

NBER WORKING PAPER SERIES

WHAT IS NEWSWORTHY? THEORY AND EVIDENCE

Luis Armona
Matthew Gentzkow
Emir Kamenica
Jesse M. Shapiro

Working Paper 32512
<http://www.nber.org/papers/w32512>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2024, Revised October 2024

We thank our dedicated research assistants for their contributions to this project. We thank Julia Cagé, Ben Enke, Kris Nimark, Jon Roth, and seminar audiences at the Philadelphia Fed, Microsoft Research, and the Political Economy at Brown Conference for helpful comments. We acknowledge support from the Department of Defense at early stages of this project, and from the Stanford Institute for Economic Policy Research (SIEPR). Wharton Research Data Services (WRDS) was used in preparing this article. This service and the data available thereon constitute valuable intellectual property and trade secrets of WRDS and/or its third-party suppliers. BLS.gov cannot vouch for the data or analyses derived from these data after the data have been retrieved from BLS.gov. Health and Retirement Study (HRS) Public Release data files were used in preparing this article. The HRS (Health and Retirement Study) is sponsored by the National Institute on Aging (grant number NIA U01AG009740) and is conducted by the University of Michigan. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w32512>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2024 by Luis Armona, Matthew Gentzkow, Emir Kamenica, and Jesse M. Shapiro. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

What is Newsworthy? Theory and Evidence

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

NBER Working Paper No. 32512

May 2024, Revised October 2024

JEL No. C44, D83, L82

ABSTRACT

We introduce a model in which a benevolent news outlet decides whether to report the realization of a state to a consumer, who pays a cost to receive it. A simple statistical rule, called a proper scoring rule, describes when the outlet should be more likely to report the realization. Using data from the US television news, we show that a particular scoring rule successfully predicts many salient features of news reporting. We show how to use this rule as a control variable to discipline tests of reporting bias, and we show that controlling for it matters 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; Cummings et al. 2024) 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; Bursztyrn 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, 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 outlet reports the realization, the consumer takes the *ex-post* optimal action. If the outlet does not report the realization, 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); see also Shoemaker and Vos (2009) and Harcup and O’Neill (2017).

²More than a quarter of US adults still report obtaining news daily from network television, more than from cable

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 US 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 proposing a particular scoring rule, the continuous ranked probability score (CRPS), as a default measure of newsworthiness. The CRPS is widely used in forecast evaluation (Bröcker 2012) and is appealing for several reasons, including that it can be applied to both continuous and discrete random variables. Importantly, because a scoring rule depends only on the prior and the realization, the CRPS does not depend on the data on news reports.

Despite being highly disciplined, the measure of newsworthiness captures key qualitative features of the data well. These include intuitive features, such as the fact that broadcasts are more likely to report on the stock market when the market moves by a large amount. These also include more subtle features, such as the fact that small movements in the market become more newsworthy when volatility is high.

Given that our proposed measure of newsworthiness is parsimonious, grounded in a theory of benevolent reporting, and able to explain key patterns across a range of domains, we propose to use it as a control variable to discipline tests of bias in reporting. If a given pattern in reporting can be well explained by our measure of newsworthiness, we think this calls into question whether the pattern is evidence of a bias. This is especially true because our proposed measure of newsworthiness does not depend on aspects of the context, such as what is “good” or “bad” news, or which groups or places are affected by the given piece of news, that might be expected to cause bias.

We illustrate our proposed approach to disciplining tests for bias with two applications. The first is to the propensity to report negative unemployment news more than positive unemployment news. The second is 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 newsworthiness greatly lessens, or even eliminates, the estimated bias.

Our primary contribution is to propose and test a parsimonious measure of newsworthiness that is grounded in a model of benevolent, selective reporting. 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 news or online-only news sites, though less than from social media (Morning Consult 2022).

is verifiable but costly to report. Our model is deliberately stark, encoding only the instrumental value of information and the opportunity cost of transmitting it, and omitting many realistic features of news reporting. 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; 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 (see also Ozmen 2005). 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).

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

A state is realized from a state space Ω according to a commonly known prior distribution $F \in \Delta(\Omega)$, where $\Delta(\Omega)$ denotes the set of distributions on the state space. A news outlet observes the realization $\omega \in \Omega$ of the state. The outlet also observes the opportunity cost $\gamma \in \mathbb{R}_{\geq 0}$ of reporting, which is distributed according to a strictly increasing and continuous cumulative distribution function, Γ . We interpret γ as capturing, among other things, the value of reporting on other topics. The outlet then chooses whether or not to report the realization to a consumer. We denote the outlet’s action by $r \in \{0, 1\}$, where $r = 1$ denotes reporting.

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

$$L(a, \omega) + r\gamma$$

for $L : \mathcal{A} \times \Omega \rightarrow \mathbb{R}_{\geq 0}$ 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 (see, e.g., Gentzkow et al. 2015). We assume that $\min_{a \in \mathcal{A}} \mathbb{E}_{\omega \sim F}[L(a, \omega)]$ exists for all $F \in \Delta(\Omega)$, something that will be true of the class of losses that we consider in our applications.

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

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

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

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

Although we have assumed a common prior for simplicity, because the prior enters the reporting decision only via the consumer's action profile, the solution concept is preserved if the outlet knows the consumer's prior but believes it is misspecified.

Our notion of equilibrium does not allow the consumer to make inferences from whether the outlet reports or not. This will allow for a parsimonious empirical test, since we can construct all the inputs to the outlet's predicted behavior without using any data on reporting. Esponda and Vespa (2014), Martínez-Marquina et al. (2019), Enke (2020), and Jin et al. (2022), among others, find that many experimental participants do not make sophisticated inferences in similar situations. Appendix A extends the model to allow for more sophisticated updating and discusses the implications for our applications.

2.2 Characterization via Scoring Rules

A *reporting rule* $R : \Delta(\Omega) \times \Omega \rightarrow [0, 1]$ describes the probability of reporting as a function of the prior and realization. Any equilibrium strategies (a^*, r^*) induce an equilibrium reporting rule R^* that summarizes the probability that the opportunity cost will be low enough that the outlet chooses to report.

