus, we consider RELATIO to identify narrative building blocks that have to be combined in some way to yield full economic narratives. Combined statements like "workers lose jobs due to Covid 19" or "god bless america because america is best country in the world" would be considered complete CENs. Therefore, our pipeline is aimed at disregarding the noise created by incomplete narratives through the identification of causal indicators such as "because" that connect RELATIO-style narrative building blocks.

Being able to analyze the narratives attributed to a specific economic event or actor over time could offer insights on how they are viewed, how this perception changes and what actions may be taken towards them. Existing narrative extraction schemes often purely try to filter out events individually and display their factual connection as Directed Acyclic Graphs or temporal flow charts. In this paper however, we differentiate between narratives and factual description of events. That is, the sentence "XYZ died of Covid 19." is not considered a narrative, as it displays an event but does not contain an opinion or a sense-making element. On the other hand, the statement "XYZ died of Covid 19 because he was unvaccinated" is a narrative, as it is a counterfactual statement that is based on an assumption, opinion or personal world view rather than pure facts.

We propose a pipeline to identify such narratives by analyzing texts for causal



structures. This pipeline identifies causal linkages between smaller statements, which can be seen as events, to provide the user with a collection of causally connected events to draw narratives from. For this purpose, first, we explain in detail, how our definition of a narrative differs from the literature in Section 2. We describe the data we use in Section 3, present each natural language processing task that is later used in our pipeline individually in Section 4 and combine them to propose our pipeline in Section 5. We present our results in Section 6 and provide an outlook to future research in Section 7. In Section 8, we conclude the paper.

## 2 Definition of an economic narrative

Different definitions of narratives coexist in several fields of research. Often times, the definitions used in applied work are not theoretically motivated and appear to be tailored to fit a specific research agenda. In some studies about narratives, the term is not defined at all. However, not subscribing to any specific definition of narratives allows for any type of verbal or textual information to be called and analyzed as a narrative. This is unfortunate because narratives possess specific characteristics that constitute their persuasiveness and differentiate them from other forms of text.

We leverage the definition put forth by Roos and Reccius (2021) according to which a collective economic narrative (CEN) is "*a sense-making story about some economically relevant topic that is shared by members of a group, emerges and proliferates in social interaction, and suggests actions*" (Roos and Reccius, 2021, p. 13). The most fundamental part of this definition is the *sense-making story*. A story can be viewed as depicting a temporal sequence of (possibly independent) events without claiming any specific connection between them. For example, "Inflation started soaring and then my colleagues all quit" is a story (Forster, 1927). When those events are logically connected with each other, the story becomes a plot: "Inflation started soaring and then my colleagues all quit because our boss would not raise their salaries accordingly" is a plot. In this case, the causal connection provides the sense-making glue between the events. It allows the recipient of the narrative to deduce that it was the rise in inflation that ultimately caused the colleagues to quit. Similar to Andre et al. (2022), we exploit this sense-making function of causal reasoning as our primary way of identifying CENs. It is important to note, however, that an explicit causal connection is not necessarily required for a CEN. Sense-making connections can

also be implied. An example is the sentence "Inflation shot up and the central bank raised interest rates.", which is a sense-making story to recipients with the proper background knowledge and economic education. In this case, the sense-making aspect of the narrative interacts with the belief system shared by most economically literate recipients.

Other elements of the definition by Roos and Reccius (2021), such as the emergence-property of narratives, are also quite demanding from an empirical standpoint because they cannot necessarily be detected or analyzed in a single text. This emergence-property for instance may entail the existence of proto-narratives – different versions and drafts of a narrative – that coexist, circulate and combine before a consensus is reached and a full CEN emerges in a group. This process cannot be captured by looking at single texts at a time. Instead, for this purpose, dynamic methods have to be used that can capture and represent narrative developments over time.

A distinction has to be made between personal narratives that people use to make sense of their own lives and collective narratives that serve a function for groups. The psychological literature tends to focus on the function of narratives in the decision-making of individuals. When faced with a decision, narratives thin out the vast space of possible actions that could be taken by favoring the option that best fit a simple and plausible narrative. While decision-making is obviously important from a macroeconomic standpoint, a narrative must transcend the individual and guide collective action in order to become relevant for the macroeconomy. In contrast to personal narratives, CENs can spur *collective* action. For example, "I need to renegotiate my wage because the soaring inflation lowers the purchasing power of my income" is a personal narrative insofar as it is employed by an individual to (i) analyze the current state of affairs, (ii) make predictions about the future and (iii) decide on a course of action. Only once it is shared and believed in within a larger group – of example a trade union – it emerges as a CEN and may affect macroeconomic dynamics such as wage-setting to a significant degree.

### 3 Data

The aim of our analysis is to identify CENs that are shared by a group of people and have the potential to become economically relevant. Like other scholars before us (Müller et al., 2022; Ter Ellen et al., 2021; Larsen and Thorsrud, 2019),

we choose media data as our unit of study in this regard. We do so mainly for two reasons: i) journalistic texts are a pretty accurate proxy for debates that move society and ii) news articles are of high relevance for many market participants, especially for central banks. That is, narratives found there can not only potentially be interpreted as CENs, but also provide an excellent starting point for further research.

### 3.1 Narratives and the media

Journalists play a crucial role in the dynamics of narrative formation (Shiller, 2017). In order to meet the interest of their readership, journalists follow Twitter debates, evaluate letters to the editor, analyse the traffic on their online portals, track the reporting of their competitors - and thus know which topics move the public and which do not. If an already existing narrative gets big enough (for example, about villains and victims of inflation), it is likely that it finds its way into editorial conferences and thus (in some way) into reporting.

