

Narrative Persuasion: A Brief Introduction

Encyclopedia of Experimental Social Science

Kai Barron

WZB Berlin

Tilman Fries

LMU Munich

May 2024

“Narratives may be the last resort of economic theorists. But they are probably the life stuff of those whose behavior they study.”

~ Jerome Bruner (1986, p. 43)

1. Introduction

In the field of economics, there has been a recent surge in interest in how stories, interpretations, and explanations shape how individuals think about and interact with the world around them (Eliaz and Spiegler, 2020; Roos and Reccius, 2024; Schwartzstein and Sunderam, 2021; Shiller, 2019). This development represents a break with notions of information processing that traditionally focused on studying how Bayesian rational individuals learn from hard information. For example, economists traditionally root persuasion in expertise: individuals may change their beliefs in response to claims made by others, but only if they believe that this claim is grounded in some superior knowledge of the facts that they do not possess (Crawford and Sobel, 1982).

In contrast, more recent accounts explore how individuals who do not possess superior knowledge can still persuade others by providing (potentially skewed) interpretations of the available facts. They might do this by attaching coherent explanations to the claims they make, ensuring that the explanation seems plausible given the available facts (Schwartzstein and Sunderam, 2021). Or, they might provide an account of the facts that implies a hopeful future, relying on the idea that individuals might prefer to believe rose-tinted versions of reality (Eliaz and Spiegler, 2020).

This article seeks to trace the early development of this “narrative persuasion” literature, with a particular focus on the nascent body of experimental work that provides empirical insights into how, when, and why narratives persuade.

2. Conceptualization of Narratives in Economics¹

Despite the rapid growth in interest in incorporating the concept of a “narrative” into economic models, there does not yet exist a commonly accepted definition of what the term narrative means in economics.² This is relatively unsurprising given the recency of economists’ interest

¹ This section borrows elements from Appendix A of Barron and Fries (2023), which has now been superseded by an updated version of the paper.

² The following examples selected from important early contributions to this literature serve to illustrate this point. Morag and Loewenstein (2023) view a narrative as “a story [that] places selected events on a timeline and establishes causal links between them” (p. 2). Shiller (2020) argues that economic narratives are “stories that offer interpretations of economic events, or morals, [or] hints of theories about the economy [that] go viral just as diseases do” (p. 792). Similar to Shiller (2020), Bénabou et al. (2020) offer a fairly broad conceptualization of a moral narrative as “... any signal, story, or heuristic that can potentially alter an agent’s beliefs about the tradeoff between private benefits and social costs” (p. 1). Recently, two approaches to formalizing the notion of a narrative have begun to gain attention. Eliaz and Spiegler (2020) model narratives as causal structures that connect events, while Schwartzstein and Sunderam (2021) view narratives as likelihood functions, emphasizing the importance of the narrative fitting the facts well. Below, we discuss these two formal approaches. For an

in the concept and the multifaceted nature of narratives, which are applied to a wide array of scenarios in everyday usage. Given this multiplicity of meanings of the concept in everyday usage, it seems natural to explore the usefulness of different conceptualizations for understanding (economic) behavior—even when the concept distillation process has reached maturity, several formalizations may survive and prove useful in parallel. Therefore, we consider it constructive to contribute to the process of organizing this “conceptual landscape”. This taxonomical exercise will help to foster clear-sighted thinking about narratives in economics, providing a clearer perspective on how these overlapping ideas relate to one another. This article will take an initial step in this direction by providing a short discussion of two questions that are central to this exercise.

2.1. What do Different Conceptualizations of Narratives Share?

One core shared feature that is present across most of the working definitions of the term “narrative” used in the literature to date is that a narrative involves “sense-making”. Specifically, the concept is used to refer to providing an explanation or interpretation of a collection of events. This collection of events can take many forms.

To illustrate this, consider the following examples. First, consider a sequence of historical events, such as those leading up to World War II or those leading up to the 2007 Financial Crisis (see, e.g., Gennaioli and Shleifer 2018). Here, a narrative explaining the causes of World War II or the Financial Crisis weaves a causal path through the preceding events. Second, consider the rise of depression amongst teenagers along with the other contemporaneous changes in society in the last twenty years, such as the rise of social media usage (see, e.g., Braghieri, Levy, and Makarin, 2022). Here, a narrative might posit that the widespread diffusion of social media is causally responsible for the rise in depression amongst teenagers. Third, consider the existence of cultural differences around the world. Here, a narrative might propose that differences in weather patterns provide an explanation for some of the differences between Southern and Northern Europeans.

