

WORKING PAPER · NO. 2024-67

What is Newsworthy? Theory and Evidence

Luis Armona, Matthew Gentzkow, Emir Kamenica, and Jesse M. Shapiro

JUNE 2024

What is Newsworthy? Theory and Evidence

Luis Armona, Matthew Gentzkow, Emir Kamenica, and Jesse M. Shapiro

NBER Working Paper No. 32512

May 2024

JEL No. C44,D83,L82

ABSTRACT

We study newsworthiness in theory and practice. We focus on situations in which a news outlet observes the realization of a state of the world and must decide whether to report the realization to a consumer who pays an opportunity cost to consume the report. The consumer-optimal reporting probability is monotone in a proper scoring rule, a statistical measure of the amount of “news” in the realization relative to the consumer’s prior. We show that a particular scoring rule drawn from the statistics literature parsimoniously captures key patterns in reporting probabilities across several domains of US television news. We argue that the scoring rule can serve as a useful control variable in settings where a researcher wishes to test for bias in news reporting. Controlling for the score greatly lessens the appearance of bias in our applications.

Luis Armona
Harvard University
larmona@hks.harvard.edu

Matthew Gentzkow
Department of Economics
Stanford University
579 Jane Stanford Way
Stanford, CA 94305
and NBER
gentzkow@stanford.edu

Emir Kamenica
University of Chicago
Booth School of Business
5807 S. Woodlawn Ave.
Chicago, IL 60637
emir.kamenica@chicagobooth.edu

Jesse M. Shapiro
Department of Economics
Harvard University
Littauer Center
Cambridge, MA 02138
and NBER
jesse_shapiro@fas.harvard.edu

1 Introduction

A large social science literature hypothesizes that news coverage is systematically biased, for example towards reporting negative events more than positive ones (Harrington 1989; Soroka 2006; Aday 2010; Lengauer et al. 2012; Sacerdote et al. 2020) or towards reporting on “sensational” topics (Grabe et al. 2001).¹ Such hypotheses are important in light of evidence that news coverage impacts public policy and other consequential outcomes (Eisensee and Strömberg 2007; Snyder Jr and Strömberg 2010; Durante and Zhuravskaya 2018; Chahrour et al. 2021; Bursztyn et al. 2023).

A convincing test for bias requires a benchmark of unbiased reporting. Developing such a benchmark is challenging. Mechanical benchmarks, such as completely unselective reporting, or reporting events in proportion to their occurrence, do not seem appealing: imagine a weather report that includes every conceivable piece of information about the weather, or one that devotes equal attention to snowy and sunny days in Los Angeles. Similar challenges arise in defining benchmarks of optimal reporting in domains outside of news reporting, such as emergency alerts (Lim et al. 2019), vehicle collision warnings (Fu et al. 2019), smartphone beeps (Ely 2017; Li et al. 2023), and scientific journals (Frankel and Kasy 2022).

In this paper, we develop, operationalize, and test a model of benevolent, selective reporting. In the model, a news consumer takes an action whose payoff depends on an unknown state. A news outlet learns the state and can verifiably report its realization to the consumer at a known, random opportunity cost. If the news outlet reports the realization, the consumer takes the ex-post optimal action. If the outlet does not report the realization, then the consumer takes the ex-ante optimal action based on their prior.

The model encodes two key forces that seem essential to any theory of benevolent reporting. The first is selectivity: because reporting incurs an opportunity cost, not all realizations will be reported. The second is dependence on the prior: the more different are the ex-post and ex-ante optimal actions, the more valuable it is to report the realization.

Our main theoretical result is that the consumer-optimal probability of reporting, as a function of the realization and the prior, is monotone in a proper scoring rule. Proper scoring rules play a role in many areas of statistics and economics, including in the assessment of probabilistic forecasts (Gneiting and Raftery 2007), the assessment of expertise (Foster and Hart 2023), the valuation of information (Frankel and Kamenica 2019), and the design of incentives in surveys and laboratory experiments (Danz et al. 2022). In our setting, we can think of a scoring rule as a measure of how much “news” is in a given realization, relative to the consumer’s prior.

We study how well the model predicts what is reported on nightly national television network news broadcasts from 1968 through 2013. For much of this period, network (broadcast) television was the dominant source of news in the US (Mayer 1993; Kohut et al. 2012).² The rigid, segmented

¹Other examples include demographic bias (Gilliam Jr et al. 1996; Gruenewald et al. 2009), political bias (Efron 1971), and fatigue bias (Moeller 2002).

²More than a quarter of US adults still report obtaining news daily from network television, more than from cable news or online-only news sites, though less than from social media (Morning Consult 2022).

structure of television broadcasts makes it especially clear what is and is not reported.

We obtain summaries of nightly television news broadcasts from the Vanderbilt Television News Archive (VTNA; 2017). For each of a set of domains—the stock market, the labor market, the weather, and foreign wars—we use a set of keywords to identify segments reporting on the domain. For each domain we select a state variable—the stock market return, the unemployment rate, the average precipitation, and the number of casualties—that summarizes the news in the domain.

We use prior research, data on realizations, and other information to estimate a time-varying prior for each state variable.

We study how well our model explains when the news outlets choose to report on a given domain. Because the set of scoring rules is large, we discipline our analysis by focusing on a particular one, the continuous ranked probability score (CRPS). The CRPS is widely used in forecast evaluation (Bröcker 2012) and is appealing for our purposes because it can be applied to both continuous and discrete random variables. Importantly, because a scoring rule depends only on the prior and the realization, CRPS does not depend on the data on news reports.

Despite being highly disciplined, the model captures key qualitative features of the data well. These include obvious features, such as the fact that broadcasts are more likely to report on the stock market when the stock return is very high or very low, or that broadcasts are more likely to report on unemployment when the unemployment rate is high. These also include more subtle features, such as the fact that extreme stock returns become less newsworthy when volatility is high, or that unemployment becomes more newsworthy when there is a large change relative to the previous unemployment rate. Our primary support for these claims is a set of simple visualizations, but as a quantitative summary, we also report estimates of the completeness (Fudenberg et al. 2022) and restrictiveness (Fudenberg et al. Forthcoming) of a model that treats the reporting probability as a monotone, smooth function of the CRPS.

Given that our model is parsimonious, grounded in a theory of benevolent reporting, and able to explain key patterns across a range of domains, we argue that it can be used as a benchmark to evaluate claims of bias in reporting. We focus on two forms of bias. Under *realization bias*, the outlet treats different realizations with the same score differently. We focus on the propensity to report negative unemployment news more than positive unemployment news. Under *domain bias*, the outlet treats different domains differently, holding constant the score. We focus on the propensity to report on US military casualties in Iraq more than those in Afghanistan. In both cases, we find that a naive estimate shows strong evidence of bias, but that controlling for CRPS greatly lessens the appearance of bias. Whether any remaining evidence of bias reflects consumer preferences, outlet preferences, or some other force, is outside the scope of our study. For our purposes, the important lesson is that accounting for how much “news” is in each realization, relative to the prior, is an essential and impactful step in evaluating claims of bias.

Our primary contribution is to propose and test a parsimonious measure of newsworthiness that is grounded in an economic model. We are not aware of prior work that operationalizes such a benchmark.

A large theoretical literature studies positive models of news reporting (Gentzkow et al. 2015). Much of this literature is concerned with how bias can arise from political or market forces. We instead focus on the question of what is reported when the news outlet is fully benevolent and news is verifiable but costly to report. What we term realization bias is closely related to what Gentzkow et al. (2015) term filtering bias. Nimark and Pitschner (2019) study a model in which, as in our model, news consumers delegate to an outlet the decision of which states to report (see also Nimark [2014] and Denti and Nimark [2022]). Nimark and Pitschner (2019) present evidence on changes in topical focus of front-page articles in US newspapers over two three-month periods, and apply their model to a beauty-contest setting with strategic interactions among information consumers. Martineau and Mondria (2023) develop and test predictions of a model in which the absence of a news report conveys information to equity investors. Nimark and Pitschner (2019) and Martineau and Mondria (2023) do not consider the role of scoring rules and do not evaluate the quantitative performance of their models in predicting what is reported.

Models of selective reporting arise in other domains as well. Our finding that a realization is more likely to be reported when it is more unusual is closely related to the characterization of optimal scientific publication in Frankel and Kasy (2022). More generally, our analysis connects to a large literature on strategic communication and information design (Bergemann and Morris 2019; Kamenica 2019), and to related themes in work on rational inattention (Maćkowiak et al. 2023).³

A large empirical literature studies the determinants of media bias (Puglisi and Snyder Jr 2015). Some of this literature proposes benchmarks for unbiased selective reporting, such as treating the same economic news symmetrically regardless of the party in power (Larcinese et al. 2011; Lott Jr and Hassett 2014). We instead propose a benchmark grounded in a model of benevolent, selective reporting. Our approach is designed to cover a wide range of situations, including those in which bias may not have a political origin.

