

Consistent Good News and Inconsistent Bad News

Rick Harbaugh, John Maxwell, and Kelly Shue*

Draft Version, June 2016

Abstract

Good news is more persuasive when it is more consistent, and bad news is less damaging when it is less consistent. We show when Bayesian updating supports this intuition so that a biased sender has “mean-variance news preferences” where more or less variance in the news helps the sender depending on whether the mean of the news exceeds expectations. We apply the result to selective news distortion of multiple projects by a manager interested in enhancing the perception of his skill. If news from the different projects is generally good, boosting relatively bad projects increases consistency across projects and provides a stronger signal that the manager is skilled. But if the news is generally bad, instead boosting relatively good projects reduces consistency and provides some hope that the manager is unlucky rather than incompetent. We test for evidence of such distortion by examining the consistency of reported segment earnings across different units in firms. As predicted by the model, managers appear to shift discretionary cost allocations to report more consistent earnings when overall earnings are above rather than below expectations. The mean-variance news preferences that we identify also apply to media bias, p-value hacking, and other situations beyond our career concerns application, and differ from standard mean-variance preferences in that more variable news sometimes helps and better news sometimes hurts.

*Harbaugh, Maxwell: Indiana University. Shue: University of Chicago and NBER. We thank conference and seminar participants at ESSET Gerzensee, Queen’s University Economics of Organization Conference, NBER Corporate Finance Meetings, HBS NOM, Indiana Kelley Business Economics, Maryland Smith Finance, MIT Sloan Accounting, Chicago Booth Finance, Toronto Rotman Business Economics, LSE Paul Woolley Conference, and UNSW Economics for helpful comments. We are also grateful to Mike Baye, Phil Berger, Bruce Carlin, Archishman Chakraborty, Alex Edmans, Alex Frankel, Simon Gervais, Eitan Goldman, Emir Kamenica, Anya Kleymenova, and Pietro Veronesi for helpful comments. We thank Tarik Umar for excellent research assistance.

1 Introduction

If a biased source can distort some of the news, what distortions are most persuasive? Suppose a skeptic wants to persuade the public that global warming is not a problem. Is it more persuasive to exaggerate studies against global warming or to downplay studies for global warming? Or suppose a manager wants to appear skilled at managing projects. If resources can be shifted across projects to affect their reported performance, is it more impressive to make the worst performing projects look less bad, or to make the best performing projects look even better? This same basic question appears in many contexts – is it more persuasive to focus on boosting the news that is more favorable to one’s cause, or instead to focus on shoring up the news that is less favorable?

To help understand this question, we consider news distortion in a sender-receiver game in which the accuracy of the news generating process is uncertain, so that the receiver uses the news to update over both the underlying state and the accuracy of the news itself. We show conditions under which such updating induces the sender to have “mean-variance news preferences” where more variance (less consistency) across multiple pieces of news hurts the sender when the mean of the news is better than the prior, and helps the sender when the mean is worse than the prior. The relative incentive to distort different pieces of news then depends on whether the overall news is generally favorable or unfavorable relative to expectations.

Applied to project performance by a manager, when news from the different projects is mostly favorable, the manager wants the news to appear more reliable so that the posterior estimate of the manager’s skill puts more weight on the news relative to the prior. Increasing the performance of any project helps raise average performance, but shoring up worse performing projects has the added benefit that it makes the news more consistent across projects. This makes the generally good news on the projects a stronger signal of the manager’s competence, so all of the good news becomes more persuasive. However, when the news is mostly unfavorable, the manager looks best by making the least bad projects look better. This makes the news less consistent and hence makes all of the bad news a weaker signal of the manager’s incompetence.

Depending on the situation, the receiver might be “naive” and not anticipate distortion, or might be “sophisticated” and rationally anticipate distortion. We allow for both of these possibilities and focus on the case where the sender can costlessly distort different news within some range as long as the mean of the news remains fixed. For instance, a manager can make some projects look better at the expense of others, or a researcher can inflate some results at the expense of others. Under these constraints, the sender’s optimal distortion strategy when the receiver is naive is also an equilibrium strategy when the receiver is sophisticated.¹ Even though a sophisticated receiver is not fooled in

¹Note that if the receiver believes that the sender might be strategic or instead might be an “honest” type who reports the true news, then highly consistent good news (or highly inconsistent bad news) can be suspicious, which mitigates the incentive to distort the news. Stone (2015) considers a related problem in a cheap talk model of binary

equilibrium, some information is still lost because the sender partially pools reports to minimize variance when the news is generally good. This contrasts with the classic “signal-jamming” result that earnings management can distort firm behavior but does not lead to information loss (e.g., Stein, 1989; Holmstrom, 1982; and Fudenberg and Tirole, 1986).

The model predicts that selective news distortion leads to lower variance when the news is generally good than when the news is generally bad. We test this prediction using the variance of corporate earnings reports for different units or segments within conglomerate firms. Since many overhead and other costs are shared by different units, managers can shift reported earnings across units by adjusting the allocation of these costs. We find evidence that managers shift costs to inflate the reported earnings of worse performing units when the firm is doing well overall. This makes it appear that all the units are doing similarly well, which is a more persuasive signal of management’s abilities than if some units do very well while others struggle. But when the firm is doing poorly, managers shift costs to inflate the reported earnings of the relatively better performing units. This makes it appear that at least some units are doing not too badly, so there is more uncertainty about management’s abilities and the overall evidence of bad performance is weaker.

Our empirical tests account for the possibility that segment earnings may be relatively more consistent during good times due to other natural factors. For example, bad times may cause higher volatility across segments. To isolate variation that is likely to be caused by strategic distortions of cost allocations, we compare the consistency of segment earnings to that implied by segment sales. Like earnings, the consistency of segment sales may vary with firm performance for natural reasons. However, sales are more difficult to distort because they are reported prior to the deduction of costs. Consistent with the model predictions, we find that segment earnings displays abnormal patterns in consistency relative to that implied by segment sales. As a direct test of the mechanism, we also compare the consistency of segment earnings in real multi-segment firms to that of counterfactual firms constructed from matched single-segment firms. In these counterfactual firms, there is neither the incentive nor ability to distort earnings across segments, and we find that the consistency of matched segment earnings does not vary with whether the firm is releasing good or bad news.

Our analysis of mean-variance news preferences contributes to the literature on “good news and bad news” (e.g., Milgrom, 1981) by showing how the impact of specific pieces of news depends on whether the overall news is good or bad. First, we show when a more precise good news signal is better than a less precise good news signal in that it moves the posterior estimate of the state more strongly in the direction of the signal. Second, when there are multiple signals, we show when greater consistency (i.e., lower variance) of the signals implies that the mean of the signals is a more precise signal of the state. Put together, these results imply that a sender wants more consistency of signals when they are good on average, and less consistency when they are bad on average. Since

signals where reporting too many favorable signals is suspicious.

the mean of the news is equally affected by the distortion of any one piece of news, but the variance of the news is affected most by the smallest and largest piece of news, this generates an incentive to selectively distort higher or lower news based on whether the overall news is good or bad.

Mean-variance preferences over the distribution of the news differ substantially from the standard model of mean-variance preferences over the distribution of the state (e.g., Meyer, 1987). First, in standard mean-variance models, lower variance is always preferred due to risk aversion, but in our model more variance is preferred when the news is bad, and the information effect we identify can be stronger than the risk aversion effect. Second, in standard mean-variance models, a higher mean is always preferred, but in our model a lower mean of the news is sometimes preferred due to a version of the “too good to be true” effect whereby very good news is inferred to be very unreliable news (Dawid, 1973; O’Hagan, 1979; Subramanyam, 1996). In our setting, this effect is even stronger than in the previous literature since raising the best news makes not just that news but all the news appear less reliable. Conversely we show that shoring up weaker news can avoid the effect by raising the mean of the news while also making the news more reliable.

Our analysis is related to the problem of “p-value hacking” in which scientists choose data cleaning and measurement methods to maximize statistical significance in classical hypothesis testing. We show how, in a Bayesian environment, artificially reducing variance increases both statistical significance as measured by the posterior probability that the true effect is above the prior, and also “economic significance” as measured by the posterior estimated effect.² Hence the problem of p-value hacking is not limited to classical hypothesis testing nor to just statistical significance. We formalize how the long-recognized strategy of adjusting outliers can persist and lead to loss of information in a strategic environment where distortion is anticipated.³ Our results also highlight that distortion is effective not just at the level of manipulating individual p-values and mean effects. For instance, if three different specifications are presented in a paper, it can be more persuasive if each specification provides a similar result, than if some results are stronger but more disparate. Finally, while our results apply to data manipulation and fraud, our predictions also apply to adjustments such as the reallocation of time and other resources across projects that may be legally and contractually permissible and may even be expected.