To characterize R^* , fix any selection $\alpha(F) \in \arg \min_{a \in \mathcal{A}} \mathbb{E}_{\omega \sim F}[L(a, \omega)]$ from the consumer's set of optimal actions. Let $\delta(\cdot)$ denote the Dirac measure on Ω , so that $\alpha(\delta(\omega))$ denotes the selected action when the consumer knows the realization. Now, for any prior $F \in \Delta(\Omega)$ and realization $\omega \in \Omega$, let

$$\rho(F, \omega) = L(\alpha(F), \omega) - L(\alpha(\delta(\omega)), \omega) \geq 0$$

denote the consumer's *regret* from choosing the *ex-ante* optimal action rather than the *ex-post*

optimal action. In equilibrium, the outlet is willing to report if and only if $\rho(F, \omega) \geq \gamma$. Because the distribution of γ is continuous, the equilibrium reporting rule R^* is given by

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

The equilibrium reporting rule is unique for a given selection $\alpha(\cdot)$.

We will show an important connection between equilibrium reporting rules and scoring rules. We first recall the standard definition of a scoring rule.

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 $\omega \in \Omega$.

Scoring rules are widely studied for their use in evaluating and improving probabilistic forecasts (Gneiting and Raftery 2007; Gneiting and Katzfuss 2014).³ 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 a forecaster 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 forecaster 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.

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 $L(\cdot, \cdot)$, $\mathcal{A}$, and $\Gamma(\cdot)$ 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 the action space $\mathcal{A}$ to be the set $\Delta(\Omega)$ of distributions on the state space, take $\Gamma(\cdot) = G(\cdot)$, take the loss to be the score, $L(a, \omega) = S(a, \omega)$, and take the selection to be $\alpha(F) = F$.

³We use negatively oriented scoring rules to simplify statements.

It is immediate that $\rho(F, \omega) = S(F, \omega)$ and hence that $R^*(F, \omega) = G(S(F, \omega))$. For the “only if” direction, fix some primitives $L(\cdot, \cdot)$, $\mathcal{A}$, and $\Gamma(\cdot)$, and some selection $\alpha(\cdot)$. Notice that the regret function has the properties that $\rho(\delta(\omega), \omega) = 0$ for all $\omega \in \Omega$, and that $E_{\omega \sim F} \rho(F', \omega) \geq E_{\omega \sim F} \rho(F, \omega)$ for all $F, F' \in \Delta(\Omega)$. This implies that the regret function is itself a normalized proper scoring rule. The result then follows because $R^*(F, \omega) = \Gamma(\rho(F, \omega))$ and $\Gamma(\cdot)$ is a strictly increasing and continuous CDF. $\square$

2.3 A Default Scoring Rule

We propose the following normalized proper scoring rule as a default for real-valued states.

Definition 3. *The **continuous ranked probability score (CRPS)** is a scoring rule S_0 such that*

$$S_0(F, \omega) = \int_{\Omega} (F(x) - \delta(x; \omega))^2 dx,$$

where $\Omega \subseteq \mathbb{R}$ and where $\delta(x; \omega)$ is an indicator for whether $\omega \leq x$.

To interpret the CRPS, imagine there is a binary state and a consumer must guess the probability that the state takes the value one (rather than zero). If the consumer’s loss is given by the squared difference between their guess and the realization, it is well known that the consumer has an incentive to state their true belief. The CRPS extends this logic to any real-valued state, by recognizing that we can describe any distribution as a set of probabilities over binary states, each indicating whether the realization is less than or equal to some value x in the domain of the state.

We propose the CRPS as a default rule for a few reasons. First, because the CRPS does not require a density, it can apply to continuous distributions, discrete distributions, and discrete-continuous distributions. Second, the CRPS has no tuning parameters, and is often available in closed form for common distributions.⁴ Third, fixing a realization and a prior, the CRPS is invariant to contextual factors that may lead to bias, such as whether a high realization of the state is “bad news” (as in the unemployment rate) or “good news” (as in the stock market return), or which population or place the state affects. These properties make the CRPS appealing as a starting point for studying news reporting and disciplining claims of bias. Appendix D reports on the fit of an alternative to the 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

⁴(Gneiting and Raftery 2007) discuss these and other properties of the CRPS.

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.⁵

We consider four state variables: the stock market return, the unemployment rate, the average precipitation, and the number of 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 B provides details on our choice of keywords for each state variable, as well as evidence on the quality of the resulting classification, and descriptive statistics on patterns of reporting over time.

3.2 State Variables and Prior Beliefs

We wish to base our analysis on prior beliefs that are reasonable, disciplined by the data, and able to vary over time as circumstances change. We therefore specify prior beliefs for each state using a flexible parametric time-series model, with a structure motivated by prior research, complexity parameters chosen by cross validation, and remaining parameters estimated 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 C. Appendix C also discusses evidence on the performance of the estimated priors, and Appendix D shows results using alternative prior specifications.

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 CBOE Volatility Index (VIX), and the risk-free interest rate r_t , for each date t .⁶ We consider those dates t for which SPX_t , r_t , and VIX_t are available. Because there is limited variation in reporting in the early part of this sample, we conduct our main analysis on data beginning in 2001. Appendix C reports results for the sample that begins in 1990.⁷

⁵We exclude data from March 1999 and February 2000 because VTNA is missing data on a substantial number of broadcasts during these months.

⁶The VIX is an options-based volatility index, quoted in units of standard deviations. The 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.

⁷The VIX began trading in 2004 but is available as a historical series beginning in 1990.

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).⁸ 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 C.1).

3.2.3 Precipitation

We obtain information on precipitation from the US 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. We match each station to its county using information from the Enhanced Master Station History Report (Vose et al. 2018).

⁸We 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 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 $(P, Q, S) = (1, 0, 3)$.

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 Kedeem et al. (1990) in modeling the average precipitation as lognormal with seasonally varying parameters (Appendix C.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 1, 2003 through December 31, 2011. 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 C.3).

4 Evidence on Determinants of Reporting

In this section, we study whether our proposed default scoring rule, the CRPS, captures sensible and relevant patterns in news reporting. For each state variable, let 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 $t \in \{1, \dots, T\}$.

4.1 Methods

For each state, we present a series of plots that compare the observed patterns in reporting to those predicted by the CRPS. The plots include, (a) a plot of the reporting frequency R_t against the log of the CRPS, (b) a plot of the log of the CRPS against the realization ω_t , and (c) a plot of the reporting frequency R_t against the realization ω_t . Each plot is a binned scatterplot with quantile-spaced bins (Cattaneo et al. 2024). We calculate the CRPS on each date t as $S_0(\hat{F}_t, \omega_t)$. We plot the log of the CRPS to facilitate comparison of magnitudes across states.

