

News Media as Suppliers of Narratives (and Information)*

Kfir Eliaz and Ran Spiegler[†]

January 26, 2026

Abstract

We present a model in which news media shape beliefs by providing information (signals about an exogenous state) and narratives (models of what determines outcomes). To amplify consumers' engagement, the media maximize their anticipatory utility. We characterize the optimal monopolistic media strategy, highlighting the synergy between false narratives and biased information. Consumer heterogeneity gives rise to a novel menu-design problem due to an “equilibrium data externality” among consumers. The optimal menu includes false narratives, and can generate polarized beliefs and choices without any information transmission. False narratives and polarization also feature in a competitive version of our model, albeit coupled with complete information. Nevertheless, competition can make some consumers objectively worse off.

*Spiegler acknowledges financial support from Leverhulme Trust grant RPG-2023-120 and the Sapir Center. Eliaz acknowledges financial support from ISF grant 455/24. We thank Nageeb Ali, Alex Clyde, Daniel Gottlieb, Nathan Hancart, Shira Dvir-Gvirtsman, Margaret Meyer, Jesse Shapiro, Adam Szeidl, Heidi Thyssen, seminar audiences, and the editor and referees of this journal, for helpful comments.

[†]Eliaz: School of Economics, Tel-Aviv University and NHH. E-mail: kfire@tauex.tau.ac.il. Spiegler: School of Economics, Tel-Aviv University and Economics Dept., University College London. E-mail: rani@tauex.tau.ac.il.

*“We’re supposed to be tellers of tales as well as purveyors of facts.
When we don’t live up to that responsibility, we don’t get read.”*
(William Blundell)

*“The masses have never thirsted after truth. Whoever can supply
them with illusions is easily their master.”* (Gustave Le Bon)

1 Introduction

Standard models of news media regard them as suppliers of information, providing noisy signals of an underlying state of Nature. A complementary view, which is absent from standard models, is that news media are a vehicle for spreading *narratives*. While this term has multiple meanings, we conceive of narratives as *qualitative accounts of what causes outcomes of interest*.

For example, when news media report about the latest development in some international conflict, they may contextualize it within a story (supported by background statistics) that assigns credit or blame for observed outcomes to one of the parties to the conflict. Likewise, news reports about poverty or discrimination may be framed by a narrative about what determines life outcomes: Is it one’s personal choices, or rather external circumstances beyond one’s control? For instance, Iyengar (1990) studies how the media shapes popular perceptions regarding the role of personal agency and external factors in escaping poverty. Loury (2020) makes a similar distinction between “development” and “bias” narratives in media discussions of ethnic discrimination.

This paper presents a stylized model of news media (in a broad sense that includes content platforms) that is based on a fusion of the two views: News outlets provide information about exogenous states as well as a narrative. Media consumers use the narrative to interpret empirical regularities and form beliefs about the mapping from states and actions to outcomes. A false narrative is a misspecified causal model, which can therefore induce distorted beliefs.

The fusion of information- and narrative-based views enables us to offer a new model of *media bias*. There is a common intuition that this phenomenon is driven in large part by consumer demand (Gentzkow and Shapiro (2010) back this intuition with empirical evidence). Yet, the standard model of consumer behavior assumes that demand for information is purely instrumental. Expected-utility maximizers weakly prefer more informative signals. Therefore,

unless there are frictions on the supply side that prevent media from providing complete and objective information, the market will provide it. Even if consumers have heterogeneous preferences, they all want more informative news.

Studies across several disciplines (psychology, political science, media studies) have provided evidence that consumer demand for news media reflects non-instrumental attitudes to beliefs (e.g., Hart et al. (2009), Van der Meer et al. (2020), Taber and Lodge (2006)). These findings have inspired models of media bias in which beliefs enter directly into consumers’ utility function (see Prat and Strömberg (2013) and Gentzkow et al. (2015) for surveys).

Our approach to non-instrumental demand for news is based on the hypothesis that people are more likely to follow a news outlet when it helps them attain “desirable beliefs”. For example, in the context of sporting events or international conflicts, consumers want to believe that their side will win (perhaps if the right action is taken).¹ Likewise, in the context of ideological debates, consumers want to believe they are “on the right side of history” (i.e., posterity will prove them right — again, if appropriate actions are taken).² In the context of reporting on social issues (police brutality, climate change), they would like to believe in the ability of policy reforms (“defunding the police”, switching to green energy) to improve social welfare. In the context of business reporting, amateur investors want to believe they can “beat the market”, and aspiring entrepreneurs want to believe they will be the next Jeff Besos.

We assume that in all these contexts, consumers are attracted to news outlets that cultivate such hopeful beliefs. Accordingly, we propose a model in which news media aim to maximize consumers’ *anticipatory utility* — i.e., their expected indirect utility from their posterior beliefs. When these are beliefs are objectively wrong, there is a gap between anticipatory utility and objective expected payoffs.

However, under the conventional assumption that news media only supply information, this objective cannot give rise to media bias. The reason is that under rational expectations, maximizing ex-ante anticipatory utility (“what I believe I will get”) is indistinguishable from maximizing conventional indirect utility of a Bayesian rational consumer (“What I will actually get”), where full

¹In discussing the popularity of patriotic coverage of the war in Afganistan and Iraq, a New York Times story (Ruthenberg (2003)) quotes MSNBC’s president Erik Sorenson: “After Sept. 11 the country wants more optimism and benefit of the doubt...It’s about being positive as opposed to being negative.”

²E.g., see Chopra et al. (2023).

information provision is optimal. Thus, even when we assume non-instrumental demand for information (based on anticipatory utility), the standard view of the media as mere information providers cannot generate media bias.

This is where our view of media as joint providers of narratives and information enters. We show that this more comprehensive approach provides a non-trivial model of media bias, such that distortion of the truth consists of biased or inaccurate reports, *together with* false narratives. Moreover, there is synergy between these two instruments: They complement each other in producing the hopeful beliefs that consumers seek.

Overview of the model

In our basic model, a representative consumer takes an action after observing a signal about a state of Nature. There is an objective stochastic mapping from states and actions to outcomes. The consumer is endowed with a vNM utility function over states, actions and outcomes, which is additively separable in the action variable. A monopolistic media outlet commits ex-ante to a “media strategy”, which consists of: (i) a Blackwell experiment (a stochastic mapping from states to signals), and (ii) a narrative, which selects a subset of the outcome’s true causes.

There are four feasible narratives. The true narrative acknowledges both states and actions as causes. The “empowering” narrative postulates that actions are the sole cause of outcomes. The “fatalistic” narrative postulates that only the state matters for the outcome. These are the analogues of Loury’s (2020) above-mentioned development and bias narratives. Finally, the “denial” narrative asserts that neither the state nor the action cause the outcome (implicitly attributing outcomes to unspecified other factors).

The representative consumer’s strategy is a stochastic mapping from signals to actions. We interpret the strategy as the long-run aggregate behavior of many identical consumers, each making a one-shot decision. Together with the media strategy, it induces a long-run empirical joint distribution over states, actions and outcomes. A narrative produces a subjective conditional belief over outcomes, by “fitting” it to this long-run joint distribution. For example, the empowering narrative interprets the empirical correlation between actions and outcomes as a causal quantity — i.e., it attributes the long-run variation in outcomes entirely to variation in actions. Once the consumer adopts a narrative, his strategy prescribes actions that maximize expected utility with respect to

the narrative-induced belief. In equilibrium, this strategy is consistent with the empirical long-run distribution. The need for an equilibrium definition of consumer response to a given narrative is typical of models of decision making under misspecified models (e.g., Esponda and Pouzo (2016), Spiegel (2016), Eliaz and Spiegel (2020)).

The media’s problem is to find a strategy and an equilibrium consumer strategy that maximize the consumer’s ex-ante expected anticipatory utility. Incorporating equilibrium responses into the choice of a media strategy is in the spirit of the information-design literature (e.g., Kamenica and Gentzkow (2011), Bergemann and Morris (2019)). However, in standard models, equilibrium effects arise in multi-agent settings with payoff externalities. In contrast, in our model equilibrium effects arise because false narratives induce misspecified beliefs.

Overview of the results

Our account of news media raises a number of questions: Will the media provide accurate, unbiased information? If not, what is the structure of media inaccuracy/bias, and which narratives will media peddle? Our analysis of the basic model in Section 3 addresses these questions.

We begin with a complete analysis of an example that imposes structure on our model’s primitives. One story behind this example is that the consumer is an aspiring entrepreneur who considers a costly investment and dreams about making it big. An outcome indicates whether the entrepreneur is the first to develop a new product, and the state of Nature indicates whether there are positive returns from being the first. Objectively, these two variables are negatively correlated: Positive returns from being the first are associated with lower chances of attaining this goal. This specification is a running example in our paper. The optimal media strategy consists of the empowering narrative and optimistically biased information (correctly reporting good news, sometimes misrepresenting bad news). The lesson is that when feeding consumers’ fantasies of material success (which, according to our assumption, is what attracts them to the news outlet), the media peddles a narrative that “your life outcome is entirely up to you”, coupled with optimistically selective reporting about opportunities for business success. We aptly call this example “*the American dream*”.

The features of the optimal media strategy in this example are robust in the

following sense: For any action-separable utility function, if a media strategy gives higher anticipatory utility than the rational-expectations benchmark, then it must involve the empowering narrative. Also, it must provide information that induces different behavior from the benchmark (as long as the benchmark involves state-contingent actions). Thus, *there is synergy between false narratives and biased information*. We also present a result that clarifies why the negative state-outcome correlation in the example is necessary for the false narrative’s emergence. Specifically, when the consumer’s payoff is supermodular in the state and the outcome, and when these variables are affiliated given any action, the media cannot outperform the rational-expectations benchmark.

The role of consumer heterogeneity. Media consumers typically have diverse tastes, and hence, may be attracted to different news outlets, which may be differentiated in the information they report and the narrative they peddle. What is the effect of such diversity on the narratives and quality of information that consumers will be exposed to? What is the role of the media market structure? We address these questions in Sections 4-5, which introduce preference heterogeneity. Section 4 envisages our monopolist as a gatekeeper or platform that restricts the entry of news outlets. Formally, the platform chooses a menu of media strategies, aiming to maximize aggregate anticipatory utility. Each consumer selects a media strategy from the menu to maximize his individual anticipatory utility.

At first glance, it may appear that incentive-compatibility of consumers’ choices from the menu is moot, because all parties have a common objective: Maximizing consumers’ anticipatory utility. However, this is not the case, thanks to of an “*equilibrium data externality*” among consumer types. When evaluating a combination of a Blackwell experiment and a narrative, a consumer’s conditional belief over states is generated by the specific Blackwell experiment. However, his conditional belief over outcomes is determined by how the narrative interprets the *aggregate* distribution over relevant variables, which reflects the choices of *all* consumer types. Consequently, changes in the behavior of one segment of the consumer population can change how another segment evaluates media strategies. Dealing with this externality in the context of a menu design problem is a methodological novelty of our paper.

We show that under a mild condition on the preference distribution, the optimal menu *must* include a false narrative. Thus, thanks to the equilibrium

data externality, consumer heterogeneity encourages the emergence of false narratives. In addition, without loss of optimality, the menu includes a narrative that admits actions as causes of outcomes.

We then turn to the American dream example, where consumers differ in their cost of investment. For tractability, we restrict the domain of feasible Blackwell experiments: Signals are binary, and good news are always reported in the good state. Under the optimal menu, low-cost consumers choose the empowering narrative coupled with optimistically biased information. Intermediate-cost types choose the true narrative coupled with more informative signals. All these consumers invest whenever they receive a good signal. The third market segment consists of high-cost consumers, who choose the fatalistic or denial narratives (coupled with an arbitrary signal function), and never invest.

One of the first two consumer segments in this characterization may be empty. In particular, the false narratives that cater to extreme consumer types can chip away at the middle segment in a way that — thanks to the equilibrium data externality — reinforces this “poaching” effect. This effect can be stark: When the negative state-outcome correlation is sufficiently strong, consumer types below a cost threshold receive no information and always invest (egged on by the empowering narrative), while types above the threshold select the denial narrative and never invest. Thus, a heterogeneous population of consumers trying to make sense of the same aggregate data can end up holding highly polarized beliefs and taking opposite actions based on *no* information, just because they select different, self-serving false narratives peddled by different news outlets.

Section 5 explores the role of market structure by examining a competitive version of the heterogeneous-consumers model. The monopolistic gatekeeper is now gone; each media provider is “small”, in the sense that it takes the joint distribution over all variables as given, without internalizing the equilibrium data externality. We show that as in the monopoly case, false narratives and polarized beliefs arise in competitive equilibrium, albeit coupled with full information. While competition leads to better information, we show that it can make some consumer types objectively worse off relative to monopoly.

2 A Model

We begin by introducing the primitives of our model. There are four relevant variables, all taking finitely many values: A *state* of Nature t , an *action* a taken by a representative consumer, a *signal* s that the consumer observes before taking the action, and an *outcome* y . The state t is drawn from some exogenous distribution. The outcome y is determined according to some exogenous distribution conditional on a and t .

The consumer has a vNM utility function $u(t, a, y) = v(t, y) - C(a)$, where C is referred to as the consumer’s cost function.³ A monopolistic news media outlet (referred to as “the media”) commits ex-ante to a pair (I, N) , where I is a *signal function*, which is a Blackwell experiment assigning a distribution over signals s to each state t ; and N is a *narrative*, which is a subset of the two direct causes of y . The four possible narratives are (with an abuse of notation): The *true narrative* $N^* = \{t, a\}$; the “*empowering*” narrative $N^a = \{a\}$; the “*fatalistic*” narrative $N^t = \{t\}$; and the “*denial*” narrative $N^\emptyset = \emptyset$, which implicitly attributes y to unspecified other factors.

The consumer’s strategy is a (possibly stochastic) mapping from signals s to actions a . We think of this strategy as a description of long-run behavioral patterns by an infinite sequence of individual consumers: Each consumer makes a one-shot choice of action given a narrative-based interpretation of historical observations of the four variables, as we describe below. The long-run distribution p induced by the two parties’ strategies can be factorized as follows:⁴

$$p(t, s, a, y) = p(t)p(s | t)p(a | s)p(y | t, a) \quad (1)$$

The first and last terms on the R.H.S are exogenous; the second term is given by the media’s signal function I ; and the third term is given by the consumer’s strategy. The factorization reflects the causal structure underlying the data-generating process, which can be described by the following directed acyclic

³In a previous version of this paper (Eliaz and Spiegler (2024)), we also considered other classes of additively separable utility functions.

⁴We use p to notate every marginal and conditional distribution that is induced by the joint distribution over t, s, a, y .

graph (DAG):

$$\begin{array}{ccc} t & \rightarrow & s \\ \downarrow & & \downarrow \\ y & \leftarrow & a \end{array}$$

In this graphical representation, borrowed from the Statistics/AI literature on probabilistic graphical models (Pearl (2009)), a node represents a variable, and an arrow represents a direct causal relation. For example, the link $s \rightarrow a$ means that s is a direct cause of a . The DAG represents N^* by including the links $t \rightarrow y$ and $a \rightarrow y$. The three false narratives $N^a, N^t, N^\emptyset$ can be represented by DAGs that omit at least one of these links into y , while maintaining the true causal relations among t, s, a . For example, N^a omits the link $t \rightarrow y$, producing the DAG $t \rightarrow s \rightarrow a \rightarrow y$.