2 A Model of Benevolent, Selective Reporting

2.1 Notation and Equilibrium

In each period t a state is realized from a state space Ω according to a commonly known prior distribution $F_t \in \Delta(\Omega)$. A news outlet observes the realization $\omega_t \in \Omega$ of the state, as well as the opportunity cost $\gamma_t \in \mathbb{R}_{\geq 0}$ of making a report, which is distributed according to the strictly increasing and continuous cumulative distribution function Γ . The outlet then chooses whether or not to report the realization to a consumer, a decision that we denote by $r_t \in \{0, 1\}$.

After receiving the outlet’s report, or learning that there is no report, the consumer chooses an action a_t from an action space $\mathcal{A}$. After the consumer chooses their action, the consumer learns the state ω_t and the opportunity cost γ_t , and both the outlet and the consumer realize the loss

$$L(\omega_t, a_t) + r_t \gamma_t$$

³Ozmen (2005) considers a model of verifiable strategic communication with a potentially uninformed sender.

for $L : \Omega \times \mathcal{A} \rightarrow \mathbb{R}_{\geq 0}$ a scalar-valued loss function. Although we treat the news outlet as sharing the same objective as the consumer, we note that a concern for consumer welfare can also arise from market forces.⁴ We assume that $\min_{a \in \mathcal{A}} \mathbb{E}_F[L(\omega, a)]$ exists for all $F \in \Delta(\Omega)$.

We focus on stationary pure strategies. A stationary pure strategy for the outlet is a mapping from the set of possible realizations, priors, and opportunity costs to a reporting decision. A stationary pure strategy for the consumer is a mapping from the set of possible realizations, priors, and reporting decisions to an action.

We assume that, in an equilibrium, the consumer chooses an *ex-post* optimal action when the realization is reported, and an *ex-ante* optimal action under the prior when the realization is not reported. The outlet chooses an optimal reporting strategy given the (known) loss from the consumer's equilibrium action. Formally, our solution concept is as follows.

Definition 1. A pair of stationary pure strategies $a^* : \Omega \times \Delta(\Omega) \times \{0, 1\} \rightarrow \mathcal{A}$ and $r^* : \Omega \times \Delta(\Omega) \times \mathbb{R}_{\geq 0} \rightarrow \{0, 1\}$ is an **equilibrium** if and only if

$$a^*(\omega, F, 1) \in \arg \min_{a \in \mathcal{A}} L(\omega, a), \quad a^*(\omega, F, 0) \in \arg \min_{a \in \mathcal{A}} \mathbb{E}_{x \sim F}[L(x, a)], \text{ and}$$

$$r^*(\omega, F, \gamma) \in \arg \min_{r \in \{0, 1\}} [r(L(\omega, a^*(\omega, F, 1) + \gamma) + (1 - r)L(\omega, a^*(\omega, F, 0)))]$$

for all $\omega \in \Omega$, $F \in \Delta(\Omega)$, and $\gamma \in \mathbb{R}_{\geq 0}$.

Although we have assumed a common prior for simplicity, the solution concept is preserved if the outlet knows the consumer's prior but believes it is misspecified, provided the outlet still knows the consumer's strategy and the realization.

Our notion of equilibrium does not allow the consumer to make inferences from whether the outlet reports or not. Enke (2020) finds that many experimental participants do not make such inferences. Our notion of equilibrium allows for a parsimonious empirical test, since we can construct all the inputs to the outlet's predicted behavior without using any data on reporting. In our applications, we find that replacing the consumer's prior with their posterior conditional on non-reporting tends, if anything, to worsen the model's fit to observed reporting behavior (see Section 4.1 and Appendix Table A1).

2.2 Characterization via Scoring Rules

Towards a characterization of equilibrium, say that a *reporting rule* $R : \Delta(\Omega) \times \Omega \rightarrow [0, 1]$ is a function that describes the probability of reporting as a function of the prior and realization. Any equilibrium strategies (a^*, r^*) induce an equilibrium reporting rule. To characterize it, let $L^*(F, \omega) = L(\omega, a^*(\omega, F, 0)) - L(\omega, a^*(\omega, F, 1))$ denote the regret of a consumer with strategy a^*

⁴Suppose, for example, that the news outlet publicly commits to a reporting strategy $r(\omega_t, F_t)$, that each of a continuum of otherwise identical consumers visits the outlet with probability proportional to a constant minus the expected loss from its reports, and that the outlet receives advertising revenue proportional to the size of its audience. In this case a revenue-maximizing outlet will choose its reporting strategy to minimize the expected loss, just as in the model we consider.

and prior F when realization ω is not reported. In an equilibrium with consumer strategy a^* , the outlet is willing to make a report if and only if $L^*(F, \omega) \geq \gamma$. Any equilibrium with consumer strategy a^* then has equilibrium reporting rule R^* given by

$$R^*(F, \omega) = \Gamma(L^*(F, \omega))$$

for all $F \in \Delta(\Omega)$ and $\omega \in \Omega$.

We will show an important connection between equilibrium reporting rules and scoring rules. To do this, we first define scoring rules.

Definition 2. A *scoring rule* is a mapping $S : \Delta(\Omega) \times \Omega \rightarrow \mathbb{R}_{\geq 0}$. A *proper scoring rule* is a scoring rule such that $\mathbb{E}_{\omega \sim F'} S(F', \omega) \geq \mathbb{E}_{\omega \sim F} S(F, \omega)$ for any $F, F' \in \Delta(\Omega)$. A *normalized scoring rule* is a scoring rule such that $S(\delta(\omega), \omega) = 0$ for all $\omega \in \Omega$, for $\delta(\cdot)$ the Dirac measure on Ω .

Scoring rules are widely studied for their use in evaluating and improving probabilistic forecasts (Gneiting and Raftery 2007; Gneiting and Katzfuss 2014).⁵ In a typical setting, the forecaster (say, a weather agency) reports a predictive distribution F for some state (say, temperature) about whose distribution the forecaster has private information. The score evaluates the predictive distribution by comparison to the realization ω of the state, penalizing distributions under which the realization was unlikely. A proper scoring rule ensures that an agent wishing to minimize their expected penalty wants to report their true belief about the distribution of the state. A normalized scoring rule assigns no penalty if the agent reports a distribution that puts all mass on the realization.

In our setting, by contrast, the prior distribution F of the state is commonly known, and the decision is whether to report the realization ω . The score nevertheless plays a central role, because it gauges whether the realization is similar to what the consumer would have expected under the prior, and therefore how costly it is to omit a report and let the consumer choose based on the prior. In this sense, the score measures the amount of “news,” relative to the prior, in a given realization.

To build intuition, consider the particular scoring rule that we will use in our applications:

Definition 3. For Ω an ordered set, the *continuous ranked probability score (CRPS)* is a normalized scoring rule S_0 such that

$$S_0(F, \omega) = \int_{\Omega} (F(x) - \mathbf{1}(x \geq \omega))^2 dx$$

for $\mathbf{1}(\cdot)$ the indicator function.

Intuitively, CRPS measures how different the prior is from a point mass distribution at the realization of the state. The more different these are, the more “news” is in the given realization, relative to the prior. CRPS is known to be proper on real spaces (Gneiting and Raftery 2007).

⁵We use negatively oriented scoring rules to simplify statements.

Our main theoretical result is that any equilibrium reporting rule is monotone in some normalized proper scoring rule.

Proposition 1. *For any given state space Ω and reporting rule $R : \Delta(\Omega) \times \Omega \rightarrow [0, 1]$, there exist primitives $L(\cdot, \cdot), \mathcal{A}, \Gamma(\cdot)$ such that $R(\cdot, \cdot)$ is an equilibrium reporting rule of the resulting game if and only if*

$$R(F, \omega) = G(S(F, \omega))$$

for all $F \in \Delta(\Omega), \omega \in \Omega$, for some strictly increasing and continuous $G : \mathbb{R}_{\geq 0} \rightarrow [0, 1]$ and some normalized proper scoring rule $S(\cdot, \cdot)$.

Proof. Fix a state space Ω . For the “if” direction, assume there are $G(\cdot), S(\cdot, \cdot)$ such that $R(F, \omega) = G(S(F, \omega))$. Take $\mathcal{A} = \Delta(\Omega)$, $\Gamma(\cdot) = G(\cdot)$, and $L(\omega, a) = S(a, \omega)$. Then we can take $a^*(\omega, F, 0) = F$, so that $L^*(F, \omega) = S(F, \omega)$ and hence $R^*(F, \omega) = G(S(F, \omega))$. For the “only if” direction, fix some primitives $L(\cdot, \cdot), \mathcal{A}, \Gamma(\cdot)$ and an optimal consumer strategy a^* . The result then follows because $R^*(F, \omega) = \Gamma(L^*(F, \omega))$, noticing that the regret $L^*(F, \omega) = L(\omega, a^*(\omega, F, 0)) - L(\omega, a^*(\omega, F, 1))$ is a normalized proper scoring rule, and recalling that $\Gamma(\cdot)$ is a strictly increasing and continuous CDF. $\square$

2.3 Evaluating Bias

To define what it means for a reporting rule to be biased, we need a benchmark of unbiased reporting. One such benchmark is optimal reporting for a consumer whose preferences treat different realizations or domains symmetrically. Suppose a researcher defines a set of preferences (i.e., loss functions and action spaces) that are symmetric. Proposition 1 then implies that there is a corresponding set of scoring rules such that, if reporting probabilities are monotone in a scoring rule in this set, observed reporting does not provide evidence of bias.

