

From Variance to Invariance: Qualitative Content Analysis for Narrative Graph Annotation

Junbo Huang^{**}, Max Weinig^{†*}, Ulrich Fritsche[†], Ricardo Usbeck[‡]

* Equal contribution.

^{*}Department of Computer Science, University of Hamburg

[†]Department of Socioeconomics, University of Hamburg

[‡]AI and Explanability Group, Leuphana University of Lüneburg

{junbo.huang, max.weinig, ulrich.fritsche}@uni-hamburg.de

ricardo.usbeck@leuphana.de

Abstract

Narratives in news discourse play a critical role in shaping public understanding of economic events, such as inflation. Annotating and evaluating these narratives in a structured manner remains a key challenge for Natural Language Processing (NLP). In this work, we introduce a narrative graph annotation framework that integrates principles from qualitative content analysis (QCA) to prioritize annotation quality by reducing annotation errors. We present a dataset of inflation narratives annotated as directed acyclic graphs (DAGs), where nodes represent events and edges encode causal relations. To evaluate annotation quality, we employed a 6×3 factorial experimental design to examine the effects of narrative representation (six levels) and distance metric type (three levels) on inter-annotator agreement (Krippendorff's α), capturing the presence of human label variation (HLV) in narrative interpretations. Our analysis shows that (1) lenient metrics (overlap-based distance) overestimate reliability, and (2) locally-constrained representations (e.g., one-hop neighbors) reduce annotation variability. Our annotation and implementation of graph-based Krippendorff's α are open-sourced. The annotation framework and evaluation results provide practical guidance for NLP research on graph-based narrative annotation under HLV.

Keywords: Narrative Graph Annotation, Qualitative Content Analysis, Human Label Variation

1. Introduction

In recent years, there has been growing interest in the study of narratives in economics (Shiller, 2017; Roos and Reccius, 2023) and social sciences (Polletta et al., 2011; Beckert and Bronk, 2018). Narratives are of central importance in these fields because they “[...] give meaning to life’s activities, to sense truth, and to create the commitment to act” (Tuckett, 2011, xvii), and they may trigger shifts or turning points that drive economic and societal dynamics (Roos and Reccius, 2023). Within the literature they are commonly defined as a “sequence of causally linked events” (Akerlof and Shiller, 2010) and are formally represented as Directed Acyclic Graphs (DAGs) (Eliaz and Spiegler, 2020).

As an important application of Natural Language Processing (NLP), identifying and extracting narratives from large text corpora is a central objective in many interdisciplinary research areas. In this paper, we present an inflation narrative dataset constructed from news articles, with a strong focus on ensuring a transparent and rigorous approach to **narrative graph annotation and evaluation**.

To begin with, we identify three major challenges in narrative annotation and its evaluation from an interdisciplinary perspective. First, **annotating narratives presents additional challenges compared to standard sequence classification or labeling tasks in NLP**. Narrative annotation involves

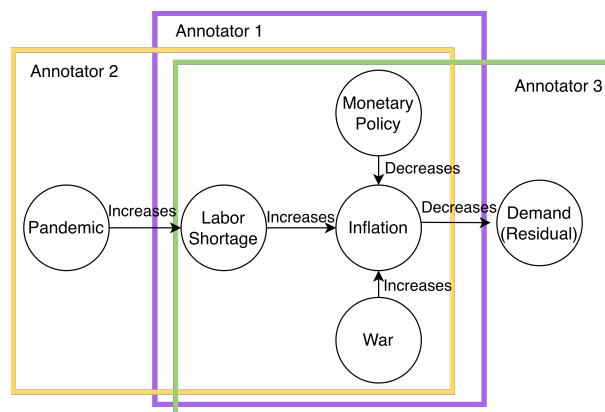


Figure 1: Example of narrative graph annotation under human label variation.

identifying and linking events across larger text segments. This requires interpretive judgments that are often subjective and context-dependent, making consistency and reproducibility more difficult to achieve with conventional annotation frameworks. Integrating qualitative methodologies such as Qualitative Content Analysis (QCA), widely used in the humanities and social sciences, can address this challenge by providing a systematic procedure for interpreting and validating complex narrative structures. Yet their adoption in NLP remains limited.

Second, as in Fig. 1, **graph-based representations introduce substantial variation across an-**

notators. Unlike flat representations of narratives, such as assigning a single label to a sequence, graph-based structures provide greater degrees of freedom in how events are connected. Annotators may differ in identifying relevant events, choosing relation types, or determining graph granularity, all of which can lead to divergent yet plausible annotations. Research on human label variation (HLV) indicates that “there exists genuine human variation in labeling due to disagreement, subjectivity in annotation or multiple plausible answers” (Plank, 2022).

Third, **there is no established consensus on how to assess inter-annotator agreement (IAA) for narrative graph annotations.** Existing graph distance metrics, such as graph edit distance, capture different structural or semantic dimensions of similarity, and their relevance often depends on the specific analytic goals or application domains. This lack of standardization complicates the interpretation of agreement scores, particularly in the presence of HLV. To address this, we compute multiple IAA measures based on diverse distance metrics and narrative representations, following the comparative evaluation approach proposed by Passonneau (2006).

To our knowledge, no previous work has jointly addressed these challenges in narrative annotation and evaluation. **Our contributions are three-fold:** (1) we propose a QCA-based methodology for narrative graph annotation; (2) we develop an graph-based annotation evaluation framework that accounts for HLV, and (3) we identify a reliable graph representation that balances contextual completeness with annotation consistency, capturing core narrative elements while maintaining robust reliability. We ask the following questions.

- RQ1:** How can QCA enhance the reliability and methodological rigor of narrative graph annotation with the presence of HLV?
- RQ2:** How does the choice of narrative representation (categorical vs. graph-based) and distance metric (lenient vs. moderate vs. strict) affect reliability score?
- RQ3:** Which graph representation best captures reliable and commonly-perceived narrative elements while maintaining contextual completeness?

2. Preliminaries

We briefly introduce narrative representation, QCA, and the computation of Krippendorff’s alpha for IAA.

2.1. Inflation Narratives as Graphs

We conceptualize news discourse narratives as structured representations of events and their causal linkages. Conversely, early work in economics and social sciences often relied on broad definitions, treating narratives as “stories that offer interpretations of economic events, morals, or hints of theories about the economy” (Shiller, 2017), and operationalized them through topics or sentiment-based features (Müller et al., 2022; Ter Ellen et al., 2022; Weinig and Fritsche, 2025). While these approaches capture general trends, they offer limited insight into the structural complexity of narratives.

More recent methods aim to structurally identify core narrative elements (e.g., agents, relations, actions, outcomes) to better represent narratives (Gehring and Grigoletto, 2023; Ash et al., 2024; Gueta et al., 2025). However, most still overlook interpretive variation, offering limited insight into the subjective nature of narrative interpretation, where multiple plausible readings of the same text coexist.

Furthermore, these approaches lack alignment with the formal definition of narratives grounded in the DAG framework (Pearl, 2009), as introduced by Eliaz and Spiegler (2020). Following this idea, they define narratives, in line with the literature studies definitions of narratives (Abbott, 2002), as a series of events that are causally linked. This formalization enables narratives to be represented as graphs, where **nodes correspond to events and edges encode directed causal relations**¹.