Given an objective full-support distribution p and the pair (I, N) , the consumer forms the following belief over t and y conditional on the signal realization s and an action a :

$$\tilde{p}(t, y \mid s, a) = p_I(t \mid s)p_N(y \mid t, a) \quad (2)$$

where $p_I(t \mid s)$ is the objective posterior probability of t conditional on s , which is induced by the signal function I via Bayes' rule; and $p_N(y \mid t, a)$ is the perceived probability of y conditional on t and a , which is shaped by the narrative N . Specifically,

$$\begin{array}{ll} p_{N^*}(y \mid t, a) = p(y \mid t, a) & p_{N^a}(y \mid t, a) = p(y \mid a) \\ p_{N^t}(y \mid t, a) = p(y \mid t) & p_{N^\emptyset}(y \mid t, a) = p(y) \end{array}$$

The interpretation is that the narrative N makes sense of the long-run distribution p by imposing a particular explanation for what causes variation in outcomes. The belief $p_N(y \mid t, a)$ is a systematic, narrative-based distortion of the objective conditional outcome distribution. Thus, the media affects the consumer's beliefs via two channels: (i) the signal function given by I , which determines the consumer's conditional belief over states; and (ii) the narrative N , which determines the consumer's conditional belief over outcomes.

More concretely, our interpretation of the second channel is that in addition to the signal s , the media also provides the *statistical data* described by $p_N(y \mid t, a)$ and frames it as a *causal* quantity. For example, when peddling the empowering narrative N^a , the media quotes statistical data about the historical

correlation between a and y and pitches it as a causal effect of a on y .

Importantly, when the narrative N is false, $p_N(y \mid t, a)$ is not invariant to the consumer's strategy, namely the long-run consumer average behavior given by $(p(a \mid s))_{a,s}$. To see why, elaborate $p_N(y \mid t, a)$ for each of the false narratives:

$$p_{N^a}(y \mid t, a) = \sum_{s', t'} p(s' \mid a) p(t' \mid s') p(y \mid t', a) \quad (3)$$

$$p_{N^t}(y \mid t, a) = \sum_{s', a'} p(s' \mid t) p(a' \mid s') p(y \mid t, a') \quad (4)$$

$$p_{N^\emptyset}(y \mid t, a) = \sum_{t'} p(t') \sum_{s', a'} p(s' \mid t') p(a' \mid s') p(y \mid t', a') \quad (5)$$

The terms $p(s' \mid a)$ and $p(a' \mid s')$ involve the consumer's strategy. In other words, long-run consumer behavior affects narrative-based perception of the mapping from actions to consequences (given a signal), which in turn affects the consumer's subjectively optimal decisions. If we view the long-run distribution p as a *steady state*, we need an equilibrium notion of the consumer's subjective optimization.

Because $p_N(y \mid t, a)$ may involve conditioning on null events, we make use of a full-support perturbation. Specifically, let $\varepsilon > 0$ be arbitrarily small, and let σ^* be an exogenous full-support strategy (i.e., $p(a \mid s) > 0$ for every a, s). The consumer's endogenous, deliberate behavior is given by a strategy σ , such that the relation between $(p(a \mid s))_{a,s}$ and σ is

$$p(a \mid s) \equiv (1 - \varepsilon) \cdot \sigma(a \mid s) + \varepsilon \cdot \sigma^*(a \mid s) \quad (6)$$

The interpretation is that a fraction ε of the consumer population follows σ^* , while a fraction $1 - \varepsilon$ makes subjectively optimal choices with respect to their beliefs, as described by the following definition.⁵

⁵This is essentially the notion of personal equilibrium by Spiegel (2016), which coincides with Berk-Nash equilibrium (Esponda and Pouzo (2016)) when the consumer's subjective model is defined by N . The only difference is in the treatment of full-support perturbations, which we modify here for convenience.

Definition 1 (Equilibrium) Fix $\varepsilon > 0$ and (I, N) . A consumer strategy σ is an equilibrium with respect to (I, N) and ε if, whenever $\sigma(a | s) > 0$,

$$a \in \arg \max_{a'} V_{I,N}(s, a') = \sum_{t,y} p_I(t | s) p_N(y | t, a') u(t, a', y) \quad (7)$$

Fix $\varepsilon > 0$. The media chooses (I, N) and a consumer strategy σ to maximize

$$U(I, N) = \sum_t p(t) \sum_s p(s | t) \sum_a p(a | s) V_{I,N}(s, a) \quad (8)$$

subject to the constraint that σ is an equilibrium with respect to (I, N) and ε . We will focus on the $\varepsilon \rightarrow 0$ *limit* of the solutions to this problem, and refer to the (I, N) component of such a limit as an *optimal media strategy*. (The only purpose of the full-support perturbation is to make expressions such as (7) well-defined; our characterizations of optimal media strategies are invariant to σ^* . Thus, unless close attention to the perturbation is required, we analyze the limit directly and take it for granted that it can be justified by an arbitrary perturbation.)

The media’s objective function (8) is the consumer’s *expected anticipatory utility*. The interpretation is that anticipatory utility drives the consumer’s demand for news media. The higher his anticipatory utility, the greater his media engagement. We do *not* regard U as a measure of the consumer’s “true welfare”. Rather, it is a proxy for his media engagement, which is what the media cares about. Our task is to characterize the media’s optimal strategy.

Private/Public interpretations of a and y

According to one interpretation of our model, a represents a *private* action that an individual media consumer takes, and y is a personal outcome of his choice. For example, a can represent a career decision or a dietary choice, in which case y represents earnings or health outcomes, respectively. The data that the consumer relies on to form beliefs is *aggregate*, reflecting the historical choices and outcomes of other consumers.

An alternative interpretation is that a represents a *public* choice (such as economic or foreign policy), and y represents a public outcome (economic growth, national security). According to this interpretation, the media consumer is a representative *voter*, and the probability $p(a | s)$ is the frequency with which society selects a political leadership that implements a . This is a reduced-form

representation of a democratic process, such that society’s choice matches what the representative voter deems optimal.

The necessity of false narratives for media bias

Suppose that the media is restricted to providing the true narrative N^* . This reduces the model to standard information provision by a sender who can commit ex-ante to a Blackwell experiment. The sender faces a Bayesian receiver whose indirect utility from a posterior belief μ over t is

$$\max_a \sum_t \mu(t) \sum_y p(y \mid t, a) u(t, a, y)$$

This is a conventional indirect utility function. Since it is a maximum over linear functions of μ , it is convex in μ . Therefore, it is (weakly) optimal for the sender to commit to a fully informative signal — i.e., $p(s = t \mid t) \equiv 1$.

It follows that in our model, given the media’s objective of maximizing the consumer’s ex-ante anticipatory utility, the media has no strict incentive to provide partial or biased information unless it also peddles a false narrative. Throughout the paper, we refer to the maximal anticipatory utility attained by the true narrative and complete information as the *rational-expectations benchmark*.

A revelation principle

Our model departs from the canonical information-design framework (see Bergemann and Morris (2019)), since it allows the designer to influence the subjective model that the receiver holds. Nevertheless, the assumption that the consumer always correctly perceives $p(t, s, a)$ ensures that the standard revelation principle in the information-design literature can be adapted to the present setting, which simplifies our analysis.

Remark 1 *For any ε and without loss of optimality, we can let the set of signals coincide with the set of feasible actions, and restrict attention to equilibria in which $\sigma(a = s \mid s) \equiv 1$.*

The proof follows the footsteps of Theorem 1 in Bergemann and Morris (2016) — adapted to the single-player setting — and is therefore omitted. The proof involves manipulating the signal function given by $(p(s \mid t))_{t,s}$ and the consumer’s behavior given by $(p(a \mid s))_{a,s}$. In general, when the consumer forms

beliefs according to a misspecified model N , such changes may affect $p_N(y \mid t, a)$, which could violate the revelation principle. The reason the principle holds in our setting is that the manipulation of $(p(s \mid t))_{t,s}$ and $(p(a \mid s))_{a,s}$ in the proof leaves $(p(t, a))_{t,a}$ unchanged. Even when the consumer adopts a false narrative, he correctly perceives the joint distribution $p(t, s, a)$. By expressions (3)-(5), this means that $p_N(y \mid t, a)$ remains unchanged as well, regardless of how t and a are jointly distributed with s . This enables the standard Bergemann-Morris proof to go through.⁶

3 Analysis

In this section we analyze the media's optimal strategy. We begin with a specification that serves as a running example in the paper. We then show that the qualitative features of this example hold more generally.

3.1 “The American Dream”

In this example, the variables t and a take values in $\{0, 1\}$. By the revelation principle, we can assume $s \in \{0, 1\}$ as well. The variable y takes non-negative real values. The exogenous components of the data-generating process are $p(t = 1) = \frac{1}{2}$ and $E(y \mid t, a) = a \cdot f_t$, where $f_0 > f_1 > 0$. The consumer's payoff function is $u(a, t, y) = ty - ca$, where $0 < c < \min\{f_1, f_0 - f_1\}$. The action a represents a private decision whether to engage in a costly economic activity. The outcome y indicates the extent to which the activity attains an objective. The state t represents the returns from such attainment. Higher attainment is associated with lower returns, reflecting background equilibrium effects.⁷

For a concrete story, the consumer is an aspiring entrepreneur who decides whether to develop a new product. Suppose $y \in \{0, 1\}$, where $y = 1$ represents being the first to succeed. The state $t = 1$ means there is demand for the product — in which case, more competitors flock to the market, thus lowering the entrepreneur's chances of being the first. In an alternative story, the consumer

⁶One minor distinction is that the obedience constraint is applied to the deliberate strategy σ , whereas the consumer's subjective-expected-utility calculations in the proof involve the conditional probabilities $p(a \mid s)$ given by (6). However, this does not affect the proof.

⁷A more elaborate specification would model these forces explicitly, incorporating how consumers' decisions contribute to the equilibrium effects. Since this would add complexity without altering the main qualitative insight, we chose not to do so. In this sense, we perform a partial equilibrium analysis.

is a high school student (or his parent) who decides whether to exert costly effort at school (private tutoring, extracurricular activities). The outcome y represents the prestige of the college the student manages to enter. The state represents the college wage premium. A higher premium makes colleges more selective (hence the negative correlation between t and y).

Under both stories, the media provides information about the fundamentals represented by t , as well as a narrative about what drives the outcome y . The narrative determines whether people attribute personal material outcomes to internal factors under their control or to external factors beyond their control. Thus, our example captures in stylized form the forces described by Iyengar (1990) and Loury (2020), as mentioned in the Introduction: The media may use a combination of false narratives and biased information to shape popular perceptions about the role of personal agency and external circumstances in determining life outcomes.

By the revelation principle, we can restrict attention to binary signals and an equilibrium in which the consumer always plays $a = s$.

Rational-expectations benchmark

Suppose the media offers the true narrative N^* . As we saw in Section 2, it is optimal to couple this narrative with a fully informative signal. When $t = 0$, the consumer knows that $ty = 0$, and therefore plays $a = 0$. When $t = 1$, he knows that $E(y \mid t = 1, a) = af_1$. Since $c < f_1$, the consumer plays $a = 1$. It follows that the rational-expectations benchmark in this example is $\frac{1}{2}(f_1 - c)$.

□

Narratives that omit the link $a \rightarrow y$

Under the narratives N^t and $N^\emptyset$, the consumer believes that his action has no effect on y , and therefore prefers to take the costless action $a = 0$. This means that in any equilibrium, $a = 0$ with certainty for every t . Since $y = 0$ whenever $a = 0$, it follows that $E(y) = 0$. Therefore, the consumer's anticipatory utility is necessarily zero, which is below the rational-expectations benchmark. It follows that the media will necessarily offer a narrative that acknowledges a as a cause of y . □

The empowering narrative

Under the narrative N^a and any signal function I ,

$$E_{N^a}(ty \mid a, s) = p(t = 1 \mid s)E(y \mid a) \quad (9)$$

Observe that although the consumer believes that only a causes y , he cares about t because his net payoff is positive only when $t = 1$.

The consumer's subjective payoff from $a = 0$ is zero regardless of s , because $E(y \mid a = 0) = 0$. Let us now turn to his payoff from $a = 1$ for each s . Applying the revelation principle, we guess an equilibrium in which $a \equiv s$ on the equilibrium path. We also guess that the obedience constraints (described below) hold, and verify the guess at the end of the derivation. Under the guess,

$$E(y \mid a = 1) = \sum_t p(t \mid a = 1) f_t = \sum_t p(t \mid s = 1) f_t \quad (10)$$

Denoting $q_t = p(s = 1 \mid t)$ and plugging

$$\begin{aligned} p(t = 1 \mid s = 1) &= \frac{q_1}{q_1 + q_0} \\ p(t = 1 \mid s = 0) &= \frac{1 - q_1}{2 - q_1 - q_0} \end{aligned}$$

and (10) in (9), we obtain

$$E_{N^a}(ty \mid s = 1, a = 1) = \frac{q_1}{q_1 + q_0} \cdot \frac{q_1 f_1 + q_0 f_0}{q_1 + q_0} \quad (11)$$

and

$$E_{N^a}(ty \mid s = 0, a = 1) = \frac{1 - q_1}{2 - q_1 - q_0} \cdot \frac{q_1 f_1 + q_0 f_0}{q_1 + q_0} \quad (12)$$

The obedience constraints (which we check later) require (11) and (12) to be weakly above and below c , respectively.

In the guessed equilibrium, when $s = 0$, the consumer plays $a = 0$ and gets zero payoffs. The consumer's anticipatory utility is $p(s = 1) \cdot [E_{N^a}(ty \mid s = 1, a = 1) - c]$, which is equal to

$$\frac{q_1 + q_0}{2} \cdot \left(\frac{q_1}{q_1 + q_0} \cdot \frac{q_1 f_1 + q_0 f_0}{q_1 + q_0} - c \right) \quad (13)$$

Observe that when the media offers a fully informative signal ($q_t \equiv t$), this expression coincides with the payoff from N^* . Thus, if the false narrative N^a outperforms the true narrative, it must be coupled with incomplete information. We now proceed to calculate the optimal $I = (q_0, q_1)$ that accompanies N^a . The following claim simplifies the problem.

Claim 1 *Under N^a , it is optimal to set $q_1 = 1$.*

Thus, if the optimal signal function has a bias, it must be an optimistic one (i.e., sometimes reporting $s = 1$ in $t = 0$). The simple proof of this claim (like all proofs in this paper) is in the Appendix. The claim reduces the consumer's anticipatory utility into

$$\frac{1}{2} \left[\frac{f_1 + q_0 f_0}{1 + q_0} - c(1 + q_0) \right] \quad (14)$$

The unique value of q_0 that maximizes this expression is

$$q_0 = \min \left\{ 1, \sqrt{\frac{f_0 - f_1}{c}} - 1 \right\} > 0 \quad (15)$$

We have thus constructed a media strategy consisting of the empowering narrative and optimistically biased information, which outperforms the rational-expectations benchmark.

Note that (14) is equal to $p(s = 1) \cdot [E_{N^a}(ty \mid s = 1, a = 1) - c]$, which is positive since it is above the rational expectations benchmark. Therefore, the R.H.S. of (11) exceeds c , thereby satisfying the obedience constraint for $s = 1$. Note also that $q_1 = 1$ implies that (12) is zero, such that playing $a = 0$ when $s = 0$ is optimal for the consumer. We have thus confirmed that the obedience constraints hold under the media strategy we have derived. $\square$

Thus, the optimal media strategy involves the narrative N^a coupled with positively biased information: Always sending a good signal in the good state, and sending it with positive probability in the bad state.⁸

In terms of the story behind the example, the false narrative N^a claims that attainment of a career or business objective depends entirely on one's initiative. The accompanying signal function has an optimistic bias, claiming that returns from attaining the objective are high even when they are not. The media exaggerates the attractiveness of the external environment, and suppresses — via the empowering narrative — the negative effect that good fundamentals have on the chances of a successful outcome. Therefore, we find it apt to refer to the media in this example as peddling the “American dream”.⁹

⁸When $c > \min\{f_1, f_0 - f_1\}$, no media strategy can outperform the rational-expectations benchmark (which induces zero payoffs when $c > f_1$).