At the same time, this reporting leads to topics, frames<sup>1</sup> and narratives becoming entrenched in public debate.

We know from agenda-setting theory (McCombs and Shaw, 1972) that the media have a decisive influence on what we discuss in a society. When it comes to evaluate the general economic situation, media are even the most important source for many citizens (Lischka, 2015; Blinder and Krueger, 2004). So, while public debates regularly find their way into editorial conferences, it is primarily the media themselves that determine the popularity of a topic and/or a certain narrative. This dynamic leads us to see media texts as a natural and very promising textual basis for finding sufficiently popular narratives on inflation.

---

<sup>1</sup>In communication science frames are often described as the selection of "some aspects of a perceived reality [to] make them more salient in a communicating text [...]" (Entman, 1993, p.52). Just like narratives, frames are judgemental and have the potential to establish a social view that influences decision-making processes. Compared to narratives, however, frames tend to be static in the sense that frames focus on specific topics, whereas narratives cover a longer period of time. Following Müller et al. (2018) we propose the interpretation that "a frame is to a narrative what a still is to a movie" (p.559). More information on media frames and their link to public opinion in Scheufele (1999); De Vreese (2005); Scheufele and Tewksbury (2007); Matthes (2014).

### 3.2 Central banks and the media

Since central banks have recognised forward guidance as a central element of their monetary policy, media reports have increasingly become the focus of their attention. Forward guidance describes central banks' efforts to manage inflation expectations – and ultimately actual inflation – through targeted communication. It is "intended to correct faulty expectations, and thereby reduce misallocations of resources" (Blinder et al., 2008, p.22). Used wisely, it has the potential to move financial markets (Blinder et al., 2008), improve the predictive power of monetary policy (Haldane and McMahon, 2018), guide inflation expectations (Armantier et al., 2016; Binder and Rodrigue, 2018; Eusepi and Preston, 2010), and influence consumer behaviour (Armantier et al., 2015).

The media play a crucial role in this process (Berger et al., 2011). As a central bank wants consumers and decision-makers to interpret the economic situation correctly, it needs journalists to understand their analyses, and communicate them to the public in an understandable and captivating way. A central bank will hardly succeed in setting a certain narrative - for example, that rising commodity prices are responsible for inflationary dynamics - if most of the business press presents a different interpretation. Therefore, central banks make use of several communication channels (Monthly Bulletin, speeches, interviews, monthly press conferences etc.) to directly influence media coverage (Conrad and Lamla, 2007; Haldane and McMahon, 2018). However, despite all their efforts, they do not always succeed in getting their own narrative featured in the press: In a Speech at the 148th Baden-Baden Entrepreneurs' Talk, ECB Director Isabelle Schnabel expressed her displeasure that "many supposed experts and the media are again rousing people's fears without explaining the reasons behind the price movements" (Schnabel, 2021). The quote impressively illustrates how important the business press is for the ECB's forward guidance<sup>2</sup>. Above all, it shows how crucial it is for a central bank to recognise any (unwanted) narratives in the press. It is not surprising that central bank researchers are increasingly involved in research on narratives, especially narratives in the press (Ter Ellen et al., 2021; Nyman et al., 2021; Kalamara et al., 2020).

---

<sup>2</sup>The influence of media coverage on inflation expectations is well documented. See e.g. Larsen et al. (2021); Coibion et al. (2018b,a); Lamla and Lein (2014).

### 3.3 The media as data

Both, the high potential of media data to contain CENs and its relevance for the economy leads us to use the business press as our unit of investigation.

The medium that fits our research interest most is the Financial Times (FT). The weekday Financial Times newspaper is one of the world's leading business, politics and world-affairs newspapers. If there are any economic narratives circulating in a society, we are likely to find them in the Financial Times - either because they originated there or because they are reproduced in its reporting.

Our analysis corpus consists of all FT articles published in the English-language edition between 1/1/2010 and 18/07/2019 that depict economic uncertainty to some extent. The search term that narrows the corpus to economic uncertainty is based on the research of Baker et al. (2016). The keyword is of the form *econom\** OR *uncertain\** and covers all articles that contain both patterns at least once.

In order to only capture narratives about the inflation, we further reduce the text base to those paragraphs in which the pattern *inflation* occurs. We define a paragraph as the sentence in which the pattern *inflation* appears as well as the six sentences before and the six sentences after it. As a further restriction we filter for paragraphs that also contain at least one causal indicator (see Table 1). Our text base comprises a total of 18,375 unique paragraphs. In this collection, we search for cause-and-effect statements that match our definition of an economic narrative.

## 4 Natural language processing sub-tasks

In the following sections we explain the process of some common tasks of natural language processing that are all incorporated into our pipeline. While Semantic Role Labeling, Named Entity Recognition and K-Means clustering are already used by the RELATIO-method Ash et al. (2021) which we cover in Section 4.5, in our pipeline, we perform an additional layer of Named Entity Recognition paired with Coreference Resolution to improve the method's consistency and make the results interpretable out of their textual context. After defining these components of our pipeline in this section, we will explain how we combine

them with additional post-processing and filtering to create a narrative extraction technique in Section 5.