We will study two kinds of bias in our applications. The first, *realization bias*, arises when consumer preferences treat different realizations—for example, good news and bad news—asymmetrically. Suppose, for example, that under the prior F , the state is distributed normally around a mean μ . Say that a decision problem is unbiased with respect to good and bad news if the associated regret is symmetric in the sense that $L^*(F, \omega) = L^*(F, \omega')$ for all ω, ω' with $|\omega - \mu| = |\omega' - \mu|$. CRPS is an example of a scoring rule that is consistent with preferences in this class, as it assigns the same score to a realization k units above the mean of distribution F as to a realization k units below the mean. An outlet’s reporting is unbiased in this sense if it does not favor bad news relative to good news when the two types of news share the same value of CRPS.

The second type of bias, *domain bias*, arises when consumer preferences treat two equally surprising realizations differently depending on whether they concern one domain or another. Suppose we consider two separate instances of our model corresponding to two domains—for example, casualties in Iraq and casualties in Afghanistan. Say that preferences toward these domains are unbiased if the consumer’s regret function $L^*(\cdot, \cdot)$ is the same across the two domains. CRPS applied separately to each domain is an example of a scoring rule that is consistent with preferences in this

class. An outlet’s reporting is unbiased in this sense if it is equally likely to report realizations in the two domains if the two realizations share the same value of CRPS.

3 Data and Estimation of Prior Beliefs

3.1 News Reports

We take our data on news reporting from the Vanderbilt Television News Archive (VTNA; 2017). The VTNA contains detailed information, including a hand-coded title and abstract, for individual segments (stories) aired during evening news broadcasts across the major television networks in the United States beginning in September 1968. We extract from the VTNA the television network, date of broadcast, duration, title, and abstract of each news segment appearing from September 1968 to December 2013 on the “Big Three” (ABC, CBS, NBC) half-hour evening news broadcasts.⁶

We consider four state variables: the stock market return, the unemployment rate, precipitation, and US military casualties. For each state variable, we construct a set of keywords. We consider a segment to report on a state if its title or abstract contains at least one of these keywords. We consider a broadcast to report on a state if at least one segment in the broadcast reports on the state. We calculate, for each date t and state, the fraction R_t of broadcasts that report on the state, and the total number of minutes in segments that report on the state. For each state, we include in our analysis only those dates on which the state variable is realized (e.g., days where financial markets are open, days where the BLS announces unemployment rates). Appendix A provides details on our choice of keywords for each state variable, as well as evidence on the quality of the resulting classification.

3.2 State Variables and Prior Beliefs

We now describe our sources for data on each state variable. In each case, we parameterize the prior as a time-series model. We choose the complexity of the prior using cross validation and estimate the remaining parameters by maximum likelihood. We describe the procedure for stock market returns in detail. For other state variables we give a high-level overview in the text, with additional details in Appendix B.

3.2.1 Stock Market Returns

We obtain from the Wharton Research Data Services (2024) the closing value SPX_t of the S&P 500 index, the closing value VIX_t of the VIX, and the risk-free interest rate r_t , for each date t .⁷ We consider those dates t for which SPX_t , r_t , and VIX_t are available. Because there is limited

⁶We exclude data from March 1999 and February 2000 because VTNA is missing data on a substantial number of broadcasts during these months. We also exclude data from 74 dates on which some segments have missing metadata, such as segment title or duration.

⁷The daily interest rate r_t is computed as the daily continuously compounded yield to maturity on day t for a Treasury bill maturing in four weeks, as implied by the average nominal price of bid-ask quotations received that day for the bill.

variation in reporting in the early part of this sample, we conduct our main analysis on data beginning in 2001. Appendix Figure A1 reports results for the sample that begins in 1990.⁸

We define our state variable ω_t to be the daily return of the S&P, expressed in log differences:

$$\omega_t = \log(SPX_t) - \log(SPX_{t-1}).$$

Similar to Black and Scholes (1973), we assume that the consumer believes that stock prices follow a geometric Brownian motion, so that over the course of a day of trading, the state variable is normally distributed. In the spirit of Bekaert and Hoerova (2014) we estimate a prior of the following form:

$$\begin{aligned}\omega_t &= \alpha_0 + r_t + \alpha_1 \sigma_t^2 + \epsilon_t \\ \epsilon_t &\sim N(0, \sigma_t^2) \\ \sigma_t^2 &= \alpha_2 + \sum_{p=1}^P \beta_p \epsilon_{t-p}^2 + \sum_{s=1}^S \gamma_s \sigma_{t-s}^2 + \sum_{q=1}^Q \delta_q VIX_{t-q}^2\end{aligned}$$

Under this prior, the drift term in returns depends on the risk-free interest rate as well as a market price for risk (α_1), while the variance follows a GARCH process that is allowed to depend on the VIX, with the α 's β 's, γ 's, δ 's, P , Q , and S representing parameters. For a fixed value of (P, Q, S) , we estimate the remaining parameters of the prior via maximum likelihood. We select the values of (P, Q, S) via k -fold h -block cross-validation (Racine 2000).⁹ We let $\hat{\theta}_t$ denote the estimated mean and variance of the distribution of ω_t on date t .

3.2.2 Unemployment Rate

We obtain from the FRED (2019) database the seasonally adjusted unemployment rate for each month t from 1948 through 2013. We consider the date within each month on which the Bureau of Labor Statistics (2019) announces the unemployment rate for the preceding month, usually the first Friday of the month. We define the state variable ω_t to be the seasonally adjusted unemployment rate announced in month t . We assume that under the consumer's prior F_t in month t the mean of the unemployment rate follows a seasonally autoregressive (SAR) process in first differences (Montgomery et al. 1998), while its variance follows a GARCH process (Appendix B.1).

⁸The VIX began trading in 2004 but is available as a historical series beginning in 1990.

⁹We implement this cross-validation procedure as follows. Fix some value of (P, Q, S) . For each of k contiguous blocks, we estimate the remaining parameters via maximum likelihood on the data excluding both the block itself and the h observations on either side of the block. We then compute the log-likelihood $-\sum_{t \in k} \frac{1}{2} \left(\log(2\pi\sigma_t^2) + \frac{\epsilon_t^2}{\sigma_t^2} \right)$ of the estimated model on the held-out block. We choose the value $(P, Q, S) \in \{0, \dots, 5\} \times \{0, \dots, 5\} \times \{0, \dots, 5\}$ that maximizes the log-likelihood across blocks. We set $k = 5$ and $h = 60$. This procedure selects $(P, Q, S) = (1, 3, 0)$.

3.2.3 Precipitation

We obtain from the U.S. Historical Climatology Network (Williams et al. 2006; National Centers for Environmental Information 2018), for a subset of US weather stations with high-quality historical data, information on the level of precipitation. We obtain from the Enhanced Master Station History Report (Vose et al. 2018) the county associated with each station’s geographic location.

For each county and date, we compute the average precipitation level across all weather stations in the county on the given date. We then take as our state variable the population-weighted average precipitation ω_t across counties on date t , using data on population from the National Bureau of Economic Research (2019), for dates from January 1, 1971 through December 31, 2013.

We follow Kedem et al. (1990) in modeling the average precipitation as lognormal with seasonally varying parameters (Appendix B.2).

3.2.4 US Military Casualties

We obtain from the Defense Casualty Analysis System (2019) data on the number ω_t of deaths of US military service members in active duty, reserve, or in the national guard on each date t from January 1st, 2003 through December 31st, 2013. Of all deaths in this dataset, 60.6% occurred in Iraq and 31.1% occurred in Afghanistan. We restrict attention to casualties in these two locations.

We assume that, in each of Iraq and Afghanistan, under the consumer’s prior F_t on date t the state variable follows a zero-inflated negative binomial distribution with parameters that depend on past casualties (Yang et al. 2013; see Appendix B.3).

3.3 Descriptive Statistics

Appendix Figure A2 presents descriptive statistics on patterns of reporting over time. Appendix Figure A3 reports on the calibration of the prior distributions $\hat{F}_t$, which we find tend to understate the probability of tail realizations.

4 Evidence on Determinants of Reporting

4.1 Methods

For each state variable, the data consist of the share R_t of broadcasts that report on the state, the realization ω_t , and the estimated prior $\hat{F}_t$, on a finite set of dates $t \in \{1, \dots, T\}$. Let $\bar{R}$ be the sample mean of R_t .

Our main findings will be via simple plots that compare the observed patterns in reporting to those predicted by CRPS. As a quantitative summary, we also report estimates of the completeness (Fudenberg et al. 2022) and restrictiveness (Fudenberg et al. Forthcoming) of a set of reporting functions.