⁹The political-economics implications of popular perceptions of the role of personal choices

The synergy between false narratives and biased signals

Biased information is necessary for N^a to beat the rational-expectations benchmark. Suppose the media provides full information. This means that t and s are perfectly correlated ($s \equiv t$). The revelation principle means that $a \equiv s$ on the equilibrium path, such that a and t are perfectly correlated, too. But this means that omitting t as an explanatory variable for y does not lead to erroneous beliefs: $p(y | a)$ coincides with $p(y | t, a)$. In turn, this implies that the consumer effectively has rational expectations and perfectly monitors t , which gives the rational-expectations benchmark. Therefore, incomplete information is necessary for N^a to enhance the consumer's anticipatory utility.

The reason that the combination of N^a and optimistically biased information outperforms the benchmark is that it produces a *correlation-neglect* effect. As expression (9) makes explicit, the consumer believes that t and y are independent conditional on (s, a) . In reality, t and y are *negatively* correlated. By neglecting this correlation, the consumer attains a more optimistic belief about the product ty conditional on $s = a = 1$. This effect is non-null only when $p(a = 1 | t = 0) > 0$, which only happens when information is biased.

3.2 Generalizing the Example

The “American dream” example has three noteworthy features. First, the empowering narrative emerges as optimal. Second, it distorts consumer behavior away from the rational-expectations benchmark. Finally, the combination of these two effects works by leveraging an underlying negative correlation between t and y (conditional on $a = 1$). We now present three results that generalize these observations.

Proposition 1 *If the media can outperform the rational-expectations benchmark, then N^a is part of an optimal media strategy.*

Thus, the empowering narrative N^a is an essential feature of media strategies that beat the rational-expectations benchmark. The logic behind the result is as follows. Because u is action-separable, a false narrative can have an effect on ex-ante anticipatory utility only when it induces a belief that distorts the joint distribution of (t, y) . By definition, the fatalistic narrative N^t cannot do that.

in life outcomes have been studied by Piketty (1995) and Alesina and Angeletos (2005).

In principle, the denial narrative $N^\emptyset$ can attain such a distortion. However, this effect can be replicated (and possibly improved upon) by N^a coupled with *no information*. As the latter point demonstrates, the synergy between the two components of media strategies plays a key role in the proof of this general result.

The next result addresses the consumer behavior that the optimal media strategy induces. We say that the payoff function and the exogenous data-generating process form a *regular environment* if, under the true narrative coupled with complete information, the consumer has a unique best-reply which is a one-to-one function of the state. That is, in regular environments different states prescribe different unique actions under rational expectations.

Proposition 2 *Suppose the environment is regular. If an optimal media strategy outperforms the rational-expectations benchmark, then its induced conditional distribution $(p(a | t))_{t,a}$ is different from that benchmark.*

Thus, when the media deviates from the rational-expectations benchmark, it necessarily induces changes in consumer behavior.¹⁰ To see the role of regularity in this result, consider the payoff specification of Section 3.1, and modify the data-generating process by assuming $E(y | t, a) = 1 - t$ for every t, a . Under rational expectations, the consumer's optimal action is $a = 0$ for every t , and the rational-expectations payoff is 0 (because $a = 0$ and $ty = 0$ with probability one). Using similar arguments as in Section 3.1, it can be shown that it is optimal for the media to provide N^a coupled with no information (or, equivalently, $N^\emptyset$ with an arbitrary signal function). The consumer responds by playing $a = 0$. His anticipatory payoff is $\frac{1}{4}$, beating the rational-expectations benchmark, although the behavior is the same. Thus, without regularity, an optimal media strategy can outperform the benchmark without any effect on consumer behavior.

Proposition 2 has implications for the consumer's objective welfare (i.e., his expected utility according to the true data-generating process). In a regular environment, the optimal media strategy lowers the consumer's true welfare.

¹⁰Proposition 2 does not claim that the media necessarily employs biased signals. We cannot rule out the possibility that it is optimal for the media to accompany N^a with full information, anticipating that the consumer's subjective best-reply will involve mixing (the revelation principle does not guarantee that $a \equiv s$ in *all* equilibria).

However, this may not be true in irregular environments: The consumer may entertain the illusion of higher expected payoffs at no objective cost.

Our final result clarifies why the negative correlation between t and y was required for the false narrative to emerge in the American dream example. In preparation for this result, suppose t and y take non-negative real values, and that v (the gross payoff from t, y) is a *supermodular* function.¹¹ We say that $p(y | t, a)$ satisfies the monotone likelihood ratio property (MLRP) if for every a , $t' > t$ and $y' > y$,

$$\frac{p(y' | t', a)}{p(y | t', a)} \geq \frac{p(y' | t, a)}{p(y | t, a)}$$

This is a familiar property, which implies that t and y are *affiliated* given a .

Proposition 3 *Suppose that v is supermodular and that $p(y | t, a)$ satisfies MLRP for every a . Then, the media strategy consisting of the true narrative N^* coupled with complete information is optimal.*

This result establishes that when t and y are positively correlated (in the MLRP sense) given a , the media cannot beat the rational-expectations benchmark. Thus, the negative-correlation-neglect aspect of the American dream example is not accidental. The proof makes subtle use of the revelation principle, combined with standard results on the supermodular order for bivariate distributions.

An example: Matching the state

We now present a simple example that illustrates the generality of the forces behind the American dream example. Let $t, a, y \in \{0, 1\}$. The state t is uniformly distributed. Let $u(t, y, a) = \mathbf{1}[t = y]$ — i.e., the consumer wishes to match the outcome to the state. Note that $C(a) = 0$ for every a . Denote $\theta_{ta} = p(y = 1 | t, a)$. Assume $\theta_{11} = \theta_{00} > \frac{1}{2}$ and $\theta_{01} > \theta_{11}$. As a result, the rational-expectations optimal strategy is $a = t$, such that the rational-expectations payoff is $\frac{1}{2}\theta_{11} + \frac{1}{2}(1 - \theta_{00}) = \frac{1}{2}$. Note that $\theta_{11} > 1 - \theta_{00}$ — i.e., the consumer’s optimal objective payoff is higher in state $t = 1$. Thus, while different from the American dream example in many respects, it shares two key features: In both examples, $t = 1$ is a “good” state; and an outcome that is considered good in $t = 1$ is actually more likely in $t = 0$.

¹¹In the “American dream” example, $v(t, y) = ty$, hence it is supermodular with respect to the natural order on t and y .

We do not derive the optimal media strategy, but merely settle for showing that the media can outperform the rational-expectations benchmark with the empowering narrative, coupled with the following biased signal function: When $t = 1, s = 1$ with probability one; and when $t = 0, s = 1$ with probability $\varepsilon > 0$. Since the obedience constraints hold with slack under rational expectations, by continuity they continue to hold with slack if ε is sufficiently small. Hence, we can take it for granted that the consumer always plays $a = s$.

It follows that the consumer's anticipatory utility is

$$p(s = 1)[p(t = 1 | s = 1)p(y = 1 | a = 1) + p(t = 0 | s = 1)p(y = 0 | a = 1)] \\ + p(s = 0)p(y = 0 | a = 0)$$

which equals

$$\left(\frac{1}{2} + \frac{1}{2}\varepsilon\right) \cdot \left[\frac{1}{1+\varepsilon} \left(\frac{1}{1+\varepsilon}\theta_{11} + \frac{\varepsilon}{1+\varepsilon}\theta_{01}\right) + \frac{\varepsilon}{1+\varepsilon} \left(1 - \frac{1}{1+\varepsilon}\theta_{11} - \frac{\varepsilon}{1+\varepsilon}\theta_{01}\right)\right] \\ + \left(\frac{1}{2} - \frac{1}{2}\varepsilon\right) \cdot (1 - \theta_{00})$$

This expression is higher than the rational-expectations payoff if

$$\theta_{00} + \theta_{01} - 2\theta_{11} + \varepsilon\theta_{00} - \varepsilon\theta_{01} > 0$$

Since $\theta_{11} = \theta_{00}$ and $\theta_{01} > \theta_{11}$, this inequality holds if ε is sufficiently small. Thus, a combination of the empowering narrative and a signal with a small optimistic bias beats the rational expectations benchmark.¹²

4 Heterogeneous Consumers

Media consumers often hold varied preferences, leading them to gravitate toward distinct news outlets that may differ both in the information they provide and the narratives they promote. This section explores how taste heterogeneity impacts the narratives and quality of information that consumers are exposed to.

To introduce preference heterogeneity into our model, we let the supply side consist of multiple media strategies that consumers can choose from, each

¹²It can be shown that the other false narratives are strictly inferior to the empowering narrative.

according to his preferences. We consider a monopolistic media *platform* acting as a gatekeeper that restricts the entry of media providers (each represented by a distinct media strategy). The monopolist’s objective is to maximize consumers’ aggregate anticipatory utility — reflecting the continued assumption that this corresponds to maximizing their platform engagement.

Formally, the platform commits ex-ante to a menu M of pairs (I, N) . The set of consumer types is $\Theta = [0, 1]$. Types are distributed according to a continuous and strictly increasing *cdf* G with full support. Let u_θ be type θ ’s payoff function. Each type θ selects a pair $(I_\theta, N_\theta) \in M$ and a signal-dependent action $a_\theta(s)$ to maximize his ex-ante anticipatory utility. The platform’s objective is to maximize consumers’ *aggregate* ex-ante anticipatory utility.

The platform faces a “second-degree discrimination” problem, which arises because it cannot prevent consumers from freely choosing their favorite media strategy on the menu. As we will see, the behavior of consumers who follow one news outlet may affect the beliefs of consumers following different outlets. This effect, which we refer to as an “equilibrium data externality”, is what makes the menu-design problem non-trivial.

The menu design problem

To formally describe the design problem, we begin with how consumers evaluate alternatives. Fix some profile of consumer types’ media-strategy choices and signal-dependent actions, $(I_\theta, N_\theta, (a_\theta(s)))_{\theta \in \Theta}$. As in Section 2, let $\varepsilon > 0$ be arbitrarily small and let σ^* be a fixed full-support signal-dependent action distribution. The perturbation’s role continues to be to ensure that consumers’ subjective beliefs do not involve conditioning on null events; and as before, the exact structure of σ^* is irrelevant.

Aggregate consumer behavior is given by $(p(a \mid t))_{a,t}$, where

$$p(a \mid t) = \int_{\theta} \sum_s p_{I_\theta}(s \mid t) \cdot \{(1 - \varepsilon) \cdot \mathbf{1}[a_\theta(s) = a] + \varepsilon \cdot \sigma^*(a \mid s)\} dG(\theta) \quad (16)$$

and $(p_{I_\theta}(s \mid t))_{s,t}$ is the Blackwell experiment given by I_θ . Denote $\mathbf{a} \equiv (a(s))_s$. Given $(p(a \mid t))_{a,t}$, consumer type c ’s ex-ante evaluation of any $(I, N, \mathbf{a})$ is

$$U_\theta(I, N, \mathbf{a}) = \sum_s p_I(s) \sum_{t,y} p_I(t \mid s) p_N(y \mid t, a(s)) u_\theta(t, a(s), y) \quad (17)$$

In this formula, $p_I(s)$ and $p_I(t \mid s)$ are induced by the objective prior

probability $p(t)$ and the Blackwell experiment I . The conditional probability $p_N(y | t, a)$ is as defined in Section 2, based on $p(t, a, y) = p(t)p(a | t)p(y | t, a)$, with $p(a | t)$ representing *aggregate* consumer behavior as in (16). Thus, although different consumer types may select different media outlets, they do not live in isolated islands: They all belong to the same society, and the news media they consume offer narratives that interpret the *same* aggregate data arising from the choices of all consumers. It follows that when N_θ is a false narrative, the anticipatory payoff that type θ gets from his chosen triplet $(I_\theta, N_\theta, (a_\theta(s)))$ is affected by the choices made by *all the other types*, since these determine the joint aggregate distribution $p(t, a, y)$. This is the *equilibrium data externality* we referred to above.

Effectively, the platform's problem is to design a menu of (I, N) pairs that maximizes consumers' aggregate anticipatory utility, subject to the constraint that each consumer type selects the pair — and subsequent signal-dependent actions — that maximize his own anticipatory utility. Formally, the platform chooses a profile of *triplets* $(I_\theta, N_\theta, \mathbf{a}_\theta)_{\theta \in \Theta}$ to maximize

$$\int_{\theta} U_{\theta}(I_{\theta}, N_{\theta}, \mathbf{a}_{\theta}) dG(\theta)$$

subject to the constraints that for every $\theta \in \Theta$: (i) the triplet $(I_\theta, N_\theta, \mathbf{a}_\theta)$ maximizes U_θ over the set $\{I_\theta, N_\theta, \mathbf{a}_\theta\}_{\theta \in \Theta}$; and (ii) $\mathbf{a}_\theta$ maximizes $U_\theta(I_\theta, N_\theta, \mathbf{a})$ given (I_θ, N_θ) . The two constraints in this problem constitute an equilibrium requirement, as in Definition 1, since U_θ is defined with respect to $(p(a | t))$, which itself is induced by consumers' choices.¹³ From now on, we take the equilibrium aspect of the platform's problem for granted, and usually refrain from explicit equilibrium terminology. Throughout this section, we analyze the $\varepsilon \rightarrow 0$ limit of the solution to the platform's problem.¹⁴

Consumers' choice of (I, N) from the platform's menu involves comparing different models, and therefore different beliefs. We do *not* envisage consumers as rationally and explicitly deliberating over this problem like rational scientists. Instead, our interpretation is that consumers are attracted to the news outlets that make them more hopeful about the future (conditional on taking certain actions). Consumers try various news media, experience the anticipatory feeling

¹³We rule out mixing by consumers. This entails no loss of generality, as the set of consumer types who would mix has measure zero.

¹⁴In proofs, we typically analyze the limit directly, and make explicit reference to the full-support perturbation only when necessary.

each generates, and gravitate toward the news outlet that offers the better illusion — even if systematic observations of the past can refute these illusions. We believe this behavioral account is plausible for news media consumption.

4.1 A General Result

This sub-section presents a result that partially characterizes optimal menus, under minimal structure on the primitives. First, we allow for arbitrary sets of feasible signal functions I , as long as they include full information and no information. Second, t and a take real values in finite sets. Let a^* denote the lowest feasible action. Third, type θ 's payoff function takes the form $u_\theta(t, a, y) = v(t, y) - C(a, \theta)$, where C is continuous in θ and strictly increasing in both arguments. To avoid uninteresting knife-edge cases, we assume: (i) $a^* = \arg \max_a [v(t, y) - C(a, 1)]$ for every t ; and (ii) $a^* < \arg \max_a [v(t, y) - C(a, 0)]$ for some t . In other words, under rational expectations and complete information, high-cost types always take the least costly action, whereas low-cost types sometimes take a costly action.

Under these restrictions, we can state the following result.

Proposition 4 *In the $\varepsilon \rightarrow 0$ limit, any optimal menu must include a false narrative. Furthermore, there is an optimal menu that includes N^* or N^a , such that a positive measure of types play $a > a^*$ with positive probability.*

The comparison with the homogenous-consumers case is stark. According to Proposition 3, there is a substantial class of environments for which *no* media strategy beats the rational-expectations benchmark in the homogeneity case. In contrast, false narratives are *always essential* to the platform's strategy in the heterogeneity case.