## 4.1 Semantic Role Labeling

Semantic Role Labeling is the process of assigning tokens to semantic roles that describe their effect within the sentence in relation to a specific verb (Jurafsky and Martin, 2020, pp.405 ff.). More specifically, in the RELATIO-method (Ash et al., 2021), which use as a part of our pipeline, Semantic Role Labeling is used to determine *who* did *what* to *whom*. In this case, *what* represents the verb and *who* as well as *whom* represent the active and receiving part of this action, which we call *agent* and *patient*.

The Semantic Role Labeling in this paper will be based on the state-of-the-art approach by AllenNLP (Gardner et al., 2018), which is also able to handle verb negations, allowing us to incorporate these into our statements and narratives.

## 4.2 Named Entity Recognition

After splitting each sentence into its semantic roles, we only want use the roles relevant to us, as these statements will represent the most important information of any event: *who did what to whom?*. We consider other roles to be unwanted noise, as anything beyond the description of the event would add minor details or linguistic mannerisms of the author, which might differ from text to text, to the statement. The statements thus only become needlessly large and distinct from another, which would greatly increase the efforts needed to identify common statements shared by multiple authors. Condensing short statements, many of which are overlapping, into a set of logical narratives is already a complex task. Hence, doing so to distinct and longer statements would be even harder. In the latter case, little mistakes from probabilistic models in the pipeline would yield an even larger amount of false statements. Hence, we will only use smaller and compact statements for our pipeline and leave narrative extraction using longer statements for future research. This compact form does however still contain unwanted noise, as many of the words applied to each semantic role may be representing the same underlying latent entity, but have a slightly different wording. For instance the terms "short term rate" or "short term interest rate"

may both represent the same entity and can be generalized by denoting them and similar entities simply as "interest rate".

Named Entity Recognition is the task of checking whether a semantic building block represents a real world entity, may it be a person, a specific building or an organization (Jurafsky and Martin, 2020, pp.164 ff.). In our context, it checks if an agent or patient itself is an entity or if it refers to a latent entity that is not specifically named. The Named Entity Recognition in this paper is performed by Spacy's Named Entity Recognition (Honnibal and Montani, 2017).

### 4.3 Coreference Resolution

Ultimately, instead of looking at each text in detail, we want to exclusively compare narratives from anywhere in the corpus to another without needing to read their respective contexts. This yields the problem that some sentence building blocks like pronouns cannot simply be interpreted out of context. For instance, a human reader may figure out that the "it" in the sentence "The fed tries to combat it by raising interest rates" represents the word "inflation". Our pipeline also needs to detect and resolve such coreferences, or else the resulting collection of narratives would still need a lot of manual work to interpret.

This automatic replacement of pronouns is called Coreference Resolution (Jurafsky and Martin, 2020, pp.445 ff). In this paper, we use Spacy's Coreference Resolution model (Honnibal and Montani, 2017). The nouns within each sentence are detected and linked to nouns and pronouns in other sentences, based on which entity they represent. These linked nouns and pronouns are called a Coreference chain. The chain is then resolved by choosing one of the chain elements to replace all other chain elements in the original texts with. We will describe how we choose our element to resolve the chains with in Section 5.2.

### 4.4 K-Means clustering

As a large number of distinct entities is not easily interpretable, reducing them to a reasonably small number of latent entities is a crucial dimension reduction step. We cluster the resulting patients and agents that cannot be linked to entities by Named Entity Recognition. This way, one latent entity represents a larger

number of agents and patients. In the context of natural language processing, clustering is performed in combination with word embeddings (Mikolov et al., 2013), which transform a word into a vector based on its semantic meaning. The word embeddings, each representing one word, are clustered using K-Means clustering. The word with an embedding closest to the resulting cluster mean, is chosen to represent said cluster. Thus, each non-recognized agent or patient is assigned to a latent entity that is determined by one of the cluster means.

The K-Means algorithm, given vectors  $x_i \in \mathbb{R}^d$  for  $i = 1, \dots, n$ , assigns each vector to one cluster, based on the distance to the cluster's mean. Thus the training objective is to minimize the distance of all our vectors to their respective closest cluster means

$$\arg \min_{S=\{S_1, S_2, \dots, S_k\}} \sum_{i=1}^k \sum_{x \in S_i} \|x - \mu_i\|$$

where  $\mu_i \in \mathbb{R}^d$  denotes the mean of cluster  $S_i$ . The only parameter of this algorithm is  $k$ , the number of clusters to use (Lloyd, 1982). To make our clusters distinct enough, we chose  $k = 300$  clusters for this paper.

## 4.5 RELATIO

RELATIO is a method developed by Ash et al. (2021) to identify narratives in large collections of text. However, the statements identified by RELATIO are narratives in a much broader sense than the CENs that we consider in this paper. To avoid confusion between the two definitions of a narrative, we will call the results of the RELATIO-method "statements" instead of narratives. Ash et al. (2021) extracts statements from a text, which are formed by three narrative building blocks called "agent", "verb" and "patient" as explained in Section 4.1. Although these building blocks do not suffice our definition of a narrative, they do provide important groundwork for us to be able to extract more complex narratives that conform more closely to the concept of a CEN.