To this end, we grounded our narrative representation as DAGs. In the context of text-to-graph annotation, the annotated DAGs, in contrast to Heddaya et al. (2024), are not restricted in size. Given the presence of HLV in narrative annotation, **we address having multiple plausible annotations not as an issue, but a signal to reveal the commonly-perceived story within a news article.**

2.2. Qualitative Content Analysis

Qualitative methods commonly used in social sciences and humanities have the ability to capture rich and complex information, including manifest content and latent meaning. One widely used empirical method for text analysis is QCA (Selvi, 2019). Unlike other interpretative methods, QCA follows a rule-based procedure that ensures a high degree of inter-subjective verifiability, making it a natural candidate for narrative annotation under HLV. Unlike formal annotation frameworks in NLP, QCA is a context-sensitive method that derives coding categories from both theoretical considerations and the

¹We use the term node instead of vertex throughout the paper for simplicity

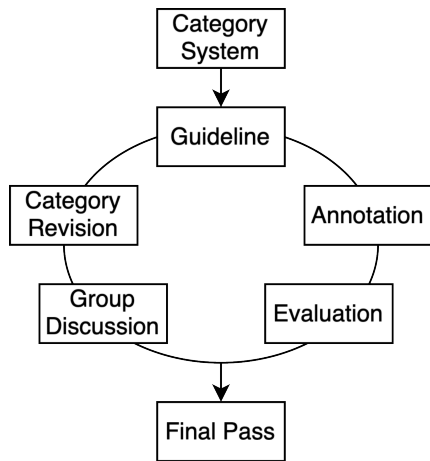


Figure 2: Flowchart of iterative QCA research procedure

textual material itself. In Fig. 2, we provide a flow diagram that explains our procedure. The central unit of analysis is the category system, which is deductively derived from existing theoretical and empirical research (Mayring, 2014). Our category system is primarily based on the work of Andre et al. (2026), who study existing backward-looking inflation narratives among households and experts in the US. They group these narratives into three supercategories (supply, demand and miscellaneous) with their corresponding subcategories.

Following Przyborski and Wohlrab-Sahr (2022)’s understanding of qualitative research as an open and circular process, we inductively extend and adjust the initial category system during the annotation process (Fig. 2). Acknowledging temporal and contextual differences in narrative structures, we iteratively develop additional narrative subcategories through regular group discussions involving annotators and the research team. Additionally, we refine our coding guidelines throughout the annotation process to reduce annotation errors and improve consistency. Finally, we arrived at 26 fine-granular subcategories (five for **demand**², 9 for **supply**³, and 12 for **miscellaneous**⁴). The final category system, including explanation of each subcategory is provided in the appendix.

To sum up, QCA is a systematic, human-centered approach to text interpretation that integrates categorical coding with iterative refinement, supporting transparent and theoretically-grounded

²Demand: Government Spending, Monetary Policy, Pent-up Demand, Demand Shift, Demand (Residuals).

³Supply: Supply Chain Issues, Transportation Costs, Labor Shortage, Wages, Energy Prices, Food Prices, Housing Costs, Supply (Residual).

⁴Miscellaneous: Pandemic, Politics, War, Inflation Expectations, Base Effect, Government Debt, Tax Increases, Exchange Rates, Medical Costs, Education Costs, Climate Crisis, Price-Gouging.

annotation of complex, context-dependent phenomena such as narratives. **In our work, we applied QCA during a pilot phase to refine both the category system and annotation guidelines. In the final annotation stage, QCA-based group discussions facilitated a shared understanding of the iteratively-refined category system among annotators.**

2.3. Krippendorff’s Alpha for Graphs

To measure IAA rigorously, we opt for the Krippendorff’s alpha (α). It is a reliability coefficient which accounts for chance agreement and can handle missing data (Krippendorff, 2011). In contrast to Cohen’s kappa, it can handle any number of annotators and does not assume equal distribution across labels (Parfenova et al., 2025).

The general form for α is:

$$\alpha = 1 - \frac{D_o}{D_e} \quad (1)$$

, where D_o is the observed disagreement among annotators per annotation item (i.e, per document) and D_e denotes the expected total disagreement within- and between-item. To capture disagreement denoted by D_o and D_e , well-defined graph distance metrics are required. For the detailed computation of D_o and D_e , we refer readers to Krippendorff (2011).

The value of α ranges from -1 to 1, where $\alpha < 0$ implies a systematic disagreement, likely caused by annotators not instructed well for the annotation task. $\alpha = 0$ suggests no agreement beyond chance, and $\alpha > 0$ the observed agreement is beyond chance-level (e.g., annotators understand the annotation task well and have some degree of agreement). It is worth mentioning that values assigned to items can be categorical, ordinal, interval, or ratio, and any distance metric can be applied. However, to our best knowledge, no publicly available implementation of Krippendorff’s α exists for graph annotations. To this end, we modified the current implementation of α and extended its computation for graph features, which includes distance metrics for nodes, edges and graphs⁵. Formally, we represent narrative graphs as a set of subject-predicate-object triples, e.g., $\{\{sub_1, pred_1, obj_1\}, \{sub_2, pred_2, obj_2\}\}$. The following distance metrics are considered.

2.3.1. Lenient Distance Metric

Node/Edge Overlap Distance As shown in Eq. 2, given two sets of nodes or edges (a set of strings),

⁵We open-source the implementation of Krippendorff’s α for graph under <https://pypi.org/project/krippendorff-graph/>

n_1 and n_2 , node/edge overlap distance measures if these two sets overlap. The measurement of set overlap is a lenient measurement of distance.

$$d_{lenient}(n_1, n_2) = \begin{cases} 0, & \text{if } |n_1 \cap n_2| > 0 \\ 1, & \text{otherwise} \end{cases} \quad (2)$$

where $|\cdot|$ denotes the cardinality of a set.

Graph Overlap Distance Similar to node/edge overlap distance, graph overlap distance (Eq. 3) measures if two graphs (a set of triples), g_1 and g_2 , overlap.

$$d_{lenient}(g_1, g_2) = \begin{cases} 0, & \text{if } |g_1 \cap g_2| > 0 \\ 1, & \text{otherwise} \end{cases} \quad (3)$$

2.3.2. Moderate Distance Metric

Node/Edge Jaccard Distance To capture partial similarity, we adopt the Jaccard distance, which measures the proportion of shared nodes or edges between two sets (Eq. 4). It provides a graded measure of agreement, reducing overestimation from simple overlap.

$$d_{moderate}(n_1, n_2) = 1 - \frac{|n_1 \cap n_2|}{|n_1 \cup n_2|} \quad (4)$$

The Jaccard distance ranges from 0 (identical) to 1 (no overlap).

Graph Edit Distance Eq. 5 defines this standard similarity measure between two graphs.

$$d_{strict}(g_1, g_2) = \min_{(e_1, \dots, e_k) \in \mathcal{P}(g_1, g_2)} \sum_{i=1}^k c(e_i) \quad (5)$$

$\mathcal{P}(g_1, g_2)$ represents the set of edit operations required to transform graph g_1 into g_2 . Each edit operation e_i can be a node or edge insertion, deletion, or substitution. The cost function $c(e_i)$ assigns a constant cost of 1 to each edit, and the total distance corresponds to the minimal number of edits needed to make the two graphs identical.

Because annotated graphs can vary substantially in size, normalization is necessary to ensure scale invariance. The raw edit distance between g_1 and g_2 is therefore normalized by the sum of their respective distances to an empty graph (g_0), yielding a bounded score between 0 and 1 (Equation 6).

$$\hat{d}_{strict}(g_1, g_2) = \frac{d_{strict}(g_1, g_2)}{d_{strict}(g_1, g_0) + d_{strict}(g_2, g_0)} \quad (6)$$

2.3.3. Strict Distance Metric

Node/Edge Exact Match Eq. 7 defines the exact match between two sets of nodes or edges. $d_{strict}(n_1, n_2)$ assigns a distance of 0 only when the two sets are identical; if any element differs between them, the distance is 1.

$$d_{strict}(n_1, n_2) = \begin{cases} 0, & \text{if } n_1 = n_2 \\ 1, & \text{otherwise} \end{cases} \quad (7)$$

Graph Exact Match For the strictest comparison, we compute graph-level exact match, which requires two graphs to be structurally identical—that is, they must contain the same set of triples (subject, relation, object) with identical labels. As shown in Equation 8, any missing, extra, or mismatched node or edge results in a maximal distance of 1.

$$d_{strict}(g_1, g_2) = \begin{cases} 0, & \text{if } g_1 = g_2 \\ 1, & \text{otherwise} \end{cases} \quad (8)$$

This metric captures complete structural equivalence between two narrative graphs and penalizes even minor discrepancies, providing a conservative lower bound on inter-annotator reliability. While it highlights exact agreement, it may underestimate consistency in cases where annotators capture the same underlying causal structure using slightly different formulations or graph sizes.