The closest approach to ours in the earnings management literature is by Kirschenheiter and Melumad (2002) who consider the incentive to smooth overall firm earnings across time so as to maximize perceived profitability. They find that firms understate sufficiently good earnings and

²Even if the researcher does not present results based on a Bayesian model, our Bayesian approach still applies if decision makers rationally update their own priors based on both the p-value and the mean effect provided by the researcher. Encouraging a joint emphasis on both the p-value and the mean effect (McCloskey and Ziliak, 1996) is hence consistent with a Bayesian approach.

³Within Babbage’s (1830) canonical typology of scientific fraud, such adjustments are “trimming” which is defined as “in clipping off little bits here and there from those observations which differ most in excess from the mean and in sticking them on to those which are too small” so as to reduce the variance while maintaining the mean.

exaggerate sufficiently bad losses. Earnings distortion across time is complicated by the firm’s need to anticipate uncertain future earnings when deciding whether to overreport or underreport current earnings, by the firm’s concern for market estimates of its profitability in each period, and by lack of a fixed end date. By considering the simpler issue of distortion across earning segments rather than time, we can focus on the underlying mechanism that is implicit in their approach – good results are more helpful when they are consistent, and bad results are less damaging when they are inconsistent. We then show that this same idea applies in a more general statistical environment with multiple pieces of news, analyze the resulting mean-variance news preferences, and apply the idea to our career concerns application and other environments.

Our focus on the variability of the news is similar to that of the Bayesian persuasion literature which analyzes ex-ante commitment to an information structure (e.g., Kamenica and Gentzkow, 2011), but we analyze ex-post distortion of the sender’s realized news. Since we consider multidimensional news, the variability of the news is important not just ex ante as in the Bayesian persuasion literature, but also ex post once the news has been realized. Within the career concerns literature, several papers follow Holmstrom (1982) in considering learning about ability from multiple signals, but without our focus on uncertainty over the accuracy of the joint data generating process.⁴ An exception is Prendergast and Stole (1996) who analyze multiple decisions by a manager where managerial ability is defined as having more accurate signals for decision-making. In our model, managerial ability affects output with some noise and the manager does not always prefer that the receiver believes the output signals are accurate.

The remainder of the paper proceeds as follows. Section 2.1 provides a simple example that shows how the consistency of performance news affects updating. In Section 2.2 we develop statistical results on consistency and precision. In Section 2.3 we use these results to show how induced preferences over the mean and variance of the news affects distortion incentives, and in Section 2.4 we consider equilibrium distortion in a sender-receiver game with rational expectations. In Section 3 we consider a range of different applications with mean-variance news preferences, and extend the model to asymmetric news weights. In Section 4 we test for distortion using our main application of segment earnings reports. Section 5 concludes the paper.

⁴We abstract away from other important issues that have been analyzed in the literature such as performance on some tasks being more observable than others (Holmstrom and Milgrom, 1991), some projects having higher returns for particular managers, or some projects being complementary with each other.

2 The model

2.1 Example

A manager has n projects where performance news x_i on each is an additive function of the manager's ability q and a measurement error ε_i , so $x_i = q + \varepsilon_i$. The prior distribution of q is given by the symmetric logconcave density f with mean μ and support on the real line. The ε_i are i.i.d. normal with zero mean and a s.d. σ_ε with non-degenerate independent prior distribution H . The manager, who may or may not know the realization of q , knows the realized values of x_i and can shift some resources to selectively boost reported performance \tilde{x}_i on one or more projects at the expense of lower reported performance on other projects. A receiver does not know q or the true $x = (x_1, \dots, x_n)$ but knows the prior distributions and sees the performance reports $\tilde{x} = (\tilde{x}_1, \dots, \tilde{x}_n)$ after the manager may have distorted them.

Suppose that there is a competitive market for the manager's talent so the manager's payoff equals the receiver's expectation of the manager's ability q given the priors and the news, $U = E[q|x]$. And suppose for this example that the receiver naively believes that the reported news is the true news, $x = \tilde{x}$, so we can focus purely on the statistical implications of different x . The receiver's posterior estimate $E[q|x]$ is a mixture of the prior and the performance news x with the weight dependent on how accurate the news is believed to be. Since the ε_i are i.i.d. normal, the news x can be summarized by the news mean $\bar{x} = \frac{1}{n} \sum_{i=1}^n x_i$, and news variance $s^2 = \sum_{i=1}^n (x_i - \bar{x})^2 / (n-1)$. Letting ϕ be the density of the standard normal distribution, the likelihood of the data x is

$$\prod_{i=1}^n \phi(x_i|q, \sigma_\varepsilon^2) = \frac{1}{(\sigma_\varepsilon \sqrt{2\pi})^n} e^{-\frac{ns^2 + n(\bar{x}-q)^2}{2\sigma_\varepsilon^2}}. \quad (1)$$

Using the assumed independence of σ_ε and q , the impact of the news on q before it is integrated with the prior for q is captured by

$$g(\bar{x} - q|s) = \int_0^\infty \frac{1}{(\sigma_\varepsilon \sqrt{2\pi})^n} e^{-\frac{ns^2 + n(\bar{x}-q)^2}{2\sigma_\varepsilon^2}} dH(\sigma_\varepsilon) \quad (2)$$

or, given the symmetry of g , by $g(q - \bar{x}|s)$. Therefore the posterior density is $f(q|x) = f(q)g(q - \bar{x}|s) / \int_{-\infty}^\infty f(q)g(q - \bar{x}|s)dq$ and the posterior estimate is

$$E[q|\bar{x}, s] = \frac{\int_{-\infty}^\infty qf(q)g(q - \bar{x}|s)dq}{\int_{-\infty}^\infty f(q)g(q - \bar{x}|s)dq}. \quad (3)$$

When the news is more consistent as measured by a lower standard deviation s , the receiver infers that the x_i are less noisy in the sense that there is more weight on lower values of σ_ε in (2). This makes g more concentrated around the news mean \bar{x} so the news mean is a more precise signal of q , and the posterior estimate of q in (3) puts more weight on the news relative to the prior. If the




Figure 1: Effects of selective news distortion on consistency and posterior estimate

news is more favorable than the prior in the sense that $\bar{x} > \mu$, then this greater weight on the news helps the manager.

To see the effect on distortion incentives, suppose there are four projects, the prior $f(q)$ for manager ability is normal with mean 0 and s.d. 2, and the prior H for the variance of project performance has density $h = 1/\sigma_\varepsilon^2$.⁵ Suppose that performance on the projects is generally good, $x = (0, 1, 2, 3)$, and the manager can shift resources to strengthen one project by one unit at the expense of another. For instance the manager could boost the best project at the expense of the worst and report $(-1, 1, 2, 4)$, or could boost the worst project at the expense of the best and report $(1, 1, 2, 2)$. Both keep the mean at $\bar{x} = 3/2$ but the former raises the original $s = \sqrt{5/3}$ to $s = \sqrt{13/3}$ while the latter lowers it to $s = \sqrt{1/3}$. Hence boosting the best project makes the news appear less precise and hence less reliable, while boosting the worst project makes the news appear more precise and hence more reliable. These effects on the apparent precision of the mean \bar{x} as an estimate of q are seen in Figure 1(a).

It would seem that more precise good news should lead to stronger updating of q , and this is seen in the right side of Figure 1(b). Since \bar{x} and s are sufficient statistics for x , the manager's utility can be written as a function of these statistics, $U(\bar{x}, s) = E[q|\bar{x}, s]$, so the manager has what we call "mean-variance news preferences." Helping the worst project lowers s and thereby makes the receiver put more weight on the news and less on the prior, so the posterior mean rises. These effects are reversed if overall performance is bad. Looking at the left side of the figure,

⁵This Jeffreys prior for H corresponds to the inverse gamma distribution with parameters $\alpha = 1, \beta = 0$ and implies $g((\bar{x} - q)/(s/\sqrt{n}))$ is the density of a standard t -distribution with $n - 1$ degrees of freedom.

suppose $x = (-3, -2, -1, 0)$ so the projects are doing poorly with $\bar{x} = -3/2$. In this case shifting resources to the best project from the worst project and reporting $(-4, -2, -1, 1)$ raises s and thereby increases the chance the overall bad outcome was due to the noisiness of the environment. The receiver then relies less on the news and more on the prior distribution of q , so the bad news hurts the posterior mean less.

These differential incentives to distort the news imply that the variance of selectively distorted news will be lower (i.e., the news will be more consistent) when it is favorable rather than unfavorable. With enough instances of such situations, distortion can then be detected probabilistically from this predicted difference. To check the generality of these results, in the following we allow for any number of data points, for different priors, for different sender preferences beyond just maximizing the posterior mean, for different pieces of news having different precision, and analyze a sender-receiver game where the receiver rationally anticipates distortion by the sender. We find that the same incentives to distort the consistency of the news remain and the same implications for distortion detection hold.