Completeness (Fudenberg et al. 2022) is a value in $[0, 1]$ that measures the fit of a theory relative to a flexible benchmark. For some reporting function $R(\omega_t, \hat{F}_t)$, we define the fit to be the

improvement in sum of squared errors relative to a constant prediction,

$$\sum_t (R_t - \bar{R})^2 - \sum_t (R_t - R(\omega_t, \hat{F}_t))^2.$$

We estimate the completeness of a smooth monotone function of the CRPS $S_0(\omega_t, \hat{F}_t)$ by taking the ratio of its fit to that of a smooth function of both the realization ω_t and the estimated prior $\hat{F}_t$.

Appendix Table A1 reports additional details and specifications, including (a) estimating separately by calendar quarter, to give a sense of the variability in the data; (b) replacing the share R_t of broadcasts that report on the state with the average number of minutes devoted to the state; (c) replacing the estimated prior $\hat{F}_t$ with the corresponding estimated posterior conditional on nonreporting, to evaluate the performance of a model that allows for more sophisticated updating by the consumer.¹⁰

Restrictiveness (Fudenberg et al. Forthcoming) is a value in $[0, 1]$ that measures the (in)ability of a theory to fit arbitrary data. We estimate the restrictiveness of a smooth monotone function of the CRPS $S_0(\omega_t, \hat{F}_t)$ by comparing its average mean squared error to that of a constant prediction when fit on 100 datasets in which we replace R_t with the output of a random smooth function. Appendix Figure A4 reports additional details and specifications.

4.2 Findings

4.2.1 Stock Market Returns

Figure 1 presents our findings on reporting on the stock market. Panel (a) shows the empirical frequency of reporting for quantiles of the stock return ω_t . Panel (b) shows the mean $\log(\text{CRPS})$ for these same quantiles. The relationship between reporting and returns in panel (a) is visually similar to that between the $\log(\text{CRPS})$ and returns in panel (b). Panel (b) reports that a monotone function of CRPS captures nearly all of the predictive information in the state and prior about the likelihood of reporting, with a completeness of 0.976, despite being highly restrictive, with a restrictiveness of 1.

Panels (c) and (d) illustrate the role of the prior in determining reporting. Panel (d) shows that when the prior standard deviation of returns is large, CRPS becomes larger when returns are less extreme and smaller when they are more extreme. This reflects the fact that, when the market is volatile, more extreme returns are less unusual, and less extreme returns are more unusual. Panel (c) shows that a qualitatively similar pattern holds for reporting.

¹⁰Appendix Table A2 reports the completeness of the more flexible reporting function with respect to an even more flexible one that can depend on any of the inputs used to estimate the parameters $\hat{\theta}_t$.

4.2.2 Unemployment Rate

Figure 2 presents our findings on reporting on the unemployment rate. Panel (a) shows the empirical frequency of reporting for quantiles of the unemployment rate ω_t (left plot) or its change $\omega_t - \omega_{t-1}$ (right plot). Panel (b) shows the mean $\log(\text{CRPS})$ for these same quantiles. The relationship between reporting and the unemployment rate or its change in panel (a) is visually similar to that between $\log(\text{CRPS})$ and the unemployment rate or its change in panel (b). The biggest difference between the two is that the reporting frequency plateaus when the change in the unemployment rate is large, whereas by construction the $\log(\text{CRPS})$ does not have such a ceiling.

Panels (a) and (b) again illustrate the role of the prior in determining reporting. The plots on the left show that larger levels of unemployment are more surprising (panel (a)) and more likely to be reported (panel (b)). The plots on the right show that larger absolute changes in the unemployment rate are more surprising (panel (a)) and are more likely to be reported (panel (b)). Larger absolute changes in the unemployment rate are surprising precisely because the prior for the current period's unemployment rate depends on the previous period's unemployment rate. Panel (b) reports that a monotone function of CRPS achieves a completeness of 0.408 with a restrictiveness of 1.

4.2.3 Precipitation

Figure 3 presents our findings on reporting on precipitation. Panel (a) shows the empirical frequency of reporting for quantiles of the average precipitation ω_t . Panel (b) shows the mean $\log(\text{CRPS})$ for these same quantiles. The relationship between reporting and average precipitation in panel (a) is visually similar to that between $\log(\text{CRPS})$ and average precipitation in panel (b). Panel (b) reports that a monotone function of CRPS achieves a completeness of 0.192 with a restrictiveness of 1.

4.2.4 US Military Casualties

Figure 4 presents our findings on reporting on US military casualties in Iraq, with a format parallel to that of Figure 3. Broadcasts on days with more casualties are more likely to have reports on US military casualties in Iraq (panel (a)), and to have higher CRPS (panel (b)). Panel (b) shows that a monotone function of CRPS achieves a completeness of 0.599 with a restrictiveness of 1. Appendix Figure A5 presents analogous plots for casualties in Afghanistan, for which estimated completeness is greater but the visual agreement between reporting probabilities and CRPS is poorer.

5 Application to Testing for Bias

In this section we show how CRPS can be used as a benchmark to evaluate different forms of bias. We focus on forms of bias that have been discussed in the literature, that are testable with our data, and that emerge as statistically significant in an analysis without controls. These filters yield

one example of realization bias and one example of domain bias.¹¹

5.1 Realization Bias

Panel (a) of Figure 5 shows estimates of a bias towards reporting on increases in the unemployment rate, a form of realization bias that favors reporting on negative news (as in, e.g., Soroka [2006]).

With no controls, we estimate that reporting the unemployment rate is 5.4 percentage points more likely when the unemployment rate has risen than when it has fallen or remained unchanged. Superficially, this contrast is evidence of a realization bias in favor of reporting negative news.

However, when we control for a smooth monotone function of the CRPS, the estimated bias falls to 2.4 percentage points and becomes statistically insignificant. One reason that controlling for CRPS reduces the estimated bias is that both CRPS and the probability of reporting tend to be smaller when there is no change in the unemployment rate (Figure 2). More subtly, changes in unemployment are positively skewed, with unemployment rising more quickly than decreasing.¹² More extreme changes in unemployment result in both a higher likelihood of reporting and a higher CRPS (Figure 2). Variation in CRPS thus accounts for many of the patterns that drive the appearance of a bias in favor of reporting negative news.

5.2 Domain Bias

Panel (b) of Figure 5 shows estimates of a bias towards reporting on casualties in Iraq, a form of domain bias (see, e.g., Jacobson [2010]).

With no controls, we estimate that reporting on casualties in Iraq is 10.1 percentage points more likely than reporting on casualties in Afghanistan. Superficially, this contrast is evidence of a domain bias in favor of reporting about Iraq.

However, when we control for the difference in a smooth monotone function of CRPS, the estimated bias falls to -0.9 percentage points and becomes statistically insignificant. Appendix Figure A7 provides an intuition behind this finding. It shows that when the difference in $\log(\text{CRPS})$ is close to zero, the probability of reporting is similar between casualties in Iraq and Afghanistan. However, casualties in Iraq tend to have a higher CRPS, due to higher variation in the number of casualties. Hence, a naive comparison of reporting probabilities concludes there is a bias in favor of reporting on casualties in Iraq, whereas our analysis suggests that reporting is instead driven by newsworthiness.

6 Conclusion

A parsimonious model of benevolent, selective reporting reproduces many key patterns in US television nightly news broadcasts. We think the model provides a useful benchmark or control variable for future studies of biased reporting in the news and other domains.

¹¹Appendix Figure A6 details the full set of forms of bias that we considered.

¹²For example, the median increase in unemployment (0.2) is twice the median decrease in unemployment (-0.1).

References

- S. Aday. Chasing the Bad News: An Analysis of 2005 Iraq and Afghanistan War Coverage on NBC and Fox News Channel. *Journal of Communication*, 60(1):144–164, 2010.
- G. Bekaert and M. Hoerova. The VIX, the Variance Premium and Stock Market Volatility. *Journal of Econometrics*, 183(2):181–192, 2014.
- D. Bergemann and S. Morris. Information Design: A Unified Perspective. *Journal of Economic Literature*, 57(1):44–95, 2019.
- S. Bird, E. Klein, and E. Loper. *Natural Language Processing with Python: Analyzing Text with the Natural Language Toolkit*. O’Reilly Media, Inc., 2009.
- F. Black and M. Scholes. The Pricing of Options and Corporate Liabilities. *Journal of Political Economy*, 81(3):637–654, 1973.
- J. Bröcker. Evaluating Raw Ensembles with the Continuous Ranked Probability Score. *Quarterly Journal of the Royal Meteorological Society*, 138(667):1611–1617, 2012.
- Bureau of Labor Statistics. Release Calendar. <https://web.archive.org/web/20190111080349/https://www.bls.gov/bls/histreleasedates.pdf> and https://web.archive.org/web/20190902022135/https://www.bls.gov/schedule/2019/09_sched.htm, 2019. Accessed in September 2019.
- L. Bursztyń, A. Rao, C. Roth, and D. Yanagizawa-Drott. Opinions as Facts. *The Review of Economic Studies*, 90(4):1832–1864, 2023.
- M. D. Cattaneo, R. K. Crump, M. H. Farrell, and Y. Feng. On Binscatter. *American Economic Review*, 114(5):1488–1514, 2024.
- R. Chahrour, K. Nimark, and S. Pitschner. Sectoral Media Focus and Aggregate Fluctuations. *American Economic Review*, 111(12):3872–3922, 2021.
- V. Chernozhukov, I. Fernandez-Val, and A. Galichon. Improving Point and Interval Estimators of Monotone Functions by Rearrangement. *Biometrika*, 96(3):559–575, 2009.
- D. Danz, L. Vesterlund, and A. J. Wilson. Belief Elicitation and Behavioral Incentive Compatibility. *American Economic Review*, 112(9):2851–2883, 2022.
- Defense Casualty Analysis System. Names of Fallen. <https://web.archive.org/web/20200130165026/https://dcas.dmdc.osd.mil/dcas/pages/main.xhtml>, 2019. Accessed in September 2019.
- T. Denti and K. Nimark. Attention Costs, Economies of Scale and Markets for Information. 2022. URL https://web.archive.org/web/20230316210939/https://www.kris-nimark.net/pdf/papers/DN_slides.pdf. Accessed in May 2024.