3. Narrative Graph Annotation

We detail the dataset sampling and filtering criteria, and outline the narrative annotation task as well as the annotator recruitment process. Furthermore, we explain how the annotation guidelines and category system were iteratively refined through a QCA-based pilot study to minimize annotation errors. Note that this QCA-based annotation methodology is best suited when annotation quality is the top priority, especially for annotation tasks where human label variation (i.e., disagreement, subjectivity and multiple plausible answers) is of interest. Additionally, the discussion around reproducibility should be focused on the annotation methodology as well as evaluation methodology, rather than the dataset itself, which is unfortunately not open-sourced due to licensing restrictions.

3.1. Data Source

We use the English news corpus from the Dow Jones Newswires Machine Text Feed and Archive (DJN)⁶ database. The dataset includes content

⁶Unfortunately DJN is a licensed corpus. Therefore we can not open-source the original text. We nevertheless open-sourced the annotation and anal-

from the following areas: market-moving M&A, exclusives, and earnings news; full-text feeds from Dow Jones sources (Newsires, Barron's, MarketWatch); global company news; central bank, macroeconomic, political, FX, commodities, and energy news; and third-party press release wires (BusinessWire, PR Newswire, Globe Newswire, among others). Notably, the corpus includes content from one of the largest US newspapers, *The Wall Street Journal*.

3.2. Tasks Description

Our annotation process is two-fold. We begin with a document-level classification task. This is followed by an extraction task, which involves highlighting event spans and relations.

3.2.1. Task 1: Narrative Identification

To determine whether a document addresses the cause(s) of inflation, Task 1 is a classification task with the following three classes:

1. Inflation-cause-dominant: The document discusses the cause(s) of inflation.
2. Inflation-related: The document mentions inflation but not inflation cause(s).
3. Non-inflation-related: The document neither mentions causes of inflation nor inflation.

3.2.2. Task 2: Narrative Extraction

The second annotation task consists of two steps. First, annotators mark relevant event spans using the derived category system. Second, annotators identify the causal relation between the marked events (i.e., *Increases* or *Decreases*). Although annotation is performed at the span level, the resulting representation is aggregated at the document level, abstracting from the exact span boundaries. Appendix 8.3 provides an example with annotations.

3.3. Annotator Recruitment

Given the complexity of the task and the importance of domain-specific knowledge, we recruited students with academic training in economics to carry out the data annotation. In total, seven students participate in the study, three in the pilot phase (one cisgender female and two cisgender males⁷) and four in the final annotation phase (two cisgender females and two cisgender males). All annotators have studied economics, at least during their bachelor's degrees, and five out of the seven are enrolled

yses: <https://github.com/semantic-systems/LREC2026-From-Variance-to-Invariance>.

⁷All genders are self-identified.

in graduate programs with a substantial economics component at the time of the study. Four of them have experience with QCA through their academic coursework. All annotators were between 24 and 30 years old, and two out of seven reported having a migration background. The hourly wage is €14.37 for the pilot study, and €18.30 for the final study. Further details are provided in Appendix 8.1.

3.4. Pilot Study

We began with a simple assumption: **fewer annotation errors should lead to higher IAA**, even accounting for inevitable differences in interpretation. To refine our setup and minimize errors, we conducted a pilot study⁸ within a QCA framework.

We identified three main sources of annotation errors: (i) **under-specified guidelines**, which left annotators uncertain about decision boundaries; (ii) **misunderstandings of the category system**, where annotators misinterpreted correct definitions or failed to apply them as intended; and (iii) **ambiguities in class definitions within the category system itself**, where the wording allowed multiple reasonable interpretations even for careful readers. Rather than explicitly measuring annotation errors or separating them from inherent human variation, we focused on iterative improvement of the annotation setup. After each annotation round, annotators participated in group discussions to resolve misunderstandings, clarified ambiguous definitions, and aligned on task interpretations, guideline usage, and the category system. The pilot study targeted four design questions:

- **Q1 (Sampling strategy):** Does the initial *temporally equal-weighted stratified sampling* approach⁹ yield enough articles primarily focused on *causes* of inflation?
- **Q2 (Annotation workflow):** How does performing Task 2 immediately after Task 1 for each document compare to completing Task 1 for all documents before selecting a subset for Task 2?
- **Q3 (Category System):** Which categories in the category system are ambiguous or difficult to apply consistently?
- **Q4 (Span annotation):** How does document length affect annotation quality in identifying relevant spans, and can pre-annotation (e.g., automated highlighting of candidate spans) improve accuracy and reduce annotator fatigue?

We have the following findings:

⁸Conducted from August to September 2024.

⁹Samples are drawn by dividing the corpus into time-based strata and randomly selecting documents within each stratum.

Q1: Annotators reported that temporally equal-weighted stratified samples produced mostly non-inflation-related documents. This bias likely reflects the tendency for inflation-related coverage to surge during price hikes. To address this, we revised the strategy to sample only from *inflation-peak years*¹⁰ between 1990 and 2023 (identified using Consumer Price Index data) (Figure 4). More descriptive statistics of the sampled documents are provided in the Appendix 8.2. While this increased the volume of inflation-related documents, many were still not cause-dominant. We then applied LLaMA-3.1¹¹ for zero-shot classification as a filtering step, which raised the average proportion of documents annotated with *inflation-cause-dominant* from ~10% to ~40% across annotators.

Q2: Increasing the share of cause-dominant documents did not guarantee sufficient annotation overlap for Task 2 annotations. Under the original *sequential labeling* design, annotators can proceed to Task 2 if they label a document from Task 1 as cause-dominant. With a small annotator pool, this leads to sparse Task 2 coverage: only 31.3% of 99 pilot documents have three annotations. To resolve this, we adopted a *batch labeling* procedure: first, all documents are fully annotated for Task 1. Then the most frequently cause-dominant documents are selected for Task 2.

Q3: Annotators repeatedly reported ambiguity in distinguishing between the *Pent-up Demand* and *Demand (Residual)* categories, as well as between *Government Spending* and *Government Debt*. These ambiguities informed targeted discussion and revisions to the coding guideline. Moreover, annotators identified cases where a finer-grained category system was required, as many such events frequently appeared in the corpus. Specifically, we introduced additional event categories *Transportation Costs*, *Wages*, *Food Prices*, *Housing Costs*, *Trade Balance*, *Exchange Rates*, *Medical Costs*, *Education Costs*, and *Climate Crisis*, which are novel relative to the original category system by Andre et al. (2026). We also renamed the labels *Energy Crisis* and *War in Ukraine* to *Energy Prices* and *War*, enhancing temporal generalizability. Finally, we introduced the *Wages* category to capture a broader range of labor market dynamics beyond *Labor Shortage*, while retaining the latter to support a more fine-grained classification.