In each of these examples, it is possible to extract a collection of events or variables that can be captured in a database. The narrative provides an explanation that makes sense of this data. Each of the examples is analogous to a particular data structure—the first corresponds to a time series, the second to a panel, and the third (arguably) to a cross-section.³ Taken together, this

extensive discussion of economists’ uses of the term *narrative* and the historical context from which the recent contributions have emerged, see Roos and Reccius (2024).

³ This analogy between the narrative-constructor and the econometrician highlights some features of narratives. First, just as the econometrician must select the relevant variables for her empirical specification, the sense-making task of the narrative-constructor often involves selecting a subset of variables from a large (possibly infinite) set of variables. Second, like the econometrician who conducts statistical analyses in the presence of unobserved variables, an individual constructing a narrative is often missing information and can only work with the events they know about. Third, similar to how econometric models can be used to forecast the impact of a policy, an individual who constructs a narrative to make sense of existing data may then use it to forecast future events.

discussion suggests the following broad definition of a narrative as a starting point, which may later be refined and subdivided into different components and types of narratives:

“A causal explanation that makes sense of a collection of events.”

2.2. How do Narratives Differ from One Another?

When thinking about different ways to conceptualize narratives, it is helpful to identify key features that distinguish the various everyday uses of the term from one another. This section provides a discussion of some of these features to help disentangle different ways we may wish to conceptualize narratives. To do this, we first discuss features that may be useful for classifying different *types* of narratives. Second, we discuss some of the *constituent components* of narratives that may play an important role when comparing narratives of same broad type. These components may, for example, be important when thinking about which narratives are persuasive.

Classifying Narratives into Broad Types:

First, it is important to consider the implicit objective of the narrative. Narratives may seek to explain the causes of one particular outcome of interest (e.g., a narrative describing the causes of World War II or the 2007 Financial Crisis) or to explain an entire collection of events, without assigning primacy to any single event (e.g., a narrative of the rise and spread of major world religions throughout history). We refer to the first type as *single-outcome-of-interest* narratives, and the second as *multiple-outcomes-of-interest* narratives.

Second, narratives may be general or specific. For example, one narrative may focus on explaining the causes of a particular recession, while another may propose an explanation for common causes of recessions in general. Similarly, one narrative might propose an explanation for high inflation observed in the US in late 2021 and early 2022, while another might propose an explanation for common causes of inflation more generally (see Andre et al., 2023, for further discussion of this example). The key difference is whether the narrative seeks to explain a particular collection of events or to identify a common repeating pattern in a set of similar collections of events. We refer to these two types as *singular narratives* (single event) or a *generic narratives* (repeating events).⁴

In addition to these two type-classifications, narratives can be distinguished and categorized along several other dimensions. For instance, narratives may be classified by thematic content, distinguishing between narratives of the physical and social world, with sub-categories such as

⁴ An alternative label that one may apply here is to refer to *generic narratives* as *mental models*. While the concept of a *mental model* is also not yet precisely defined in economics, it can be broadly construed as referring to the mental representations that individuals form to understand processes in the world around them. A fairly crude way to distinguish narratives from mental models would be to think of narratives as answering “why” questions (why did X happen?), while mental models answer “how” questions (how does X work?). An example of a mental model is provided by Andre et al., (2024) who study how different actors in financial markets understand how financial markets operate.

cultural, economic, and political narratives. One can also distinguish narratives that are predominantly *quantitative*, focusing on statistical data to tell a story, from those that are more *qualitative*, focusing more on personal stories (e.g., ethnographic studies). A full discussion of these and other dimensions of narratives goes beyond the scope of this article.

Components of Narratives that May Matter for Persuasion:

In thinking about different ways to conceptualize narratives, breaking them down into their constituent components is a central exercise. In particular, it is informative to think about which components may be important for determining which narratives people find compelling. The following provides a very brief overview of some of these components, while Section 3 continues the discussion by outlining the progress made in the literature in formalizing a subset of these components and incorporating them into economic models.