The equilibrium data externality drives this effect. To see why, recall the American Dream example. Let $f_1 = \frac{1}{2}$, and suppose there are two consumer types, L and H , characterized by $c_L \approx 0$ and $c_H \approx 1$, with equal shares in the population. Under the true narrative and full information, low-cost consumers obey the signal and earn $\frac{1}{2}(\frac{1}{2} - c_L) > 0$, whereas high-cost consumers choose $a = 0$ and earn zero. Now imagine that the platform introduces an additional media strategy consisting of the fatalistic narrative (coupled with an arbitrary signal function). Suppose consumers self-select as follows: Low-cost consumers

stick to the true narrative, while high-cost consumers now switch to the fatalistic narrative.

We now check that this self-selection is subjectively optimal, and that it induces higher overall anticipatory utility. For low-cost consumers, the true narrative (coupled with full information) continues to induce $a = t$ as an optimal response, and generate an anticipatory utility of $\frac{1}{2}(\frac{1}{2} - c_L)$. For any consumer who adopts the fatalistic narrative, the optimal response is $a = 0$ and the induced anticipatory utility is $p(t = 1) \cdot E(y = 1 \mid t = 1)$. Under the assumption that only low-cost consumers play $a = 1$ at $t = 1$, $E(y = 1 \mid t = 1) = \frac{1}{2}f_1$. It follows that the anticipatory utility from the fatalistic narrative is $\frac{1}{8}$. For low-cost consumers, this is inferior to the true narrative, while the opposite is the case for high-cost consumers.

This confirms that the self-selection we assumed is subjectively optimal for all consumers. Moreover, since now high-cost consumers earn a strictly positive anticipatory utility, the new menu improves on the original one. Importantly, this improvement necessarily arises because the equilibrium externality is one-directional: The choices made by consumers who adopt the true narrative exert a positive externality on consumers who opt for the fatalistic narrative; while consumers who adopt the true narrative are immune to any externality.

This example also illustrates a key difference between the heterogeneous-population case and the homogenous-population model analyzed in Section 3. For the fatalistic narrative to generate a positive anticipatory utility for any consumer, there must be a positive fraction of the consumer population who play $a = 1$ at $t = 1$. If all consumers are of type H , none of them will find this profitable. In contrast, when there are also low-cost consumers in the population, the narratives N^* and N^a can drive them to play $a = 1$ at $t = 1$, thus generating a positive externality on high-cost consumers who adopt N^t .

The result's second part states that without loss of optimality, the platform's menu includes a narrative that includes actions as a cause of outcomes. Unlike the first part, here, we cannot say categorically whether the menu must include such a narrative, nor whether it is distinct from the false narrative implied by the result's first part.

4.2 Revisiting the “American Dream”

We now turn back to the “American dream” example of Section 3.1, extending it by introducing consumer heterogeneity. Specifically, we identify the consumer type θ with the cost parameter c . We impose two restrictions that are consistent with those made at the beginning of the previous sub-section. First, let $f_1 < 1$. Second, the domain of feasible signal functions is restricted: $s \in \{0, 1\}$, and $\Pr(s = 1 \mid t = 1) = 1$. This restriction entailed no loss of generality in the representative-consumer case of Section 3.1. This is no longer the case here. The restriction also means that we cannot apply the revelation principle. Accordingly, we will *not* take it for granted that consumers’ actions mimic the signal they receive. We impose this restriction for tractability, as a result of several non-standard sources of complexity. First, since different types can select different narratives having non-linear effects on their beliefs, there is no obvious single-crossing-like argument that would impose order on the incentive constraints. Second, the equilibrium data externality is *global*: When we change the (I, N) that one interval of types selects, this affects how *every* type evaluates the false narratives on the menu. Therefore, we cannot reduce the problem to checking local incentive constraints. Finally, since mechanism-design with multi-dimensional instruments is typically difficult to solve, our domain restriction simplifies the analysis by reducing the dimensionality of the platform’s instruments (which consist of Blackwell experiments *and* narratives).

Thus, in what follows, each signal function I is identified with q , which is the probability of submitting $s = 1$ when $t = 0$. The probability of $t = 1$ conditional on s under I is thus $p_q(t = 1 \mid s) = s/(1 + q)$. In particular, when the consumer observes the signal $s = 0$, he infers that $t = 0$ and therefore $ty = 0$ with probability one. Therefore, we can take it for granted that all consumer types play $a = 0$ and earn zero payoffs when receiving the signal $s = 0$. This observation enables us to identify $\mathbf{a}$ with the action taken at $s = 1$, denoted $a(1)$, and simplify $U_c(q, N, a(1))$ into

$$\begin{aligned} U_c(q, N, a(1)) &= p_q(s = 1) \cdot [p_q(t = 1 \mid s = 1)E_N(y \mid t = 1, a(1)) - ca(1)] \\ &= \frac{1 + q}{2} \cdot \left[\frac{1}{1 + q} E_N(y \mid t = 1, a(1)) - ca(1) \right] \\ &= \frac{1}{2} E_N(y \mid t = 1, a(1)) - \frac{c(1 + q)}{2} a(1) \end{aligned} \tag{18}$$

Likewise, consumers' aggregate state-dependent behavior can be simplified into

$$p(a = 1 \mid t = 1) = \int_0^1 a_c(1) dG(c) \quad p(a = 1 \mid t = 0) = \int_0^1 q_c a_c(0) dG(c)$$

We can now restate the platform's problem: Choose a profile $(q_c, N_c, a_c(1))_{c \in C}$ that maximizes

$$\int_0^1 U_c(q_c, N_c, a_c(1)) dG(c)$$

subject to the constraints that for every c , $U_c(q_c, N_c, a_c(1)) \geq U_c(q_{c'}, N_{c'}, a_{c'}(1))$ for every $c' \in C$; and that $a_c(1)$ maximizes U_c given (q_c, N_c) . The latter constraint can be written as follows:

$$\begin{aligned} & \frac{1}{1+q} \cdot E_{N_c}(y \mid t = 1, a = a_c(1)) - ca_c(1) \\ & \geq \frac{1}{1+q} \cdot E_{N_c}(y \mid t = 1, a = 1 - a_c(1)) - c(1 - a_c(1)) \end{aligned}$$

Recall that the constraints are equilibrium conditions, because $(a_c(1))_{c \in C}$ affects $E_N(y \mid t = 1, a)$ when N is false.

Proposition 5 *In the $\varepsilon \rightarrow 0$ limit, the platform maximizes its objective function with a menu that has the following structure. There exist $c^{**} \in (0, 1]$ and $c^* \in [0, c^{**}]$ such that:*

- (i) *All consumer types in $[0, c^*]$ choose (q^a, N^a) with $q^a > 0$ and play $a \equiv s$.*
- (ii) *All consumer types in $(c^*, c^{**}]$ choose (q^*, N^*) and play $a \equiv s$. Moreover, if $c^* = 0$, then $q^* = 0$; and if $c^* > 0$, then $q^* < q^a$.*
- (iii) *All consumer types in $(c^{**}, 1]$ choose one narrative $N \in \{N^t, N^\emptyset\}$, coupled with an arbitrary q , and always play $a = 0$. If $c^* = 0$, then $N = N^t$.*

This result demonstrates that differentiation among consumer types plays out differently here than in the homogenous-consumers case of Section 3.1. In that environment, the same narrative (N^a) serves any consumer type; the differentiation between them is done entirely through the signal function.¹⁵ In contrast, differentiation between types in the heterogeneous case is carried out by offering a menu of narratives.

¹⁵In principle, the menu-design problem can be defined for a homogenous population, where identical consumers can select different pairs (I, N) . Nevertheless, it can be shown that in the “American dream” example, the degenerate menu of Section 3.1 is optimal.

Specifically, each of the narratives that admit a as a cause of y is coupled with a unique signal function. The reason is that thanks to our restricted domain of signal functions, different media strategies that share the *same* narrative are Blackwell-ordered, and consumers will never strictly prefer the Blackwell-dominated ones among them. The empowering narrative, which caters to low-cost consumers, is accompanied by a biased signal. The true narrative, which caters to mid-cost consumers, is accompanied by a less biased signal (a fully unbiased one if the first segment is empty). The narrative that omits a as a cause of y (without loss of optimality, there is at most one) is accompanied by an arbitrary signal function, and induces the action $a = 0$ with certainty. Thus, we have a proliferation of narratives, which lead to polarized beliefs and behavior.

The possible emergence of the upper market segment, which adopts a narrative that omits a as a cause of y , is a consequence of the equilibrium data externality. As we saw in Section 3.1, these narratives never prevail in the homogenous-consumers case. In contrast, under consumer heterogeneity, high-cost types can “free ride” on low-cost types, who sometimes play $a = 1$ and thus generate a positive anticipatory utility for any consumer who adopts N^t or $N^\emptyset$. These narratives are self-serving for high-cost types because they rationalize shirking.

The equilibrium data externality can chip away at the middle market segment of consumers who adopt N^* . To see why, imagine that the denial and fatalistic narratives are added to a menu that originally excludes them. As upper-middle cost types are attracted to one of these narratives, the fraction of consumers who play $a = 1$ at $t = 0$ among those who ever play $a = 1$ increases (because low-cost consumer types receive a more biased signal than mid-cost types). Since $f_0 > f_1$, this makes the empowering narrative *even more attractive* for lower-middle cost types. As such consumers switch to the empowering narrative, the overall frequency of consumers who play $a = 1$ at $t = 0$ increases, which makes the denial narrative *even more attractive* for upper-middle cost types. Thus, the effects of the equilibrium data externality on the two sides of the middle market segment are *mutually reinforcing*. When the negative correlation between t and y is sufficiently strong, this force can all but *eliminate* the middle segment, as the following result demonstrates.

Proposition 6 *If f_0 is sufficiently large, then the platform maximizes its objective with the menu $\{(1, N^a), (1, N^\emptyset)\}$, such that $c^* = c^{**} < 1$ (these are the cutoffs defined in Proposition 5). Consumers with $c < c^*$ choose $(1, N^a)$ and always play $a = 1$; whereas consumers with $c > c^*$ choose $(1, N^\emptyset)$ and always play $a = 0$.*

Thus, when f_0 is large, news media never provide *any* information to *any* consumer, and only false narratives prevail. Consumer behavior is highly polarized: High- c consumers always play $a = 0$ whereas low- c consumers always play $a = 1$. Both market segments are non-empty. What generates this polarization is the *different narratives* that the two consumer segments adopt: Low- c consumers opt for the empowering narrative while high- c consumers opt for the denial narrative.

Proposition 6 does not rely on any assumptions about the distribution of consumer types. For specific distributions, the same pattern arises even when the gap between f_0 and f_1 is moderate, as the following result illustrates.

Claim 2 *Let $f_0 = 1$, $f_1 = \frac{1}{2}$, and $c \sim U[0, 1]$. There is an optimal menu as described in Proposition 6, with $c^* = \frac{3}{11}$.*

This claim, as well as Proposition 6, provide examples in which the platform's optimal menu involves no information provision. This is not a general feature. Indeed, for other specifications the optimal menu can involve full information transmission. For illustration, suppose that $f_0 < 2f_1$, and that G is almost entirely concentrated around some $c \in (f_0 - f_1, f_1)$. This is a small perturbation of a homogenous-population model in which $(0, N^*)$ is the optimal media strategy. It can be shown that in the perturbed case, $c^* = 0$ — i.e., N^a is not on the menu. The reason is that there are too few low- c types to make N^a attractive (they would find N^a superior to N^* only when there are enough consumers who adopt N^a and play $a = 1$ at $t = 0$). As a result, $(0, N^*)$ is the only media strategy on the menu that ever generates $a = 1$. Consumer types below some cutoff $\bar{c}$ select $(0, N^*)$. Types above $\bar{c}$ adopt the narrative N^t (which, without loss of generality, is also accompanied by fully informative signals). Thus, Proposition 5 does not imply a clear-cut conclusion regarding the prevalence of media bias in the American-dream example: That will depend on the consumer type distribution.

A feature that *is* robust in the American-dream example is the proliferation of narratives that generate polarized beliefs. By comparison, if the platform were restricted to offering the true narrative, it would provide full information to all consumer types, thus producing homogenous beliefs. Consumers' different preferences would still lead to different behavior, but there would be no belief polarization. This is not the case under the optimal menu, even when all consumers receive full information.

This distinction has objective welfare implications. Under the true narrative with full information, consumer types above some threshold $\hat{c}$ play $a = 0$ and earn zero anticipatory utility. Under the optimal menu with full information provision, these same types act the same but earn positive anticipatory utility because they adopt the false fatalistic narrative (which generates positive anticipatory utility by free riding on the low-cost types' adoption of the true narrative and action choice $a \equiv t$). As a result, consumer types slightly below $\hat{c}$ are attracted to the false narrative and also play $a = 0$. The cutoff type $\bar{c}$ induced by the optimal menu, which separates consumers who play $a \equiv t$ from those who always play $a = 0$, lies below $\hat{c}$. In terms of objective welfare, the types $c \in (\bar{c}, \hat{c})$ are worse off under the optimal menu relative to the rational-expectations benchmark, because they play $a = 0$ at $t = 1$ even though $f_1 > c$.

4.3 Summary

In the representative-consumer version of the American-dream example, the optimal media strategy was a bundle consisting of the empowering narrative and optimistically biased information. In contrast, consumer preference heterogeneity gives rise to a *proliferation of narratives*, some (or all) of which are false. High-cost consumers adopt a narrative that gives them license not to take the costly action. Low-cost types, who sometimes do take it, exert a positive externality on high-cost types, who now enjoy a positive anticipatory utility. In terms of the phenomena highlighted by Iyenger (1990) and Loury (2020), consumers hold different worldviews regarding the role of personal agency and external factors in determining their material success, because they self-servingly take these worldviews from different news outlet.

Thus, the heterogenous-consumer version of our example predicts a coexistence of *conflicting narratives* that cater to different segments of the consumer

population. Depending on the consumer type distribution, media bias can be extreme, which exacerbates narrative-driven belief polarization and helps crowding out the true narrative.

5 A Competitive Media Market

In this section we continue to study the heterogeneous consumer population, under the same restrictions on primitives as in Section 4.1, but with a different market structure: We remove the monopolistic gatekeeper and analyze a “perfectly competitive” news-media market. The key difference from the monopoly case is that competitive media providers do *not* internalize the equilibrium data externality that played a key role in the previous section.

Definition 2 (Competitive equilibrium) *A profile $(I_\theta, N_\theta, \mathbf{a}_\theta)_{\theta \in \Theta}$ is a competitive equilibrium if for every $\theta \in \Theta$, $(I_\theta, N_\theta, \mathbf{a}_\theta)$ maximizes U_θ over all possible triples $(I, N, \mathbf{a})$; where U_θ is defined as in (17) and calculated taking as given the aggregate distribution $(p(a | t)_{a,t})$ that is induced by $(\mathbf{a}_\theta)_{\theta \in \Theta}$.*

Unlike the monopoly case, here each media strategy targets a consumer type and maximizes his anticipatory utility. As in Section 4.1 — and under the same restrictions on primitives — we can describe some general features of competitive equilibrium.

Proposition 7 *(i) A competitive equilibrium includes at least one false narrative. (ii) If $(I_\theta, N_\theta, \mathbf{a}_\theta)_{\theta \in \Theta}$ is a competitive equilibrium, then so is $(I^*, N_\theta, \mathbf{a}_\theta)_{\theta \in \Theta}$, where I^* is a perfectly informative signal function.*

Part (i) of this result establishes that market competition does not eliminate false narratives. The culprit, as before, is the equilibrium data externality: High-cost consumer types can adopt the fatalistic or denial narratives, and enjoy an anticipatory payoff thanks to low-cost types’ choice of narratives that induce costly actions.