Q4: Annotators reported that annotating long documents (some documents are above 1000 words)

¹⁰Inflation-peak years include 1990, 1991, 1999, 2001, 2007, 2008, 2021, 2022 and 2023

¹¹<https://huggingface.co/meta-llama/Llama-3.1-8B>

made it harder to identify precise spans, with a tendency to either overlook relevant segments or select overly broad spans. On average, each document required approximately 15 minutes to annotate in Task 2, which is longer than the initially-estimated 10 minutes, further motivating the need to reduce cognitive load during annotation. To address this, we introduced a pre-annotation step using the Gliner model¹² (Zaratiana et al., 2024). Gliner is a BERT-based, zero-shot, similarity-driven entity extraction model trained on large news corpora. At inference time, it computes distance metrics between entity-type embeddings and mention embeddings, enabling entity extraction for arbitrary sets of entity types without task-specific fine-tuning. Annotators were free to ignore or modify the automatically highlighted spans, which served purely as navigational aids rather than prescriptive labels. In the final annotated graphs, 41.08% of nodes were manually added by annotators, reflecting both the model's bias toward named entities and its imperfect coverage. Annotators were explicitly encouraged to challenge Gliner's suggestions and to mark additional spans when relevant content was missed. According to annotator feedback, this pre-annotation strategy effectively reduced the cognitive effort while preserving independent interpretive judgment.

4. Annotation Evaluation

For Task 1, we present a standard annotation evaluation practice with a single IAA score (Krippendorff's α). For Task 2, we employed a 6×3 factorial experimental design to evaluate the effects of two independent variables on Krippendorff's α : (1) narrative representation and (2) distance metric type. This design yielded 18 distinct reliability scores, enabling a systematic comparison of how both factors influence annotation reliability.

4.1. Task 1: Narrative Identification

Turning to Task 1, which involves classification across three labels, we note that no missing or invalid values were recorded by the four annotators. Krippendorff's alpha ($\alpha = 0.668$) is computed with a nominal distance metric, where distance between two annotations is 0 if and only if they are identical. The acquired α indicates a moderate agreement.

Instead of developing an optimal model for identifying narratives mentioning causes of inflation, the goal of Task 1 is to select a subset of the annotated sample ($N = 488$) for which annotators consistently agree that inflation causes are explicitly reported by the media. Among the 488 documents, there

¹²https://huggingface.co/EmergentMethods/gliner_large_news-v2.1

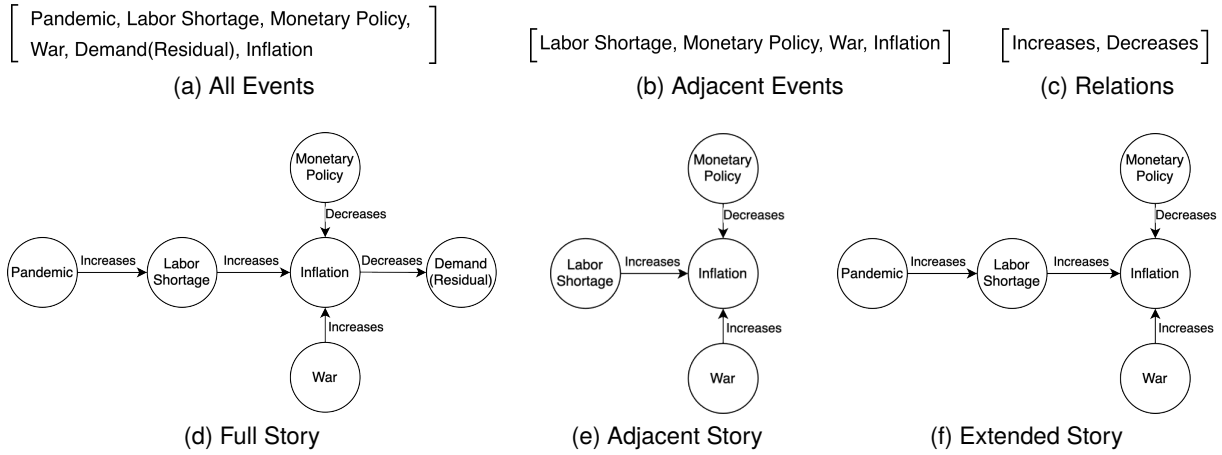


Figure 3: Visual example of six types of narrative representations. (a-c) represent the categorical representation of narratives, where (a) includes all annotated events causing and affected by *Inflation*, (b) includes only events directly causing *Inflation* and (c) includes a set of relations (*Increases* or *Decreases*) from (e). (d-f) represent the graph-based representation of narratives, where (d) includes all annotated events and their relations *from* and *to* *Inflation*, (e) includes only annotated events directly causing *Inflation*, as well as their relations, and (f) includes only annotated events directly and indirectly causing *Inflation*, as well as their relations.

are 227 instances where all four annotators assigned the same label from the label set: 33 for *inflation-cause-dominant*, 37 for *inflation-related*, and 157 for *non-inflation-related*. 261 documents depict disagreement. Of these 261 disagreed documents, 202 instances have a unique dominant label via majority voting (71 for *inflation-cause-dominant*, 84 for *inflation-related*, and 47 for *non-inflation-related*). In cases without a clear majority, it is assigned to *inflation-related*. We concatenated these 33 documents labeled as *inflation-cause-dominant* by all four annotators and 71 majority-voted *inflation-cause-dominant* documents. These 104 documents are the basis for Task 2.

4.2. Task 2: Narrative Extraction

As Task 2 is formulated as a graph-based annotation task, we employed a 6×3 factorial experimental design to examine the effects of narrative representation (six levels) and distance metric type (three levels) on narrative annotation reliability. Since only a single set of IAA measurements was available, we report only descriptive statistics, instead of providing significance tests¹³.

The first factor, narrative representation (see Fig. 3), spans six variants capturing different levels of abstraction and structural complexity. Three are categorical representations: (1) **All Events**, a set of node labels in a multi-hop subgraph reachable *from* and *to* the target node *Inflation*, includ-

ing all events identified by annotators¹⁴; (2) **Adjacent Events**, a set of node labels in a one-hop subgraph reachable *to* *Inflation*, limited to annotated events causing a change in *Inflation*¹⁵; and (3) **Relations**, a set of edge labels in Adjacent Story (described below), representing annotators' perceived direction of change in *Inflation*¹⁶. The remaining three are graph-based representations: (4) **Full Story**, a multi-hop subgraph encompassing all annotated events and relations in triple forms $\{\{e_i, r_{ij}, e_j\}, \{e_j, r_{jk}, e_k\}\}$; (5) **Adjacent Story**, a one-hop subgraph of *Full Story*, restricted to events and their relations directly causing a change in *Inflation*; and (6) **Extended Story**, a multi-hop subgraph extended from *Adjacent Story*, including events and their relations indirectly causing a change in *Inflation*.

The second factor, distance metric type, consists of three types: (1) **Lenient metric**, based on partial overlaps between annotations (Section 2.3.1); (2) **Moderate metric**, which captures graded similarity by considering the proportion of shared elements between annotations (Section 2.3.2), and (3) **Strict metric**, which requires exact matches in a set of nodes, edges or triples (Section 2.3.3).

It is important to note that IAA scores for nodes and edges are not directly comparable to those for graphs, since categorical features and graph features measure different aspects of narratives, and they require different distance metrics. On the other hand, a comparison of IAA scores between cate-

¹³Typically, a 6×3 factorial design requires at least 15 samples per group to enable ANOVA.

¹⁴A set of node labels in *Full Story*.

¹⁵A set of node labels in *Adjacent Story*.