First, many types of narratives can be classified into common structures. In literary studies, this may involve categorizing narratives into familiar story arcs with standard archetypal characters. This idea extends to causal narratives, where narratives can be classified according to common causal structures represented by directed acyclic graphs.

Second, given a particular structure, narratives may differ in important details. When individuals are faced with assessing different narratives and deciding which to adopt, these details may crucially influence their decision. These include:

- **Coherence:** How well does the narrative fit the individuals' existing factual knowledge?
- **Optimism:** Does the narrative promise a rose-tinted view of the world?
- **Familiarity:** Does the narrative follow a structure that the individual is familiar with?
- **Complexity:** Is the narrative simpler than the alternatives?
- **Conformity:** Does the narrative conform with the individuals pre-existing views?

In the next section, we discuss two key contributions to formalizing narratives and integrating them into economic models and summarize early empirical studies that build on these formalizations.

3. Narrative Persuasion Experiments

3.1 Narratives as Likelihood Functions

Schwartzstein and Sunderam (2021) develop a theoretical model in which narratives are treated as likelihood functions that map data into beliefs. They consider a problem where a persuader and a decision maker (DM) observe a signal (data), with the persuader providing a narrative to the DM about how to interpret the signal. The DM evaluates the persuader's

narrative by comparing it to some default narrative that the DM would use to interpret the signal in the absence of the persuader’s narrative. Schwartzstein and Sunderam assume that the DM adopts the received narrative if and only if it has a higher likelihood of having generated the signal than the default narrative. They then develop a theory of persuasion where the persuader anticipates needing to surpass the default narrative in terms of likelihood fit to successfully persuade the DM.

Building on this theoretical work, Barron and Fries (2024) conduct an experiment to investigate the mechanics of narrative persuasion. The study tests whether people do actually follow a likelihood-based adoption rule, and also asks whether human persuaders anticipate this adoption rule when constructing their narratives.

The experiment is framed as an investment game where “advisors” try to persuade “investors” about the future success prospects of a company. Both advisors and investors observe data about the company. Investors wish to accurately assess the company’s future success prospects, while advisors have their own objectives and may wish to bias the investors’ beliefs. The way that an advisor may try to influence an investor’s beliefs is by proposing a narrative that describes a possible data-generating process responsible for producing the company data.

The study documents several findings. First, it shows that people are susceptible to narrative persuasion. Investors beliefs are biased by advisors’ narratives, even when they know that the advisor is likely to hold misaligned incentives. Second, investors assess narratives based on their likelihood fit; narratives that are more consistent with the company data are more persuasive. Third, human advisors anticipate the importance of fit for how investors assess narratives. This yields predictable features in the narratives advisors construct as they try to exploit investors’ fit-based adoption rule. Finally, the study shows that the presence of a competing narrative can influence how advisors construct narratives. For example, when an advisor competes with a narrative that fits the data well, she has less flexibility in constructing her own persuasive narrative because she wants it to fit better than the alternative. This results in her being less ambitious in her attempts to use her narrative to bias the beliefs of the investor.

These results are robust to various changes in the experimental environment. Most notably, investors can still be persuaded even if they know that advisors do not have *any* more information than they do, which is in line with Schwartzstein and Sunderam's (2021) framework. This highlights a key difference from traditional approaches that have typically emphasized information asymmetries as the foundation of persuasion.

3.2 Narratives as Directed Acyclic Graphs (DAGs)

A series of theoretical papers, starting with Spiegler (2016), use the machinery of Directed

Acyclical Graphs (DAGs) to develop a theoretical framework that puts misperceived causal relations at the root of nonrational expectations. Since narratives typically impose a causal structure that connects events, in their work on competing narratives, Eliaz and Spiegler (2020) focus on studying how individuals might select between causal structures (DAGs) and the implications of this.