As a simple measure of the bivariate relationship between the reporting frequency and the CRPS, we report the coefficient and standard error from a linear regression of the reporting frequency R_t on the log of the CRPS. The coefficient can be interpreted as the change in the probability of reporting associated with a 100 log-point increase in newsworthiness, as measured by the CRPS. The linear regression restricts the functional form of the relationship, and does not afford a comparison of the predictive power of the CRPS to that of other models. Appendix D therefore reports estimates of the completeness (Fudenberg et al. 2022) and restrictiveness (Fudenberg et al. 2023) of a set of reporting rules, including a smooth, monotone function of the CRPS.

4.2 Findings

4.2.1 Figure 1: Stock Market Return

Panel (a) shows that reporting is more likely on dates when the CRPS is high. The remaining panels unpack this finding by looking separately at the roles of the realization and the prior, the two determinants of the CRPS.

Panels (b) and (c) illustrate the role of the realization. Panel (b) shows that reporting is more likely on dates when the return is greater in absolute magnitude. Panel (c) shows that this pattern is precisely what is predicted by the CRPS.

Panels (d) and (e) illustrate the role of the prior. These panels separate the plots in panels (b) and (c) according to whether the prior standard deviation is high or low. Less extreme returns are more newsworthy, and more extreme returns are less newsworthy, when the market is volatile. This pattern is visible in news reporting (panel (d)) and is predicted by the CRPS (panel (e)).

4.2.2 Figure 2: Unemployment Rate

Panel (a) shows that reporting is more likely on dates when the CRPS is high. Panels (b) and (c) show, respectively, that both the frequency of reporting, and the log of the CRPS, tend to be higher on dates when the unemployment rate is higher.

Panels (d) and (e) illustrate the role of the prior. These panels repeat the plots in panels (b) and (c) but replace the unemployment rate with the change in the unemployment rate relative to its previous level. The previous level of the unemployment rate, in turn, affects the current prior. Panel (d) shows that reporting is more likely when the change in the unemployment rate is greater in absolute value. Panel (e) shows that this pattern is predicted by the CRPS, because larger changes tend to be associated with larger surprises, relative to the prior.

4.2.3 Figure 3: Precipitation

Panel (a) shows that reporting is more likely on dates when the CRPS is high. Panels (b) and (c) show that both the frequency of reporting, and the log of the CRPS, tend to be higher on dates when precipitation is either very high or very low.

4.2.4 Figure 4: US Military Casualties

Panels (a) and (b) show that reporting is more likely on dates when the CRPS is high, though this pattern is more pronounced for casualties in Iraq (panel (a)) than for casualties in Afghanistan (panel (b)). Panels (c) and (d) for Iraq, and (e) and (f) for Afghanistan, show that both the frequency of reporting, and the log of the CRPS, tend to be higher on dates when casualties are greater.

5 Application to Testing for Bias

We propose to use the CRPS as a default control variable in studies of bias in reporting. In this section we elaborate on this proposal and illustrate it with applications.

5.1 Concepts and Methods

Suppose that a researcher hypothesizes a form of reporting bias and selects a variable that proxies for it in the data. For example, the bias might be a bias towards reporting negative news, and the variable might be an indicator for whether the realization of a state on a given day is “bad.” A natural way to test for reporting bias is to assess whether the proxy variable is associated with reporting. If it is, this is consistent with, though not proof of, the presence of a bias.

We propose to refine this test by controlling for newsworthiness. If controlling for newsworthiness substantially reduces the estimated relationship between the proxy variable and reporting, this casts doubt on interpreting the relationship as evidence of a bias. If, on the other hand, controlling for newsworthiness does not much affect the estimated relationship between the proxy variable and reporting, then this is consistent with, though still not proof of, the presence of a bias.

We propose the CRPS as an appealing default proxy for newsworthiness. The CRPS requires no inputs other than the realization and the prior, and captures the intuitive idea that more surprising realizations are more newsworthy. Fixing the realization and prior, the CRPS does not depend on details of the context, such as whether a high value of the state is “bad” or “good” news, that might proxy for causes of bias. And, as we have shown, the CRPS does a good job of capturing many salient features of real-world reporting decisions.

Because our theory does not prescribe a specific functional relationship between the CRPS and reporting, we propose to estimate this relationship flexibly. In the applications below, we control for a smooth, monotone function of the CRPS. We estimate this function via kernel regression with Epanechnikov kernel, bandwidth chosen via 10-fold block cross-validation (Racine 2000), and the CRPS entering with a log transformation. We enforce monotonicity via rearrangement following Chernozhukov et al. (2009).

To select our applications, we began with a broad set of biases that are discussed in the literature and are testable in our data. In each case we selected a proxy variable and assessed its relationship to reporting, without controlling for the CRPS. In two cases, this relationship was statistically significant. We focus on these two applications in what follows. Appendix E details the full set of cases that we considered.

5.2 Application to Negativity Bias in Economic Reporting

Our first application is to negativity bias, the hypothesis that reporting is more likely when the news is “bad.” We study negativity bias in the context of reporting on the unemployment rate (as in, e.g., Soroka [2006]). As a proxy for negativity bias, we construct an indicator x_t for whether the unemployment rate on a given date t is strictly greater than on the previous date on which it

was reported. An ordinary least squares regression of the frequency R_t of reporting on the proxy x_t implies that reporting is a statistically significant 5.4 percentage points more likely when the unemployment rate has risen (see Figure 2(f)).

There are two reasons to suspect that some of the apparent negativity bias can be explained by newsworthiness as we have defined it. First, the unemployment rate is more newsworthy when it is different from its previous value (see Figure 2(e)). The indicator x_t is equal to one only when the unemployment rate is strictly greater than its previous value, which can occur only when the current and previous values are different. Therefore, the proxy variable conflates negativity with change, and change is newsworthy.

Second, the unemployment rate is more newsworthy when it is more different from its previous value (see again Figure 2(e)). This matters because large increases in the unemployment rate are more common than large decreases, but are still rare enough to be surprising. This fact is visible from Figure 2(e), which shows that quantiles of the change in the unemployment rate include more large increases than large decreases. Because increases in unemployment tend to be larger than decreases, increases tend to be more newsworthy.

For both of these reasons, we expect that controlling for newsworthiness will reduce the estimated negativity bias. That is what we find. When we control for a smooth monotone function of the CRPS, the estimated bias falls from 5.4 to 2.4 percentage points and becomes statistically insignificant (see again Figure 2(f)). In this application, then, newsworthiness accounts for much of the apparent bias, and after accounting for newsworthiness, there is no statistical evidence of a bias towards reporting negative news.