In contrast, part (ii) means that competition does eliminate media bias. The reason is that the maximization of type θ ’s anticipatory utility takes p as given, without internalizing the effect of the behavior induced by (I_θ, N_θ) on p_N . Thus, for any given N , the maximization of U_θ with respect to I and

$\mathbf{a}$ is reduced to a conventional problem of maximizing a convex function of posterior beliefs. It follows that a fully informative signal maximizes U_θ (as in the rational-expectations benchmark). It is the unique maximizer if it induces $a \equiv t$ as the unique best-reply.

We now present a complete characterization of competitive equilibria for our “American dream” specification. As before, consumer types are identified by their cost parameter c . Unlike Section 4.2, here we need *not* restrict the set of feasible signal functions, except for the purely expositional restriction to binary signals that take the values 0 or 1.

Proposition 8 *There is an essentially unique competitive equilibrium in the “American dream” setting. Specifically, there is $\bar{c} \in (0, 1)$ uniquely given by the equation $\bar{c} = f_1(1 - G(\bar{c}))$, such that $N_c = N^*$ for every $c < \bar{c}$; and $N_c = N^t$ for every $c > \bar{c}$.*

By essential uniqueness, we mean that there could be other menus of media strategies that implement the same profile of beliefs and actions. When a consumer chooses N^t , the exact signal function is irrelevant for his beliefs and actions. Also, we could replace N^* with N^a , and consumers’ beliefs would be identical.

The competitive-equilibrium menu has the same structure as in the monopoly case when it leads to $c^* = 0$ (indeed, the cutoff $\bar{c}$ is the same as the one mentioned toward the end of Section 4.2). Consumers who always play $a = 0$ select the narrative N^t . The reason is that $N^\emptyset$ beats N^t when $\mathbb{E}(y) > \mathbb{E}(y \mid t = 1)$ — which can only happen if $p(a = 1 \mid t = 0) > 0$. However, this is never the case in competitive equilibrium, because consumers who sometimes play $a = 1$ have fully informed, correct beliefs and therefore play $a \equiv t$.

Revisit our numerical example from Section 4.2, where $f_0 = 1$, $f_1 = \frac{1}{2}$, and $c \sim U[0, 1]$. In this case, we have $\bar{c} = \frac{1}{3}$, which is above the cutoff $c^* = \frac{3}{11}$ of the monopoly case. Thus, under this specification, competition weakly improves informativeness and objective welfare for *all* consumer types, and strictly so for the types in $(\frac{3}{11}, \frac{1}{3})$. However, the following claim shows that this feature is not robust.

Claim 3 *There exist (f_0, f_1, G) for which a positive measure of consumer types in the “American dream” example are objectively better off under monopoly than in competitive equilibrium.*

The intuition for why market competition need not be objectively beneficial for all consumer types is as follows. In competitive equilibrium, the value of f_0 is irrelevant for the structure of equilibrium and consumers’ objective welfare. In contrast, when f_0 is large enough, it encourages a monopolistic platform to introduce N^a into the menu, coupled with maximally biased information, which leads many consumers to adopt this narrative and play $a = 1$ when $t = 0$. When f_1 is low enough, some of these same consumers commit the opposite error — always playing $a = 0$ — in competitive equilibrium. For suitable primitives, the latter error is objectively worse than the former.

6 Discussion

This paper continues a research program on the role of causal narratives in economic and political interactions. Eliaz and Spiegler (2020) presented a modeling framework that formalizes narratives as directed acyclic graphs (building on Spiegler (2016)), where agents’ adoption of narratives is based on the anticipatory utility they generate.¹⁶ The present paper merges this framework with the traditional information-design agenda, showing how narrative and information provision can be viewed as two dimensions of a broader belief-manipulation strategy. Indeed, a key message of this paper is that there are synergies between these two dimensions. Using this framework, the paper offers a new approach to modeling the market for news, focusing on the role of media as suppliers of narratives. Two additional methodological innovations are the screening problem that arises under consumer heterogeneity (where the narrative instrument brings a new kind of externality into the information-design problem), and our conception of a competitive media market.

We conclude the paper with a few comments on our modeling approach, and a discussion of related literature (especially on the phenomenon of demand-driven media bias).

Non-instrumental demand for news

Our model assumes that consumers’ demand for news is entirely non-instrumental, rather than involving a mixture of instrumental and non-instrumental motives.

¹⁶Eliaz and Spiegler (2020) and Eliaz et al. (2025) apply this framework to political competition. Levy et al. (2022b) offer a related approach, focusing on dynamics and multi-dimensional policies. Recent empirical and experimental approaches to causal economic narratives include Ash et al. (2021), Andre et al. (2022), Charles and Kendall (2022), Macaulay and Song (2023) and Ambuehl and Thysen (2023).

We make this assumption for several reasons. First, it obviously enables a sharper analysis. Second, as we saw, the distinction between instrumental and non-instrumental demand is irrelevant when the media is restricted to the true narrative. Thus, demand for news in our model already has a heavy dose of rationality. Finally, we believe that a model in which consumers evaluate media strategies according to a weighted average of material and anticipatory utility (as in Brunnermeier and Parker (2005)) would deliver similar qualitative results while making the analysis less transparent.¹⁷

The media’s anticipation of equilibrium effects

In solving its problem, the media takes into account the consumer’s equilibrium response to the media strategy. This naturally raises the question of whether the media knowingly anticipates equilibrium effects. One interpretation is that the media is not aware of them a priori. Instead, it reacts to past data about consumer engagement, possibly using algorithmic learning. The equilibrium effects that shape consumers’ media engagement will be reflected in the learning process. At any rate, our methodology is in essence the same as in the multi-agent information design literature (e.g., Bergemann and Morris (2019)), which evaluates information structures according to agents’ equilibrium responses. And as in that literature, our media can select among equilibria when its strategy induces multiple equilibria. The key difference is that the equilibrium notion in our model deviates from rational expectations.

6.1 Related Literature

The media-studies literature has examined the effects of narrative employment by the media. Iyengar (1990) provides experimental evidence on how two different narratives used by the media when reporting on poverty affect subjects’ perceptions. The “thematic narrative” describes poverty as part of a national trend or against the background of some public policy. The “episodic narrative” focuses on the poverty of specific individuals or families. Subjects exposed to the thematic narrative tend to attribute poverty to government or society at large, while those exposed to the episodic narrative attribute poverty to the poor. More recently, Schmidt (2025) investigates how economic journalists

¹⁷A previous version of this paper (Eliaz and Spiegler (2024)) analyzes a variant of our model with a mixed consumer population, some of whom know the true model N^* and therefore have purely instrumental demand for news.

explain post-Covid inflation causes and persistence compared to professional economists: The evidence suggests that relative to professionals, journalists are more likely to attribute inflation to specific entities such as the ECB and corporations.

Within the economics literature, the phenomenon of media bias has been studied from various points of view: See Prat and Strömberg (2013) and Gentzkow et al. (2015) for comprehensive reviews. Our paper contributes to a theoretical strand in this literature that tries to explain media bias as a demand-based phenomenon arising from consumers' non-instrumental demand for information. The basic idea that consumers derive intrinsic utility from beliefs or from the news they consume is related to findings in disciplines outside economics. For example, a meta-study by Hart et al. (2009) finds that when participants are faced with a choice between information that supports their prior beliefs and information that may challenge it, they exhibit a preference for the former. Within the context of news media, Van der Meer et al. (2020) find evidence that participants are more likely to view news that confirm their prior beliefs than news that oppose them.

Mullainathan and Shleifer (2005) attempt to model this phenomenon. They formalize both states of Nature and news as points along an interval. When a consumer confronts news, he incurs a cost that increases in the distance between the news and the mean of his prior belief. Media's strategic choices are thus reduced to a Hotelling-style model, where the consumer's psychological cost is analogous to a transportation cost in the standard Hotelling model.

Gentzkow et al. (2015) present a model in which consumers' utility has two additively separable components. The first component is a standard material expected-utility term that employs the consumer's posterior beliefs, which are obtained conventionally via Bayesian updating. This component treats beliefs in the usual instrumental manner. The second component is a function of the consumer's prior belief and the distribution of signals, such that if the prior leans in the direction of one state, then the function increases in the frequency of the signal whose label coincides with that state's label. This captures the idea that people like consuming news that support their prior beliefs. Note that this non-standard component does not reflect any belief updating. In particular, if the media always sends a signal that coincides with the state the consumer deems more likely (such that effectively the signal is entirely uninformative), the non-instrumental term reaches its maximal possible level given the consumer's

prior belief.

Thus, both Mullainathan and Shleifer (2005) and Gentzkow et al. (2015) assume that the hedonic effect of news is orthogonal to belief updating. We view this dissociation as a limitation of the existing models (that we are aware of). We believe that even when people appear to behave as if they dislike a clash between news and their *prior* beliefs, this reflects their prediction that the news will generate undesirable *posterior* beliefs. Therefore, incorporating this forward-looking motive into models of non-instrumental demand for information has value.

Against this background, our model introduces two innovations. To our knowledge, it is the first model of news media as suppliers of narratives in addition to information. It also appears to be among the first models (along with Herrera and Sethi (in press)) in which the hedonic aspect of media consumers' beliefs is fully integrated with Bayesian updating. Eliaz and Spiegel (2006) is a precedent for this aspect of our model. In that paper, we studied demand for information — represented by prior-dependent preferences over Blackwell experiments — driven by maximization of expected utility from (correctly specified) Bayesian posterior beliefs. Since that model allows for non-convex utility from beliefs, it accommodates demand for information that is non-increasing in Blackwell informativeness. Lipnowski and Mathevet (2018) examine optimal information provision for agents with such preferences.

The assumption that news consumers seek hopeful narratives may appear to be at odds with the common notion that news media exhibit a “negativity bias” (e.g., see Robertson et al. (2023)). We believe, however, that the two ideas are orthogonal. First, what often attracts consumers to negative news is their element of drama or sensationalism (e.g., a collapsing bridge). Second, when we measure negativity of a news piece by the prevalence of “negative words”, we may fail to capture its message that bad outcomes are a consequence of wrong decisions (which a false narrative like N^a in our model conveys). Third, it is not clear that media consumers invariably regard bad things that happen to *other people* as bad news.

There is also a sense in which negativity bias complements our perspective. When news exaggerate how bad things are, they only strengthen the apparent solution offered by a false empowering narrative. Therefore, news media may have an incentive to magnify today's problems in order to amplify the illusory gain provided by the false narrative (e.g., think of a cable news channel that

exaggerates a national-security risk and offers military intervention as a means for removing it). This effect can be captured by a modification of our model, in which news media maximize consumers’ anticipatory utility *relative* to the perceived status quo, rather than absolute anticipatory utility.

Our assumption of Bayesian updating rules out non-Bayesian responses to information due to motivated reasoning. Taber and Lodge (2006) show that when subjects are confronted with information that questions their prior beliefs, they try to discredit it. Thaler (2023) studies experimentally the supply of information to agents whose belief updating exhibits motivated reasoning. Brunnermeier and Parker (2005) offer a model of motivated reasoning, in which decision-makers distort beliefs to maximize a combination of anticipatory and material utility. However, Spiegler (2008) shows that augmenting this model with an information-acquisition stage does not produce a taste for biased or noisy information.

The idea that misspecified models can be used to manipulate agents’ beliefs has been studied in other contexts. Eliaz et al. (2021a) analyze a cheap-talk model in which the sender provides not only information but also statistical data (or, equivalently, a model) that enables the receiver to interpret the information. Eliaz et al. (2021b) characterize the maximal distortion of perceived correlation between two variables that a causal model can generate in Gaussian environments. Schwartzstein and Sunderam (2021) and Aina (2023) study persuasion problems in which the sender proposes models, formalized as likelihood functions, and the receiver chooses among them according to how well they fit historical data. Szeidl and Szucs (2024) present a model in which the sender can use “propaganda” to alter the receiver’s perception of the sender’s motives. Finally, our paper is related to a small literature on strategic communication with agents whose inference from signals departs from the standard Bayesian, rational-expectations model (e.g., Hagenbach and Koessler (2020), Levy et al. (2022a), de Clippel and Zhang (2022)).

References

- [1] Aina, C. (2023). Tailored stories. Working paper.
- [2] Alesina, A., & Angeletos, G. (2005). Fairness and redistribution. *American economic review* 95, 960-980.

- [3] Ambuehl S., & Thysen, H. (2023). Competing causal interpretations: An experimental study. Working paper.
- [4] Andre, P., Roth, C., Haaland, I., & Wohlfart, J. (2021). Inflation narratives. Working paper.
- [5] Ash, E., Gauthier, G., & Widmer, P. (2021). Text semantics capture political and economic narratives. Center for Law & Economics Working Paper Series, 2108-01720.
- [6] Bergemann, D., & Morris, S. (2016). Bayes correlated equilibrium and the comparison of information structures in games. *Theoretical Economics* 11(2), 487-522.
- [7] Bergemann, D., & Morris, S. (2019). Information design: A unified perspective. *Journal of Economic Literature*, 57(1), 44-95.
- [8] Brunnermeier, M., & Parker, J. (2005). Optimal expectations. *American Economic Review* 95, 1092-1118.
- [9] Kendall, C. W., & Charles, C. (2022). Causal narratives. National Bureau of Economic Research, WP no. 30346.
- [10] Chopra, F., Haaland, I., & Roth, C. (2023). The demand For news: Accuracy concerns versus belief confirmation motives. NHH Dept. of Economics Discussion Paper.
- [11] de Castro, L. (2009). Affiliation and dependence in economic models (No. 1479). Discussion paper.
- [12] de Clippel, G., & Zhang, X. (2022). Non-bayesian persuasion. *Journal of Political Economy* 130(10), 2594-2642.
- [13] Eliaz, K., Galperti, S., and Spiegler, R. (2025). False narratives and political mobilization. *Journal of the European Economic Association* 23(3), 983-1027.
- [14] Eliaz, K., & Spiegler, R. (2006). Can anticipatory feelings explain anomalous choices of information sources? *Games and Economic Behavior* 56(1), 87-104.

- [15] Eliaz, K., & Spiegler, R. (2020). A model of competing narratives. *American Economic Review* 110(12), 3786-3816.
- [16] Eliaz, K., & Spiegler, R. (2024). News media as suppliers of narratives (and information). arXiv preprint arXiv:2403.09155.
- [17] Eliaz, K., Spiegler, R., & Thysen, H. C. (2021a). Strategic interpretations. *Journal of Economic Theory* 192, 105192.
- [18] Eliaz, K., Spiegler, R., and Weiss, Y. (2021b). Cheating with models. *American Economic Review: Insights* 3(4), 417-434.
- [19] Esponda, I., & Pouzo, D. (2016). Berk–Nash equilibrium: A framework for modeling agents with misspecified models. *Econometrica* 84(3), 1093-1130.
- [20] Gentzkow, M., & Shapiro, J. M. (2010). What drives media slant? Evidence from US daily newspapers. *Econometrica* 78(1), 35-71.
- [21] Gentzkow, M., Shapiro, J. M., and Stone, D. F. (2015). Media bias in the marketplace: Theory. In *Handbook of media economics* (Vol. 1, pp. 623-645). North-Holland.
- [22] Hart, W., Albarracin, D., Eagly, A. H., Brechan, I., Lindberg, M. J., & Merrill, L. (2009). Feeling validated versus being correct: a meta-analysis of selective exposure to information. *Psychological bulletin* 135(4), 555-588.
- [23] Herrera, H., & R. Sethi (in press). Identity-based elections. *Journal of Political Economy: Microeconomics*.
- [24] Iyengar, S. (1990). Framing responsibility for political issues: The case of poverty. *Political Behavior* 12(1), 19-40.
- [25] Kamenica, E., & Gentzkow, M. (2011) Bayesian persuasion. *American Economic Review* 101(6), 2590-2615.
- [26] Levy, G., Morreno de Barreda, I., & Razin, R. (2022a). Persuasion with correlation neglect: A full manipulation result. *American Economic Review: Insights* 4(1), 123-138.
- [27] Levy, G., Razin, R., & Young, A. (2022b). Misspecified politics and the recurrence of populism. *American Economic Review* 112(3), 928-962.