2.2 Consistency, precision, and persuasion

We are interested in when more consistent news is more persuasive. To do this, we first show when greater consistency of the news as represented by a lower standard deviation s implies the mean \bar{x} of the news is a more precise signal of q , then show when a more precise signal of q is more persuasive in that it implies stronger updating in the direction of the signal, and then connect these results. We say news is more consistent if the variance of the news is smaller, and we say a signal is more precise if its density is less variable in the uniform variability (UV) order.⁶ Looking back at Figure 1(a), notice that the ratio $g(q - 3/2|s = \sqrt{1/3})/g(q - 3/2|s = \sqrt{13/3})$ is strictly increasing below the mode and strictly decreasing thereafter. So in this case greater consistency as ordered by s leads to greater precision as ordered by uniform variability. Using the definition of g from (2), the following property shows that this relation holds more generally. All proofs are in the Appendix.

Property 1 (Consistency implies precision) *Suppose for a given q that $x_i = q + \varepsilon_i$ for $i = 1, \dots, n$ where i.i.d. $\varepsilon_i \sim N(0, \sigma_\varepsilon^2)$ and σ_ε^2 has independent non-degenerate distribution H . Then $g(\bar{x} - q|s') \succ_{UV} g(\bar{x} - q|s)$ for $s' > s$.*

This result establishes that more consistent news makes the mean of the news a more precise signal of q in the strong sense of making it uniformly less variable. We now show generally when

⁶Following Whitt (1985), we say $g(\bar{x} - q|s') \succ_{UV} g(\bar{x} - q|s)$ if, for $s' > s$, the ratio $g(\bar{x} - q|s)/g(\bar{x} - q|s')$ is strictly quasiconcave with an internal maximum. Note that the uniform variability order implies second order stochastic dominance, but is not implied by it.

ordering of a signal y by uniform variability orders the effect on the posterior estimate for good and bad news.⁷

Property 2 (Precision implies persuasion) *Suppose $g(q - y|\rho)$ is a symmetric quasiconcave density with support on the real line where $g(q - y|\rho') \succ_{UV} g(q - y|\rho)$ for $\rho' > \rho$, and $f(q)$ is independent, symmetric, and logconcave with support on the real line. Then $E[q|y, \rho'] > E[q|y, \rho]$ if $y < \mu$; $E[q|y, \rho'] = E[q|y, \rho]$ if $y = \mu$; and $E[q|y, \rho'] < E[q|y, \rho]$ if $y > \mu$.*

The symmetry and quasiconcavity conditions ensure that the posterior is updated toward the news (Chambers and Healy, 2012).⁸ The additional logconcavity and uniform variability conditions, which are both likelihood ratio conditions, ensure that more precise news results in greater updating towards the news.⁹

Connecting these two results, we can apply Property 1 and let \bar{x} and s take the roles of y and ρ in Property 2.

Proposition 1 *Suppose for a given q that $x_i = q + \varepsilon_i$ for $i = 1, \dots, n$ where i.i.d. $\varepsilon_i \sim N(0, \sigma_\varepsilon^2)$ and σ_ε^2 has independent non-degenerate distribution H , and $f(q)$ is independent, symmetric, and logconcave with support on the real line. Then $\frac{d}{ds}E[q|\bar{x}, s] > 0$ if $\bar{x} < \mu$; $\frac{d}{ds}E[q|\bar{x}, s] = 0$ if $\bar{x} = \mu$; and $\frac{d}{ds}E[q|\bar{x}, s] < 0$ if $\bar{x} > \mu$.*

This proposition shows that more consistent news as measured by a lower s is more persuasive in the sense of moving the posterior estimate $E[q|\bar{x}, s]$ away from the prior and in the direction of the mean of the news.

2.3 Mean-variance news preferences

Proposition 1 establishes that, if $U(\bar{x}, s) = E[q|\bar{x}, s]$ as in the example, then the sender's preferences generally have the shape of Figure 1(b) where the impact of s flips based on the size of the mean \bar{x} relative to the prior μ . To analyze the resulting distortion incentives, it is helpful to think more generally of sender preferences over the news that have these same properties. We will consider

⁷Most of the related literature considers expectations of convex or concave functions of the state, e.g., SOSD results for concave $u(q)$. As we show in Section 3.4, the effects of news precision on the posterior estimate of $u(q)$ can be ordered for all news only if u is linear. The closest result we know of for linear u is by Hautsch, Hess, and Müller (2012) who consider a normal prior and normal news of either high or low precision, with a noisy binary signal of this precision.

⁸As Chambers and Healy show, surprisingly strong conditions are necessary to ensure that seemingly good news is really good news. For instance, Milgrom's standard MLR results on when news y' is more favorable than y do not rule out $y' > y > E[q]$ but $E[q] > E[q|y'] > E[q|y]$, i.e., two pieces of seemingly good news can be ranked by which is better news, yet both can actually be bad. See Finucan (1973) and O'Hagan (1979) for related results.

⁹Logconcavity of f is equivalent to $f(q - a) \succ_{MLR} f(q)$ for any $a > 0$. Uniform variability is, for $\rho' > \rho$, equivalent to $g(y - q|\rho) \succ_{MLR} g(y - q|\rho')$ for $y < q$ and $g(y - q|\rho') \succ_{MLR} g(y - q|\rho)$ for $y > q$.

general mean-variance news preferences $U : \mathbb{R} \times \mathbb{R}_+ \rightarrow \mathbb{R}$ by a sender such that, denoting partial derivatives by subscripts,

$$\begin{aligned} U_s(\bar{x}, s) &> 0 \text{ for } \bar{x} < \mu \\ U_s(\bar{x}, s) &= 0 \text{ for } \bar{x} = \mu \\ U_s(\bar{x}, s) &< 0 \text{ for } \bar{x} > \mu \end{aligned} \tag{4}$$

for all $(\bar{x}, s) \in \mathbb{R} \times \mathbb{R}_+$.¹⁰ Clearly U satisfies these conditions if U is any strictly increasing function of $E[q|\bar{x}, s]$, and in Section 3 we provide other situations where U satisfies these conditions. These are preferences over the mean and variance (or standard deviation) of the news x due to the effects of Bayesian updating, not preferences over the mean and variance of the state q due to risk aversion as in traditional mean-variance models (e.g., Meyer, 1987). We discuss this distinction further in Section 3.4. Note that we do not restrict the sign of $U_{\bar{x}}(\bar{x}, s)$ and, in Section 3.2, we consider the issue of “too good to be true” news preferences where $U_{\bar{x}}(\bar{x}, s)$ is not monotonic.

To see the implications of (4) for selective news distortion, note that for any j ,

$$\frac{d\bar{x}}{dx_j} = \frac{1}{n}, \quad \frac{ds}{dx_j} = \frac{x_j - \bar{x}}{(n-1)s} \tag{5}$$

so every piece of news has the same effect on \bar{x} , but the effect on the variance is increasing in the size of x_j relative to \bar{x} . Since a lower s helps when $\bar{x} > \mu$ and hurts when $\bar{x} < \mu$, the marginal gain is higher from increasing lower news in the former case, and from increasing higher news in the latter case. In particular, exaggerating the best news increases \bar{x} and also increases s , so the effects on the posterior estimate counteract each other if $\bar{x} > \mu$ but reinforce each other if $\bar{x} < \mu$. And improving the worst news increases \bar{x} but also decreases s , so the effects on the posterior estimate reinforce each other if $\bar{x} > \mu$ but counteract each other if $\bar{x} < \mu$. The next result follows.

Proposition 2 *For U satisfying (4) and $x_i < x_j$, $\frac{d}{dx_i}U(\bar{x}, s) < \frac{d}{dx_j}U(\bar{x}, s)$ if $\bar{x} < \mu$; $\frac{d}{dx_i}U(\bar{x}, s) = \frac{d}{dx_j}U(\bar{x}, s)$ if $\bar{x} = \mu$; and $\frac{d}{dx_i}U(\bar{x}, s) > \frac{d}{dx_j}U(\bar{x}, s)$ if $\bar{x} > \mu$.*

With these marginal incentives to distort better and worse news, we can now analyze the sender’s choice of how to distort the news.