- R. Durante and E. Zhuravskaya. Attack When the World Is Not Watching? US News and the Israeli-Palestinian Conflict. *Journal of Political Economy*, 126(3):1085–1133, 2018.
- E. Efron. *The News Twisters*. Manor Books, New York, 1971.
- T. Eiseensee and D. Strömberg. News Droughts, News Floods, and US Disaster Relief. *The Quarterly Journal of Economics*, 122(2):693–728, 2007.
- J. C. Ely. Beeps. *American Economic Review*, 107(1):31–53, 2017.
- B. Enke. What You See Is All There Is. *The Quarterly Journal of Economics*, 135(3):1363–1398, 2020.
- D. Foster and S. Hart. “Calibeating”: Beating Forecasters at Their Own Game. *Theoretical Economics*, 18(4):1441–1474, 2023.
- A. Frankel and E. Kamenica. Quantifying Information and Uncertainty. *American Economic Review*, 109(10):3650–3680, 2019.
- A. Frankel and M. Kasy. Which Findings Should Be Published? *American Economic Journal: Microeconomics*, 14(1):1–38, 2022.
- FRED. Civilian Unemployment Rate (UNRATE). <https://web.archive.org/web/20190913025936/https://fred.stlouisfed.org/series/UNRATE>, 2019. Original data from the U.S. Bureau of Labor Statistics. Accessed in September 2019.
- E. Fu, S. Sibi, D. Miller, M. Johns, B. Mok, M. Fischer, and D. Sirkin. The Car That Cried Wolf: Driver Responses to Missing, Perfectly Performing, and Oversensitive Collision Avoidance Systems. In *2019 IEEE Intelligent Vehicles Symposium (IV)*, pages 1830–1836. IEEE, 2019.
- D. Fudenberg, J. Kleinberg, A. Liang, and S. Mullainathan. Measuring the Completeness of Economic Models. *Journal of Political Economy*, 130(4):956–990, 2022.
- D. Fudenberg, W. Gao, and A. Liang. How Flexible Is That Functional Form? Quantifying the Restrictiveness of Theories. *Review of Economics and Statistics*, Forthcoming.
- M. Gentzkow, J. M. Shapiro, and D. F. Stone. Media Bias in the Marketplace: Theory. In *Handbook of Media Economics*, volume 1, pages 623–645. Elsevier, 2015.
- F. D. Gilliam Jr, S. Iyengar, A. Simon, and O. Wright. Crime in Black and White: The Violent, Scary World of Local News. *Harvard International Journal of Press/Politics*, 1(3):6–23, 1996.
- T. Gneiting and M. Katzfuss. Probabilistic Forecasting. *Annual Review of Statistics and Its Application*, 1:125–151, 2014.
- T. Gneiting and A. E. Raftery. Strictly Proper Scoring Rules, Prediction, and Estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007.

- M. E. Grabe, S. Zhou, and B. Barnett. Explicating Sensationalism in Television News: Content and the Bells and Whistles of Form. *Journal of Broadcasting & Electronic Media*, 45(4):635–655, 2001.
- J. Gruenewald, J. Pizarro, and S. M. Chermak. Race, Gender, and the Newsworthiness of Homicide Incidents. *Journal of Criminal Justice*, 37(3):262–272, 2009.
- D. E. Harrington. Economic News on Television: The Determinants of Coverage. *Public Opinion Quarterly*, 53(1):17–40, 1989.
- G. C. Jacobson. A Tale of Two Wars: Public Opinion on the US Military Interventions in Afghanistan and Iraq. *Presidential Studies Quarterly*, 40(4):585–610, 2010.
- E. Kamenica. Bayesian Persuasion and Information Design. *Annual Review of Economics*, 11: 249–272, 2019.
- B. Kedem, L. S. Chiu, and G. R. North. Estimation of Mean Rain Rate: Application to Satellite Observations. *Journal of Geophysical Research: Atmospheres*, 95(D2):1965–1972, 1990.
- A. Kohut, C. Doherty, M. Dimock, and S. Keeter. In Changing News Landscape, Even Television Is Vulnerable. *Pew Internet & American Life Project*, 1:1–85, 2012. <https://web.archive.org/web/20230314211900/https://www.pewresearch.org/wp-content/uploads/sites/4/legacy-pdf/2012-News-Consumption-Report.pdf>. Accessed in November 2023.
- V. Larcinese, R. Puglisi, and J. M. Snyder Jr. Partisan Bias in Economic News: Evidence on the Agenda-Setting Behavior of US Newspapers. *Journal of Public Economics*, 95(9-10):1178–1189, 2011.
- G. Lengauer, F. Esser, and R. Berganza. Negativity in Political News: A Review of Concepts, Operationalizations and Key Findings. *Journalism*, 13(2):179–202, 2012.
- T. Li, J. K. Haines, M. F. R. De Eguino, J. I. Hong, and J. Nichols. Alert Now or Never: Understanding and Predicting Notification Preferences of Smartphone Users. *ACM Transactions on Computer-Human Interaction*, 29(5):1–33, 2023.
- J. R. Lim, B. F. Liu, and M. Egnoto. Cry Wolf Effect? Evaluating the Impact of False Alarms on Public Responses to Tornado Alerts in the Southeastern United States. *Weather, Climate, and Society*, 11(3):549–563, 2019.
- J. R. Lott Jr and K. A. Hassett. Is Newspaper Coverage of Economic Events Politically Biased? *Public Choice*, 160(1-2):65–108, 2014.
- B. Maćkowiak, F. Matějka, and M. Wiederholt. Rational Inattention: A Review. *Journal of Economic Literature*, 61(1):226–273, 2023.

- C. Martineau and J. Mondria. News Selection and Asset Pricing Implications. 2023. URL <https://doi.org/10.31235/osf.io/ame2f>. Accessed in May 2024.
- W. G. Mayer. Poll Trends: Trends in Media Usage. *The Public Opinion Quarterly*, 57(4):593–611, 1993.
- S. D. Moeller. *Compassion Fatigue: How the Media Sell Disease, Famine, War and Death*. Routledge, New York, 2002.
- A. L. Montgomery, V. Zarnowitz, R. S. Tsay, and G. C. Tiao. Forecasting the U.S. Unemployment Rate. *Journal of the American Statistical Association*, 93(442):478–493, 1998.
- Morning Consult. National Tracking Poll No. 2202053. 2022. https://web.archive.org/web/20220228130226/https://assets.morningconsult.com/wp-uploads/2022/02/25132637/2202053_crosstabs_MC_ENTERTAINMENT_MEDIA_AND_NEWS_ANCHOR_TRUST_Adults_v1_SH.pdf. Accessed in November 2023.
- National Bureau of Economic Research. US Intercensal Population by County and State 1970-on. <https://web.archive.org/web/20201112014741/https://www.nber.org/research/data/us-intercensal-population-county-and-state-1970>, 2019. Original data from the U.S. Census Bureau. Accessed in September 2019.
- National Centers for Environmental Information. United States Historical Climatology Network Daily Temperature, Precipitation, and Snow Data. <https://web.archive.org/web/20170507031704/https://www1.ncdc.noaa.gov/pub/data/ghcn/daily/>, 2018. Accessed in June 2018.
- K. P. Nimark. Man-Bites-Dog Business Cycles. *American Economic Review*, 104(8):2320–2367, 2014.
- K. P. Nimark and S. Pitschner. News Media and Delegated Information Choice. *Journal of Economic Theory*, 181:160–196, 2019.
- D. Ozmen. *Information Transmission and Recommender Systems*. Yale University Dissertation, 2005.
- R. Puglisi and J. M. Snyder Jr. Empirical Studies of Media Bias. In *Handbook of Media Economics*, volume 1, pages 647–667. Elsevier, 2015.
- J. Racine. Consistent Cross-Validatory Model-Selection for Dependent Data: *h_v*-Block Cross-Validation. *Journal of Econometrics*, 99(1):39–61, 2000.
- B. Sacerdote, R. Sehgal, and M. Cook. Why Is All COVID-19 News Bad News? *National Bureau of Economic Research Working Paper*, 28110, 2020.