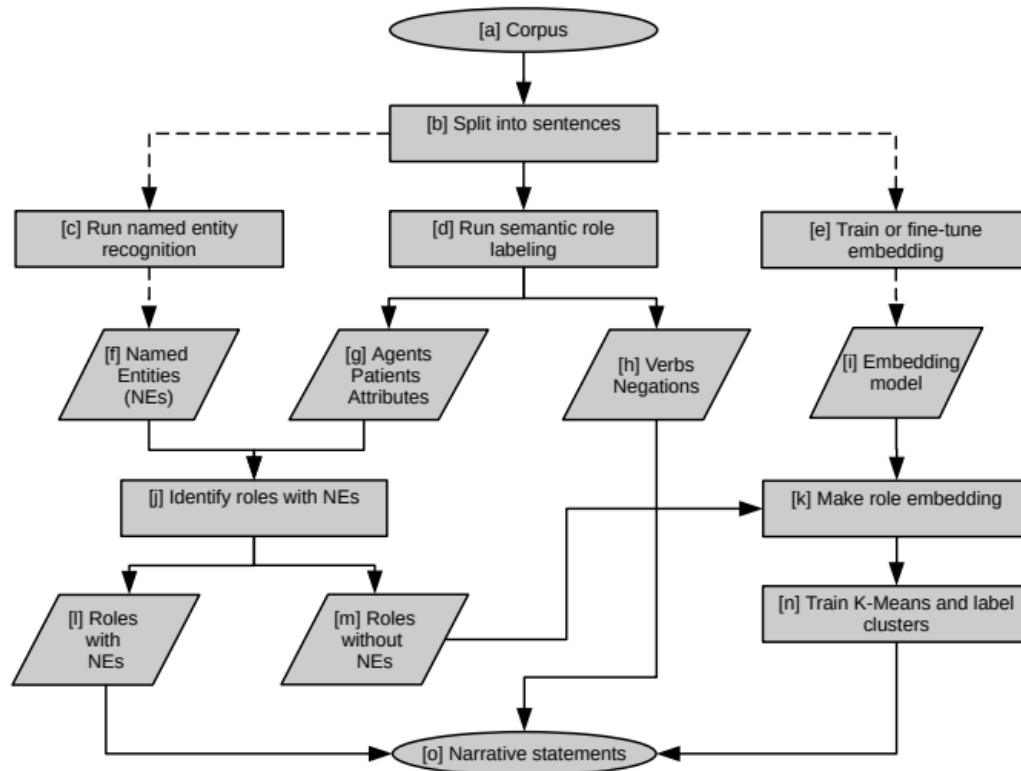


Figure 1: Flowchart of the RELATIO-method (Ash et al., 2021)

The RELATIO-method is displayed as a flowchart in Figure 1. It first splits each document into parts to evaluate each sentence separately. Then, Spacy's (Honnibal and Montani, 2017) Named Entity Recognition is used to detect named entities within each sentence. Parallel to this, Semantic Role Labeling is used to distinguish agents, patients and verbs. If the agents and patients align with named entities that were found, they are directly used as an agent of patient respectively. The remaining agents and patients that cannot be identified as named entities are likely a very high dimensional set of unique terms. That is, a lot of agents or patients will represent the same latent entity, but will have slightly different wordings. To reduce the dimension of these terms, they are clustered to a smaller set of latent entities using K-Means clustering. Thus, all agents and patients are assigned an entity. The final statements can be formed by combining agents and patients as well as the verbs that were identified during Semantic Role Labeling (verb negations included, if any).

## 5 Causal RELATIOns

In this section, we present our proposed pipeline by describing, how we combine and enhance the methods defined in Section 4 and proceed beyond the scope of the RELATIO-method. We detail our filtering process to only look at relevant articles, describe how we pre- and postprocess our data to be able to detect more complex statements and discuss how we intend to look for causal connections.

### 5.1 Detecting causal articles

As defined in Section 2, we are focusing on narratives displaying (explicit) causal connection between events. Filtering the corpus we seek to analyze for causal articles or paragraphs is thus a preprocessing step that will reduce unwanted noise. Like most tasks in natural language processing, this can be done using either a lexicon-based or learning-based approach.

We did train a BERT-model (Devlin et al., 2019) to detect causal articles (which we describe in further detail in Section 7), however, in this paper, we decided to filter our data using a lexicon-based approach. As we look for causally connected events, causality not only needs to be detected within articles in general, but, ultimately, when looking at two particular statements. While a learning-based approach that is specifically trained for this task will likely do this sufficiently well (see Section 7), causal indicators such as "because" can be used to link causal statements both more easily and more transparently. Crucially, lexicon-based identification reduces the chance for misinterpretation due to a failure of the learning-based model. Therefore we decided to use a lexicon approach to filter for causal articles, as implicit causality detected by a learning-based approach would not improve our pipeline, unless the model used to connect statements later on is also learning-based. The filtering process is described in Section 3.

### 5.2 Preprocessing and detecton of statements

After the texts to analyze have been filtered from the corpus, we use further preprocessing to improve our narrative extraction. For this we utilize a combination of Named Entity Recognition and Coreference Resolution, using both Spacy's Coreference Resolution and Named Entity Recognition on our original

texts (Honnibal and Montani, 2017). We then look at each Coreference chain separately and use the Named Entity as an indicator on how to resolve the Coreference chain. If an entity detected by our Named Entity Recognition algorithm or at least 33% of its tokens (most tokens are simply words) are part of an element of the chain, the chain is resolved by inserting this entity into all of its references. The reasoning for the 33% threshold is that it enables us to identify three-part names, even if only one of its parts (e.g. only the last name of "Donald J. Trump") comes up in a chain element. If no entity is overlapping with the chain, the chain resolves using its first element, as this is likely the correct entity and later elements consist of pronouns referencing it. The texts with resolved chains are then saved as preprocessed texts.