5.3 Application to Context Bias in War Reporting

Our second application is to bias due to context. Prior work finds that the war in Iraq received more media attention in the US than the war in Afghanistan (see, e.g., Jacobson [2010]). As a measure of bias towards reporting on Iraq, we calculate the difference in the frequency R_t of reporting on casualties in each war on each date t . An ordinary least squares regression of this difference on a constant implies that reporting on casualties in Iraq is a statistically significant 10 percentage points more likely than reporting on casualties in Afghanistan (see Figure 4(g)).

There are reasons to suspect that some of the apparent bias towards reporting on casualties in Iraq can be explained by newsworthiness as we have defined it. Specifically, US casualties tended to be more variable and unpredictable in Iraq than in Afghanistan.⁹ As a result, the CRPS for casualties in Iraq tends to be greater than the CRPS for casualties in Afghanistan. This fact is visible from the rug plot along the x-axis of Figure 4(h). The binned scatterplot in Figure 4(h) also shows that the frequency of reporting tends to be similar between Iraq and Afghanistan on dates when the CRPS is similar.

We therefore expect that controlling for newsworthiness will reduce the estimated bias towards

⁹For example, during our sample period, the standard deviation of US casualties was 67.4 percent larger in Iraq than in Afghanistan.

reporting on casualties in Iraq. That is what we find. When we control for a smooth monotone function of the CRPS, the estimated bias falls from 10 to -1.1 percentage points and becomes statistically insignificant. In this application, as in our earlier one, then, our measure of newsworthiness accounts for much of the apparent bias and, after accounting for newsworthiness, there is no statistical evidence of a bias towards reporting on casualties in Iraq relative to Afghanistan.

5.4 Takeaways

In these two applications, newsworthiness can account for much of the apparent bias. These findings do not mean that news media are unbiased, but they do show the value of controlling for newsworthiness in a disciplined and parsimonious way.

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. *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.
- R. Cummings, G. Fraccaroli, and N. Mahoney. Bad News Bias in Gasoline Price Coverage. *Briefing Book*, 2024. URL <https://web.archive.org/web/20240922022344/https://www.briefingbook.info/p/bad-news-bias-in-gasoline-price-coverage>. Accessed in September 2024.
- 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.
- I. Esponda and E. Vespa. Hypothetical Thinking and Information Extraction in the Laboratory. *American Economic Journal: Microeconomics*, 6(4):180–202, 2014.
- 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 US 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)*, page 1830–1836. IEEE Press, 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*, pages 1–50, 11 2023.


- 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.
- T. Harcup and D. O’Neill. What is News? News Values Revisited (Again). *Journalism Studies*, 18(12):1470–1488, 2017.
- D. E. Harrington. Economic News on Television: The Determinants of Coverage. *Public Opinion Quarterly*, 53(1):17–40, 1989.
- Health and Retirement Study. Health and Retirement Study, Public Survey Data public use dataset. Produced and distributed by the University of Michigan with funding from the National Institute on Aging (grant number NIA U01AG009740). Ann Arbor, MI, (2024). https://hrsdata.isr.umich.edu/data-products/public-survey-data?_gl=1*1rr3m3j*_ga*0DY5MjAzNzg5LjE3MjQ2Mzg5MTI.*_ga_FF28MW3MW2*MTcyNjc3MTc0NS42LjEuMTcyNjc3MTkzOS4wLjAuMA.., 2008, 2010 and 2012. Original data from Health and Retirement Study. Accessed in September 2024.
- 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.
- G. Z. Jin, M. Luca, and D. Martin. Complex Disclosure. *Management Science*, 68(5):3236–3261, 2022.
- 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.
- A. Martínez-Marquina, M. Niederle, and E. Vespa. Failures in Contingent Reasoning: The Role of Uncertainty. *American Economic Review*, 109(10):3437–3474, 2019.
- 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 US 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 US 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.
- P. J. Shoemaker and T. Vos. *Gatekeeping Theory*. Routledge, 2009.
- 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.
- University of Michigan Surveys of Consumers. University of Michigan, Survey Research Center, Surveys of Consumers. <https://data.sca.isr.umich.edu/faq.php>, 2024. Original data from University of Michigan Surveys of Consumers. Accessed in September 2024.
- 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 October 2017 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 the Stock Market Return



Notes: The unit of analysis is the outlet-date. Each plot is a binscatter. Panels (a), (b), and (c) include 95 percent uniform confidence bands based on inference clustered by month following Cattaneo et al. (2024). In panels (d) and (e), observations are split by whether the prior standard deviation is above or below its sample median. In panel (a), the x-axis variable is $\log(\text{CRPS})$ for the stock market return on the given date. In panels (b), (c), (d), and (e), the x-axis variable is the stock market return on the given date. In panels (a), (b), and (d) the spacing of bins is chosen using the IMSE-optimal bandwidth selection of Cattaneo et al. (2024). In panels (c) and (e), the spacing of bins is chosen to match that in panels (b) and (d), respectively. In panels (a), (b), and (d) the y-axis variable is an indicator for whether the given outlet reports on the stock market on the given date. In panels (c) and (e), the y-axis variable is $\log(\text{CRPS})$ for the given date. The box in panel (a) reports the coefficient from a regression of reporting on $\log(\text{CRPS})$, with standard error in parentheses based on inference clustered by month.


Figure 2: Reporting on the Unemployment Rate




Notes: The unit of analysis is the outlet-date. Each plot in panels (a), (b), (c), (d), and (e) is a binscatter with 95 percent uniform confidence band based on inference clustered by month following Cattaneo et al. (2024). In panel (a), the x-axis variable is $\log(\text{CRPS})$ for the unemployment rate on the given date. In panels (b) and (d), 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 panels (a), (b), and (d), the spacing of bins is chosen using the IMSE-optimal bandwidth selection of Cattaneo et al. (2024). In panels (c) and (e), the spacing of bins is chosen to match that in panel (b) and (d), respectively. In panels (a), (b), and (d), the y-axis variable is an indicator for whether the given outlet reports on the unemployment rate on the given date. In panels (c) and (e), the y-axis variable is $\log(\text{CRPS})$ for the given date. The box in panel (a) reports the coefficient from a regression of reporting on $\log(\text{CRPS})$, with standard error in parentheses based on inference clustered by month. In panel (f), we report the coefficient and confidence interval from a regression of reporting on an indicator for whether the unemployment rate is larger than the previous period, with and without a control for newsworthiness as defined in Section 5, with a 100-replicate bootstrap-based p-value for the equality of the two coefficients, clustered by month.