¹⁶A set of relation labels in *Adjacent Story*.

| Narrative Representation | | Distance Metric Type | | | Δ_{lm} | Δ_{ms} | Δ_{ls} | $\mu(\cdot)$ | $std(\cdot)$ |
|--------------------------|-------------|----------------------|--------------|--------------|---------------|---------------|---------------|------------------|------------------|
| | | Lenient | Moderate | Strict | | | | | |
| Category | All Events | 0.868 | 0.481 | 0.244 | 0.387 | 0.237 | 0.624 | 2.764 | 1.809 |
| | Adj. Events | 0.717 | 0.424 | 0.285 | 0.293 | 0.139 | 0.432 | 2.112 | 1.452 |
| | Relations | 0.642 | 0.513 | 0.483 | 0.129 | 0.030 | 0.159 | 1.220 | 0.415 |
| Graph | Full Story | 0.678 | 0.320 | 0.127 | 0.358 | 0.193 | 0.551 | 2.783 | 2.070 |
| | Adj. Story | 0.702 | 0.441 | 0.202 | 0.261 | 0.239 | 0.500 | 2.143 | 1.567 |
| | Ext. Story | 0.702 | 0.356 | 0.184 | 0.346 | 0.172 | 0.518 | 2.573 | 1.886 |

Table 1: Effects of narrative representation and distance metric type on Krippendorff’s α for narrative extraction task. Δ denotes the difference in α between lenient, moderate and strict metrics. $| \cdot |$ indicates the average set cardinality (i.e., the number of events, relations, or triples per annotation). $\mu(\cdot)$ and $std(\cdot)$ denote mean and standard deviation operation.

gorical representations or between graph representations describes the influence of representation granularity to annotation variability.

A closer examination of Table 1 reveals three key observations regarding the effects of narrative representation and distance metric type on IAA.

First, **lenient metrics overestimate reliability in narrative annotation for both categorical and graph-based representation.** For example, the notably high α for *All Events* (0.861) with the lenient metric, followed by a sharp decline under the moderate metric (0.481, $\Delta_{ls}=0.387$) and strict metric (0.244, $\Delta_{ls}=0.624$), indicates that overlap-based lenient metrics tend to overestimate agreement for dense, high-coverage annotations, which increase incidental overlaps between annotators. As a result, lenient metric masks the difference in how annotators identify events and interpret causal relations. This decline of α is observed across all narrative representations, indicating stricter metrics reveals and penalizes more types of dissimilarity, e.g., mismatched or missing nodes and edges. Therefore, we encourage reporting multiple reliability scores across distance metrics of varying granularity in narrative graph annotation task, offering an informative overview of the dataset.

Second, **locally-constrained representations yields higher annotation consistency (or lower annotation variability).** Annotators show higher agreement and consistency when the annotation scope is limited to local structures. The *Relations* category yields the highest moderate (0.513) and strict (0.483) α , while having the smallest declines ($\Delta_{lm}=0.129$, $\Delta_{ms}=0.030$, $\Delta_{ls}=0.159$), indicating that annotators more consistently identify directional relations than events or graphs. Similarly, the *Adjacent Story* achieves the most balanced agreement among graph-based representations (lenient: 0.702, moderate: 0.441, strict: 0.202). In other words, its relatively high reliability score experiences a less-pronounced decline. Notably, higher annotation consistency is also reflected in lower standard deviation of annotation cardinality

(i.e., the number of nodes, edges or triples per annotation) for *Relations* (0.415) and *Adjacent Story* (1.567). On the contrary, full annotation coverage introduces the largest mean cardinality and cardinality spread (*All Events*: $\mu=2.764$, $std=1.809$; *Full Story*: $\mu=2.783$, $std=2.070$), which exhibits the largest lenient to strict evaluation gaps (Δ_{ls}). This highlights a trade-off between coverage of narrative graph annotation and inter-annotator consistency.

Third, **Adjacent Story provides the most reliable and consistent graph representation.** Among graph-based variants, *Adjacent Story* achieves the highest agreement across all distance metrics (lenient: 0.702, moderate: 0.441, strict: 0.202) while maintaining smaller or comparable gaps in reliability scores (e.g., $\Delta_{lm}=0.261$, $\Delta_{ms}=0.239$) compared to *Full Story* and *Extended Story*. In short, constraining graphs to immediate neighbors of the target node yields the best overall balance between contextual coverage and reliability robustness in our annotations, suggesting that *Adjacent Story* can be considered a reliable and consistent graph-based representation for evaluating narrative extraction datasets.

4.3. Where Do Annotators Disagree?

To investigate label-specific disagreement in Task 1, we conduct an analysis of variance (ANOVA) to test whether the proportion of documents with full annotator agreement differs across majority-voted labels. In other words, we examine whether annotators achieve higher consensus for certain labels than for others. As shown in Fig. 6, articles with the majority label *non-inflation-related* exhibit a substantially higher rate of full agreement ($\mu=0.770$, $std=0.422$) compared to those labeled *inflation-related* ($\mu=0.205$, $std=0.405$) and *inflation-cause-dominant* ($\mu=0.317$, $std=0.468$). The ANOVA results indicate a statistically significant difference across labels ($df=2$, $F=91.701^{***}$, $p<0.001$, $\eta^2=0.274$). Post-hoc pairwise Student’s *t*-tests further show significant differences between *inflation-*

related and non-inflation-related ($t=-13.348^{***}$, $p<0.001$), between inflation-related and inflation-cause-dominant ($t=-2.035^*$, $p<0.05$), and between non-inflation-related and inflation-cause-dominant ($t=8.290^{***}$, $p<0.001$). These results suggest that annotators have greater difficulty reaching agreement on inflation-related and inflation-cause-dominant instances than on non-inflation-related ones. For Task 2, we analyze disagreement at the triple level by identifying, for each article, the least-agreed triples—that is, those annotated least frequently across annotators. For each article, we first compute the union of all annotated triples and then determine which triples have the lowest annotation frequency. Across 104 articles (each annotated by four annotators), the five least-agreed triples are: (*Supply (residual), Increases, Inflation*), (*Monetary Policy, Increases, Inflation*), (*Energy Prices, Increases, Inflation*), (*Monetary Policy, Decreases, Inflation*), and (*Wages, Increases, Inflation*). These findings should be interpreted as preliminary. Future work should develop more systematic and fine-grained methods to quantify and characterize annotator disagreement.

5. Related work

5.1. Integrating QCA in NLP

The integration of QCA as a methodological framework for data annotation remains limited in NLP. Yu et al. (2011) provided epistemological arguments for the compatibility of computational and qualitative research methods, and the importance of iterative refinement of category system. Birhane et al. (2022) applied QCA at a meta-research level, providing an analysis of research culture and values in ML papers. Another line of work attempts to automate QCA by training models to perform the coding process on textual data (Liew et al., 2014; Grandjeit et al., 2020; Fischer and Biemann, 2024; Overney et al., 2024). While this approach leverages computational efficiency, it differs from a **human-centered** QCA workflow, where iterative interpretation and discussion guide the annotation process. To our knowledge, no prior work has applied QCA as a methodological framework for data annotation in machine learning. This may be attributed to the inherently time-intensive process of QCA, which contrasts with the fast-paced, benchmark-driven culture typical of machine learning research.

5.2. Evaluation on Narrative Graph Annotations

The most similar work to our annotation evaluation framework is in NLP concerns using suitable distance metrics that account for the complex phenomena of semantic and pragmatic annotation. *Measur-*

ing Agreement on Set-valued Items (MASI) metric extends Jaccard similarity by assigning weights to different type of edit operations, applied in coreference resolution (Passonneau, 2004) and document summarization (Passonneau, 2006). Building on this foundation, we introduce a family of distance metrics (lenient, moderate, and strict) that capture different granularities of reliability to better capture the complex phenomenon of narrative annotation. *Labeled Attachment Score (LAS)* and *Unlabeled Attachment Score (UAS)* are distance metrics used in dependency parsing (Buchholz and Marsi, 2006). While LAS provides a stricter, more fine-grained measure by requiring both the correct head and dependency label, UAS is a more lenient metric that only considers structural correctness. We draw a parallel between LAS/UAS and our proposal of providing lenient vs. moderate vs. strict metrics for narrative graph evaluation. However, unlike LAS/UAS which are used to assess parser performance, our metrics are designed to evaluate annotation quality in terms of providing multiple IAA scores under different conditions, a critical step in understanding human variation in complex tasks like narrative understanding.