- J. M. Snyder Jr and D. Strömberg. Press Coverage and Political Accountability. *Journal of Political Economy*, 118(2):355–408, 2010.
- S. N. Soroka. Good News and Bad News: Asymmetric Responses to Economic Information. *Journal of Politics*, 68(2):372–385, 2006.
- Vanderbilt Television News Archive. <https://web.archive.org/web/20170601191609/https://tvnews.vanderbilt.edu>, 2017. Accessed in June 2017.
- R. S. Vose, S. McNeill, K. Thomas, and E. Shepherd. Enhanced Master Station History Report. <https://web.archive.org/web/20221226063945/https://www.ncei.noaa.gov/access/homr/file/list>, 2018. NOAA National Climatic Data Center. Accessed in June 2018.
- Wharton Research Data Services. SPX, Treasury Riskfree, and VIX Series. <https://web.archive.org/web/20191015172939/https://wrds-www.wharton.upenn.edu>, 2024. Accessed in September 2019 and May 2024.
- C. Williams, R. Vose, D. Easterling, and M. Menne. United States Historical Climatology Network Daily Temperature, Precipitation, and Snow Data. <https://web.archive.org/web/20170302223347/http://cdiac.ornl.gov/epubs/ndp/ushcn/usa.html>, 2006. ORNL/CDIAC-118, NDP-070. Carbon Dioxide Information Analysis Center, Oak Ridge National Laboratory, US Department of Energy, Oak Ridge, Tennessee. Accessed in April 2018.
- M. Yang, G. K. Zamba, and J. E. Cavanaugh. Markov Regression Models for Count Time Series with Excess Zeros: A Partial Likelihood Approach. *Statistical Methodology*, 14:26–38, 2013.


Figure 1: Reporting on Stock Returns



(a) Frequency of reporting vs. stock market return



(b) log(CRPS) vs. stock market return




(c) Frequency of reporting, by prior std. deviation



(d) log(CRPS), by prior std. deviation


Notes: The unit of analysis is the outlet-date. Each panel depicts the mean of the variable on the y-axis within quantiles of the variable on the x-axis. Panels (a) and (b) include pointwise 95 percent confidence intervals clustered by date. In panels (c) and (d), observations are split by whether the standard deviation of priors is above or below the median and the sample is restricted to dates after January 1, 2001. In each panel, the x-axis variable is the stock market return on the given date. In panels (a) and (c) the number of quantiles is chosen using the IMSE-optimal bandwidth selection of Cattaneo et al. (2024). In panels (b) and (d), the number of quantiles is chosen to match the number of quantiles in panels (a) and (c), respectively. In panels (a) and (c) the y-axis variable is an indicator for whether the given outlet reports on the stock market on the given date. In panels (b) and (d) the y-axis variable is the log(CRPS) for the given date.

Figure 2: Reporting on the Unemployment Rate




Notes: The unit of analysis is the outlet-date. Panels (a) and (b) each depict the mean of the variable on the y-axis within quantiles of the variable on the x-axis, along with pointwise 95 percent confidence intervals clustered by date. In panels (a) and (b), the x-axis variable is the unemployment rate on the given date (left plot) or the change relative to the previous unemployment rate (right plot). In panel (a) the number of quantiles is chosen using the IMSE-optimal bandwidth selection of Cattaneo et al. (2024). In panel (b), the number of quantiles is chosen to match the number of quantiles in panel (a). In panel (a) the y-axis variable is an indicator for whether the given outlet reports on the unemployment rate on the given date. In panel (b) the y-axis variable is the log(CRPS) for the given date.

Figure 3: Reporting on Precipitation




(a) Frequency of reporting vs. precipitation




(b) log(CRPS) vs. precipitation

Notes: The unit of analysis is the outlet-date. Panels (a) and (b) each depict the mean of the variable on the y-axis within quantiles of the variable on the x-axis, along with pointwise 95 percent confidence intervals clustered by date. In panels (a) and (b), the x-axis variable is the average precipitation on the given date. In panel (a) the number of quantiles is chosen using the IMSE-optimal bandwidth selection of Cattaneo et al. (2024). In panel (b), the number of quantiles is chosen to match the number of quantiles in panel (a). In panel (a) the y-axis variable is an indicator for whether the given outlet reports on precipitation on the given date. In panel (b) the y-axis variable is the log(CRPS) for the given date.

Figure 4: Reporting on US Military Casualties in Iraq




(a) Frequency of reporting vs. number of casualties




(b) $\log(\text{CRPS})$ vs. number of casualties

Notes: The unit of analysis is the outlet-date. Panels (a) and (b) each depict the mean of the variable on the y-axis within quantiles of the variable on the x-axis, along with pointwise 95 percent confidence intervals clustered by date. In panels (a) and (b), the x-axis variable is the number of military casualties on the given date. In panel (a) the number of quantiles is chosen using the IMSE-optimal bandwidth selection of Cattaneo et al. (2024). In panel (b), the number of quantiles is chosen to match the number of quantiles in panel (a). In panel (a) the y-axis variable is an indicator for whether the given outlet reports on military casualties in Iraq on the given date. In panel (b) the y-axis variable is the $\log(\text{CRPS})$ for the given date.

Figure 5: Estimates of Bias, Controlling for Newsworthiness



(a) Realization bias: Reporting on increases vs. decreases in the unemployment rate



(b) Domain bias: Reporting on Iraq vs. Afghanistan

Notes: Plots show the estimated bias in reporting frequency when reporting on unemployment (panel (a)) and military casualties (panel (b)) under different model specifications. In each panel we estimate the bias without (“naive”) and with (“control for newsworthiness”) a control for CRPS. In panel (a), we estimate the bias towards reporting on negative news using a linear regression, where the dependent variable is the share of broadcasts that report on the state, the estimate of bias is the coefficient on an indicator for whether the unemployment rate is larger than in the previous period, and the control is included using the CRPS kernel function defined in Appendix Table A1. In panel (b), we estimate the bias towards reporting on Iraq war casualties relative to Afghanistan war casualties using a linear regression, where the dependent variable is the difference in the share of broadcasts that report on Iraq casualties and the share of broadcasts that report on Afghanistan casualties, the estimate of bias is the coefficient on the constant term, and the control is included as the Iraq-Afghanistan difference in the kernel function defined in Appendix Table A1. P-values for differences in coefficients are indicated in brackets and are obtained via a bootstrap over calendar months with 100 replicates.

A Details on Definition of Reporting Variables

To define whether a given segment reports on a given state, we code whether the title or abstract contains a specific set of keywords. Here we provide additional details on the keywords. Appendix Table A3 reports the most common non-keyword terms associated with each state. Appendix Figure A8 reports the results of a random audit in which a human auditor unaware of our classification classified 100 randomly chosen segments based on their textual descriptions.

A.1 Stock Market Returns

We flag any segment with the keywords “stock market,” “stks.,” or “stk.” as a report on the stock market’s performance that day. We use these abbreviations because stocks are commonly abbreviated by VTNA coders as “stk.” or “stks.”.

A.2 Unemployment Rate

We flag as a report on the unemployment rate a segment meeting at least one of the following criteria:

- Contains the keywords “employment” or “unempl” (an abbreviation sometimes used for employment by VTNA coders).
- Contains the keyphrase “jobs report.”
- Contains the keyword “job” and one of the following keywords: “figure,” “statistic,” or “number” (i.e., jobs numbers, job figures, etc.).

A.3 US Military Casualties

We flag as a report on US military casualties a segment meeting all of the following criteria:

- Mention of casualties: use of the word “casualty,” “casualties,” “death,” or “kill.”
- Mention of military: use of the word “troop,” “troops,” “soldier,” “army,” “navy,” “military,” “air force,” or “marine.”
- Mention of the US: includes either “united states” or “usa” somewhere in the text or “us” followed immediately by one of the military keywords (e.g., includes “us troops” in the text). This is important to avoid reporting on theaters of war unrelated to the US (such as reporting on combat in Israel/Palestine).
- No mention of civilians: does not include the word “civilian.”

We further filter these segments based on whether they reference a particular theater of war:

- Mention of Iraq: use of the word “iraq” or “baghdad.”
- Mention of Afghanistan: use of the word “afghanistan” or “kabul.”

A.4 Precipitation

We flag as a report on precipitation a segment with the word “weather” in the segment title. The word “weather” is a keyword explicitly used by VTNA coders in the segment titles to identify weather segments. Weather reporting may depend on factors besides precipitation (e.g., a heat wave or tornado). Appendix Figure A9 reproduces the results in Figure 3, requiring a weather event to include one of the following words that directly relate to precipitation: “snow,” “snowy,” “snowstorm,” “snowstorms,” “snowfall,” “flood,” “flooding,” “floods,” “rain,” “drought,” “blizzard,” “hurricane,” “ice,” “storm,” “storms,” or “stormy.” These account for 67.7% of our regularly coded weather reports.