Figure 3: Reporting on Precipitation


(a) Frequency of reporting vs. $\log(\text{CRPS})$



(b) Frequency of reporting vs. precipitation




(c) $\log(\text{CRPS})$ vs. precipitation



Notes: The unit of analysis is the outlet-date. Each plot is a binscatter with 95 percent uniform confidence band based on inference clustered by month following Cattaneo et al. (2024). In panel (a), the x-axis variable is $\log(\text{CRPS})$ for the average precipitation on the given date. In panels (b) and (c), the x-axis variable is the average precipitation on the given date. In panels (a) and (b), the spacing of bins is chosen using the IMSE-optimal bandwidth selection of Cattaneo et al. (2024). In panel (c), the spacing of bins is chosen to match that in panel (b). In panels (a) and (b), the y-axis variable is an indicator for whether the given outlet reports on precipitation on the given date. In panel (c), the y-axis variable is $\log(\text{CRPS})$ for the given date. The box in panel (a) reports the coefficient from a regression of reporting on $\log(\text{CRPS})$, with standard error in parentheses based on inference clustered by month.

Figure 4: Reporting on US Military Casualties



Notes: The unit of analysis is the outlet-date. Each plot in panels (a), (b), (c), (d), (e), and (f) is a binscatter with 95 percent uniform confidence band based on inference clustered by month following Cattaneo et al. (2024). In panels (a) and (b), the x-axis variable is $\log(\text{CRPS})$ for the number of military casualties on the given date. In panels (c), (d), (e), and (f), the x-axis variable is the number of military casualties on the given date. In panels (a), (b), (c), and (e), the spacing of bins is chosen using the IMSE-optimal bandwidth selection of Cattaneo et al. (2024). In panels (d) and (f), the spacing of bins is chosen to match that in panel (c) and (e), respectively. In panels (a), (b), (c), and (e), the y-axis variable is an indicator for whether the given outlet reports on military casualties on the given date. In panels (d) and (f), the y-axis variable is $\log(\text{CRPS})$ for the given date. The box in panel (a) reports the coefficient from a regression of reporting on $\log(\text{CRPS})$, with standard error in parentheses based on inference clustered by month. In panel (f), we report the coefficient and confidence interval from a regression of the difference in reporting between casualties in Iraq and Afghanistan on a constant, with and without a control for the difference in newsworthiness as defined in Section 5, with a 100-replicate bootstrap-based p-value for the equality of the two coefficients, clustered by month. Panel (h) shows a binscatter of the difference in the probability of reporting of military casualties between Iraq and Afghanistan (y-axis) versus the difference in $\log(\text{CRPS})$ between Iraq and Afghanistan (x-axis). The bins are ventiles and the average difference in $\log(\text{CRPS})$ between Iraq and Afghanistan in each ventile is marked on the x-axis in a rug plot.

A Analysis with Sophisticated Updating

A.1 Model, Solution Concept, and Characterization

Here we define and characterize equilibrium in a model with a sophisticated (Bayesian) consumer who makes inferences about the state from the outlet's decision not to report. We suppose that $L : \mathcal{A} \times \Omega \rightarrow [0, \bar{L}]$ where $\Gamma(\bar{L}) < 1$.

Definition 4. A pair of pure strategies $a^* : \Delta(\Omega) \times \Omega \times \{0, 1\} \rightarrow \mathcal{A}$ and $r^* : \Delta(\Omega) \times \Omega \times \mathbb{R}_{\geq 0} \rightarrow \{0, 1\}$ is a *sophisticated equilibrium* if and only if

$$\begin{aligned} a^{**}(F, \omega, 1) &\in \arg \min_{a \in \mathcal{A}} L(a, \omega), \quad a^{**}(F, \omega, 0) \in \arg \min_{a \in \mathcal{A}} \mathbb{E}_{x \sim F_0} [L(a, x)], \text{ and} \\ r^{**}(F, \omega, \gamma) &\in \arg \min_{r \in \{0, 1\}} [r(L(a^{**}(F, \omega, 1), \omega) + \gamma) + (1 - r)L(a^{**}(F, \omega, 0), \omega)] \end{aligned}$$

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

$$F_0(x; r^{**}) = \frac{F(x) \mathbb{E}_{\gamma \sim \Gamma, x' \sim F} [1 - r^{**}(F, x', \gamma) | x' \leq x]}{\mathbb{E}_{\gamma \sim \Gamma, x' \sim F} [1 - r^{**}(F, x', \gamma)]}.$$

The restriction that $\Gamma(\bar{L}) < 1$ guarantees that the denominator in the definition of F_0 is nonzero.

Any sophisticated equilibrium a^{**}, r^{**} induces a unique reporting rule $R^{**} = \Gamma(\rho(F_0(r^{**}), \omega))$ for a given selection $\alpha(\cdot)$. Under a regularity condition on the consumer's decision problem, when reporting is rare any such reporting rule is close to the equilibrium reporting rule R^* implied by the baseline model.

Proposition 2. Suppose that $L(\alpha(F), \omega)$ is continuous in the sup norm for all $\omega \in \Omega$. Then for any ω, F , and $\epsilon > 0$, there exists $\Delta \in (0, 1)$ such that the reporting probability $R^{**}(F, \omega)$ induced by any sophisticated equilibrium is within ϵ of the equilibrium reporting probability $R^*(F, \omega)$ whenever $\Gamma(\bar{L}) < 1 - \Delta$.

Proof. Pick some ω, F and some $\epsilon > 0$. Pick some sophisticated equilibrium that induces reporting probability $R^{**}(\omega, F)$. Recall that, for given selection $\alpha(\cdot)$, the model primitives imply a unique equilibrium reporting probability $R^*(\omega, F)$. Write the difference in reporting probabilities as

$$|R^{**}(\omega, F) - R^*(\omega, F)| = |\Gamma(\rho(F, \omega)) - \Gamma(\rho(F_0(r^{**}), \omega))|.$$