To provide some intuition for how DAGs can influence the inferences drawn from data, consider the example of an agent who can take an action a with the goal of increasing an output y . To figure out the marginal impact that increasing a has on y , the agent may draw on an *event database* that records realizations of various events such as a and y at different points in time. Suppose that the database records the variables a , y , and z . Now, there are different possible beliefs that the agent might have about the causal relationship between these three variables. The agent might believe in a *lever narrative* where a changes z which in turn changes y . In the language of DAGs, $a \rightarrow z \rightarrow y$. Alternatively, the agent might believe that a and z directly affect y without influencing one another, $a \rightarrow y \leftarrow z$. Such causal structures are commonly referred to as *threat narratives*. Importantly, agents who believe in the threat or lever narrative may draw different conclusions to one another about the marginal effect of a on y even if they observe the same data. This is because an essential feature of the conditional expectation of $E[y|a]$ is that its estimate depends on the causal structure imposed on the dataset.

Charles and Kendall (2024) investigate this idea empirically by using an experiment to test whether people draw different conclusions from data if they are exposed to different narratives. Participants in their experiment can access a database and are provided with a description of either a lever narrative or a threat narrative. As predicted by the DAG framework, participants are influenced by the narrative they are exposed to, biasing their actions towards the optimal action suggested by using the narrative as a lens to interpret the data. Additionally, they find that lever narratives are more influential than threat narratives; individuals are more easily persuaded by them, and they also come to mind more easily when participants are tasked with identifying the causal structure underlying a data set themselves.

Ambuehl and Thyssen (2024) examine how experimental participants choose between DAG explanations of data, aiming to analyze which features of DAGs influence their adoption. Participants observe data about the relationships between a small set of variables and are presented with different DAGs that may represent the true causal structure generating the data. Each DAG is paired with a recommended action, which is optimal if it matches the correct causal structure. The experimental design allows the researchers to identify the rules individuals use when choosing between different explanations. Among many possible rules, the experiment suggests that 40% of student participants use *data-based criteria* (as opposed

to *structure-based* or *advice-based criteria*) to rule out explanations that are inconsistent with the underlying data. This proportion is markedly lower in an additional experiment conducted with members of the general population. Other popular decision rules involve choosing the explanation that requires the lowest investment for taking the associated action or the explanation that promises the highest payoff. The latter finding is important since it is the decision-making rule that Eliaz and Spiegler (2020) assume agents use when choosing between competing narratives.

3.3 Experiments on the Persuasive Function of Arguments

One overarching theme in the narrative persuasion literature is that claims are assessed based on the accompanying justifying explanations or arguments. Therefore, we conclude this discussion of the literature by summarizing two papers that zoom in on the role of arguments or explanations in persuasion.

Hüning et al. (2022) conduct an experiment where participants are randomly assigned to group chats to discuss a rent control reform that will be voted on in a future local election. An analysis of the chat messages suggests that, as individuals are exposed to more arguments against their initial voting intention, the likelihood of their final vote not aligning with their initial intention increases. In contrast, exposure to claims that are not supported by arguments tends to entrench pre-existing voting intentions without persuading those who initially disagree. This suggests that individuals have higher standards of evaluating a claim if it contradicts their priors, but can be persuaded to change their decision regarding a relevant real-life outcome if supplied with counterarguments.

Graeber et al. (2024) conduct a social learning experiment where participants must answer a financial decision-making quiz while observing either *just choices* or *choices plus explanations* provided by other participants who took the same quiz before them. The authors find higher imitation rates in the treatment where participants receive the explanation. Participants also make better decisions when presented with an explanation. Interestingly, however, participants do not become less likely to follow an objectively wrong decision when additionally provided with an explanation. Therefore, while explanations persuade individuals to take the right decision more often, they do not curb the spread of harmful imitation.

4. Going Forward

Initial experimental studies on narrative persuasion demonstrate that theoretical frameworks have empirical relevance, indicating their potential to enhance economists' understanding of persuasion. There exist numerous avenues for future research to build on this early work and help advance this endeavor.

First, the narrative persuasion framework has already been extended in several interesting theoretical directions. For example, Aina (2024) studies how a persuader can benefit from providing a menu of narratives in anticipation of future signals with uncertain outcomes, Ichihashi and Meng (2021) study a situation where the persuader can curate a narrative as well as the data observed by the DM, Jain (2023) studies how narrative persuasion can be constrained by a strategic data gatekeeper, and Schwartzstein and Sunderam (2024) examine the implications of sharing models in a community or network. These studies extend the initial framework in naturally appealing ways and provide testable assumptions and predictions that future experimental research should build upon and test.