- [28] Lipnowski, E., & Mathevet, L. (2018). Disclosure to a psychological audience. *American Economic Journal: Microeconomics* 10(4), 67-93.
- [29] Loury, G. C. (2020). The bias narrative vs. the development narrative. City Journal, December 2020, <https://www.city-journal.org/article/the-bias-narrative-v-the-development-narrative>.
- [30] Macaulay, A., & Song, W. (2023). News media, inflation, and sentiment. *AEA Papers and Proceedings* 113, 172-176.
- [31] Meyer, M. & Strulovici, B. (2012). Increasing interdependence of multivariate distributions. *Journal of Economic Theory* 147(4), 1460-1489.
- [32] Milgrom, P. & Weber, R. (1982). A theory of auctions and competitive bidding. *Econometrica* 50, 1089-1122.
- [33] Mullainathan, S., & Shleifer, A. (2005). The market for news. *American Economic Review* 95(4), 1031-1053.
- [34] Müller, A. and Stoyan, D. (2002). Comparison methods for stochastic models and risks. John Wiley & Sons, Chichester.
- [35] Pearl, J. (2009). *Causality*. Cambridge university press.
- [36] Piketty, T. (1995). Social mobility and redistributive politics. *The Quarterly journal of economics* 110, 551-584.
- [37] Prat, A., & Strömberg, D. (2013). The political economy of mass media. *Advances in economics and econometrics*, 2, 135.
- [38] Robertson, C. E., Pröllochs, N., Schwarzenegger, K., Pärnamets, P., Van Bavel, J. J., & Feuerriegel, S. (2023). Negativity drives online news consumption. *Nature Human Behaviour* 7(5), 812-822.
- [39] Rutenberg, J. (2003). A nation at war: The news media; cable's war coverage suggests a new "Fox effect" on television journalism. *The New York Times*. April 16.
- [40] Schmidt, T. (2025). Narrating inflation: How German economic journalists explain post-covid price rises (No. 14). DoCMA Working Paper.

- [41] Schwartzstein, J., & Sunderam, A. (2021). Using models to persuade. *American Economic Review* 111(1), 276-323.
- [42] Spiegel, R. (2008). On two points of view regarding revealed preferences and behavioral economics, pp. 95-115 in *Perspectives on the Future of Economics: Positive and Normative Foundations. Handbook of Economic Methodology*, Vol. 1, ed. A. Caplin and A. Schotter, Oxford University Press.
- [43] Spiegel, R. (2016). Bayesian networks and boundedly rational expectations. *Quarterly Journal of Economics* 131, 1243-1290.
- [44] Spiegel, R. (2020). Behavioral implications of causal misperceptions. *Annual Review of Economics* 12, 81-106.
- [45] Szeidl, A., & Szucs, F. (2024), The political economy of alternative realities. Working paper.
- [46] Taber, C., & Lodge, M. (2006). Motivated skepticism in the evaluation of political beliefs. *American Journal of Political Science* 50(3), 755–69.
- [47] Thaler, M. (2023). The supply of motivated beliefs. arXiv preprint arXiv:2111.06062.
- [48] Van der Meer, T.G.L.A., Hamelers, M., & Kroon, A.C. (2020). Crafting our own biased media diets: The effects of confirmation, source, and negativity bias on selective attendance to online news, *Mass Communication and Society* 23(6), 937-967.

Appendix: Proofs

Claim 1

Rewrite (13) as

$$\frac{1}{2} \left[\frac{(q_1)^2}{q_1 + q_0} f_1 + \frac{q_1 q_0}{q_1 + q_0} f_0 - (q_1 + q_0) c \right]$$

Suppose $q_0 > q_1$. Then, the expression will evidently go up if we swap q_0 and q_1 . Therefore, it is optimal to set $q_1 \geq q_0$. Now rewrite the expression as

$$\frac{q_1}{2} \left[f_0 - \frac{1}{1 + \frac{q_0}{q_1}} (f_0 - f_1) - \left(1 + \frac{q_0}{q_1}\right) c \right]$$

The terms inside the square brackets only depends on the ratio q_0/q_1 , while the term outside them increases in q_1 . It follows that $q_1 = 1 \geq q_0$ in optimum. ■

Proposition 1

Fix $\varepsilon > 0$, such that $p(y | t, a)$ never involves conditioning on a null event. Denote $\min_a C(a) = c^{**}$. Under the narratives N^t and $N^\emptyset$, the consumer believes that a has no causal effect on y . Therefore, for every s , the consumer's deliberate strategy σ will only assign positive probability over actions that minimize C . Moreover, without loss of optimality, I is completely uninformative, such that $a \perp t$ under p — i.e., $p(a | t) = p(a)$ for every a, t . In particular, we can treat both σ and σ^* as action distributions that are independent of t . Denote $c^* = \sum_a \sigma^*(a)C(a)$.

Under N^t (coupled with no information), any equilibrium for ε induces the following expression for $U(I, N^t)$:

$$\begin{aligned} & \sum_t p(t) \sum_y p(y | t) v(t, y) - (1 - \varepsilon)c^{**} - \varepsilon c^* \\ = & \sum_t p(t) \sum_{a'} [(1 - \varepsilon)\sigma(a') + \varepsilon\sigma^*(a')] p(y | t, a') v(t, y) - (1 - \varepsilon)c^{**} - \varepsilon c^* \end{aligned}$$

Since $C(a') = c^{**}$ whenever $\sigma(a') > 0$, $U(I, N^t)$ can be rewritten as

$$\begin{aligned} & (1 - \varepsilon) \sum_t p(t) \sum_{a'} \sigma(a') \left[\sum_y p(y | t, a') v(t, y) - C(a') \right] \\ + & \varepsilon \sum_t p(t) \sum_{a'} \sigma^*(a') \sum_y p(y | t, a') v(t, y) - \varepsilon c^* \end{aligned} \quad (19)$$

Note that the second term in this expression involves only exogenous quantities.

By definition,

$$\sum_{a'} \sigma(a') \left[\sum_y p(y | t, a') v(t, y) - C(a') \right] \leq \max_a \left[\sum_y p(y | t, a) v(t, y) - C(a) \right]$$

If we replace the L.H.S. of this inequality with its R.H.S. in (19), we obtain $U(I, N^*)$ for ε . It follows that N^t cannot be part of a media strategy that outperforms N^* coupled with full information.

Now turn to $N^\emptyset$, coupled with a fully uninformative I (such that $a \perp t$ under p). Any equilibrium induces

$$\begin{aligned} U(I, N^\emptyset) &= \sum_t p(t) \sum_y p(y) v(t, y) - (1 - \varepsilon)c^{**} - \varepsilon c^* \\ &= \sum_t p(t) \sum_y (\sum_{a'} p(a') p(y | a')) v(t, y) - (1 - \varepsilon)c^{**} - \varepsilon c^* \end{aligned}$$

Since $C(a') = c^{**}$ whenever $\sigma(a') > 0$, $U(I, N^\emptyset)$ can be rewritten as

$$(1 - \varepsilon) \sum_t p(t) \sum_{a'} \sigma(a') \left[\sum_{t'} p(t') \sum_y p(y | t', a') v(t, y) - C(a') \right] \\ + \varepsilon \sum_t p(t) \sum_{a'} \sigma^{**}(a') \sum_{t'} p(t') \sum_y p(y | t', a') v(t, y) - \varepsilon c^{**} \quad (20)$$

As in the previous case, the second term in this expression involves only exogenous quantities.

By definition,

$$\sum_t p(t) \sum_{a'} \sigma(a') \left[\sum_y \sum_{t'} p(t') p(y | t', a') v(t, y) - C(a') \right] \\ \leq \max_a \sum_t p(t) \left[\sum_y \sum_{t'} p(t') p(y | t', a) v(t, y) - C(a) \right]$$

If we replace the L.H.S of this inequality with its R.H.S. in (20), we obtain $U(I, N^a)$ for ε (where I is fully uninformative). It follows that a media strategy that includes $N^\emptyset$ is weakly dominated by a media strategy that consists of N^a and a fully uninformative signal function.

We conclude that if a media strategy outperforms the rational-expectations benchmark, then there is an optimal media strategy that includes N^a . ■

Proposition 2

Assume the contrary — i.e., suppose there is a media strategy that induces the same $(p(a | t))_{t,a}$ as in the rational-expectations benchmark, yet outperforms it.

We first show that N^a is the only narrative that can be part of the strategy. The proof of Proposition 1 showed that N^t can never outperform the benchmark; and N^* cannot do so by definition. Now consider $N^\emptyset$. Under this narrative, the consumer will assign probability one to $\arg \min_a c(a)$ for every t . By assumption, this is also the consumer's behavior under rational expectations, but this contradicts the definition of regularity. This leaves N^a as the only possible narrative.

By regularity, $p(a | t)$ assigns probability one to a distinct action for each t . Let $t(a)$ be the unique state for which a is played under p . Since $t = t(a)$ whenever $p(t, s, a) > 0$, it follows that $p(y | a) = p(y | t, a)$ for every (t, a) in the support of p . Consequently, $p_{N^a}(t, y | s, a) = p(t, y | s, a)$, and therefore, the consumer's anticipatory utility under p and N^a is equal to the rational-expectations benchmark, a contradiction. ■

Proposition 3

By Proposition 1, if any media strategy outperforms the rational-expectations benchmark, it involves the narrative N^a without loss of optimality. The ex-ante anticipatory utility induced by such a media strategy is

$$\sum_s p(s) \sum_a p(a | s) \left[\sum_{t,y} p(t | s) p(y | a) (v(t, y) - C(a)) \right] \quad (21)$$

By the revelation principle, there is no loss of optimality in assuming that $a \equiv s$. Therefore, (21) can be rewritten as follows:

$$\sum_a p(a) \sum_{t,y} p(t | a) p(y | a) v(t, y) - \sum_a p(a) C(a) \quad (22)$$

Now suppose that the media strategy switches from N^a to N^* , while keeping the signal function I intact. Suppose further that we keep the consumer's strategy ($a \equiv s$) intact, even if it now violates the obedience constraints as a result of the media's deviation. The consumer's resulting anticipatory utility can be written as

$$\sum_a p(a) \sum_{t,y} p(t, y | a) v(t, y) - \sum_a p(a) C(a)$$

It should be emphasized that this would *not* be a correct expression for an arbitrary, counterfactual consumer strategy. However, It is a correct expression given the *postulated* consumer strategy.

By assumption, v is supermodular and $p(t, y | a)$ satisfies affiliation for every a . Therefore, we can apply a standard result in the literature on stochastic orders (see Theorem 3.8.2 in Müller and Stoyan (2002, p. 108) and Theorem 3.2 in De Castro (2009)), and obtain

$$\sum_{t,y} p(t, y | a) v(t, y) \geq \sum_{t,y} p(t | a) p(y | a) v(t, y)$$

for every a .¹⁸

It follows that (22) is weakly above (21). By definition, (22) is weakly below

¹⁸Alternatively, we can apply Milgrom and Weber's (1982) demonstration that affiliation implies association, and Theorem 1 in Meyer and Strulovici's (2012), which implies that a bivariate distribution that satisfies association dominates its independent counterpart according to the supermodular order.

the maximal anticipatory utility that is attainable with the true narrative N^* . Therefore, (21) is weakly below the rational-expectations anticipatory-utility benchmark. We conclude that no media strategy can beat the benchmark. ■

Proposition 4

In this proof, we repeatedly make claims that directly refer to the $\varepsilon \rightarrow 0$ limit, rather than making statements about small $\varepsilon > 0$ perturbations and then taking limits. We take this shortcut merely to shorten the proof and avoid clutter. The shortcut is legitimate, since the perturbation's only role is to ensure that $p(y \mid t, a)$ is always well-defined.

We begin by showing that the optimal menu must include a false narrative. Assume the contrary — namely, that only the true narrative N^* is offered. Then, it is optimal to couple it with the fully informative signal function, denoted I^* . By assumption, the unique best-reply for type $\theta = 1$ is to always play a^* , whereas type $\theta = 0$ necessarily plays $a > a^*$ at some t . By continuity of C , nearby types act the same.

We now show that by adding to the menu a media strategy that consists of the narrative N^t and an arbitrary signal function I , the platform necessarily increases consumers' aggregate anticipatory utility. We do so via a series of three claims regarding consumers' response to the new menu.

Claim 1: There is a neighborhood of types near $\theta = 0$ that select the media strategy (I^*, N^*) .

Proof: Assume the contrary — i.e., all types choose the media strategy that includes N^t . Then, every type's best-reply is a^* , such that in the $\varepsilon \rightarrow 0$ limit, $p(y \mid t) = p(y \mid t, a^*)$. As a result, any type θ 's anticipatory utility in the $\varepsilon \rightarrow 0$ limit is

$$\sum_t p(t) \sum_y p(y \mid t, a^*) v(y, t) - C(a^*, \theta) \quad (23)$$

By assumption, for types θ sufficiently close to zero, (23) is strictly below

$$\sum_t p(t) \max_a \left[\sum_y p(y \mid t, a) v(y, t) - C(a, \theta) \right] \quad (24)$$

Therefore, these types will strictly prefer (I^*, N^*) , a contradiction. □

Claim 2: There is a neighborhood of types near $\theta = 1$ that select (I, N^t) . The anticipatory utility of every type that selects (I, N^t) is strictly higher than

under the original, singleton menu.

Proof: By assumption, (23) is equal in the $\varepsilon \rightarrow 0$ limit to the maximal payoff for types near $\theta = 1$ under (I^*, N^*) . By Claim 1, there is a positive measure of types who adopt (I^*, N^*) and play some $a > a^*$ at some t in response to the new menu. By revealed preferences,

$$\sum_y p(y | t, a) v(y, t) > \sum_y p(y | t, a^*) v(y, t)$$

for every such (a, t) , since $C(a, \theta) > C(a^*, \theta)$. It follows that under consumers' response to the new menu and the joint distribution p that is induced by it,

$$\begin{aligned} \sum_t p(t) \sum_y p(y | t) v(y, t) &= \sum_t p(t) \sum_y \sum_{a'} p(a' | t) p(y | t, a') v(y, t) \\ &> \sum_y p(y | t, a^*) v(y, t) \end{aligned}$$

It follows that any type θ that selects (I, N^t) earns an anticipatory payoff above (23). By revealed preferences, it is also above the maximal payoff he can derive from (I^*, N^*) .

Claim 3: When every consumer type optimally selects from the new menu, aggregate consumer welfare exceeds the rational-expectations benchmark.

Proof: By revealed preference, if a type favors (I, N^t) over (I^*, N^*) , his anticipatory utility is weakly higher, and therefore weakly above the rational-expectations benchmark. Furthermore, any type's evaluation of (I^*, N^*) is invariant to other types' choices. Finally, by Claim 2, there is a positive measure of consumer types who strictly prefer (I, N^t) to (I^*, N^*) . It follows that consumers' aggregate anticipatory utility exceeds the rational-expectations benchmark.