By the definition of $\rho(\cdot)$ and the continuity of $\Gamma(\cdot)$, it suffices to show that $L(\alpha(F), \omega)$ and $L(\alpha(F_0(r^{**})), \omega)$ are sufficiently close when Δ is sufficiently large. Towards such a result, note that

$$F_0(x; r^{**}) = \frac{F(x) \mathbb{E}_{\gamma \sim \Gamma, x' \sim F} [1 - r^{**}(F, x', \gamma) | x' \leq x]}{\mathbb{E}_{\gamma \sim \Gamma, x' \sim F} [1 - r^{**}(F, x', \gamma)]},$$

where both $\mathbb{E}_{\gamma \sim \Gamma, x' \sim F} [1 - r^{**}(F, x', \gamma) | x' \leq x]$ and $\mathbb{E}_{\gamma \sim \Gamma, x' \sim F} [1 - r^{**}(F, x', \gamma)]$ are bounded below

by Δ by the definitions of Δ and r^{**} . We therefore have that

$$F_0(x; r^{**}) \in \left[F(x) \frac{\Delta}{1}, F(x) \frac{1}{\Delta} \right]$$

so that

$$|F_0(x; r^{**}) - F(x)| \leq F(x) \left(\frac{1 - \Delta}{\Delta} \right)$$

and

$$\sup_{x \in \Omega} |F_0(x; r^{**}) - F(x)| \leq \frac{1 - \Delta}{\Delta}.$$

The desired result then follows by the continuity of $L(\alpha(F), \omega)$ in the sup norm. $\square$

Proposition 2 states that, when reporting is rare—for example, because there are many possible states on which to report—the outlet’s behavior is similar whether or not the consumer makes sophisticated inferences from nonreporting. Intuitively, this is because when reporting is rare, nonreporting conveys limited information about the realization of the state.


A.2 Evidence on Inferences from Nonreporting

Here we consider the scope in our applications for rational inferences from the outlet’s decision not to report. For each state variable, we estimate the empirical reporting rule $R(\omega, \hat{F})$ using the kernel methods described in Appendix D. For each state variable, we use the estimated reporting rule $\hat{R}(\omega, \hat{F})$ and the estimated prior $\hat{F}$ to numerically approximate the posterior distribution F_0 that a sophisticated consumer would assign to the state variable after learning that the outlet decided not to report on it. We calculate the CRPS based on the estimated posterior distribution $\hat{F}_0$ and compare it to CRPS based on the estimated prior distribution $\hat{F}$.

Appendix Figure A1 presents the results of this exercise. The correlation between the log of the CRPS based on the prior, and the log of the CRPS based on the posterior, is nearly one for all state variables. In practice, then, our proposed measure of newsworthiness is very similar regardless of whether consumers do or do not make inferences based on an outlet’s decision not to report on the given state.

To unpack this finding further, Appendix Figure A1 shows statistics of the standardized absolute difference in means between the prior and posterior distributions, $|\mathbb{E}_{\omega \sim \hat{F}_0}(\omega) - \mathbb{E}_{\omega \sim \hat{F}}(\omega)| / \sqrt{\text{Var}_{\omega \sim \hat{F}}(\omega)}$. We can think of the difference in means as a measure of the difference in optimal decisions under the prior and posterior for a consumer with a quadratic loss. Scaling by the prior standard deviation facilitates comparison across states. There are some dates on which the standardized difference in means is meaningful. For example, when the stock market is very volatile, reporting becomes fairly likely (Figure 1(d)), so that nonreporting indicates small returns. But, overall, the standardized differences in means tend to be small.

Figure A1: Inference from Nonreporting



Notes: The unit of analysis is the date. The diamonds report the correlation between $\log(\text{CRPS})$ calculated based on the prior distribution for the state and $\log(\text{CRPS})$ calculated based on the posterior conditional on the outlet not reporting the state, following the methods described in Appendix A. The boxplots report statistics of the distribution of the absolute difference between the prior and posterior means, standardized by dividing by the prior standard deviation. The hollow triangles report the fraction of observations excluded from the other calculations due to having a probability of nonreporting below 0.002.

B Details on Definitions 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 A1 reports the most common non-keyword terms associated with each state. Appendix Figure A2 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. Appendix Figure A3 presents descriptive statistics on patterns of reporting over time.

B.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 “stks.” or “stk.”.

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

B.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,” “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).
- 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.”

B.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 A4 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 67.6% of our regularly coded weather reports.

Table A1: Most Common Words in Broadcast News Segments Covering State Realizations

All News Segments	Reports on State Variable				
	Stock market	Unemployment	US military casualties in Iraq	US military casualties in Afghanistan	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 B for the keywords used to define reports for each state.

Figure A2: 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 A3: 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 A4: Reporting on Precipitation, Alternative Reporting Definition



Notes: The unit of analysis is the outlet-date. Each plot is a binscatter with 95 percent uniform confidence band based on inference clustered by month following Cattaneo et al. (2024). In panel (a), the x-axis variable is $\log(\text{CRPS})$ for the average precipitation on the given date. In panels (b) and (c), the x-axis variable is the average precipitation on the given date. In panels (a) and (b), the spacing of bins is chosen using the IMSE-optimal bandwidth selection of Cattaneo et al. (2024). In panel (c), the spacing of bins is chosen to match that in panel (b). In panels (a) and (b), the y-axis variable is an indicator for whether the given outlet reports on precipitation on the given date, using the alternative definition described in Appendix B.4. In panel (c), the y-axis variable is $\log(\text{CRPS})$ for the given date. The box in panel (a) reports the coefficient from a regression of reporting on $\log(\text{CRPS})$, with standard error in parentheses based on inference clustered by month.

C Details on Parameterization of Prior Distributions

C.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.¹⁰ We select the values of (P_1, Q_1, P_2, Q_2) via k -fold h -block cross-validation (Racine 2000).¹¹ We let $\hat{\theta}_t$ denote the estimated mean and variance of the distribution of ω_t in month t .

C.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 .¹² 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:

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

¹⁰We 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.

¹¹We 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 $(P_1, Q_1, P_2, Q_2) = (4, 0, 2, 0)$.

¹²We 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.¹³ We optimize the P 's and Q 's via k -fold h -block cross-validation.¹⁴ We let $\hat{\theta}_t$ denote the estimated mean and variance of the distribution of $\log(\omega_t)$ on date t .

C.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

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

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

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.¹⁵ We let $\hat{\theta}_t$ denote the estimated parameters of the distribution of ω_t on date t .

C.4 Coverage and Comparison to Survey Evidence

Appendix Figure A5 reports on the coverage of the prior distributions $\hat{F}_t$. We find that our prior distributions tend to understate the probability of tail realizations.

Appendix Figure A6 compares the evolution of features of the estimated prior distributions $\hat{F}_t$ to related concepts measured in survey data. Panel (a) of Appendix Figure A6 considers the stock return. Here, we compare the implied volatility under the prior on a twelve-month horizon to the probability assigned to an extreme movement in the stock market over that same horizon by an average respondent to the Health and Retirement Study (2008, 2010 and 2012). Both series are

¹³We initialize the mean by assuming that $\epsilon_t = 0$ on dates t prior to January 1st, 1970. This implies that $(\log(\omega_t) - \vec{x}_t \vec{\beta}_1) = 0$ 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.