After these preprocessing steps, the RELATIO-method is applied to the resulting texts. In addition to the usual pipeline, we return information about which token within a sentence represents the beginning and end of a statement. This is helpful in identifying smaller rather than larger statements. RELATIO often creates very long statements that tend to add additional noise and complicate interpretation. For instance, the sentence "economists say, the current inflation is caused by supply shortages and central banks must raise interest rates" will contain three entangled statements in the style of agent; verb; patient: *inflation; is caused; supply shortages, central banks; raise; interest rate* and *economists; say; the current inflation is caused by supply shortages and central banks must raise interest rates*. The first two statements are interesting for CEN extraction, but the third statement only repeats the entire sentence. This is because Semantic Role Labeling assigns the entire relative sentence as a patient to the verb "say". While this is factually correct, the resulting statement cannot be easily grouped alongside other statements due to its length and will only add unwanted noise to our resulting data. Instead of using all three statements, we only want to use the first two. We choose the maximum number of non-overlapping statements per sentence to represent the sentence – this task is similar to common scheduling-tasks in computer science. Thus, to filter out unwanted long statements, we use a greedy approach within sentences by always adding the current statements which stops first and does not overlap with any statement that was identified before (and thus starts at a later token than the previous statement had ended on) to our collection of filtered statements. This collection is our final statement list, which we use to display our results.

### 5.3 Causal direction

A meaningful narrative extraction pipeline should incorporate a way to detect a causal connection so that it can connect statements in a logical manner. This can be done using lexicon approaches or by analyzing the linguistic structure of the documents. While explicit causality like in the sentence "The inflation is high because of the supply shortage." are easily identifiable by lexicon approaches due to the word "because", implicit causality is not trivially detectable. However, lexicon-based methods that simply check a document for certain causal or temporal keywords are not flexible enough to detect causality in the statements "The Russo-Ukrainian war caused the gas prices rise." or "Donald Trump was a bad President. He divided the American people.", even though the causality is easily identifiable for a human.

While learning-based approaches, such as transformer-based models would be suitable for such a task, they also need to be trained correctly and carefully. By trying to detect implicit causality, such models can yield a high false-positive error rate and thus connect a lot of statements that are in reality not causally connected. This error rate can also not be easily resolved by persons checking the results manually, as such implicit causality might be hard to detect by the statements alone and the person thus would need to check the surrounding context of the statements. A learning-based model still remains a possible improvement in future research, but needs to be trained well to minimize the occurrences of false-positive errors. A lexicon approach on the other hand does not require additional training and the false-positive error is greatly reduced, as a explicit causal indicator directly implies the existence of causality, while implicit causality always relies on the context of the statements. These causal indicators can be found in Table 1 in the appendix. We look for these causal indicators in combination with two adjacent statements to combine them into a larger statement. We show examples of this in Section 6.2.

## 6 Evaluation

In this chapter, we evaluate the performance of our narrative extraction pipeline and compare it to the RELATIO-method in two steps. In Section 6.1 we show the graphical representation of our results using pre- and posprocessing compared to the original RELATIO-method. Then we show exemplary "complete"

narratives that can be found using our lexicon approach for causality-filtering in Section 6.2.

## 6.1 Results

We display our resulting statements using the graph structure of the original RELATIO in Figure 2 and Figure 3. While Figure 2 shows the original detected actors and patients, Figure 3 shows the results after clustering the non-recognized entities to reduce dimensionality. In addition to this, Figure 7 shows the results when using the original RELATIO-method without any pre- or postprocessing. We can see that the most prominent statements of Figure 7 are similar to the ones of our adjusted pipeline in Figure 3, which was expected as the goal of Coreference Resolution was not to change the important statements of the corpus. Instead, it increased the number of already prominent statements, as common entities like *fed*, *interest rate* or *inflation* appear even more often. The pre- and postprocessing are also working as intended in context of their goal described in Section 5.2. Long statements that are too long to be generalized like `commitee; continiue; anticipate economic condition include low rate resource utilization...` (see Figure 7 at the bottom) are filtered out in Figure 2 and Figure 3.

While we see some statements of interest for our research in Figure 2, such as "fed eliminate inflation", "inflation adjust interest rate", "economy gain strength" or "economy need help", we can see there are a lot of distinct statements due to slight differences in wording. Entities like "interest rate", "benchmark interest rate", "bank benchmark interest rate", "fed benchmark rate", "short term rate", "short term interest rate" all refer to either the same or a very similar entity that can be generalized as "interest rate". Due to this, we use RELATIO's clustering algorithm to reduce dimensionality and work with less, more interpretable entities.

Figure 3 shows that a smaller amount of distinct statements improves our search for inflation narratives, as the clustering to latent entities reduces noise and also increases the number of inflation statements, as it is recognized as a latent entity. The clustering process behind the dimension reduction profits from the preprocessing as well as our filter for shorter statements. On the one hand, important entities like "inflation" appear more often, so that less agents and patients have to be clustered and are instead directly assigned to an entity. On the other hand, unreasonably long terms like the patients that can be found at the

bottom of Figure 7 cannot simply be summarized into a single word. Removing them from model by filtering for shorter statements thus also increases the quality of the clustering process.

With these new statements we can filter out additional statements of interest like "oil price push inflation", "economist expect inflation", "government raise inflation", "inflation adjust unemployment".



Figure 2: Top 75 statements (39 distinct agents and patients) detected by the Causal RELATIONS pipeline without dimension reduction.