2.4 Optimal and equilibrium distortion

We analyze the sender’s optimal distortion strategy when the receiver is “naive” and does not anticipate distortion, and also the sender’s equilibrium distortion strategy when the receiver is “sophisticated” and rationally anticipates distortion. To focus on how the sender can affect the receiver’s confidence in the news by distorting its consistency, we assume that the sender’s distortions cannot change the overall mean of the news. To capture the difficulty of distorting the news,

¹⁰We focus on preferences over summary statistics of multiple signals, but the analysis also applies to preferences over one signal with known variability, $U(y, \rho)$, when the variability parameter ρ can be directly influenced.

we assume the total amount of distortion is limited. Under these two constraints, we show that the sender's optimal distortion strategy when the receiver is naive is also an equilibrium strategy when the receiver is sophisticated.

Let $\tilde{x}(x)$ be the sender's pure strategy of reporting \tilde{x} based on the sender's true news type x . The receiver estimates the posterior distribution of q given her priors f and H , the observed news \tilde{x} , and her beliefs that map \tilde{x} to the set of probability distributions over \mathbb{R}^n . In the naive receiver case, the receiver does not anticipate distortion, so receiver beliefs put all weight on $x = \tilde{x}$. In the sophisticated receiver case, the receiver's beliefs are consistent with the sender's strategy along the equilibrium path. Therefore if $\tilde{x}(x)$ is one-to-one the receiver puts all weight on $x = \tilde{x}^{-1}(\tilde{x}(x))$. If not, the receiver weights the distribution of x according to $\tilde{x}(x)$ and Bayes' rule given f and H . If the sender makes a report that is off the equilibrium path, the beliefs put all weight on whichever type is willing to deviate for the largest set of rationalizable payoffs, i.e., we impose the standard D1 refinement (Cho and Kreps, 1987).

We assume that sender distortions are subject to a constant mean constraint and a maximum distortion constraint,

$$\sum_i \tilde{x}_i - x_i = 0 \text{ and } \sum_i |\tilde{x}_i - x_i| \leq d, \quad (6)$$

where $d > 0$ is the maximum total distortion across the news. Given the constant mean constraint, receiver beliefs about the distribution of the true x can be summarized by receiver beliefs about the distribution of s which we denote by $p(s|\tilde{x})$. Therefore the sender maximizes her expected utility

$$\int_0^\infty U(\bar{x}, s) dp(s|\tilde{x}). \quad (7)$$

First consider the naive receiver case. When the news is generally unfavorable, $\bar{x} < \mu$, the sender wants to increase s as much as possible. Figure 2(a) shows the same case as Figure 1(b) with a prior of $N(0, 2)$ and $h = 1/\sigma^2$, except that $n = 2$ so the contour sets for the posterior mean can be seen directly as a function of x . Looking at the bottom left quadrant where the red line shows combinations of x_1 and x_2 that maintain the same mean $\bar{x} = -2$, the sender increases the posterior mean by moving the news away from the center where $x_1 = x_2$ and toward either edge. This increases s by maximizing the difference in the news. So if $x_1 > x_2$ the sender reports $\tilde{x} = (x_1 + d/2, x_2 - d/2)$, and if $x_1 < x_2$ the sender reports $\tilde{x} = (x_1 - d/2, x_2 + d/2)$. If $n > 2$, this same logic applies. From (5), the largest increase in s occurs when the smallest news is decreased and the largest news is increased, so the sender simply decreases the smallest news by $d/2$ and increases the largest news by $d/2$, which satisfies (6).

When the news is generally favorable, $\bar{x} > \mu$, the sender wants to decrease s as much as possible. From the upper right quadrant of Figure 2(a), for any x_1 and x_2 with the same given mean $\bar{x} = 2$, the sender wants to move inward along the red line toward the center where $x_1 = x_2$.




Figure 2: Selective news distortion for bad news and good news

Therefore if $x_1 - x_2 \geq d$ the sender reports $\tilde{x} = (x_1 - d/2, x_2 + d/2)$, if $x_2 - x_1 \geq d$ the sender reports $\tilde{x} = (x_1 + d/2, x_2 - d/2)$, and otherwise the sender reports $\tilde{x} = (\bar{x}, \bar{x})$ without having to exhaust the total distortion budget. If $n > 2$, the sender starts by squeezing in the most extreme news. As extreme news moves inward, it might bump into other news, which then is equally extreme so that this news is also moved in jointly. This continues from each side until the side's budget of $d/2$ distortion, which maintains the prior mean, is exhausted. If all the data starts out sufficiently close, the data is completely squeezed to the mean \bar{x} before the budget is exhausted. This strategy is specified in part (i) of Proposition 3 below, where notationally to capture the pooling of potentially multiple pieces of news as the news is squeezed in from either extreme, we let \bar{x}_a solve $\sum_{i=1}^a (\bar{x}_a - x_i) = d/2$ subject to $x_a \leq \bar{x}_a$, and let \bar{x}_b solve $\sum_{i=b}^n (x_i - \bar{x}_b) = d/2$ subject to $x_b \geq \bar{x}_b$.

Now consider the sophisticated receiver case and suppose that the sender follows the same strategy as in the naive receiver case. For bad news, not all reports are on the equilibrium path. As seen in Figure 2(a), if $d = 1$ then for any x such that $\bar{x} = -2$, a report along the dashed line between $(-5/2, -3/2)$ and $(-3/2, -5/2)$ should never be observed. As we show in the proof of Proposition 3, in such cases it is always the “worst type” $x_1 = x_2$ with the lowest s that is willing to deviate to any such report for the largest range of rationalizable payoffs. Therefore, by the D1 refinement, the receiver should assume that such a deviation was done by this type. Given such beliefs, even the worst type gains nothing from deviation. When $\bar{x} > \mu$, if the reports for the projects differ, a sophisticated receiver can again invert the equilibrium strategy and back out the true x , but otherwise there is some pooling. Looking at Figure 2, if $d = 1$, then for all x between

$(3/2, 5/2)$ and $(5/2, 3/2)$, the sender will report $(2, 2)$, so the receiver cannot invert the reports. In this case, the receiver will form a belief over the true x that induces a distribution over s , where s is always smaller than when the receiver is thought to be outside of the region between $(3/2, 5/2)$ and $(5/2, 3/2)$. Since the sender prefers a lower s and any other report will lead the receiver to infer the news is outside this region with a higher s , the sender has no incentive to deviate.

Following this logic, the optimal strategy when the receiver is naive is also an equilibrium strategy when the receiver is sophisticated, leading to part (ii) of Proposition 3, the proof of which is extended to $n > 2$ in the Appendix.

Proposition 3 *(i) Assume the receiver is naive. If $\bar{x} < \mu$ then the sender's optimal strategy is $\tilde{x}_1 = x_1 - d/2, \tilde{x}_n = x_n + d/2$, and $\tilde{x}_i = x_i$ for $i \neq 1, n$. If $\bar{x} > \mu$ then (a) if $\sum_i |x_i - \bar{x}| \leq d$ then $\tilde{x}_i = \bar{x}$ for all i ; (b) if not, then $\tilde{x}_i = \bar{x}_a$ for $i \leq a$, $\tilde{x}_i = \bar{x}_b$ for $i \geq b$, and $\tilde{x}_i = x_i$ for $a < i < b$. (ii) Assume the receiver is sophisticated. Then the sender's strategy in (i) is a perfect Bayesian equilibrium.*

Since the equilibrium is fully separating for $\bar{x} \leq \mu$, the receiver correctly “backs out” the true values by discounting the reported values according to the equilibrium strategy. However the equilibrium is partially pooling for $\bar{x} \geq \mu$, so some information is lost even though receiver correctly anticipates distortion.

The distortion strategy given by Proposition 3 leads to higher variance for \tilde{x} than for x when $\bar{x} < \mu$, and lower variance for \tilde{x} than for x when $\bar{x} > \mu$. By our symmetry assumptions on the prior density of q and on the news given q , the expected standard deviation of the true x is the same for any \bar{x} equidistant from the prior on either side. Therefore the reported standard deviation for \tilde{x} should on average be higher below the prior than above the prior.¹¹

Proposition 4 *The distortion strategy in Proposition 3 implies that, in expectation, $s(\tilde{x})$ is higher when $\bar{x} < \mu$ than when $\bar{x} > \mu$.*

This result is the main testable implication of the model, which we examine using data on firm segment earnings in Section 4.

3 Applications and extensions

We now consider different applications and extensions of mean-variance news preferences. Sections 3.1 to 3.3 show environments where preferences have the general fan-shape of Figure 1(b) where

¹¹In addition to the constraints we model, there might be other constraints such as only some news can be distorted, and/or distortion might be costly, with some distortions more costly than others. For any constraints or costs, the same prediction applies for the naive receiver case since the sender never benefits from a higher s when news is good or a lower s when news is bad. For a sophisticated receiver, the same intuition would appear to hold, but we do not analyze this general case.