¹⁴We 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 $(P_1, Q_1, P_2, Q_2) = (4, 1, 2, 1)$.

¹⁵We 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 31, 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 $(P, Q) = (3, 30)$.


elevated during periods of elevated market risk, such as the May 2010 “flash crash,” the September 2011 decline due to the European debt crisis, and the selloffs in June 2012 and November 2012.

Panel (b) of Appendix Figure A6 considers the unemployment rate. Here, we compare the implied mean change under the prior over a twelve-month horizon to the difference in the fraction of respondents to the University of Michigan Surveys of Consumers (2024) reporting that they expect the unemployment rate will increase and the fraction reporting that they expect the rate will decrease over that same horizon. The two series comove closely.

Figure A5: 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 A6: Comparing Estimated Prior Beliefs to Survey Reports

(a) Stock returns




(b) Unemployment rate



Notes: Panel (a) shows the standard deviation in the stock return over the next twelve months implied by our estimated prior ("Prior standard deviation"), as well as the average probability that respondents to the Health and Retirement Study (2008, 2010 and 2012) report assigning to the scenario that the stock market changes by more than 20 percent over the next twelve months ("Consumer forecast"). We assign each probability to the month in which the respondent answers the survey; not all months are covered by the survey. Panel (b) shows the expected change in the unemployment rate over the next twelve months implied by our estimated prior ("Prior forecast"), as well as the difference in the fraction of respondents to the University of Michigan Surveys of Consumers (2024) reporting that the unemployment rate will increase over the next twelve months and the fraction reporting that it will decrease over the next twelve months ("Consumer forecast"), excluding respondents who do not provide a forecast.

Figure A7: Reporting on the Stock Market Return, Sample Beginning in 1990



Notes: The figure is analogous to Figure 1 but includes dates beginning in 1990. The unit of analysis is the outlet-date. Each plot is a binscatter. Panels (a), (b), and (c) include 95 percent uniform confidence bands based on inference clustered by month following Cattaneo et al. (2024). In panels (d) and (e), observations are split by whether the prior standard deviation is above or below its sample median. In panel (a), the x-axis variable is $\log(\text{CRPS})$ for the stock market return on the given date. In panels (b), (c), (d), and (e), the x-axis variable is the stock market return on the given date. In panels (a), (b), and (d) the spacing of bins is chosen using the IMSE-optimal bandwidth selection of Cattaneo et al. (2024). In panels (c) and (e), the spacing of bins is chosen to match that in panels (b) and (d), respectively. In panels (a), (b), and (d), the y-axis variable is an indicator for whether the given outlet reports on the stock market on the given date. In panels (c) and (e), the y-axis variable is $\log(\text{CRPS})$ for the given date. The box in panel (a) reports the coefficient from a regression of reporting on $\log(\text{CRPS})$, with standard error in parentheses based on inference clustered by month.

D Estimating Completeness and Restrictiveness

D.1 Completeness

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 rule $R(\omega_t, \hat{F}_t)$, we define the fit to be the improvement in sum of squared errors relative to a constant prediction,

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

where $\bar{R}$ is the sample mean of the reporting frequency R_t . We define the completeness to be the ratio of the fit of the given reporting rule to the fit of a more flexible benchmark.

We consider three reporting rules. The first reporting rule is a smooth, monotone function of the CRPS $S_0(\omega_t, \hat{F}_t)$. This reporting rule captures the default scoring rule that we propose. The second rule function is a smooth function of the realization ω_t . This reporting rule captures a model in which reporting depends only on the realization and not on the prior. The third reporting rule is a smooth function of the realization ω_t and of the estimated prior $\hat{F}_t$. This reporting rule captures a model in which reporting depends flexibly on both the realization and the prior, and is the flexible benchmark against which we measure completeness.

We estimate each reporting rule using a kernel regression with Epanechnikov kernel and bandwidth chosen via 10-fold block cross-validation (Racine 2000). When appropriate, we enforce monotonicity via rearrangement following Chernozhukov et al. (2009), and we incorporate the CRPS with a log transformation. For the flexible benchmark, we make the estimation problem finite-dimensional by representing the estimated prior $\hat{F}_t$ as a function of a finite-dimensional vector $\hat{\theta}_t$ of estimated parameters.

For each reporting rule, our estimate of completeness is the ratio of the fit of the rule to the fit of the flexible benchmark, truncated at zero and one. Appendix Figure A8 reports these estimates.

In the case of stock returns, a smooth, monotone function of the CRPS achieves a completeness of 1, meaning that it achieves the same fit as the flexible benchmark. In the case of unemployment, the completeness is 0.407, meaning that a smooth monotone function of the CRPS achieves 0.407 of the fit of the flexible benchmark. In the case of precipitation, the completeness is 0.194. In the case of US military casualties in Iraq and Afghanistan, the completeness is 0.657 and 0.872 respectively.

Appendix Table A2 reports additional specifications. Each panel corresponds to a different state variable. Within each panel, the first row corresponds to our baseline calculation in Appendix Figure A8. The second row in each panel reports the range of estimates when we estimate completeness separately by calendar quarter. This specification gives a sense of the variability in the data. The third row in each panel reports estimates when we replace the share R_t of broadcasts that report on the state with the average number of minutes devoted to the state. The fourth row in each panel reports estimates when we replace the estimated prior $\hat{F}_t$ used to calculate the CRPS

with a nonparametric marginal prior estimated with the empirical CDF of previous realizations. The fifth row in each panel reports estimates when we replace the estimated prior $\hat{F}_t$ with one with greater dispersion. The sixth row in each panel reports estimates when we replace the CRPS with the negative log density, which is an alternative scoring rule.

Across many specifications we find that a smooth, monotone function of the CRPS captures a meaningful portion of the predictive power of a model that can depend more flexibly on the realization and the prior. A separate question is whether the realization and the prior are themselves good predictors of reporting behavior. To inform this question, Appendix Table A3 reports the completeness of a smooth function of the realization and prior with respect to a flexible function of the realization and the inputs used to estimate the prior parameters $\hat{\theta}_t$. We find that completeness is meaningful, but also well below one in most cases.

D.2 Restrictiveness

Our estimates of completeness show that a smooth, monotone function of the CRPS captures a meaningful portion of the predictive power of a more flexible model. To measure differences in the flexibility of different models, we turn to estimating restrictiveness.