Figure 3: Top 75 statements (24 distinct agents and patients) detected by the Causal RELATIONS pipeline including dimension reduction.

We can also see statements that seem to stem from conflicting narratives such, that is "fed fight inflation" and "fed tolerate inflation". These two statements also show, why it is important to extract the entire narrative, rather than parts of the statements. Without further information we cannot tell, if these two statements are conflicting or whether they display the same narrative but show different perspectives, such as "The fed tolerates inflation in the livestock sector while it fights inflation in the energy sector". To identify these narratives, we need to analyze causal connections to relate the incomplete statements. We aim to display such narrative chains as graphs similar to the statements above. For now, we will display some exemplary CENs extracted from the text data in the following section.

## 6.2 CENs as chains of causally linked statements

The ultimate goal of the narrative extraction pipeline is to identify CENs in a fully automated way. Naturally, the requirement for a causal, sense-making link between narrative building blocks reduces the set of narratives that are extracted by the pipeline compared to RELATIO. This narrative structure also makes for more comprehensive and linguistically more complex narratives, as it increases the total number of featured tokens. Figure 6 visualizes two examples of a CEN as two causally linked statements that both consist of an AGENT, a PATIENT and VERB, where the latter specifies the relationship between the two former elements.

The CEN displayed in Figure 4 identifies a causal relation between the price-setting of fast-food restaurants and government subsidies for certain crops. Both halves of the CEN are identified separately and then connected through the causal indicator "because". In this case, coreference resolution was not utilized to attain the results because all entities and proper names in the narrative are stated and no pronouns are used.

The CEN in Figure 5 has indeed only been identified through coreference resolution. In the original text, the statement reads "*Yet funds are buying bonds because they (or their advisers) deem them low-risk.*" The coreference resolution algorithm has correctly replaced the pronoun "*they*" with the named entity "funds" and the pronoun "*them*" with the term "bonds" before the narrative was identified. Without the prior use of coreference resolution, the second half of the CEN would not have been detected by RELATIO.






Figure 4




Figure 5

Figure 6: Examples of extracted CENs. Both sides of the plot show extracted statements that our pipeline links by the means of a causal indicator (here: "because").

It is important to note that, at this point, the CENs identified with our method have to strictly conform to the structure displayed in Figure 6. This empirical strategy minimizes the identification of false positives, as was the primary objective of this research. However, this approach also increases the false negative rate with regards to the sense-making criterion of CENs. For example, the statement *"The last 12 months have brought growing reports of mass layoffs, unpaid wages and factory closings because of outstanding government debts."* is missed by the algorithm because its second half is a CEN, but does not conform to the AGENT, VERB, PATIENT structure. It would be very easy, however, to restate the narrative in a way that would be recognizable to the CEN detection algorithm: *"The last 12 months have brought growing reports of mass layoffs, unpaid wages and*

*factory closings because **the government has accumulates a lot of debt.***". The two statements are semantically identical, but structurally different.

## 7 Outlook

This paper is designed to show our current work in extracting narratives from texts. As this work is not yet complete, we want to designate a section to provide an outlook to future research. As the extraction of narratives is a very complex process, we do have a lot of ideas of how to improve our pipeline, as there are many ways it can be adjusted step by step.

While we are able to detect explicit causality within chains by looking for causal indicators, displaying a large amount of narratives is a complicated task, as the causal connections open up a large variety of different narratives to be found. A graphic thus cannot simply show all detected narratives, but has to display them in a coherent way. In addition to this, we want our causality detection to not be based on lexicon approaches but rather learning-based approaches such as BERT, which could work like a variation of C-BERT (Khetan et al., 2020) that detects the flow of causality between two statements. As explained in detail in Section 2, the reason for this is that lexicon approaches can necessarily only identify explicit causal cues. Because sense-making stories can exist even when events are connected temporally or in other ways, a learning-based approach is required to improve the precision of the identification algorithm. However, the C-BERT model cannot be used directly, as it only detects the direction of causality within a statement and not between two statements. Creating such a model is thus an outlook to future research.

As soon as this causal detection between statements is not based on a lexicon any more, we will be able to filter our texts more freely. While researching for this topic, we fine-tuned a BERT model (Devlin et al., 2019) to detect (even implicit) causality within paragraphs using a self-created training data set based on the Yahoo Questions-and-Answers data set (Yahoo! Webscope, 2008). As far as we know, there is no data set available that displays causal versus non-causal texts in a journalistic setting, or at least not in a scale that is sufficient for training a BERT-model, which is why we use this Yahoo data set instead. This training data set uses the answers to questions asked on Yahoo Answers as its documents. We split the data into causal and non-causal texts by filtering by the type of question. An answer is considered causal if the question contains the word "why" and is

considered non-causal if it does not contain the word "why" as well as a set of previously defined causal indicators. While a data set based on Yahoo Answers is not optimal for analyzing journalistic texts, it yields a starting point to work with. With the help of manual work, a journalistic data set can be created in future research, for instance by using a data set creation approach similar to the one of Pavllo et al. (2018).

Lastly, the clustering of entities used by RELATIO can be improved by using an embedding model that is fine tuned for economic texts rather than the general-use model proposed by Ash et al. (2021). In a similar matter, to analyze the current economic developments, we aim to use an updated version of the Financial Times data set, as our version only includes articles up until 2019. We could also improve the other aspects of the RELATIO-method or even our preprocessing in the form of Coreference Resolution or Named Entity Recognition. We are however currently content with their results and want to focus on the tasks that we are still missing to display a narrative in the form we want to. Choosing which Coreference Resolution, Semantic Role Labeling or Named Entity Recognition model we use can give our model the final touch, but we do not expect these changes to massively impact our results.