6. Conclusion

In this work, we presented a narrative graph dataset on inflation narratives. We introduced a QCA-based annotation methodology to reduce annotation errors by iteratively refining our category system and annotation guideline. An important goal of this QCA-based annotation methodology is to establish a common-ground understanding of the task. To evaluate the annotation under HLV, we employed a 6×3 factorial experimental design to examine the effects of narrative representations and distance metrics on annotation reliability. Our analysis demonstrates that both the choice of narrative representation and distance metric strongly influence IAA. Specifically, while lenient, overlap-based distance metrics overestimate reliability in narrative annotation, stricter metrics revealed commonly-perceived narrative elements among annotators. This reflects the coexistence of multiple plausible readings of the same text. Furthermore, constraining annotation to local structures, as in *Relations* and *Adjacent Story*, improves consistency in reliability while retaining essential narrative elements. Multi-hop subgraph representations, while richer in context, reduce strict agreement and highlight a trade-off between contextual completeness and reliability. These findings provide practical guidance for the design and evaluation of narrative graph annotations and underscore the value of comparing multiple reliability scores to capture the complexity of narrative interpretation.

Acknowledgments

The authors acknowledge the financial support by the Hub of Computing and Data Science (HCDS) of University of Hamburg within the Cross-Disciplinary Lab programme, as well as by the Institute of Information Systems (IIS) of the Leuphana University Lüneburg.

Ethics Statement and Broader Impacts

To our knowledge, this work does not concern any substantial ethical issue. No private or sensitive data is used or displayed. Annotators were trained and informed about the research purpose. Identity of annotators is not revealed. The dataset and evaluation methodology can benefit applications in media narrative analysis, economic forecasting, and social science research, where understanding narratives is crucial. At the same time, users should be mindful that even with structured annotations, narratives reflect human interpretation and may not fully capture all perspectives or nuances in news reporting.

Limitations

6.1. Data, Annotators and Annotation Methodology

In terms of data, the dataset used in this study is both limited in size and constrained by licensing, which may hinder reproducibility and broader adoption. Its relatively small scale also reduces statistical power and limits the generalizability of findings across domains or narrative genres.

In terms of annotators, on the one hand, the recruitment criteria could be improved by including a more diverse range of sociodemographic backgrounds. On the other hand, the limited number of annotators restricts our ability to perform rigorous statistical testing. Future work should include a larger and more diverse participant pool to support more robust experimental designs and statistically sound comparisons.

By deductively introducing our initial category system, rather than developing it solely inductively from the data, we predefine interpretations of events, which may lead to a higher level of disagreement among our annotators. Additionally, establishing shared understanding among annotators required ongoing negotiation of theoretical and methodological perspectives. This process is context-dependent and time-intensive, increasing the difficulty to achieve standardization (Stamann and Janssen, 2019).

6.2. Localization of Disagreements

Although we conduct quantitative analyses to identify where disagreement occurs, our approach to localizing disagreement remains relatively coarse-grained. For Task 1, disagreement is measured at the document level through majority labels and full-agreement proportions, which does not capture more nuanced patterns such as partial agreement, systematic pairwise disagreement, or uncertainty concentrated around specific linguistic constructions. As a result, the analysis may obscure fine-grained sources of ambiguity.

For Task 2, disagreement is measured with annotation frequency of triples. While this provides an initial indication of which triples are less consistently identified, it does not distinguish between different types of disagreement (e.g., structural differences between graphs vs. semantic differences between nodes or edges).

Future work should therefore adopt more systematic methods for measuring disagreement.

6.3. Narrative Representation and Distance Metrics

In terms of narrative representation, DAGs effectively capture the semantic structure of narratives, such as events and their causal relationships, but fall short in representing narrative discourse features. These include the manner in which a story is told, e.g., aspects of tone, emotion, and cultural framing. Naturally, the discussion of what the distance metrics are measuring arises when we speak about feature representation. Current evaluations of graph annotations do not consider the semantic similarity of node labels. This means that concepts with similar meanings but different labels are treated as completely different. This could result in annotators' actual agreement being underestimated. Moreover, weighted edit-distance as in [Passonneau \(2006\)](#) can further granulate the notion of similarity, in the context of narrative graph. Future work should further explore the theory-driven measure of narrative similarity.

6.4. Representing Core Narrative Elements

A key question in narrative annotation is which representations effectively capture core narrative elements. Drawing from literary structuralism, the smallest functional units in narratives are events that cause a change of state. As [Abbott \(2002\)](#) note, narratives consist of constituent events, without which the story would fundamentally change or would not make sense. They are crucial to the plot's development, driving the narrative forward.

While our results suggest that graph-based representations, such as *Adjacent Story*, reflect some level of shared understanding among annotators, they do not provide a direct measure of core narrative elements. Addressing this limitation requires a more precise operationalization of what constitutes a “core” event within a narrative. Recent work in NLP has begun to explore this direction: [Antoniak et al. \(2024\)](#) proposed methods to extract salient elements from Reddit stories, while [Huang and Usbeck \(2024\)](#) highlighted the importance of extracting “core story” and distinguished constituent events from supplementary events, which served a decorative purpose in narratives. Future work should build on these insights to define and evaluate core narrative elements more rigorously.

7. Bibliographical References

- H. Porter Abbott. 2002. *The Cambridge introduction to narrative*. Cambridge University Press, Cambridge, UK; New York, NY, USA.
- George A. Akerlof and Robert J. Shiller. 2010. *Animal spirits: how human psychology drives the economy, and why it matters for global capitalism*. Princeton University Press, Princeton, N.J Woodstock.
- Peter Andre, Ingar Haaland, Christopher Roth, Mirko Wiederholt, and Johannes Wohlfart. 2026. [Narratives about the macroeconomy*](#). *The Review of Economic Studies*, page rdag014.
- Maria Antoniak, Joel Mire, Maarten Sap, Elliott Ash, and Andrew Piper. 2024. [Where do people tell stories online? story detection across online communities](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2024, Bangkok, Thailand, August 11-16, 2024*, pages 7104–7130. Association for Computational Linguistics.
- Elliott Ash, Germain Gauthier, and Philine Widmer. 2024. [Relatio: Text Semantics Capture Political and Economic Narratives](#). *Political Analysis*, 32(1):115–132.
- Jens Beckert and Richard Bronk, editors. 2018. *Uncertain futures: imaginaries, narratives, and calculation in the economy*, first edition edition. Oxford University Press, Oxford New York, NY.
- Abeba Birhane, Pratyusha Kalluri, Dallas Card, William Agnew, Ravit Dotan, and Michelle Bao. 2022. [The values encoded in machine learning research](#). In *FAccT '22: 2022 ACM Conference on Fairness, Accountability, and Transparency, Seoul, Republic of Korea, June 21 - 24, 2022*, pages 173–184. ACM.
- Sabine Buchholz and Erwin Marsi. 2006. [Conll-x shared task on multilingual dependency parsing](#). In *Proceedings of the Tenth Conference on Computational Natural Language Learning, CoNLL 2006, New York City, USA, June 8-9, 2006*, pages 149–164. ACL.
- Kfir Eliaz and Ran Spiegler. 2020. [A Model of Competing Narratives](#). *American Economic Review*, 110(12):3786–3816.
- Tim Fischer and Chris Biemann. 2024. [Exploring large language models for qualitative data analysis](#). In *Proceedings of the 4th International Conference on Natural Language Processing for Digital Humanities*, pages 423–437, Miami, USA. Association for Computational Linguistics.
- Kai Gehring and Matteo Grigoletto. 2023. [Analyzing Climate Change Policy Narratives with the Character-Role Narrative Framework](#). *SSRN Electronic Journal*.
- Philipp Grandeit, Carolyn Haberkern, Maximiliane Lang, Jens Albrecht, and Robert Lehmann. 2020. [Using BERT for qualitative content analysis in psychosocial online counseling](#). In *Proceedings of the Fourth Workshop on Natural Language Processing and Computational Social Science*, pages 11–23, Online. Association for Computational Linguistics.
- Almog Gueta, Amir Feder, Zorik Gekhman, Ariel Goldstein, and Roi Reichart. 2025. [Can LLMs learn macroeconomic narratives from social media?](#) In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 57–78, Albuquerque, New Mexico. Association for Computational Linguistics.
- Mourad Heddaya, Qingcheng Zeng, Alexander Zentefis, Rob Voigt, and Chenhao Tan. 2024. [Causal micro-narratives](#). In *Proceedings of the 6th Workshop on Narrative Understanding*, pages 67–84, Miami, Florida, USA. Association for Computational Linguistics.
- Junbo Huang and Ricardo Usbeck. 2024. [Narration as functions: from events to narratives](#). In *Proceedings of the 6th Workshop on Narrative Understanding*, pages 1–7, Miami, Florida, USA. Association for Computational Linguistics.
- Klaus Krippendorff. 2011. [Computing krippendorff's alpha-reliability](#).