Restrictiveness (Fudenberg et al. 2023) is a value in $[0, 1]$ that measures the (in)ability of a theory to fit arbitrary data. To construct arbitrary data, we create 100 datasets in which we replace R_t with the output of a random smooth function. We define the random smooth function as a kernel function with Epanechnikov weights. We seed the kernel function by choosing reporting probabilities uniformly randomly at points on a grid. We choose the grid 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. We choose the kernel bandwidth on each dimension as either the bandwidth chosen for the estimation of completeness, or the smallest bandwidth that permits a kernel value at each support point, whichever is larger.

We estimate each reporting rule on the arbitrary data, and compute its average mean squared error across the 100 datasets. We estimate restrictiveness as the ratio of average mean squared error of the given reporting rule to that of the constant prediction, truncated at one. Appendix Figure A8 reports these estimates.


In all cases, a smooth, monotone function of the CRPS has restrictiveness of nearly one, much greater than that of a smooth function of the realization and prior. We therefore find that a smooth, monotone function of the CRPS is much less flexible than a smooth function of the realization and prior.

Figure A8: Estimates of Completeness and Restrictiveness


(a) Stock returns




(b) Unemployment rate




(c) Precipitation



(d) US military casualties in Iraq



(e) US military casualties in Afghanistan



Notes: Each panel corresponds to a state. The left column of plots depicts estimates of completeness as defined in Fudenberg et al. (2022), following the details in Appendix D.1. The right column of plots depicts estimates of restrictiveness as defined in Fudenberg et al. (2023), following the details in Appendix D.2.

Table A2: Completeness Under Alternative Specifications

	Newsworthiness	Realization Only	Realization + Prior
Panel (a): Stock market			
Baseline	1	1	1
[Min Quarter, Max Quarter]	[1, 1]	[1, 1]	
Reporting in minutes	0.887	0.857	1
Nonparametric marginal prior	1	1	—
Dispersed prior	1	1	1
Log score	0.094	1	1
Panel (b): Unemployment			
Baseline	0.407	0.993	1
[Min Quarter, Max Quarter]	[0.250, 0.543]	[0.908, 0.990]	
Reporting in minutes	0.341	0.931	1
Nonparametric marginal prior	0.581	0.993	—
Dispersed prior	0.398	0.993	1
Log score	0.051	0.993	1
Panel (c): Precipitation			
Baseline	0.194	0.229	1
[Min Quarter, Max Quarter]	[0.071, 0.371]	[0.145, 0.703]	
Reporting in minutes	0.285	0.318	1
Nonparametric marginal prior	0.207	0.229	—
Dispersed prior	0.220	0.229	1
Log score	0.092	0.229	1
Panel (d): US military casualties in Iraq			
Baseline	0.657	0.674	1
[Min Quarter, Max Quarter]	[0.506, 0.669]	[0.575, 0.740]	
Reporting in minutes	0.641	0.699	1
Nonparametric marginal prior	0.384	0.674	—
Dispersed prior	0.725	0.674	1
Log score	0.055	0.674	1
Panel (e): US military casualties in Afghanistan			
Baseline	0.872	1	1
[Min Quarter, Max Quarter]	[0.348, 1]	[0.869, 1]	
Reporting in minutes	0.229	0.742	1
Nonparametric marginal prior	0.895	1	—
Dispersed prior	0.871	1	1
Log score	0.014	1	1

Notes: The table displays estimates of completeness as defined in Fudenberg et al. (2022), following the details in Appendix D.1. Each panel represents a different setting. Each column corresponds to a different reporting rule: a smooth monotone function of the CRPS $S_0(\omega_t, \hat{F}_t)$ (CRPS), a smooth function of the realization ω_t (realization only), and a smooth function of both the realization ω_t and the estimated prior $\hat{F}_t$ (realization+prior). Each row corresponds to a different specification. The row “Baseline” reports our baseline estimate. 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 “Nonparametric marginal prior” replaces the baseline CRPS with one calculated based on a prior estimated as the empirical CDF that uniformly weights all preceding realizations within the timeframe used to train the baseline prior, excluding dates for which all previous realizations are identical to the current realization. The row “Dispersed prior” replaces the baseline prior $\hat{F}_t$ with an alternative that increases the scale parameter of the corresponding distribution family by 25 percent, fixing the location parameter. The row “Log score” replaces the CRPS with the negative log of the density (or mass) at the realization.

Table A3: Completeness of Smooth Function of Realization and Prior

	Realization + Prior
Stock market	0.296
Unemployment	0.164
Precipitation	0.722
US military casualties in Iraq	1
US military casualties in Afganistan	1


Notes: The table displays estimates of completeness (as defined in Fudenberg et al. 2022) for a smooth function of the realization ω_t and the estimated prior $\hat{F}_t$ (as defined in Appendix D.1), with respect to a more flexible reporting rule that can depend on the realization and on any of the inputs used to estimate the parameters θ_t . Each row corresponds to a different setting. The more flexible reporting rule is a bagged regression tree, estimated via the random forest algorithm with hyperparameters chosen via 10-fold block cross-validation. All random forest covariate sets include the realization. For the stock market, we additionally include 5 lags of the realization, the risk-free rate, and the VIX. For unemployment, we additionally include 6 lags of the monthly first difference of the realization, and 5 annual lags of the monthly change in the realization. For casualties, we additionally include 30 lags of the realization. For precipitation, we additionally include 5 lags of the realization, and the seasonal proxy $[\cos(\frac{2\pi t}{365}), \sin(\frac{2\pi t}{365})]$.

E Additional Empirical Results


Appendix Figure A9 shows alternative binned scatterplots for cases where the x-axis variable has low-dimensional support. Appendix Figure A10 shows the full set of forms of bias that we considered.

Figure A9: Alternative Plots with Saturated Bin Structure


(a) Frequency of reporting and $\log(\text{CRPS})$ vs. unemployment rate change



(b) Frequency of reporting and $\log(\text{CRPS})$ vs. number of casualties in Iraq




(c) Frequency of reporting and $\log(\text{CRPS})$ vs. number of casualties in Afghanistan



Notes: The unit of analysis is the outlet-date. Each plot is a binscatter with 95 percent pointwise confidence intervals based on inference clustered by month following Cattaneo et al. (2024). In panel (a), we repeat the right-side plots from Figure 2. In panels (b) and (c), we repeat the plots from Figure 4. In all cases the spacing of bins is chosen to be saturated following Cattaneo et al. (2024).

Figure A10: 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 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.