B Details on Parameterization of Prior Distributions

B.1 Unemployment Rate

We assume that under the consumer's prior F_t in month t the state variable is normally distributed as follows:

$$\begin{aligned}\Delta\omega_t &= \beta_{0,1} + \sum_{s=1}^{P_1} \beta_{s,1}(\Delta\omega_{t-s} - \beta_{0,1}) + \sum_{s=1}^{Q_1} \gamma_{s,1}(\Delta\omega_{t-12\cdot s} - \beta_{0,1}) + \epsilon_t \\ \epsilon_t &\sim N(0, \sigma_t^2) \\ \sigma_t^2 &= \beta_{0,2} + \sum_{s=1}^{P_2} \beta_{s,2}\epsilon_{t-s}^2 + \sum_{s=1}^{Q_2} \gamma_{s,2}\sigma_{t-s}^2.\end{aligned}$$

where $\Delta\omega_t = \omega_t - \omega_{t-1}$. Under this prior, the mean of the unemployment rate follows a seasonally autoregressive (SAR) process in first differences (Montgomery et al. 1998), while its variance follows a GARCH process, with the $\beta's$, $\gamma's$, $P's$, and $Q's$ representing parameters. For a fixed value of (P_1, Q_1, P_2, Q_2) , we estimate the remaining parameters of the prior via maximum likelihood.¹³ We select the values of (P_1, Q_1, P_2, Q_2) via k -fold h -block cross-validation (Racine 2000).¹⁴ We let $\hat{\theta}_t$ denote the estimated mean and variance of the distribution of ω_t in month t .

B.2 Precipitation

For each county c and date t , we compute the average precipitation level $\omega_{c,t}$ across all weather stations in the county on the given date. We then take as our state variable the population-weighted average precipitation ω_t across all counties c on date t .¹⁵ We follow Kedem et al. (1990) in modeling the average precipitation as lognormal with seasonally varying parameters. Specifically, we assume that under the consumer's prior F_t on date t the log of the state variable, $\log(\omega_t)$, is normally distributed as follows:

$$\begin{aligned}\log(\omega_t) &= \vec{x}_t\vec{\beta}_1 + \sum_{s=1}^{P_1} \alpha_{s,1} \left(\log(\omega_{t-s}) - \vec{x}_{t-s}\vec{\beta}_1 \right) + \sum_{s=1}^{Q_1} \gamma_{s,1}\epsilon_{t-s} + \epsilon_t \\ \epsilon_t &\sim N(0, \sigma_t^2) \\ \sigma_t^2 &= \exp \left(\vec{x}_t\vec{\beta}_2 \right) + \sum_{s=1}^{P_2} \alpha_{s,2}\epsilon_{t-s}^2 + \sum_{s=1}^{Q_2} \gamma_{s,2}\sigma_{t-s}^2.\end{aligned}$$

¹³We initialize the mean by assuming that $\epsilon_t = 0$ in months t prior to January 1948. This implies that $(\Delta\omega_t - \beta_{0,1}) = 0$ in months t prior to January 1948. We initialize the variance by assuming that σ_t^2 in months t prior to January 1948 is equal to the unconditional variance of our estimates of ϵ_t over the period 1948 through 2013.

¹⁴We consider $(P_1, Q_1, P_2, Q_2) \in \{1, \dots, 6\} \times \{0, \dots, 5\} \times \{1, \dots, 5\} \times \{0, 1, 2\}$ and evaluate these via the average log-likelihood of the data in the held-out blocks. We set $k = 5$ and $h = 60$. This procedure selects $(P_1, Q_1, P_2, Q_2) = (5, 3, 2, 1)$.

¹⁵We obtain the annual population of each county c from the National Bureau of Economic Research (2019).

where $\vec{x}_t = [1, \cos(\frac{2\pi t}{365}), \sin(\frac{2\pi t}{365})]$. Under this prior, the expectation of the log average precipitation follows an ARMA model with controls for seasonality, and its variance follows a GARCH process with controls for seasonal cyclicalities, with the α 's, β 's, γ 's, P 's, and Q 's, representing parameters. For a fixed value of (P_1, Q_1, P_2, Q_2) , we estimate the remaining parameters of the prior via maximum likelihood.¹⁶ We optimize the P 's and Q 's via k -fold h -block cross-validation.¹⁷ We let $\hat{\theta}_t$ denote the estimated mean and variance of the distribution of $\log(\omega_t)$ on date t .

B.3 US Military Casualties

We assume that, in each of Iraq and Afghanistan, under the consumer's prior F_t on date t the state variable follows a zero-inflated negative binomial distribution with parameters that depend on past casualties (Yang et al. 2013):

$$\omega_t \sim ZINB(\lambda_t, \zeta_t, \phi)$$

with

$$\lambda_t = \exp(\beta_0 + \sum_{s=1}^P \beta_r \log(\omega_{t-s} + 1))$$

$$\zeta_t = \frac{\exp(\gamma_0 + \sum_{s=1}^Q \gamma_r \log(\omega_{t-s} + 1))}{1 + \exp(\gamma_0 + \sum_{s=1}^Q \gamma_r \log(\omega_{t-s} + 1))}$$

and the β 's, γ 's, P , and Q representing parameters. Under this prior, with probability ζ_t casualties are zero, and with probability $(1 - \zeta_t)$ casualties follow a negative binomial distribution with mean λ_t and dispersion ϕ . For fixed values of (P, Q) , we estimate the remaining parameters of the prior via (partial) maximum likelihood as in Yang et al. (2013). We select (P, Q) via k -fold h -block cross validation.¹⁸ We let $\hat{\theta}_t$ denote the estimated parameters of the distribution of ω_t on date t .

¹⁶We initialize the mean by assuming that $\epsilon_t = 0$ on dates t prior to January 1st, 1970. This implies that $(\log(\omega_t) - \vec{x}_t \beta_1) = 0$ on dates t prior to January 1st, 1970. We initialize the variance by assuming that σ_t^2 on dates prior to January 1st, 1970 is equal to the unconditional variance of our estimates of ϵ_t over the period 1970 through 2014.

¹⁷We choose the value $(P_1, Q_1, P_2, Q_2) \in \{1, \dots, 5\} \times \{1, \dots, 5\} \times \{1, \dots, 3\} \times \{0, 1, 2\}$ that maximizes the average log-likelihood across blocks. We set $k = 5$ and $h = 365$. This procedure selects $(P_1, Q_1, P_2, Q_2) = (4, 1, 2, 1)$.

¹⁸We consider $(P, Q) \in \{1, \dots, 30\}^2$ and evaluate these via the average log-likelihood of the data in the held-out blocks, using dates beginning on January 31st, 2003, the first date for which data on $\{\omega_{t-1}, \dots, \omega_{t-30}\}$ is available. We set $k = 5$ and $h = 100$. This procedure selects $(P, Q) = (15, 30)$.

Table A1: Completeness Under Alternative Specifications

	CRPS	State Only	State + Prior
Panel (a): Stock market			
Baseline	0.976	0.425	1
[Min Quarter, Max Quarter]	[0.818, 1]	[0.322, 0.455]	
Reporting in minutes	0.886	0.853	1
Posterior conditional on nonreporting	0.844	0.425	1
Panel (b): Unemployment			
Baseline	0.408	0.975	1
[Min Quarter, Max Quarter]	[0.278, 0.536]	[0.887, 0.973]	
Reporting in minutes	0.335	0.927	1
Posterior conditional on nonreporting	0.408	0.975	1
Panel (c): US military casualties in Iraq			
Baseline	0.599	0.616	1
[Min Quarter, Max Quarter]	[0.366, 0.590]	[0.429, 0.654]	
Reporting in minutes	0.659	0.730	1
Posterior conditional on nonreporting	0.380	0.616	1
Panel (d): Precipitation			
Baseline	0.192	0.228	1
[Min Quarter, Max Quarter]	[0.068, 0.370]	[0.132, 0.682]	
Reporting in minutes	0.273	0.308	1
Posterior conditional on nonreporting	0.202	0.228	1

Notes: The table displays estimates of completeness as defined in Fudenberg et al. (2022). Each panel represents a different setting. Each column corresponds to a different reporting function: a smooth monotone function of the CRPS $S_0(\omega_t, \hat{F}_t)$ (CRPS), a smooth function of the realization ω_t (state only), and a smooth function of both the realization ω_t and the estimated prior $\hat{F}_t$ (state+prior). We estimate each reporting function via kernel regression with bandwidth chosen via 10-fold block cross-validation. We enforce monotonicity via rearrangement following Chernozhukov et al. (2009). The estimation problem is finite-dimensional because we can represent the estimated prior $\hat{F}_t$ as a function of a finite-dimensional vector $\hat{\theta}_t$ of estimated parameters. Completeness is the ratio (truncated at 1) of the fit of the given reporting function to that of the most flexible reporting function (state+prior). Each row corresponds to a different specification. The row “Baseline” reports our baseline estimate following the details in Section 4.1. The remaining rows report variants of our baseline estimates, all using the same kernel bandwidth. The row “[Min Quarter, Max Quarter]” reports the minimum and maximum completeness when estimating separately by calendar quarter. The row “Reporting in minutes” replaces the share R_t of broadcasts that report on the state with the number of minutes devoted to the state in the average broadcast. The row “Posterior conditional on nonreporting” replaces the prior distribution on the state with the posterior distribution on the state, calculated via Bayes rule using our baseline kernel estimates of the probability of reporting as a function of the prior, and the probability as a function of the state and prior. When the prior probability of reporting given the state is less than 0.0075, we treat the prior distribution on the state as the posterior distribution.