- Jasy Suet Yan Liew, Nancy McCracken, Shichun Zhou, and Kevin Crowston. 2014. [Optimizing features in active machine learning for complex qualitative content analysis](#). In *Proceedings of the ACL 2014 Workshop on Language Technologies and Computational Social Science*, pages 44–48, Baltimore, MD, USA. Association for Computational Linguistics.
- Philipp Mayring. 2014. *Qualitative content analysis: theoretical foundation, basic procedures and software solution*. Klagenfurt.
- Henrik Müller, Tobias Schmidt, Jonas Rieger, Lena Maria Hufnagel, and Nico Hornig. 2022. [A German Inflation Narrative](#). *DoCMA Working Paper*;9. Publisher: TU Dortmund.
- Cassandra Overney, Belén Saldías, Dimitra Dimitrakopoulou, and Deb Roy. 2024. [Sensemate: An accessible and beginner-friendly human-ai platform for qualitative data analysis](#). In *Proceedings of the 29th International Conference on Intelligent User Interfaces, IUI 2024, Greenville, SC, USA, March 18-21, 2024*, pages 922–939. ACM.
- Angelina Parfenova, Andreas Marfurt, Jürgen Pfeffer, and Alexander Denzler. 2025. [Text annotation via inductive coding: Comparing human experts to LLMs in qualitative data analysis](#). In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 6471–6484, Albuquerque, New Mexico. Association for Computational Linguistics.
- Rebecca J. Passonneau. 2004. [Computing reliability for coreference annotation](#). In *Proceedings of the Fourth International Conference on Language Resources and Evaluation, LREC 2004, May 26-28, 2004, Lisbon, Portugal*. European Language Resources Association.
- Rebecca J. Passonneau. 2006. [Measuring agreement on set-valued items \(MASI\) for semantic and pragmatic annotation](#). In *Proceedings of the Fifth International Conference on Language Resources and Evaluation, LREC 2006, Genoa, Italy, May 22-28, 2006*, pages 831–836. European Language Resources Association (ELRA).
- Judea Pearl. 2009. *Causality: Models, reasoning, and inference*, 2 edition. Cambridge University Press, Cambridge.
- Barbara Plank. 2022. [The "problem" of human label variation: On ground truth in data, modeling and evaluation](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, pages 10671–10682. Association for Computational Linguistics.
- Francesca Polletta, Pang Ching Bobby Chen, Beth Gharrity Gardner, and Alice Motes. 2011. [The Sociology of Storytelling](#). *Annual Review of Sociology*, 37(1):109–130.
- Aglaja Przyborski and Wohlrab-Sahr. 2022. [Forschungsdesign für die qualitative Sozialforschung](#). In Nina Baur and Jörg Blasius, editors, *Handbuch Methoden der empirischen Sozialforschung*. Springer Fachmedien Wiesbaden, Wiesbaden.
- Michael Roos and Matthias Reccius. 2023. [Narratives in economics](#). *Journal of Economic Surveys*, page joes.12576.
- Ali Fuad Selvi. 2019. [Qualitative content analysis](#). In Jim McKinley and Heath Rose, editors, *The Routledge Handbook of Research Methods in Applied Linguistics*, 1 edition. Routledge, New York : Taylor and Francis, 2020. Series: Routledge handbooks in applied linguistics.
- Robert J. Shiller. 2017. [Narrative Economics](#). *American Economic Review*, 107(4):967–1004.
- Christoph Stamann and Markus Janssen. 2019. [Becoming able to work together as a central challenge for a successful joint practice—experiences from a qualitative content analysis interpretation group and suggestions for arranging interpretation group sessions](#). *Forum Qualitative Sozialforschung: Qualitative Social Research*, 20(3).
- Saskia Ter Ellen, Vegard H. Larsen, and Leif Anders Thorsrud. 2022. [Narrative Monetary Policy Surprises and the Media](#). *Journal of Money, Credit and Banking*, 54(5):1525–1549.
- David Tuckett. 2011. [Minding the Markets: An Emotional Finance View of Financial Instability](#). Palgrave Macmillan UK.
- Max Weing and Ulrich Fritsche. 2025. [Going viral: Inflation narratives and the macroeconomy](#). WiSo-HH Working Paper Series 86, Universität Hamburg, Fakultät für Wirtschafts- und Sozialwissenschaften, WiSo-Forschungslabor, Hamburg.
- Chong Ho Yu, Angel Jannasch-Pennell, and Samuel DiGangi. 2011. [Compatibility between text mining and qualitative research in the perspectives of grounded theory, content analysis, and reliability](#). *Qualitative Report*, 16(3):730–744.
- Urchade Zaratiana, Nadi Tomeh, Pierre Holat, and Thierry Charnois. 2024. [Gliner: Generalist model](#)

for named entity recognition using bidirectional transformer. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, NAACL 2024, Mexico City, Mexico, June 16-21, 2024, pages 5364–5376. Association for Computational Linguistics.

8. Appendices

8.1. Further Annotator Recruitment Information

The students involved in the pilot phase are student assistants already affiliated with our research group. They are compensated according to the standard student assistant payment at the time of the study: €13.52 per hour (without a bachelor's degree) and €15.21 per hour (with a bachelor's degree or higher).

For the final study, we employ four additional students under short-term service contracts. To ensure fair compensation and promote high-quality annotation, we provide payment above standard student rates. Each annotator receives a total of €912.60. Given 488 documents annotated in Task 1 (expected time: 4 minutes per document) and 104 documents in Task 2 (expected time: 10 minutes per document), the total estimated workload amounts to approximately 50 hours. This corresponds to an effective hourly wage of €18.30, exceeding both the standard student assistant wage at our institution and the country-level minimum wage by over €6 per hour.

8.2. Descriptive Statistics of the Annotated Dataset

The dataset consists of 488 documents, with an average length of 609 words (min: 112, max: 1774, median: 541). Articles are attributed with their publication year, as shown in Fig. 4. The sampling strategy focuses on inflation-peak years: 1990, 1991, 1999, 2001, 2007, 2008, 2021, 2022 and 2023 (only until January). Fig. 5 shows an increase of word counts through time.