We also have ideas on how to proceed further, beyond the scope of this pipeline. While we are currently working on detecting and displaying narratives, we are also asking the question on how to monitor a narrative. For this, we are contemplating the idea of using a combination of this narrative extraction model in combination with change detection methods, which we already worked on (Rieger et al., 2022), to link change points in time to narrative shifts. We are also working on alternative perspectives to narrative detection, such as the detection of entity-related events constituting (new) narratives in a time line by focusing on entities rather than causal connections (Benner et al., 2022). As motivated in Section 2, a method based on change detection may also be able to take account of the dynamic aspects of narrative emergence. Such dynamic aspects are fundamental to CENs, but they elude cross-sectional methods by construction.

## 8 Conclusion

We present a pipeline that aims to extract economic narratives from a corpus of journalistic texts. We identify the definitions of narratives in the literature to be

too broad for our understanding, as most of the definitions subsume common phrases or simple factual descriptions under the umbrella term "narrative". Instead, our definition puts the focus on statements that establish a sense-making connection between events and named entities that is based on a personal assessment, a world view or an opinion. Our pipeline uses the RELATIO-method as a groundwork and adapts it to detect complete CENs rather than smaller statements or narrative building blocks. We add Coreference Resolution in combination with Named Entity Recognition to our pipeline to improve the pipelines consistency and to make it interpretable out of context. By filtering the text for causal indicators, we remove unwanted noise and are able to detect complete narratives more easily by detecting causal links between pairs of RELATIO statements.

We evaluate the performance of our pipeline on a data set generated from articles of the Financial Times that is filtered for the word "inflation" and causal indicators. By equipping the RELATIO-method with our pre- and postprocessing we are able to remove unwanted and large statements and to include new statements that were found using Coreference Resolution. Consequently, our pipeline is able to detect meaningful economic narratives, for which we provide two examples.

The automatic detection of complete narratives that display a sense-making, causal connection between events still requires further research. This is mainly because of the inherent complexity of language and the often central role of subtext in establishing meaning in texts. However, we show that our adjustments improve the detection and consistency of the RELATIO-method for our purpose and provide some examples of narratives that we were able to manually extract using the causality detection algorithm.

## References

- Andre, P., Haaland, I., Roth, C., and Wohlfart, J. (2022). Narratives about the macroeconomy. CEBI Working Paper Series, (18/21).
- Armantier, O., Bruine de Bruin, W., Topa, G., Van Der Klaauw, W., and Zafar, B. (2015). Inflation expectations and behavior: Do survey respondents act on their beliefs? International Economic Review, 56(2):505–536.
- Armantier, O., Nelson, S., Topa, G., Van der Klaauw, W., and Zafar, B. (2016). The price is right: Updating inflation expectations in a randomized price information experiment. Review of Economics and Statistics, 98(3):503–523.
- Ash, E., Gauthier, G., and Widmer, P. (2021). Text semantics capture political and economic narratives. arXiv preprint.
- Baker, S. R., Bloom, N., and Davis, S. J. (2016). Measuring economic policy uncertainty. The quarterly journal of economics, 131(4):1593–1636.
- Bénabou, R., Falk, A., and Tirole, J. (2018). Narratives, imperatives, and moral reasoning. NBER Working Paper, (24798).
- Benner, N., Lange, K.-R., and Jentsch, C. (2022). Named entity narratives. Ruhr Economic Papers, 962.
- Berger, H., Ehrmann, M., and Fratzscher, M. (2011). Monetary policy in the media. Journal of Money, Credit and Banking, 43(4):689–709.
- Binder, C. and Rodrigue, A. (2018). Household informedness and long-run inflation expectations: Experimental evidence. Southern Economic Journal, 85(2):580–598.
- Blei, D. M., Ng, A. Y., and Jordan, M. I. (2003). Latent dirichlet allocation. J. Mach. Learn. Res., 3:993–1022.
- Blinder, A. S., Ehrmann, M., Fratzscher, M., De Haan, J., and Jansen, D.-J. (2008). Central bank communication and monetary policy: A survey of theory and evidence. Journal of economic literature, 46(4):910–45.
- Blinder, A. S. and Krueger, A. B. (2004). What does the public know about

economic policy, and how does it know it? Princeton University's Industrial Relations Section Working Papers, 875.

Bruner, J. (1991). The narrative construction of reality. Critical Inquiry, 18(1):1–21.

Coibion, O., Gorodnichenko, Y., and Kamdar, R. (2018a). The formation of expectations, inflation, and the phillips curve. Journal of Economic Literature, 56(4):1447–91.

Coibion, O., Gorodnichenko, Y., and Kumar, S. (2018b). How do firms form their expectations? new survey evidence. American Economic Review, 108(9):2671–2713.

Conrad, C. and Lamla, M. J. (2007). An den lippen der ezb–der kof monetary policy communicator. KOF Analysen, 2007(4):33–45.

De Vreese, C. H. (2005). News framing: Theory and typology. Information design journal+ document design, 13(1):51–62.

Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers), pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.

Entman, R. M. (1993). Framing: Towards clarification of a fractured paradigm. McQuail's reader in mass communication theory, 390:397.

Eusepi, S. and Preston, B. (2010). Central bank communication and expectations stabilization. American Economic Journal: Macroeconomics, 2(3):235–71.