$U_s > 0$ for \bar{x} below the prior and $U_s < 0$ for \bar{x} above the prior.¹² Section 3.4 combines our model based on Bayesian updating with a traditional mean-variance model based on risk aversion. An extension to weighted means and weighted standard deviations is given in Section 3.5. This extension is used in our test of earnings management in Section 4. For each case, we focus on the underlying distortion incentives when the receiver is naive, though the analysis can be extended in the same manner as above to equilibrium distortion with a sophisticated receiver.

3.1 Posterior probability

Rather than maximizing their estimated skill, a manager might want to maximize the estimated probability that they are competent so as to attain a promotion or avoid a demotion (Chevalier and Ellison, 1999). This can be modeled as maximizing the posterior probability that q is sufficiently high. In the Appendix we establish Property 3 which is an equivalent to Property 2 for the posterior probability $F(q|y, \rho)$ rather than the posterior estimate $E[q|y, \rho]$. Letting $y = \bar{x}$ and $\rho = s$, and focusing on the posterior probability that q exceeds the prior μ , gives the following result.

Result 1 *The posterior probability satisfies $\frac{d}{ds} \Pr[q > \mu|\bar{x}, s] > 0$ if $\bar{x} < \mu$; $\frac{d}{ds} \Pr[q > \mu|\bar{x}, s] = 0$ if $\bar{x} = \mu$; and $\frac{d}{ds} \Pr[q > \mu|\bar{x}, s] < 0$ if $\bar{x} > \mu$.*

This establishes that $U = \Pr[q > \mu|\bar{x}, s]$ satisfies (4), so the predictions regarding selective news distortion are the same as those for maximizing estimated skill. Figure 3(a) shows the same situation as Figure 1 except the manager wants to maximize the probability that his skill q is above the prior which is normalized to zero. Focusing on the posterior probability rather than posterior estimate can be seen as emphasizing “statistical significance” rather than “economic significance”, and distortion can be seen as “p-value hacking.” Note that if f is uninformative and $h = 1/\sigma_\varepsilon^2$, then the posterior distribution of q is the t-distribution with $n - 1$ degrees of freedom, so the probability that $q < 0$ is given by $T_{n-1}(\bar{x}/s)$ and the indifference curves in the figure are linear. Hence this result generalizes the t-distribution case where an increase in s helps or hurts depending on the sign of \bar{x} .

3.2 Too good to be true

Can news be so good that it is no longer credible? Dawid (1973) and O’Hagan (1979) show that an increase in a single piece of news y can be “too good to be true” in that it decreases $E[q|y]$. In particular, they show that $\lim_{y \rightarrow \infty} E[q|y] = \mu$ if the prior f has thinner tails than the signal g , e.g., f is normal and g is the t -distribution. In this case, as y increases it becomes very unlikely based

¹²Contrary to our assumptions, in some situations U might depend on details of the performance news rather than on \bar{x} and s , e.g., x_i might affect U differently because a manager’s compensation depends on how particular units perform.

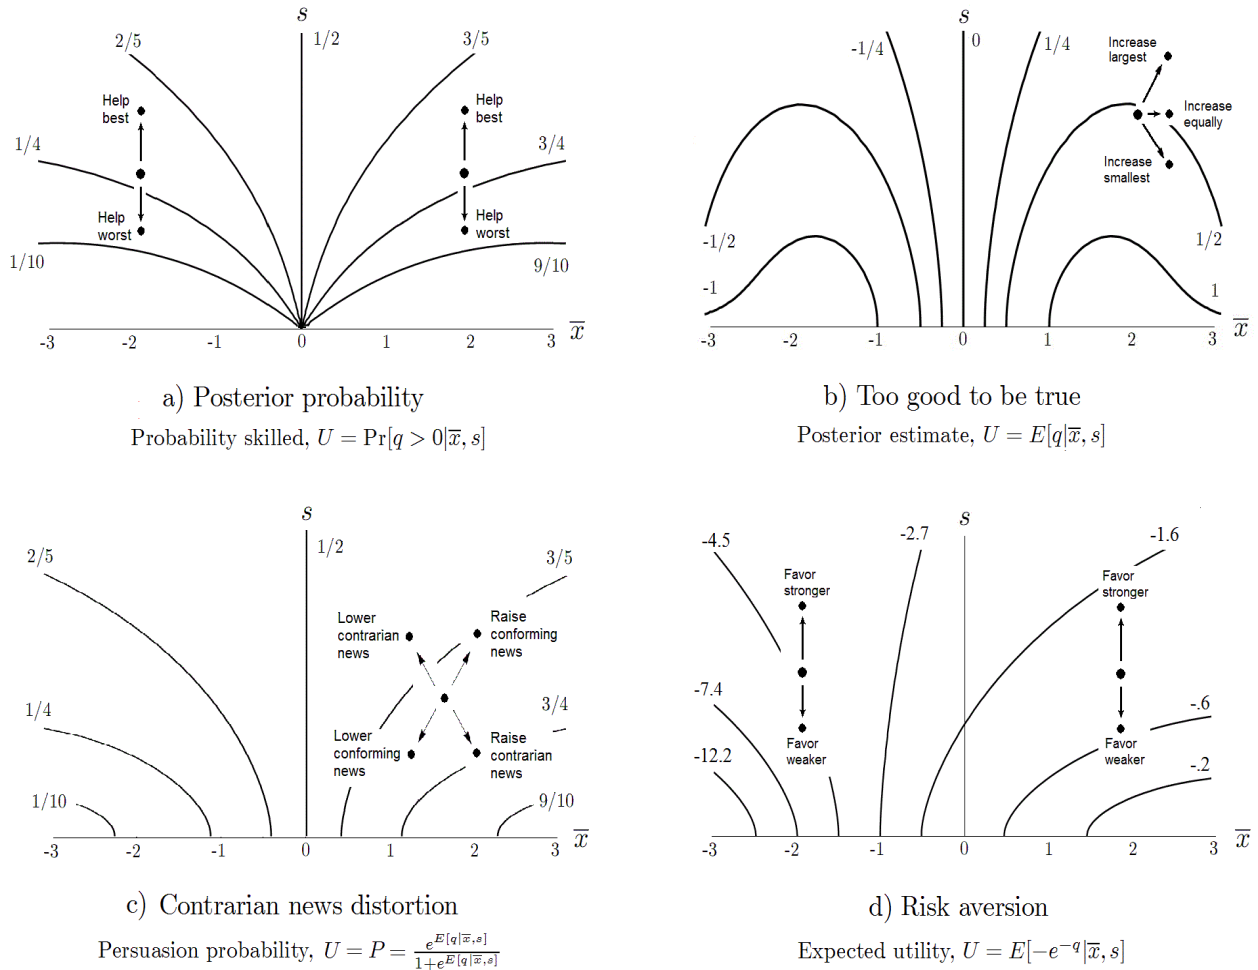


Figure 3: Applications with mean-variance news preferences $U(\bar{x}, s)$

on the prior that the true value is as extreme as the signal indicates, so the signal is believed to be just noise rather than informative of the true state. In particular, Subramanyam (1996) shows that if f is normal and g is normal with uncertain variance, which includes the t -distribution case, that as $y > \mu$ increases, $E[q|y]$ is first increasing and then decreasing.

Applied to our environment with $y = \bar{x}$, these standard results imply that, as \bar{x} increases with a fixed s , the news can eventually become too good to be true.¹³ This effect is aggravated or mitigated when an individual x_i changes, depending on its position relative to the mean. An increase in $x_i > \bar{x}$ not only raises \bar{x} but has the additional effect that the tails of the news distribution become fatter as s rises. However, for $x_i < \bar{x}$, the two effects counteract each other so the too good to be true

¹³If H is such that g is the t -distribution then news must eventually be too good to be true as O'Hagan shows. But if g is instead logconcave, which is also possible for mixtures of normals, then higher news must always be better by an application of Milgrom's good news result.

effect is mitigated and potentially avoided.

Based on these differential effects, it is possible to increase \bar{x} and avoid the too good to be true problem entirely through selective distortion. Suppose the total distortion constraint is $\sum_i |\tilde{x}_i - x_i| \leq d$ for some given $d > 0$. If the sender reports $\tilde{x}_1 = x_1 + (d/2 + \varepsilon)$ and $\tilde{x}_n = x_n - (d/2 - \varepsilon)$ for $0 \leq \varepsilon \leq d/2$, then s falls discontinuously for any such ε while \bar{x} increases continuously as ε increases from zero. Therefore, by the continuity of $E[q|\bar{x}, s]$ in s and \bar{x} , if $\frac{d}{ds}E[q|\bar{x}, s] < 0$ it is always possible to choose an ε that raises $E[q|\bar{x}, s]$ even in the range where $\frac{d}{d\bar{x}}E[q|\bar{x}, s] < 0$, with the only exception being the zero measure case where $s = 0$. These two results, and the equivalents for unfavorable news, are stated in the following proposition.