We have thus constructed a menu that improves on the rational-expectations benchmark, a contradiction.

We now show that without loss of optimality, the menu includes N^* or N^a . Assume the contrary — i.e., the only offered narratives are N^t or $N^\emptyset$. Since every consumer type has the same evaluation of these narratives and responds to them with the action a^* , we can assume without loss of generality that the menu consists of a single narrative $N \in \{N^t, N^\emptyset\}$, coupled with an arbitrary

signal function. Suppose first that $N = N^t$. Then, type θ 's anticipatory utility is arbitrarily close to (23). If the platform deviates to $\{(N^*, I^*)\}$, where I^* is fully informative, then type θ 's anticipatory utility is (24), which is a weak improvement.

Now suppose that $N = N^\emptyset$. Since I is fully uninformative, $a \perp t$ under p . It follows that type θ 's anticipatory utility is

$$\sum_t p(t) \sum_y \sum_{t'} p(t') p(y | t', a^*) v(t, y) - C(a^*, \theta)$$

If the platform deviates to $\{(N^a, I)\}$, then, type θ 's anticipatory utility is

$$\begin{aligned} & \sum_t p(t) \max_a \left[\sum_y p(y | a) v(t, y) - C(a, \theta) \right] \\ &= \sum_t p(t) \max_a \left[\sum_y \sum_{t'} p(t') p(y | t', a) v(t, y) - C(a, \theta) \right] \end{aligned}$$

where the equality makes use of the fact that $p(t' | a) \equiv p(t')$. Therefore, $\{(N^a, I)\}$ is weakly better than $\{(N^\emptyset, I)\}$. ■

Proposition 5

The proof proceeds stepwise.

Step 1: *Under any menu, every consumer type $c > 0$ chooses $a = 0$ with certainty in response to the signal $s = 0$. Moreover, every consumer type $c > 0$ who chooses the narratives N^t or $N^\emptyset$ plays $a = 0$ for every s . Without loss of optimality, the menu includes at most one of these two narratives.*

By our restriction on the set of signal functions, $\Pr(t = 0 | s = 0) = 1$ under any media strategy. Therefore, the consumer understands that $ty = 0$ with probability one, regardless of a . As a result, any consumer type with $c > 0$ will optimally choose $a = 0$, regardless of the narrative he adopts.

As we have shown before, a consumer type $c > 0$ who chooses N^t or $N^\emptyset$ believes that a has no effect on y , and therefore prefers not to incur the cost c of playing $a = 1$.

Finally, note that N^t and $N^\emptyset$ not only induce the same action choices, but the payoff they induce is *type-independent*. Clearly, if one of these narratives is not chosen by any consumer type, then removing it from the menu does not affect consumers' behavior or payoffs. If both narratives are chosen with

positive frequency, then they must confer the same anticipatory payoff. Since they also induce the same action choices, removing one of them from the menu will not affect any externality they might exert on other types, and therefore has no effect on the platform's objective. $\square$

Step 2: *Without loss of optimality, each narrative is coupled with a unique q .*

Assume the contrary — i.e., the optimal menu contains two pairs (q, N) and (q', N) with $q' < q$. This means that the signal function given by q' Blackwell-dominates the signal function given by q (recall that $\Pr(s = 1 \mid t = 1) = 1$ under both functions). Any consumer type c who compares the two pairs will weakly prefer (q', N) . The reason is that both pairs share the same narrative N , hence they both induce the same $p_N(y \mid t, a)$. This reduces the comparison between the pairs to a standard comparison between signal functions by an expected-utility maximizer. Thus, whenever $q' < q$, every consumer type weakly prefers (q', N) to (q, N) .

Suppose $N \in \{N^t, N^\emptyset\}$. Regardless of the signal function that accompanies it, the narrative N induces $a = 0$ as the unique best-reply for every $c > 0$. Its induced anticipatory payoff is also invariant to q . Therefore, removing (q, N) from the menu does not affect any consumer's behavior or anticipatory utility.

Now suppose $N \in \{N^*, N^a\}$. If (q', N) and (q, N) induce the same action choices, then they also induce the same payoff to the consumer who adopts them, and they exert the same externality on any other consumer type who adopts another narrative. Therefore, removing (q, N) from the menu does not affect any consumer's behavior or anticipatory utility. Finally, if (q', N) and (q, N) induce different choices, then the set of consumer types who are indifferent between (q', N) and the Blackwell-dominated pair (q, N) has measure zero. This means that removing (q, N) from the menu will have no effect on the aggregate anticipatory utility that the menu generates. $\square$

Step 3: *Under any optimal menu, a positive measure of consumer types play $a \equiv s$. All these types necessarily choose N^a or N^* .*

Assume that under an optimal menu, almost all consumer types always play $a = 0$. Then, regardless of the media strategy they choose, their anticipatory utility is zero. This is obviously the case for consumer types who choose N^* or N^a because these narratives induce the correct belief that $a = 0$ causes $y = 0$ with certainty. As to types who choose N^t , they estimate the conditional expectation $E(y \mid t) = p(a = 1 \mid t) \cdot f_t = 0$ for every t . Finally, types who choose

$N^\emptyset$ form the correct belief that $E(y) = 0$ (because $a = 0$ with probability one by assumption, and $E(y \mid a = 0) = 0$). It follows that all types earn zero anticipatory utility. However, if the platform offers the singleton menu consisting of the media strategy $(0, N^*)$, every type $c < f_1$ will earn $\frac{1}{2}(f_1 - c) > 0$, a contradiction.

It follows that under an optimal menu, a positive measure of consumer types sometimes play $a = 1$. By Step 1, this means that these consumer types play $a = s$ for every s and cannot choose N^t or $N^\emptyset$. $\square$

Step 4: Any optimal menu induces a cutoff $c^{**} \in (0, 1]$ such that every type $c > c^{**}$ chooses N^t or $N^\emptyset$, while every type $c < c^{**}$ chooses N^* or N^a . If $c^{**} < 1$, then every $c < c^{**}$ plays $a \equiv s$.

By Step 1, every type that selects N^t or $N^\emptyset$ always plays $a = 0$. Therefore, the anticipatory payoff that each of these narratives induces is type-independent. In contrast, for $N \in \{N^*, N^a\}$, the anticipatory utility from (I, N) is $\max\{0, p(t = 1)E_N(y \mid t = 1, a = 1) - cp_I(s = 1)\}$, which is decreasing in c . Therefore, if the menu offers a choice between a narrative in $\{N^t, N^\emptyset\}$ and a narrative in $\{N^*, N^a\}$, then if type c opts for the former, so does every $c' > c$. This establishes the cutoff c^{**} . By Step 1, every $c > c^{**}$ always plays $a = 0$. Therefore, by Step 3, $c^{**} > 0$.

Finally, suppose $c^{**} < 1$ — i.e., a narrative in $\{N^t, N^\emptyset\}$ is on the menu. Since we have established that a positive measure of types play $a \equiv s$ (and therefore they always play $a = 1$ at $t = 1$), the narratives N^t and $N^\emptyset$ both induce strictly positive anticipatory utility. By contrast, a type that selects N^* or N^a and always plays $a = 0$ earns zero anticipatory utility, hence he does not choose optimally from the menu, a contradiction. It follows that every $c < c^{**}$ plays $a \equiv s$. $\square$

Step 5: Suppose both N^a and N^* are selected by a positive measure of consumers. Then, there is $c^* \in (0, c^{**})$, such that every $c < c^*$ chooses (q^a, N^a) , whereas every $c \in (c^*, c^{**})$ chooses (q^*, N^*) . Furthermore, $q^a > q^*$.

First, we present expressions for the ex-ante anticipatory utility (derived from (18)) that a consumer type c obtains from the pairs (q, N^*) and (q, N^a) when he responds to them by playing $a \equiv s$:

$$U_c(q, N^*, 1) = \frac{1}{2}E(y \mid t = 1, a = 1) - \frac{1+q}{2}c = \frac{1}{2}f_1 - \frac{1+q}{2}c \quad (25)$$

and

$$\begin{aligned} U_c(q, N^a, 1) &= \frac{1}{2}E(y \mid a = 1) - \frac{1+q}{2}c \\ &= \frac{1}{2}[p(t = 0 \mid a = 1)f_0 + p(t = 1 \mid a = 1)f_1] - \frac{1+q}{2}c \end{aligned} \quad (26)$$

Since $f_0 > f_1$, it is immediate that $U_c(q, N^a, 1) \geq U_c(q, N^*, 1)$.

By the definition of c^{**} , every $c < c^{**}$ chooses (q^*, N^*) or (q^a, N^a) and plays $a \equiv s$. Note that if $q^a = q^* = 0$, then (25) and (26) are identical for every c . Thus, removing one of these pairs will not affect aggregate consumer utility.

Suppose that $q^a = q^* > 0$. Therefore, $p(t = 0 \mid a = 1) > 0$. Since $f_0 > f_1$, it follows from (25) and (26) that $U_c(q, N^a) > U_c(q, N^*)$ for every c . It follows that no consumer type $c < c^{**}$ will choose (q^*, N^*) , a contradiction.

Now suppose $q^* \neq q^a$. Then, q^* or q^a are strictly positive, hence $p(t = 0 \mid a = 1) > 0$. Since $f_0 > f_1$, it follows that $U_c(q^*, N^*) \geq U_c(q^a, N^a)$ only if $q^a > q^*$. Therefore, we must have $q^a > q^*$ in order to have a positive measure of consumers who choose (q^*, N^*) . Note that if $U_c(q^*, N^*) > U_c(q^a, N^a)$, then $U_{c'}(q^*, N^*) > U_{c'}(q^a, N^a)$ for every $c' > c$. It follows that if both (q^*, N^*) and (q^a, N^a) are chosen by a positive measure of consumers, then the set of types who choose (q^*, N^*) lies above the set of types who choose (q^a, N^a) .

Note that this ordering of the types that select N^a and N^* does *not* rely on the optimality of the menu, and therefore holds for any menu that induces positive measures of consumers who adopt each of these two narratives. $\square$

Step 6: *If $c^* = 0$, then $q^* = 0$, and types above c^{**} choose N^t . If $c^* = c^{**}$ and there is no alternative optimal menu that would induce $c^* = 0$, then $q^a > 0$.*

Suppose $c^* = 0$ and yet $q^* > 0$. Then, by (25)-(26), $U_c(q, N^a) > U_c(q, N^*)$ for every c . If the platform replaces (q^*, N^*) with (q^*, N^a) , it increases the payoff of every type that selected the original pair and now selects the new pair.

If types above c^{**} (who formerly always played $a = 0$) now switch to the new pair, then by revealed preferences their payoff increases. At the same time, they do not affect $p(t \mid a = 1)$, and therefore do not affect $U_c(q^*, N^a)$ for any c . Finally, they exert a positive externality on types above c^{**} who do not switch, because it increases $p(a = 1 \mid t)$ at any t . Thus, every type weakly benefits from the deviation, even when equilibrium responses are taken into account.

If types who formerly chose (q^*, N^*) now switch to a pair that induces always playing $a = 0$, then again by revealed preferences, their payoff increases.

However, this switch does not affect $p(t \mid a = 1)$, since all consumers who ever play $a = 1$ face the same signal function given by q^* . Therefore, the switch exerts no externality on types who now select (q^*, N^a) .

It follows that replacing (q^*, N^*) with (q^*, N^a) is profitable for the platform. Let us now calculate the anticipatory utility from N^t and $N^\emptyset$ for any type c when $c^* = 0$:

$$\begin{aligned} U_c(q, N^t) &= p(t = 1)p(a = 1 \mid t = 1)f_1 = \frac{1}{2}G(c^{**})f_1 \\ U_c(q, N^\emptyset) &= p(t = 1) \sum_t p(t)p(a = 1 \mid t)f_t = \frac{1}{4}G(c^{**})f_1 \end{aligned}$$

where the last equality follows from the fact that $q^* = 0$, such that $p(a = 1 \mid t = 0) = 0$. Therefore, $U_c(q, N^t) > U_c(q, N^\emptyset)$.

Now suppose $c^* = c^{**}$. If $q^a = 0$, then the pair $(0, N^a)$ is equivalent to $(0, N^*)$. Therefore, we can replicate the menu with an equivalent menu that induces $c^* = 0$ and sets $q^* = 0$. $\square$

This completes the proof. $\blacksquare$

Proposition 6

Throughout the proof, we take it for granted that f_0 is arbitrarily large.

Consider a monopolistic menu M in which $0 \leq c^* < c^{**} \leq 1$. Suppose first that $c^{**} = 1$. Then, type 1 earn an anticipatory utility of

$$\frac{1}{2}f_1 - \frac{1 + q^*}{2} < \frac{1}{2}f_1$$

Alternatively, suppose $c^{**} < 1$. Recall that all types $c \geq c^{**}$ earn the same anticipatory utility. Type c^{**} is indifferent between the menus that cater to the types on his two sides. Therefore, all types $c \geq c^{**}$ earn an anticipatory utility of

$$\frac{1}{2}f_1 - \left(\frac{1 + q^*}{2} \right) c^{**} < \frac{1}{2}f_1$$

Note that if $c^* > 0$, each consumer type $c \in (c^*, c^{**}]$ prefers (q^*, N^*) to (q^a, N^a) , i.e.,

$$\frac{1}{2}f_1 - \left(\frac{1 + q^*}{2} \right) c \geq \frac{1}{2}E(y \mid a = 1) - \frac{1 + q^a}{2}c$$

where $E(y \mid a = 1)$ is computed with respect to consumers' equilibrium strategies. This implies that $E(y \mid a = 1) \leq f_1 + c < f_1 + 1$.

It follows every consumer type earns an anticipatory utility that is bounded from above by $f_1 + 1$. As a result, this is also an upper bound on the aggregate anticipatory utility achieved by M . This is a very loose bound, but it suffices for our purposes.

Now suppose the platform deviates to the menu $M' = \{(1, N^a), (1, N)\}$, where $N \in \{N^\emptyset, N^t\}$. Let $\hat{c}$ be the threshold type such that types below $\hat{c}$ choose $(1, N^a)$, while types above it choose $(1, N)$. The expected anticipatory utility of any type $c \leq \hat{c}$ is $\frac{1}{4}f_0 + \frac{1}{4}f_1 - c$, while the expected anticipatory utility of every type $c \geq \hat{c}$ is $\frac{1}{4}f_0 + \frac{1}{4}f_1 - \hat{c}$ (because as before, all types $c \geq \hat{c}$ earn the same anticipatory utility). It follows that the aggregate anticipatory utility from M' is at least $\frac{1}{4}f_0 + \frac{1}{4}f_1 - 1 > f_1 + 1$. Hence, the optimal monopolistic menu will not include N^* .

We have thus established that the optimal menu must induce $c^* = c^{**}$. Therefore, we can restrict attention to two cases: (i) $M = \{(q^a, N^a)\}$ and $c^{**} = 1$; and (ii) $M = \{(q^a, N^a), (1, N)\}$ and $c^{**} < 1$, where $N \in \{N^t, N^\emptyset\}$ (such that type c^{**} is indifferent between the two menus). In both cases, we can write the aggregate anticipatory utility induced by the menu as

$$\int_0^{c^{**}} U_c(q^a, N^a) dG(c) + (1 - G(c^{**})) U_{c^{**}}(q^a, N^a)$$

where

$$U_c(q^a, N^a) = \frac{1}{2} \left[\frac{q^a}{1 + q^a} f_0 + \frac{1}{1 + q^a} f_1 \right] - \frac{1 + q^a}{2} c$$

Note that for every c ,

$$\frac{\partial}{\partial q} U_c(q, N^a) = \frac{f_0 - c(q + 1)^2 - f_1}{2(q + 1)^2} \geq \frac{f_0 - f_1 - 4}{2(q + 1)^2} > 0$$

Therefore, aggregate anticipatory utility is maximized at $q^a = 1$.