Forster, E. (1927). Aspects of the Novel. Penguin Books Ltd.

Gardner, M., Grus, J., Neumann, M., Tafjord, O., Dasigi, P., Liu, N. F., Peters, M., Schmitz, M., and Zettlemoyer, L. (2018). AllenNLP: A deep semantic natural language processing platform. In Proceedings of Workshop for NLP Open Source Software (NLP-OSS), pages 1–6, Melbourne, Australia. Association for Computational Linguistics.

- Haldane, A. and McMahon, M. (2018). Central bank communications and the general public. In AEA papers and proceedings, volume 108, pages 578–83.
- Honnibal, M. and Montani, I. (2017). spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing.
- Jurafsky, D. and Martin, J. H. (2020). Speech and Language Processing. Third Edition.
- Kalamara, E., Turrell, A., Redl, C., Kapetanios, G., and Kapadia, S. (2020). Making text count: economic forecasting using newspaper text. Journal of Applied Econometrics.
- Khetan, V., Ramnani, R. R., Anand, M., Sengupta, S., and Fano, A. E. (2020). Causal-bert : Language models for causality detection between events expressed in text. arXiv.
- Lamla, M. J. and Lein, S. M. (2014). The role of media for consumers' inflation expectation formation. Journal of Economic Behavior & Organization, 106:62–77.
- Larsen, V. and Thorsrud, L. A. (2019). Business cycle narratives. CESifo Working Paper.
- Larsen, V. H., Thorsrud, L. A., and Zhulanova, J. (2021). News-driven inflation expectations and information rigidities. Journal of Monetary Economics, 117:507–520.
- Lischka, J. A. (2015). What follows what? relations between economic indicators, economic expectations of the public, and news on the general economy and unemployment in germany, 2002-2011. Journalism & Mass Communication Quarterly, 92(2):374–398.
- Lloyd, S. P. (1982). Least squares quantization in pcm. IEEE Transactions on Information Theory, 28:129–137.
- Matthes, J. (2014). Framing. Nomos Verlagsgesellschaft mbH & Co. KG.

- McCombs, M. E. and Shaw, D. L. (1972). The agenda-setting function of mass media. Public opinion quarterly, 36(2):176–187.
- Meer, D. (2022). Überlegungen zum begriff des narrativs – kommunikation von nachhaltigkeit am beispiel des european green deals. Technical report.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G., and Dean, J. (2013). Distributed representations of words and phrases and their compositionality. In Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 2, NIPS’13, page 3111–3119, Red Hook, NY, USA. Curran Associates Inc.
- Müller, H., Schmidt, T., Rieger, J., Hufnagel, L. M., and Hornig, N. (2022). A german inflation narrative. how the media frame price dynamics: Results from a rollinglda analysis. DoCMA Working Paper.
- Müller, H., von Nordheim, G., Boczek, K., Koppers, L., and Rahnenführer, J. (2018). Der wert der worte–wie digitale methoden helfen, kommunikations- und wirtschaftswissenschaft zu verknüpfen. Publizistik, 63(4):557–582.
- Nyman, R., Kapadia, S., and Tuckett, D. (2021). News and narratives in financial systems: exploiting big data for systemic risk assessment. Journal of Economic Dynamics and Control, 127:104119.
- Pavlo, D., Piccardi, T., and West, R. (2018). Quootstrap: Scalable unsupervised extraction of quotation-speaker pairs from large news corpora via bootstrapping. Proceedings of the International AAAI Conference on Web and Social Media, 12(1).
- Rieger, J., Lange, K.-R., Flossdorf, J., and Jentsch, C. (2022). Dynamic change detection in topics based on rolling ldas. In Proceedings of the Text2Story’22 Workshop, pages 5–13.
- Roos, M. and Reccius, M. (2021). Narratives in economics. Ruhr Economic Papers, 922.
- Scheufele, D. A. (1999). Framing as a theory of media effects. Journal of communication, 49(1):103–122.
- Scheufele, D. A. and Tewksbury, D. (2007). Framing, agenda setting, and prim-



ing: The evolution of three media effects models. Journal of communication, 57(1):9–20.

Schnabel, I. (2021). New narratives on monetary policy—the spectre of inflation. speech at the 148th.

Shiller, R. J. (2017). Narrative economics. American economic review, 107(4):967–1004.

Ter Ellen, S., Larsen, V. H., and Thorsrud, L. A. (2021). Narrative monetary policy surprises and the media. Journal of Money, Credit and Banking.

Tuckett, D. and Nikolic, M. (2017). The role of conviction and narrative in decision-making under radical uncertainty. Theory & Psychology, 27(4):501–523.

Yahoo! Webscope (2008). Yahoo! answers comprehensive questions and answers.

## Appendix

In Table 1 we present the causal indicators that we used to filter for causal texts and determine the direction of a causal effect between two statements. This causal direction can be forward (event A caused event B), backward (event B caused event A) or simply associative (events A and B are associated, but not necessarily the cause of one another).

causal indicator	because	thus	therefore	after
direction	forward	forward	forward	backward
causal indicator	if	hence	consequently	resulting
direction	backward	forward	forward	forward
causal indicator	reason	due	then	since
direction	associative	backward	forward	backward

Table 1: List of causal indicators and the most likely direction of their causal effect. The direction might be forward, backward or associative.




Figure 7: Top 75 statements (48 distinct agents and patients) of the original RELATIO on filtered texts, not including dimension reduction.