Result 2 (i) If $\frac{d}{d\bar{x}}E[q|\bar{x}, s] \geq 0$ then $\frac{d}{dx_i}E[q|\bar{x}, s] > 0$ for all $x_i < \bar{x}$, and if $\frac{d}{d\bar{x}}E[q|\bar{x}, s] \leq 0$ then $\frac{d}{dx_i}E[q|\bar{x}, s] < 0$ for all $x_i > \bar{x}$. (ii) For any $d > 0$, there almost surely exists a distortion \tilde{x} such that $\tilde{\bar{x}} > \bar{x}$ and $E[q|\tilde{\bar{x}}, \tilde{s}] > E[q|\bar{x}, s]$, and an alternative distortion \tilde{x}' such that $\tilde{\bar{x}}' < \bar{x}$ and $E[q|\tilde{\bar{x}}', \tilde{s}'] < E[q|\bar{x}, s]$.

This result is shown in Figure 3(b) where the environment is the same as Figure 1(b) except the prior has lower variance so that, as \bar{x} increases and becomes less reliable, the posterior $E[q|\bar{x}, s]$ converges more quickly to the prior in the pictured range. As seen on the right side of the figure, increasing all the x_i keeps s the same and $E[q|\bar{x}, s]$ falls as the data becomes less believable relative to the prior, but if x is selectively distorted with increases in the smaller data points this is avoided. By the same logic, even in the range of “too bad to be true” on the left side of the figure, it is possible to reduce $E[q|\bar{x}, s]$ further by selective reduction of x that focuses on reducing s by reducing the largest data points.

The same analysis extends to posterior probabilities. Dawid (1973) shows that not only does the mean revert to the prior when f has thinner tails than g , but the entire posterior distribution reverts to the prior distribution, so $\lim_{\bar{x} \rightarrow \infty} \Pr[q > a|\bar{x}, s] = 1 - F(a)$ for any a .¹⁴ The same selective distortion strategy used for the posterior mean above can then also be used to avoid the too good to be true problem for the posterior probability.

3.3 Contrarian news distortion: seeding doubt and promoting consensus

The literature on news bias has focused on distortions that push a scalar news variable in the source’s favored direction at some reputational or other cost (e.g., Gentzkow and Shapiro, 2006).

¹⁴The basic model of Student (1908) with an uninformative prior already incorporates a version of the too good to be true idea due to its use of the same data to estimate both the mean and the standard deviation. In particular, letting $t(x)$ be the t-value, direct calculations show that $\lim_{x_i \rightarrow \infty} t(x) = 1$, so if $\bar{x} > 0$ and s is small enough that $t(x) > 1$, raising any x_i eventually undermines the reliability of all the data so much that significance decreases. Note that if $h = 1/\sigma_\varepsilon^2$ and, counter to our assumption, f is uninformative, then $f(q|x)$ follows the t-distribution. So an informative f strengthens the too good to be true effect for raising any x_i by making it hold for all x such that $\bar{x} > 0$.

If there are multiple pieces of news, then the consistency of the news also becomes a factor that the source can manipulate.¹⁵ For instance, opponents of action on climate change are claimed to exaggerate evidence against the scientific consensus as part of a strategy of “seeding doubt” (e.g., Oreskes and Conway, 2010),¹⁶ while proponents are claimed to make the consensus appear stronger by downplaying opposing evidence. These are not the only distortion strategies available – opponents could instead focus on downplaying evidence for the consensus, while proponents could instead focus on exaggerating outliers in the direction of the consensus. Given that the preponderance of scientific studies support climate change, our model implies that seeding doubt and promoting consensus are indeed the best strategies for each side.

Applying our model to such situations, we define news as *contrarian* relative to other news if it is on the prior’s side of the mean of the news and *conforming* otherwise. That is, for $\bar{x} > \mu$ we say news x_i is *contrarian* if $x_i < \bar{x}$ and *conforming* if $x_i > \bar{x}$, and for $\bar{x} < \mu$ we say news x_i is *contrarian* if $x_i > \bar{x}$ and *conforming* if $x_i < \bar{x}$.¹⁷ For $\bar{x} > \mu$ distorting contrarian news $x_i < \bar{x}$ downward increases s and also lowers \bar{x} , while distorting contrarian news upward decreases s and also raises \bar{x} . So both sides – those who want a higher estimate and who want a lower estimate – get a double effect from focusing on distorting contrarian news in their favored direction. In contrast, distorting conforming news always creates a trade-off of either making the mean of the news more favorable but the consistency less favorable, or making the mean of the news less favorable but the consistency more favorable. If $P(\bar{x}, s)$ is the probability that the audience is persuaded to one side, which could be a function of $E[q|\bar{x}, s]$ or, as in Section 3.1, of $\Pr[q > \mu|\bar{x}, s]$, we have the following result by application of Proposition 2.

Result 3 *Suppose the persuasion probability $P(\bar{x}, s)$ satisfies (4). For either side of a debate, $U = P$ or $U = 1 - P$, distorting contrarian news is more effective than distorting conforming news.*

For instance, following a standard random utility model based on uncertainty in voter preferences, suppose the probability that voters are persuaded to take action on global warming is $P = e^{E[q|\bar{x}, s]} / (1 + e^{E[q|\bar{x}, s]})$, as shown in Figure 3(c). Since the news mean is above the prior in the figure, opponents want to make contrarian evidence more damaging and supporters want to make it less damaging, and neither side benefits as much from distorting conforming news. Given the

¹⁵Chakraborty and Harbaugh (2010) consider multidimensional news but without uncertainty over the news generating process. Their focus is on the implicit opportunity cost of pushing one dimension versus another.

¹⁶Internal memos from Exxon indicate an explicit strategy to “emphasize the uncertainty in scientific conclusions” regarding climate change. NYT 11/7/2015.

¹⁷Recall that “news” in our model comes from the same data generating process so that the credibility of all the news rises and falls with its consistency. Data from different processes is modeled as contributing to the prior. This makes the question of whether a given analysis really follows standard methods, and hence has the spillover effects we analyze, of particular importance and hence a likely area of controversy.

definition of contrarian news, the same would hold if the news mean was below the prior. In general, the model implies that debates are likely to focus on the exact meaning of the most contrarian evidence, and such evidence is a good place to look for signs of distortion.

3.4 Risk aversion

In a standard mean-variance model with a location-scale distribution $f(q)$, higher variance in $f(q)$ for a fixed mean $E[q]$ lowers $E[u(q)]$ for concave u (risk aversion) and raises $E[u(q)]$ for convex u (risk seeking). In our approach, we assume risk neutrality and instead show higher variance in the news distribution $g(\bar{x} - y|s)$, which need not increase variance in the posterior distribution $f(q|\bar{x}, s)$,¹⁸ raises $E[q|\bar{x}, s]$ when the news is unfavorable and lowers it when the news is unfavorable.

To see how these two different approaches interact, suppose the sender is a firm, q is the firm's true value, and the receiver is an undiversified investor with utility $u(q)$. The investor's valuation of the asset, and the payoff to the firm U , are increasing in the investor's expected utility $E[u(q)|\bar{x}, s]$, e.g., $U = E[u(q)|\bar{x}, s]$. Since \bar{x} and s are sufficient statistics for x , and since we are taking the prior $f(q)$ as given, the investor and hence the firm must have "mean-variance" utility over the news in the sense that no other information matters, but we are no longer assured that $U(\bar{x}, s)$ satisfies (4). For a risk averse investor, the risk aversion and information effects work together if the news is good so less variance is always preferred, but counteract each other if the news is bad so more or less variance may be preferred. If the investor is risk seeking, e.g., due to option value or other considerations, the opposite pattern holds – more variance is always preferred if the news is bad, but more or less variance may be preferred if the news is good.

In particular, Property 4 in the Appendix shows that for $\rho' > \rho$, $\int_{-\infty}^a F(q|\rho', y)dq > \int_{-\infty}^a F(q|\rho, y)dq$ for all a if $y > \mu$, and $\int_{-\infty}^a F(q|\rho', y)dq < \int_{-\infty}^a F(q|\rho, y)dq$ for all a if $y < \mu$. The former result establishes that $F(q|\rho, y) \succ_{SOSD} F(q|\rho', y)$ if $y > \mu$ which, together with Property 2 and Proposition 1, implies part (i) of the following for concave u . The latter result establishes the equivalent result for the increasing convex order and similarly implies part (ii) for convex u . Together parts (i) and (ii) imply part (iii), as already established directly in Proposition 1.