It remains to show that case (ii) prevails, and that $N = N^\emptyset$. Any consumer who adopts N^t or $N^\emptyset$ always plays $a = 0$. At the same time, the anticipatory utility that (q^a, N^a) induces only involves the conditional probability $p(y \mid a = 1)$. Therefore, $U_c(q^a, N^a)$ is invariant to the fraction of consumers who select N^t or $N^\emptyset$. It follows that if the platform adds both N^t or $N^\emptyset$ to the menu and allow consumers to self-select, this will maximize the platform's objective function.

To see why every consumer type ranks $N^\emptyset$ above N^t , let us calculate the

anticipatory utility that each of these narratives (coupled arbitrarily with $q = 1$):

$$U_c(1, N^t) = \frac{f_1 G(c^{**})}{2} < \frac{(f_0 + f_1) G(c^{**})}{4} = U_c(1, N^\emptyset)$$

Therefore, the menu $\{(1, N^a), (1, N^\emptyset)\}$ is optimal. Assume that $c^{**} = 1$. Then, $U_1(1, N^\emptyset) = (f_0 + f_1)/4$, which is clearly above $U_1(1, N^a)$. Therefore, $c^{**} < 1$. ■

Claim 2

The proof proceeds stepwise, taking the characterization in Proposition 5 as a starting point.

Step 1: $c^* = c^{**}$

Assume that $c^{**} > c^* \geq 0$. The payoffs induced by (q^*, N^*) and (q^a, N^a) at some c are

$$\begin{aligned} U_c(q^*, N^*) &= \frac{1}{4} - \frac{1 + q^*}{2} c \\ U_a(q^a, N^a) &= \frac{1}{4} [2 - p(t = 1 \mid a = 1)] - \frac{1 + q^a}{2} c \end{aligned}$$

In the proof of Step 5 of Proposition 5, we showed that $q^a > q^*$. Since $c \sim U[0, 1]$, we can write

$$p(t = 1 \mid a = 1) = \frac{c^{**}}{c^{**} + c^* q^a + (c^{**} - c^*) q^*} = \frac{c^{**}}{c^{**}(1 + q^*) + c^*(q^a - q^*)}$$

At c^* , the indifference between (q^*, N^*) and (q^a, N^a) can be written as follows:

$$\frac{1}{2} c^*(q^a - q^*) = \frac{1}{4} \left[1 - \frac{c^{**}}{c^{**}(1 + q^*) + c^*(q^a - q^*)} \right]$$

Observe that if we slightly raise c^* and lower q^a such that q^a is still above q^* and $c^*(q^a - q^*)$ remains unchanged, then the indifference condition continues to hold, as long as we keep c^{**} fixed. In this way, $p(t = 1 \mid a = 1)$ remains unchanged. This modified consumer action profile is an equilibrium and it is strictly profitable for the media. To see why, note first that c^{**} is unchanged because by construction, $p(a = 1)$ and $p(a = 1 \mid t = 1)$ are both unchanged, hence the payoff from N^t or $N^\emptyset$ is unchanged. Since the payoff from (q^*, N^*) is by definition invariant to $(p(a \mid t))$, the indifference at c^{**} continues to hold. Thus, the set of types who always play $a = 0$ and their utility are unaffected.

Now consider the infra-marginal types $c < c^*$. These types are now better off thanks to the decrease in q^a , and since $p(a = 1 \mid t = 1)$ is unchanged. The types who chose and continue to choose (q^*, N^*) are unaffected by definition. Therefore, the new equilibrium is an improvement, a contradiction.

It follows that we can restrict attention to menus M that include N^a and exclude N^* . By the same argument as at the end of the proof of Proposition 6, the menu must also include N^t or $N^\emptyset$. Therefore, there are only two cases to examine:

- (i) $M = \{(q^a, N^a), (q^t, N^t)\}$, all consumer types in $[0, c^*]$ choose (q^a, N^a) and play $a = s$, and all consumer types $c > c^*$ choose (q^t, N^t) and play $a = 0$; and
- (ii) $M = \{(q^a, N^a), (q^\emptyset, N^\emptyset)\}$, all consumer types in $[0, c^*]$ choose (q^a, N^a) and play $a = s$, and all consumer types $c > c^*$ choose $(q^\emptyset, N^\emptyset)$ and play $a = 0$. $\square$

Step 2: *Completing the characterization when M includes N^t*

Aggregate utility under $M = \{(q^a, N^a), (q^t, N^t)\}$ is

$$\int_0^{c^*} U_c(q^a, N^a) dc + \int_{c^*}^1 U_c(q^t, N^t) dc$$

where

$$U_c(q^a, N^a) = \frac{1}{4} \left[2 - \frac{1}{1 + q^a} \right] - \frac{1 + q^a}{2} c$$

and

$$\begin{aligned} U_c(q^t, N^t) &= p(ty = 1) = p(t = 1) \cdot p(y = 1 \mid t = 1) \\ &= \frac{1}{2} \cdot p(a = 1 \mid t = 1) \cdot \frac{1}{2}(2 - 1) = \frac{1}{4} c^* \end{aligned}$$

Thus, the objective function can be written as

$$\begin{aligned} &\int_0^{c^*} \left\{ \frac{1}{4} \left[2 - \frac{1}{1 + q^a} \right] - \frac{1 + q^a}{2} c \right\} dc + (1 - c^*) \cdot \frac{1}{4} c^* \\ &= c^* \cdot \frac{1}{4} \left[2 - \frac{1}{1 + q^a} \right] - \frac{1 + q^a}{2} \cdot \frac{1}{2} (c^*)^2 + (1 - c^*) \cdot \frac{1}{4} c^* \end{aligned}$$

The cutoff c^* satisfies

$$\frac{1}{4} \left[2 - \frac{1}{1 + q^a} \right] - \frac{1 + q^a}{2} c^* = \frac{1}{4} c^*$$

Plugging this equation into the objective function, we obtain

$$\frac{2q^a + 1}{(2q^a + 3)^2}$$

The optimal value of q^a is $\frac{1}{2}$, yielding an aggregate utility of $\frac{1}{8}$. $\square$

Step 3: *Completing the characterization when M includes $N^\emptyset$*

Aggregate utility under $M = \{(q^a, N^a), (q^\emptyset, N^\emptyset)\}$ is

$$\int_0^{c^*} U_c(q^a, N^a) dc + \int_{c^*}^1 U_c(q^\emptyset, N^\emptyset) dc$$

where

$$U_c(q^a, N^a) = \frac{1}{4} \left[2 - \frac{1}{1 + q^a} \right] - \frac{1 + q^a}{2} c$$

and

$$\begin{aligned} U_c(q^\emptyset, N^\emptyset) &= p(t = 1) \cdot p(y = 1) \\ &= p(t = 1) \cdot [p(t = 1) \cdot p(y = 1 \mid t = 1) + p(t = 0) \cdot p(y = 1 \mid t = 0)] \\ &= \frac{1}{2} \cdot \left[\frac{1}{2} \cdot p(a = 1 \mid t = 1) \cdot \frac{1}{2}(2 - 1) + \frac{1}{2} \cdot p(a = 1 \mid t = 0) \cdot \frac{1}{2}(2 - 0) \right] \\ &= \frac{1}{2} \cdot \left[\frac{1}{2} \cdot c^* \cdot \frac{1}{2}(2 - 1) + \frac{1}{2} \cdot c^* q^a \cdot \frac{1}{2}(2 - 0) \right] \\ &= \frac{c^*}{4} \left[\frac{1}{2} + q^a \right] \end{aligned}$$

Thus, the objective function can be written as

$$\begin{aligned} &\int_0^{c^*} \left\{ \frac{1}{4} \left[2 - \frac{1}{1 + q^a} \right] - \frac{1 + q^a}{2} c \right\} dc + (1 - c^*) \cdot \frac{c^*}{4} \left[\frac{1}{2} + q^a \right] \\ &= c^* \cdot \frac{1}{4} \left[2 - \frac{1}{1 + q^a} \right] - \frac{1 + q^a}{2} \cdot \frac{1}{2} (c^*)^2 + (1 - c^*) \cdot \frac{c^*}{4} \left[\frac{1}{2} + q^a \right] \end{aligned}$$

The cutoff c^* satisfies

$$\frac{1}{4} \left[2 - \frac{1}{1 + q^a} \right] - \frac{1 + q^a}{2} c^* = \frac{c^*}{4} \left[\frac{1}{2} + q^a \right]$$

Plugging this equation into the objective function, we obtain

$$\frac{3}{4} (2q^a + 1)^2 \frac{2q^a + 3}{(6q^a + 5)^2 (q^a + 1)}$$

This expression is monotonically increasing in q^a , hence the optimal value of q^a is 1, yielding an aggregate utility of approximately 0.139. $\square$

Since the menu characterized by Step 3 yields a higher payoff than the one characterized by Step 2, the optimal menu includes the denial narrative, and sets $q^a = 1$.

The only remaining case is $c^* = 0$. By Proposition 5, this means that all types $c < c^{**}$ choose the media strategy $(0, N^*)$. However, recall that this pair is equivalent to $(0, N^a)$ for all types. Steps 2 and 3 established that this pair is inferior to $(1, N^a)$. $\blacksquare$

Proposition 7

We provide a proof of part (i), since part (ii) is explained in the main text. Suppose there is a unique competitive equilibrium in which the only item is (I^*, N^*) . By part (ii), we can let I^* be fully informative. Let $p(a, t)$ be the joint distribution induced by this equilibrium. The expected anticipatory utility of type $\theta = 1$, given the joint distribution p induced by the competitive equilibrium is

$$\sum_t p(t) \sum_y p(y | t, a^*) v(y, t) - C(a^*, 1)$$

Let

$$a(t, \theta) = \arg \max_a \left[\sum_y p(y | t, a) v(y, t) - C(a, \theta) \right]$$

By assumption, types θ close to zero get an expected anticipatory utility of

$$\sum_t p(t) \left[\sum_y p(y | t, a(t, \theta)) v(y, t) - C(a(t, \theta), \theta) \right] > \sum_t p(t) \sum_y p(y | t, a^*) v(y, t) - C(a^*, \theta)$$

where $a > a^*$ in at least one state t . Since $C(a(t, \theta), \theta) \geq C(a^*, \theta)$, with a strict inequality for at least one t , it follows that for each such type θ ,

$$\sum_y p(y | t, a(t, \theta)) v(y, t) > \sum_t p(t) \sum_y p(y | t, a^*) v(y, t)$$

Let us now calculate the anticipatory utility that the media strategy (I^*, N^t)

generates for type $\theta = 1$:

$$\begin{aligned}
& \sum_t p(t) \sum_y p(y | t) v(y, t) - C(a^*, 1) \\
&= \sum_t p(t) \sum_y \sum_{a'} p(a' | t) p(y | t, a') v(y, t) - C(a^*, 1) \\
&> \sum_y p(y | t, a^*) v(y, t) - C(a^*, 1)
\end{aligned}$$

By continuity, this inequality holds for types in the neighborhood of $\theta = 1$. It follows that (I^*, N^t) . It follows that given p , (I^*, N^t) delivers higher anticipatory utility than (I^*, N^*) for consumer types near $\theta = 1$, contradicting the definition of competitive equilibrium. ■

Proposition 8

First, by Proposition 7, there is no loss of generality in letting I_c be fully informative for every c . Therefore, all consumers play $a = 0$ when $t = 0$. This means that $p(t = 0 | a = 1) = 0$, such that the formulas for U_c under N^* and N^a coincide. Thus, from now on, we assume for convenience that the only narrative that can induce $a = 1$ with positive probability is N^* . Let us denote by γ the fraction of consumers who play $a = 1$ when $t = 1$.

We now show that N^t weakly outperforms $N^\emptyset$ for every consumer type. The anticipatory utility under N^t is

$$p(t = 1)E(y | t = 1) = \frac{1}{2}\gamma f_1$$

The anticipatory utility under $N^\emptyset$ is

$$\begin{aligned}
p(t = 1)E(y) &= \frac{1}{2} \cdot \left[\frac{1}{2}E(y | t = 1) + \frac{1}{2}E(y | t = 0) \right] \\
&= \frac{1}{4}E(y | t = 1) \\
&= \frac{1}{4}\gamma f_1
\end{aligned}$$

Note that $E(y | t = 0) = 0$ because all consumers play $a = 0$ when $t = 0$.

Thus, the only narratives we need to consider are N^* and N^t . Moreover, we can assume that any consumer who adopts N^* will play $a = 1$ when $t = 1$, because otherwise he would get zero payoffs. A consumer of type c will prefer N^* if $\frac{1}{2}(f_1 - c) > \frac{1}{2}\gamma f_1$. Therefore, there is a unique cutoff $\bar{c}$, such that all $c < \bar{c}$ choose N^* and play $a = t$, while all $c > \bar{c}$ choose N^t and always play $a = 0$. Plugging $\gamma = G(\bar{c})$, we obtain the implicit equation for $\bar{c}$. ■

Claim 3

We assume throughout that f_1 is arbitrarily small while f_0 is arbitrarily large. By Proposition 8, a competitive equilibrium has the property that there exists a cutoff type $\bar{c} = f_1(1 - G(\bar{c}))$ such that all types above this cutoff choose $a = 0$ obtaining an objective payoff of zero. By Proposition 6, under the optimal monopolistic menu, types below c^{**} always choose $a = 1$ while types above c^{**} always choose $a = 0$.

We first argue that $c^{**} > \bar{c}$. To see why, note that in the monopolistic menu, type c^{**} is indifferent between $(1, N^a)$ and $(1, N^\emptyset)$, i.e.,

$$\frac{f_0 + f_1}{4} - c^{**} = \frac{f_0 + f_1}{4} G(c^{**})$$

so that

$$c^{**} = \frac{f_0 + f_1}{4} (1 - G(c^{**}))$$

If $\bar{c} \geq c^{**}$, then

$$f_1(1 - G(c^{**})) \geq \frac{f_0 + f_1}{4} (1 - G(c^{**}))$$

However, since f_0 is arbitrarily large while f_1 is arbitrarily small, it must be that $G(c^{**}) > G(\bar{c})$, in contradiction to $\bar{c} \geq c^{**}$.

Consider a type $c \in (\bar{c}, \bar{c} + \varepsilon)$, where $\varepsilon > 0$ is arbitrarily small. In competitive equilibrium, this type always chooses $a = 0$ and earns an objective payoff of zero. In contrast, under monopoly, this type always chooses $a = 1$ and earns an objective payoff of $\frac{1}{2}f_1 - c$. If $G(c)$ satisfies $G(\bar{c}) > \frac{1}{2}$, then $\frac{1}{2}f_1 > \bar{c}$. To see why, note that

$$\frac{1}{2}f_1 > \bar{c} \iff \frac{\bar{c}}{1 - G(\bar{c})} > 2\bar{c}$$

It then follows that all types in $(\bar{c}, \bar{c} + \varepsilon)$ get a strictly positive objective payoff under monopoly, but they get an objective payoff of zero in a competitive equilibrium.¹⁹ ■

¹⁹To see that there exist parameter values for f_1 and distributions G satisfying $G(\bar{c}) > \frac{1}{2}$, let $f_1 = 0.01$ and $G(\theta) = \theta^{0.05}$. Then $\bar{c} \approx 0.0026$ and $G(\bar{c}) \approx 0.74$.