8.3. Annotation Example

Snippet of a sample article:

ADB Exec Calls For Tighter Monetary Policies In Asia

KUALA LUMPUR (AP)—Asia's growth is being threatened by spiraling inflation due to higher food and fuel cost, an Asian Development Bank executive warned Sunday, as he called on governments to




Figure 4: Distribution of documents across publication years. The sampling strategy focuses on inflation-peak years: 1990, 1991, 1999, 2001, 2007, 2008, 2021, 2022 and 2023.

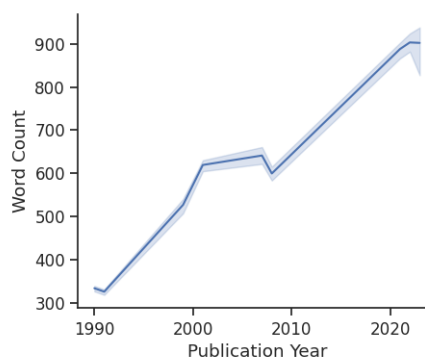


Figure 5: Average word count of documents across publication years.

tighten monetary policies to deal with the scourge.

The bank is reviewing its growth forecast of 7.6% this year for Asia, excluding Japan, amid concerns inflation will widen income inequality in the region and cause more people to plunge into poverty, said its managing director general Rajat M.Nag.

...

Annotator 1:

[('Energy Prices', 'Increases', 'Inflation'), ('Monetary Policy', 'Decreases', 'Inflation'), ('Food Prices', 'Increases', 'Inflation')]

Annotator 2:

[('Energy Prices', 'Increases', 'Inflation'), ('Monetary Policy', 'Decreases', 'Inflation'), ('Food Prices', 'Increases', 'Inflation')]

Annotator 3:

[('Energy Prices', 'Increases', 'Inflation'), ('Food Prices', 'Increases', 'Inflation')]

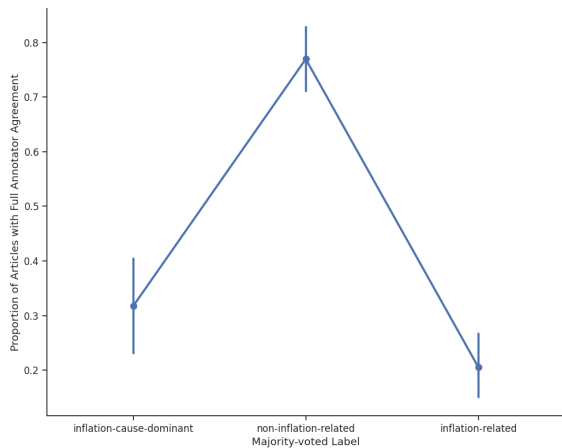


Figure 6: Proportion of Articles with Full Annotators Agreement per Label.

Annotator 4:

[('Energy Prices', 'Increases', 'Inflation'), ('Government Spending', 'Increases', 'Inflation'), ('Food Prices', 'Increases', 'Inflation')]

8.4. Annotation Guidelines

Objective: Identify inflation narratives in large text corpora using a two-stage annotation process.

Stage I: Document-level Classification

Determine whether the article addresses the cause of inflation. There are three cases:

- Inflation-cause-dominant: The text discusses the causes of inflation.
- Inflation-related: The text mentions inflation but not inflation causes.
- Non-inflation-related: The text neither mentions causes of inflation nor inflation.

What is inflation? The rate of inflation indicates how much prices in a country rise from year to year, measured in percent. It is defined as the annual growth in the general price level of goods and services. A decline in the price level is commonly referred to as deflation. For example, an inflation rate of 2% means that consumer prices rise by 2% over 12 months. A basket of typical monthly purchases that costs €1,000 at the beginning of the year costs €1,020 at the end of the year. Therefore, texts that mention solely movements in single prices (e.g. crude oil or wheat prices) are not to be labeled as “inflation-related” or “inflation-cause-dominant”. Results from this stage are reviewed to select documents for further annotation.

Stage II: Text-Span Annotation

For documents labeled as Inflation-cause-dominant, identify text spans that indicate event mentions related to inflation causes. An event mention is a specific instance in text that describes an action, occurrence, or state change. It typically includes key information like the action itself (often expressed by a verb and/or noun), involved participants, and sometimes a time or place. Each event mentioned reflects one occurrence of an event as described in the text. A selection of possible events are shown in Table 1. We distinguish between Demand, Supply, and Miscellaneous Events.

When highlighting the event spans:

- Highlight only the minimum span (single/few token if possible)
- Usually, events include verbs and nouns, and sometimes adjectives
- If you are uncertain, always make a note and discuss it in the plenary sessions

To support event detection and increase the number of matching text spans, we pre-annotated the text using a pre-trained Named Entity Recognition (NER) zero-shot language model that outperforms general-purpose models like ChatGPT. However, due to potential prediction errors, human validation is essential. Specifically, some events might not be identified correctly or may not be labeled at all.

Furthermore, we aim to identify the narrated causal structure of the narratives. Therefore, you need to establish the causal relationships between the annotated events using Label Studio. To distinguish the effect direction, we include two relational attributes: Decreases and Increases.

8.5. Category System for QCA

| Category | Subcategory | Explanation |
|----------------------|------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Demand | Government Spending | Adjustments in government spending (e.g., stimulus payments) |
| | Monetary Policy | Monetary policy by the Federal Reserve or other Central Banks |
| | Pent-up Demand | Reopening of the economy and the associated higher incomes, new spending opportunities, and optimism about the future |
| | Demand Shift | Shift of demand across sectors (particularly increases in durables) |
| | Demand (Residuals) | Increase in demand that cannot be attributed to the other channels |
| Supply | Supply Chain Issues | Disruption of global supply chains |
| | Transportation Costs | Rising shipping and freight costs, including container shortages and port congestion |
| | Labor Shortage | Shortage of workers, e.g., due to some workers dropping out of the labor force, and higher wage costs |
| | Wages | Changes in wage levels driven by labor market dynamics, such as increased labor demand or worker bargaining power |
| | Energy Prices | Higher energy prices, e.g., due to the global energy crisis, leading to shortages of oil and natural gas |
| | Food Prices | Increases in food prices, e.g., driven by rising input costs, or global agricultural disruptions |
| | Housing Costs Supply (Residual) | Rising housing and rental costs Other negative supply effects |
| Miscellaneous | Pandemic | The COVID-19 pandemic, the global pandemic recession, lockdowns, and other policy measures |
| | Politics | Policy failure, and mismanagement by policymakers, policymakers are blamed |
| | War | Armed conflict involving state or non-state actors, such as the Russian invasion of Ukraine and the associated international economic, political, and military responses |
| | Inflation Expectations | Expectations about high inflation in the coming years, making firms preemptively increase prices and workers bargain for higher wage |
| | Base Effect | Mentions that inflation is high due to a base effect, i.e., a very low inflation rate during the pandemic, leading almost mechanically to high inflation rates now |
| | Government Debt | High level of government debt |
| | Tax Increases | Tax increases, such as VAT hikes |
| | Trade Balance | Inflationary pressures linked to changes in exports, imports, or trade deficits |
| | Exchange Rates | Effects of exchange rate fluctuations on import prices and domestic inflation |
| | Medical Costs | Rising healthcare or insurance costs, including structural issues in medical pricing |
| | Education Costs | Increases in education-related expenses such as tuition fees |
| | Climate Crisis | All aspects related to the climate crisis and natural disasters, as well as related environmental and economic consequences |
| | Price-Gouging | Companies exploit opportunities to increase profits |

Table 2: Events and Explanations in the Category System.