Result 4 *Suppose U is an increasing function of $E[u(q)|\bar{x}, s]$ where u is increasing. (i) For u concave $U_s < 0$ if $\bar{x} \geq \mu$; (ii) for u convex $U_s > 0$ if $\bar{x} \leq \mu$; and (iii) for u linear $U_s \geq 0$ if $\bar{x} \leq \mu$ and $U_s \leq 0$ if $\bar{x} \geq \mu$.*

Figure 3(d) shows the case of $U = E[u(q)|\bar{x}, s]$ where u is concave with constant absolute risk aversion, $u = -e^{-q}$. In the realm of good news, smaller s both increases $E[q|\bar{x}, s]$ and lowers risk so the gains from reducing s are accentuated. In the realm of bad news, a smaller s decreases

¹⁸Moreover, the posterior $f(q|x, s)$ will not in general be a location-scale distribution, so the standard mean-variance result of Meyer (1987) still would not apply.

$E[q|\bar{x}, s]$ but it does not necessarily decrease $E[u(q)|\bar{x}, s]$. As seen in the figure, over some range the information effect dominates, and over some range the risk aversion effect dominates. If instead we assume u is convex, the positive information effect of a higher s for bad news is reinforced, but the negative information effect for good news is weakened or reversed.

3.5 Asymmetric news weights

If projects vary predictably in size, the noise terms for each project are likely to have different variances rather than be identically distributed as we have assumed so far. In this extension, we show that a weighted mean-variance model is the same statistically as the symmetric model with appropriate substitution of weighted parameters. Moreover, under natural assumptions that fit environments including our segment earnings application, the strategic implications are also the same. Since we will use this weighted model in our test in Section 4, we focus on the case of segment earnings.

Following standard accounting practice, let segment performance be measured by segment Return on Assets (ROA), i.e., $x_i = e_i/a_i$ where e_i is segment earnings and a_i is segment assets which are known, so $e_i/a_i = q + \varepsilon_i$. Suppose that the variance of ROA performance is inversely proportional to segment size, so $\varepsilon_i \sim N(0, \sigma_\varepsilon^2/a_i)$ where σ_ε^2 is distributed according to H as before. A simple justification for this assumption is each segment i is composed of a_i different subsegments with equal assets normalized to one, where subsegment earnings of the k th subsegment are $e_{ik} = q + \varepsilon_{ik}$ for $k = 1, \dots, a_i$ and ε_{ik} is i.i.d. normal with mean 0 and s.d. σ_ε . By normality, $Var[e_i] = Var[\sum_{k=1}^{a_i} e_{ik}] = a_i \sigma_\varepsilon^2$, and hence $Var[x_i] = Var[e_i/a_i] = \frac{1}{a_i^2} a_i \sigma_\varepsilon^2 = \sigma_\varepsilon^2/a_i$ as assumed.

Using segment asset shares $\frac{a_i}{A}$ as weights, the weighted mean and standard deviation of the firm's news performance are

$$\bar{x}_w = \sum_{i=1}^n \frac{a_i}{A} \frac{e_i}{a_i} \text{ and } s_w = \left(\sum_{i=1}^n n \frac{a_i}{A} \left(\frac{e_i}{a_i} - \sum_{i=1}^n \frac{a_i}{A} \frac{e_i}{a_i} \right)^2 / (n-1) \right)^{1/2} \quad (8)$$

It is straightforward to verify that $g(q - \bar{x}_w, s_w)$ is the same as (2) with \bar{x}_w and s_w^2 in place of \bar{x} and s^2 , so the same result from Property 1 for uniform variability holds with these weighted sufficient statistics. Note that increases in ROA e_i/a_i for larger segments have a bigger effect on \bar{x}_w and s_w since they are weighted more heavily. However, distortions of e_i for larger segments have proportionally less effect on segment ROA due to the larger denominator a_i . These two factors cancel each other out so we are left with the same relative incentives to distort news as before,

$$\frac{d}{de_i} \bar{x}_w = \frac{1}{A} \text{ and } \frac{d}{de_i} s_w = \frac{\frac{n}{A} (x_i - \bar{x}_w)}{(n-1) s_w}. \quad (9)$$

In particular, the effect on average ROA \bar{x}_w is the same regardless of which segment earnings are

changed,¹⁹ and the effect on s_w depends on the size of segment ROA relative to average ROA.

Therefore, the above case where $w_i = a_i/A$ and distortion is of e_i where $x_i = e_i/a_i$ is an example of the following more general result.

Result 5 *If $\varepsilon_i \sim N(0, \sigma_\varepsilon^2/w_i)$ where w_i is known, then the above asymmetric model generates the same restrictions on $U(\bar{x}_w, s_w)$ as the symmetric model does for $U(\bar{x}, s)$, and also generates the same relative distortion incentives for the sender if distortion ability is inversely proportional to w_i .*

Since earnings segments often vary substantially in size, we use this weighted model for our empirical analysis of segment earnings distortion, and in particular we consider how changes in cost allocations across segments affect earnings e_i and hence affect segment ROA $x_i = e_i/a_i$.

¹⁹Note that average ROA $\bar{x}_w = \sum_{i=1}^n \frac{a_i}{A} \frac{e_i}{a_i}$ equals the firm's overall ROA $\sum_{i=1}^n \frac{e_i}{A}$, so this just says that total ROA cannot be changed by moving earnings between segments. Also note that (9) reduces to the symmetric case of (5) for $a_i = 1$ and $A = n$.

4 Empirical test using earnings management across segments

We now turn to an empirical test of the theory. In earnings reports, managers of US public firms have discretion in how to attribute total firm earnings to business segments operating in different industries. The reporting of earnings across segments is therefore one aspect of “earnings management,” whereby a managers tries to influence the short-run appearance of the firm’s profitability, or of her own managerial ability, by adjusting reported earnings. The shifting of total firm earnings across time is a well-studied topic in the theoretical and empirical literature (e.g., Stein, 1989; Kirschenheiter and Melamud, 1992), but the shifting of earnings across segments has not received as much attention. In particular, the strategy of influencing the consistency of earnings across segments has not, to our knowledge, been analyzed theoretically or empirically.²⁰

4.1 Overview of empirical setting

Segment earnings (also known as segment profits or EBIT) are a key piece of information used by boards and investors when evaluating firm performance and managerial quality. In a survey of 140 star analysts, Epstein and Palepu (1999) find that a plurality of financial analysts consider segment performance to be the most useful disclosure item for investment decisions, ahead of the three main firm-level financial statements (statement of cash flows, income statement, and balance sheet).

Under regulation SFAS No. 14 (1976—1997) and SFAS No. 131 (1997—present), managers exercise substantial discretion over the reporting of segment earnings.²¹ Firms are allowed to report earnings based upon how management internally evaluated the operating performance of its business units. In particular, segment earnings are approximately equal to sales minus costs, where costs consist of costs of goods sold; selling, general and administrative expenses; and depreciation, depletion, and amortization. As shown in Givoly et al. (1999), the ability to distort segment earnings is primarily due to the manager’s discretion over the allocation of shared costs to different segments.²² This discretion over cost allocations approximately matches our model of strategic dis-

²⁰The literature has considered issues such as withholding segment earnings information for proprietary reasons (Berger and Hann, 2007), the effects of transfer pricing across geographic segments on taxes (Jacob, 1996), and the channeling of earnings to segments with better growth prospects (You, 2014). Our analysis of the distortion of allocations across segments is also related to the literature on the “dark side of internal capital markets,” e.g., Scharfstein and Stein (2000).

²¹Prior to SFAS No. 131, many firms did not report segment-level performance because the segments were considered to be in related lines of business. SFAS No. 131 increased the prevalence of segment reporting by requiring that disaggregated information be provided based on how management internally evaluated the operating performance of its business units.

²²GE’s 2015 10Q statement offers an example of managerial discretion over segment earnings: “Segment profit is determined based on internal performance measures used by the CEO ... the CEO may exclude matters such as charges for restructuring; rationalization and other similar expenses; acquisition costs and other related charges; technology and product development costs; certain gains and losses from acquisitions or dispositions; and litigation settlements or other charges ... Segment profit excludes or includes interest and other financial charges and income taxes according to how a particular segment’s management is measured ... corporate costs, such as shared services, employee benefits and information technology are allocated to our segments based on usage.”

tortion of news under a fixed mean and total distortion constraint. We assume that total segment earnings are approximately fixed in a period and managers have a limited amount of discretionary costs that can be flexibly allocated across segments to alter the consistency of segment earnings. Our theory predicts that managers will distort segment earnings to appear more consistent when overall firm news is good relative to expectations. When firm news is bad, managers will distort segment earnings to appear less consistent.²³

By focusing on segment earnings management within a time period rather than firm-level earnings management over time, we are able to bypass an important dynamic consideration for the management of earnings over time. The manager can only increase total firm-level earnings in the current period by borrowing from the future, which limits the manager’s ability to report high earnings again next period. In contrast, distortion of the consistency of earnings across segments in the current period does not directly constrain the manager’s ability to distort segment earnings again next period. Nevertheless, dynamic considerations may still apply to how total reported earnings this period are divided across segments. For example, investors may form expectations of segment-level growth using reported earnings for a particular segment. In this first test of the theory, we abstract away from these dynamic concerns and consider a manager who distorts the consistency of segment earnings to improve short-run perceptions of her managerial ability, e.g., to improve the manager’s probability of receiving an outside job offer.