Second, future experiments could focus on identifying relevant dimensions of individual heterogeneity. For example, what are the determinants of being skilled at narrative persuasion, and how is this skill distributed in the general population or amongst members of professions where a proclivity for storytelling might be advantageous?

Third, the narrative persuasion literature should also stay up-to-date on developments with the literature investigating how individuals construct models from data by themselves (Frechette et al., 2024). There are likely to be several overlaps in the psychological mechanisms that matter for individual model-building and the strategic transmission of models and narratives—for example, the models that come to mind more easily when learning from data alone may also be those that individuals find more persuasive when received from others.

References

- Aina, C. (2024). Tailored Stories. *Mimeo*.
- Ambuehl, S., & Thysen, H. C. (2024). Competing Causal Interpretations: An Experimental Study. *Mimeo*.
- Andre, P., I. Haaland, C. Roth, and J. Wohlfart (2023). Narratives About the Macroeconomy. *CESifo Working Paper No. 10535*.
- Andre, P., Schirmer, P., & Wohlfart, J. (2024). Mental Models of the Stock Market. *Mimeo*.
- Barron, K., & Fries, T. (2023). Narrative Persuasion. *CESifo Working Paper No. 10206*.
- Barron, K., & Fries, T. (2024). Narrative Persuasion. *Mimeo*.
- Bénabou, R., Falk, A., & Tirole, J. (2020). Narratives, Imperatives, and Moral Persuasion. *Mimeo*.
- Braghieri, L., Levy, R., & Makarin, A. (2022). Social Media and Mental Health. *American Economic Review*, 112(11), 3660–3693. <https://doi.org/10.1257/aer.20211218>
- Bruner, J. (1986). *Actual Minds, Possible Worlds*. Harvard University Press.
- Charles, C., & Kendall, C. (2024). Causal Narratives. *Mimeo*.
- Crawford, V. P., & Sobel, J. (1982). Strategic Information Transmission. *Econometrica*, 50(6), 1431–1451. <https://doi.org/10.2307/1913390>
- Eliasz, K., & Spiegler, R. (2020). A Model of Competing Narratives. *American Economic Review*, 110(12), 3786–3816. <https://doi.org/10.1257/aer.20191099>
- Frechette, G., Yuksel, S., & Vespa, E. (2024). Extracting Models from Data Sets: An Experiment. *Mimeo*.
- Gennaioli, N., & Shleifer, A. (2018). *A Crisis of Beliefs: Investor Psychology and Financial Fragility*. Princeton University Press.
- Graeber, T., Roth, C., & Schesch, C. (2024). Explanations. *Mimeo*.
- Hüning, H., Mechtenberg, L., & Wang, S. (2022). Using Arguments to Persuade: Experimental Evidence. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.4244989>
- Ichihashi, S., & Meng, D. (2021). The Design and Interpretation of Information. *Mimeo*.
- Jain, A. (2023). Informing Agents amidst Biased Narratives. *Mimeo*.
- Morag, D., & Loewenstein, G. (2023). Narratives and Valuations. *Mimeo*.
- Roos, M., & Reccius, M. (2024). Narratives in economics. *Journal of Economic Surveys*, 38(2), 303–341. <https://doi.org/10.1111/joes.12576>
- Schwartzstein, J., & Sunderam, A. (2021). Using Models to Persuade. *American Economic Review*, 111(1), 276–323. <https://doi.org/10.1257/aer.20191074>
- Schwartzstein, J., & Sunderam, A. (2024). Sharing Models to Interpret Data. *Mimeo*.
- Shiller, R. J. (2019). *Narrative Economics*. Princeton University Press.
- Shiller, R. J. (2020). Popular economic narratives advancing the longest U.S. expansion 2009–2019. *Journal of Policy Modeling*, 42(4), 791–798. <https://doi.org/10.1016/j.jpolmod.2020.03.005>
- Spiegler, R. (2016). Bayesian Networks and Boundedly Rational Expectations. *The Quarterly Journal of Economics*, 131(3), 1243–1290. <https://doi.org/10.1093/qje/qjw011>