Table A2: Completeness Under More Flexible Benchmark

	State + Prior
Stock market	0.345
Unemployment	0.186
US military casualties in Iraq	1
Precipitation	0.448


Notes: The table displays estimates of completeness as defined in Fudenberg et al. (2022) for the function of both the realization ω_t and the estimated prior $\hat{F}_t$, under a more flexible benchmark that can depend on any of the inputs used to estimate the parameters $\hat{\theta}_t$. Each row corresponds to a different setting. We estimate the more flexible reporting using a random forest whose hyperparameters are chosen via 10-fold cross-validation. All random forest covariate sets include the current value of the state. For the stock market, we additionally include 5 lags of the state, the risk-free rate, and the VIX. For unemployment, we additionally include 6 lags of the monthly first difference of the state, and 5 lags of the year-on-year change in the state in the given month of the year. For casualties, we additionally include 30 lags of the state. For rain, we additionally include 5 lags of the state, and the seasonal proxy $[\cos(\frac{2\pi t}{365}), \sin(\frac{2\pi t}{365})]$. See Appendix Table A1 for additional details.

Table A3: Most Common Words in Broadcast News Segments Covering State Realizations

All News Segments	Reports on State Variable				
	Stock market	Unemployment	Iraq military casualties	Afghanistan military casualties	Precipitation
economy	economy	economy	war	war	winter
campaign	watch	money	attacks	attacks	tomorrow's
war	money	watch	saddam	helicopter	watch
california	prices	matters	violence	crash	storm
united	interest	prices	fallujah	Wars	midwest
relations	rates	rate	unrest	pakistan	storms
vietnam	oil	campaign	al	taliban	heat
medicine	sales	rates	prisoner	bombing	severe
york	consumer	recession	interview	terrorism	cold
states	retail	vs	attack	violence	spring

Notes: The table displays the ten most frequently occurring words in segments flagged as reports for each topic, in descending order of occurrence frequency, with ties broken in ascending alphabetical order. We exclude all keywords that are used to explicitly flag segments as reports, in addition to all English stopwords, defined according to the NLTK package (Bird et al. 2009). We also manually filtered the following common words in segments that do not describe news content: “studio”, “report”, “shown”, “scenes”, “reporter”, “introduced”, “new”, “featured”, and “story”. See Appendix A for the keywords used to define reports for each state.


Figure A1: Reporting on Stock Returns, Sample Beginning in 1990



(a) Frequency of reporting vs. stock market return



(b) log(CRPS) vs. stock market return




(c) Frequency of reporting, by prior std. deviation



(d) log(CRPS), by prior std. deviation


Notes: The figure is analogous to Figure 1 but includes dates beginning in 1990. The unit of analysis is the outlet-date. Each panel depicts the mean of the variable on the y-axis within quantiles of the variable on the x-axis. Panels (a) and (b) include pointwise 95 percent confidence intervals clustered by date. In panels (c) and (d), observations are split by whether the standard deviation of priors is above or below the median. In each panel, the x-axis variable is the stock market return on the given date. In panels (a) and (c) the number of quantiles is chosen using the IMSE-optimal bandwidth selection of Cattaneo et al. (2024). In panels (b) and (d), the number of quantiles is chosen to match the number of quantiles in panels (a) and (c), respectively. In panels (a) and (c) the y-axis variable is an indicator for whether the given outlet reports on the stock market on the given date. In panels (b) and (d) the y-axis variable is the log(CRPS) for the given date.

Figure A2: Smoothed Monthly Marginal Probability of Reporting on Each State




Notes: The plot shows, for each topic, a 12-month equally weighted backward-looking moving average of the fraction of nightly news broadcasts that report on the state in each month, among all broadcasts occurring on relevant dates. For stock market reporting, the dashed portion of the series indicates dates outside of our main sample.

Figure A3: Coverage of Prior Distributions




(a) Observed data



(b) Data simulated from prior


Notes: Panel (a) shows, for each topic, the fraction of realizations that fall into each decile of the prior belief for the state variable. Panel (b) shows the analogue of panel (a) using data simulated from the prior distribution. For military casualties, because the event space is discrete, some realizations may cover multiple deciles. To correct for this, we uniformly distribute a state realization across all percentiles in the prior distribution covered by that particular count of military casualties, then aggregate to deciles.

Figure A4: Completeness and Restrictiveness of Reporting Models




Notes: The unit of analysis is the outlet-date. Each panel represents a different domain of interest. Plots depict the estimated completeness (left-column) and restrictiveness (right-column) of each of a set of reporting models, as described in Section 4.1. Estimation of completeness follows details in Appendix Table A1. Estimation of restrictiveness uses 100 datasets in which we replace R_t with the output of a random kernel function. The kernel function is seeded with reporting probabilities chosen uniformly randomly at points on a grid. The grid is chosen by taking equal-spaced points in the support of the realization ω_t and the parameters $\hat{\theta}_t$, with the number of points chosen so that the number of seed points is equal to the number of observations in the data sample. The kernel bandwidth on each dimension is selected as either the bandwidth chosen in Appendix Table A1 or the smallest bandwidth that permits a kernel value at each support point, whichever is larger. Restrictiveness is the ratio (truncated at 1) of the average mean squared error of each given reporting function to the average mean squared error of the constant prediction.

Figure A5: Reporting on US Military Casualties in Afghanistan




(a) Frequency of reporting vs. number of casualties



(b) log(CRPS) vs. number of casualties


Notes: The unit of analysis is the outlet-date. Panels (a) and (b) each depict the mean of the variable on the y-axis within quantiles of the variable on the x-axis, along with pointwise 95 percent confidence intervals clustered by date. In panels (a) and (b), the x-axis variable is the number of military casualties on the given date. In panel (a) the number of quantiles is chosen using the IMSE-optimal bandwidth selection of Cattaneo et al. (2024). In panel (b), the number of quantiles is chosen to match the number of quantiles in panel (a). In panel (a) the y-axis variable is an indicator for whether the given outlet reports on military casualties in Afghanistan on the given date. In panel (b) the y-axis variable is the log(CRPS) for the given date.

Figure A6: Naive Estimates of Reporting Bias




Notes: Each row corresponds to a separate regression and reports a naive estimate of bias along with its 95 percent confidence interval based on inference clustered by date. The naive estimate of bias is the coefficient on an independent variable of interest from a regression whose dependent variable is an indicator for reporting. For the topic “Stock market returns,” the unit of analysis is the outlet-date, and the independent variables of interest are, respectively, an indicator for negative returns, an indicator for whether the outlet reported on the stock market on the previous day, an indicator for a Republican presidential administration (estimated on the sample of dates with negative returns), and an indicator for a Democratic administration (estimated on the sample of dates with non-negative returns). For the topic “Unemployment rate,” the unit of analysis is the outlet-date, and the independent variables of interest are, respectively, an indicator for an increase in the rate relative to the previous rate, an indicator for a Republican presidential administration (estimated on the sample of dates with an increase in the rate), and an indicator for a Democratic administration (estimated on the sample of dates without an increase in the rate). For the topic “Military casualties,” the unit of analysis is an outlet-date-theater, and the independent variable of interest is an indicator for the Iraq theater.

Figure A7: Newsworthiness and the Probability of Reporting on US Military Casualties




Notes: The plot shows a binned scatter plot of the difference in the probability of reporting of military casualties between Iraq and Afghanistan (y-axis) versus the difference in the log(CRPS) between Iraq and Afghanistan (x-axis). The bins are ventiles and the average difference in the log(CRPS) between Iraq and Afghanistan in each ventile is marked on the x-axis in a rug plot.

Figure A8: Audit of Segment Classification




Notes: Panel (a) shows, for each of our main settings, the number of times (out of 20) that a human agreed with our parser's decision to categorize the segment in the given setting. Exact agreement corresponds to cases where the human flagged the exact type of news that the parser flagged (e.g., stock market news). Broad agreement corresponds to cases where the human flagged the broad type of news that the parser flagged (e.g., financial news). Panel (b) shows the number of times (out of 20) that a human agreed with our parser's decision to categorize the segment in *none* of our settings.

Figure A9: Information Factors and Reporting on Precipitation (Using Alternative Reporting Definition)



(a) Frequency of reporting vs. precipitation



(b) log(CRPS) vs. precipitation

Notes: The unit of analysis is the outlet-date. Panels (a) and (b) each depict the mean of the variable on the y-axis within quantiles of the variable on the x-axis, along with pointwise 95 percent confidence intervals clustered by date. In panels (a) and (b), the x-axis variable is the average precipitation on the given date. In panel (a) the number of quantiles is chosen using the IMSE-optimal bandwidth selection of Cattaneo et al. (2024). In panel (b), the number of quantiles is chosen to match the number of quantiles in panel (a). In panel (a) the y-axis variable is an indicator for whether the given outlet reports on precipitation on the given date. In panel (b) the y-axis variable is the log(CRPS) for the given date.