We empirically test whether segment earnings display abnormally high (low) consistency when overall firm performance is better (worse) than expected. Our analysis allows for the possibility that the consistency of segment earnings varies with firm performance for other natural reasons. For example, bad times may cause higher volatility across segments. In addition, performance across segments may be less variable during good times because good firm-level news is caused by complementarities arising from the good performance of related segments. There may also be scale effects, in that the standard deviation of news may naturally increase in the absolute values of the news. Therefore, we don’t use zero correlation between the consistency of segment earnings and overall firm performance as our null hypothesis.

Instead, we compare the consistency of reported segment earnings to a benchmark consistency of earnings implied by segment-level sales data. Like earnings, the consistency of segment sales may vary with firm performance for natural reasons. However, sales are more difficult to distort because they are reported prior to the deduction of costs. This benchmark consistency implied by segment

²³Our analysis is also motivated by anecdotal evidence that managers emphasize consistent or inconsistent segment news depending on whether overall firm performance is good or bad. For example, Walmart’s 2015 Q2 10Q highlights balanced growth following strong performance, “Each of our segments contributes to the Company’s operating results differently, but each has generally maintained a consistent contribution rate to the Company’s net sales and operating income in recent years.” In contrast, Hewlett-Packard CEO Meg Whitman highlights contrarian segment performance after sharply negative growth in five out of six segments in 2015, “HP delivered results in the third quarter that reflect very strong performance in our Enterprise Group and substantial progress in turning around Enterprise Services.”

sales leads to a conservative null hypothesis. Prior to strategic cost allocations, managers may have already distorted the consistency of segment sales through transfer pricing or the targeted allocation of effort and resources across segments. We also compare the consistency of segment earnings in real multi-segment firms to that of counterfactual firms constructed from matched single-segment firms. Unlike real multi-segment firms, the matched counterfactual firms mechanically cannot shift costs across segments to alter the consistency of reported earnings. However, the matched sample should capture natural changes in consistency that may be driven by industry trends among connected segments during good and bad times.

4.2 Data and empirical framework

We use Compustat segment data merged with I/B/E/S and CRSP for multi-segment firms in the years 1976-2014. We restrict the sample to business and operating segments (some firms report geographic segments in addition to business segments). We exclude observations if they are associated with a firm that, at any point during our sample period, contained a segment in the financial services or regulated utilities sectors, as these firms face additional oversight over their operations and accounting disclosure. In our baseline analysis, we also exclude very small segments (segments with assets in the previous year less than one-tenth that of the largest segment), although we explore how our results vary with size ratios in supplementary analysis.

We measure segment earnings as EBIT (raw earnings) scaled by segment assets (assets are measured as the average over the current and previous year). This scaled measure of earnings is also known as return on assets (ROA). We focus on this scaled measure of earnings because it is commonly used by financial analysts, investors, and corporate boards to assess performance and is easily comparable across firms and segments of different sizes. We measure firm earnings as the sum of segment EBIT divided by the sum of segment assets. This is commonly known as firm-level ROA. Due to the scaling, firm earnings are equal to the weighted mean of segment earnings, with the weight for each segment equal to segment assets divided by total firm assets. As shown earlier in Section 3.5, we can extend our model to a setting with weights over the pieces of news. The vector of news $x = (x_1, \dots, x_n)$ represents segment earnings, with weighted mean \bar{x}_w and weighted standard deviation s_w . All the main model predictions carry over to a setting with weights. In particular, \bar{x}_w remains constant if costs are shifted across segments. The intuition is that, while a shift in costs will have a greater impact on the (scaled) earnings of a smaller segment due to its smaller denominator, smaller segments have less weight, so shifting costs from one segment to another does not affect the weighted mean \bar{x}_w . This fits with our model in which managers can distort the consistency of news (as measured by s_w), holding \bar{x}_w constant.²⁴

²⁴In supplementary results, omitted for brevity, we find approximately similar results if we instead equal-weight each segment within a firm-year. Using equal weights, segment news is measured by EBIT scaled by assets within the segment, and firm news is measured as the equal-weighted mean of segment news.

We use segment sales data to construct a benchmark for how the consistency of segment earnings would vary with overall firm news in the absence of strategic cost allocations. Consider segment i in firm j in year t . Total firm earnings (unscaled) equal total sales minus total costs ($E_{jt} = Sales_{jt} - Costs_{jt}$) and segment earnings (unscaled) equal segment sales minus costs associated with the segment ($e_{ijt} = sales_{ijt} - costs_{ijt}$). For our first benchmark, we use a “proportional costs” assumption. We assume that, absent distortions, total costs are associated with segments according to the relative levels of sales for each segment. Predicted segment earnings (scaled by segment assets a_{ij}) can be estimated as:

$$\frac{\widehat{e}_{ijt}}{a_{ijt}} = \frac{1}{a_{ijt}} \left(sales_{ijt} - \frac{sales_{ijt}}{Sales_{jt}} \cdot Costs_{jt} \right). \quad (10)$$

We estimate the predicted consistency as the log of the weighted standard deviation of the predicted segment earnings:

$$\widehat{s}_{jt} \equiv \log \left(SD \left(\frac{\widehat{e}_{ijt}}{a_{ijt}} \right) \right). \quad (11)$$

Our baseline regression specification tests whether the difference between the actual standard deviation and predicted standard deviation of segment earnings depends on whether firm news exceeds expectations:

$$s_{jt} - \widehat{s}_{jt} = \beta_0 + \beta_1 I_{jt}^{goodnews} + controls + \epsilon_{jt}. \quad (12)$$

$I_{jt}^{goodnews}$ is a dummy variable for whether overall firm news exceeds expectations. *Controls* include year fixed effects and the weighted mean of the absolute values of segment sales and earnings, to account for scale effects in the average relationship between standard deviations and means in the data. Standard errors are allowed to be clustered by firm.

We refer to $s_{jt} - \widehat{s}_{jt}$ as the abnormal standard deviation of segment earnings. Our null hypothesis is $\beta_1 = 0$, i.e., that differences between the actual and predicted standard deviations of segment earnings are unrelated to whether the firm is releasing good or bad news overall. This null hypothesis allows for the possibility that we predict the consistency of segment earnings with error, but requires that the prediction error is uncorrelated with whether firm news exceeds expectations. Our model of strategic distortion of consistency predicts that $\beta_1 < 0$, i.e., that the abnormal standard deviation of segment earnings is lower when firm news is good than when firm news is bad.

We can also use industry data to improve the predictions of earnings consistency absent cost allocation distortions. Instead of assuming that total costs would be associated with segments according to the relative levels of segment sales, we can further adjust using industry averages calculated from single-segment firms in the same industry. This helps to account for the possibility that some segments are in industries that tend to have very low or high costs relative to sales.

Let γ_{it} equal the average ratio of costs to sales among single segment firms in the SIC2 industry corresponding to segment i in each year. Let $Z_{jt} \equiv \sum_i (\gamma_{it} \cdot sales_{ijt})$. Under an “industry-adjusted” assumption, total costs are associated with segments according to the relative, industry-adjusted, level of sales of each segment:

$$\widehat{\frac{e_{ijt}}{a_{ijt}}} = \frac{1}{a_{ijt}} \left(sales_{ijt} - \frac{\gamma_{it} \cdot sales_{ijt}}{Z_{jt}} Costs_{jt} \right) \quad (13)$$

We can then substitute the above definition for Equation (10) and reestimate our baseline regression specification.

Our baseline specification assumes that the receiver focuses on earnings news in terms of the level of earnings scaled by assets, otherwise known as ROA. The receiver of news may alternatively focus on performance relative to other similar firms. We can extend our analysis to the case in which receivers of earnings news focus on earnings relative to the industry mean. We measure relative segment earnings as $\frac{e_{ijt}}{a_{ijt}} - m_{it}$, where m_{it} is the value-weighted mean earnings (also scaled by assets) for the segment’s associated SIC2 industry in year t . We measure firm relative earnings as $\frac{E_{ijt}}{A_{jt}} - M_{ijt}$, where $M_{ijt} \equiv \sum_i \left(\frac{a_{ijt}}{A_{jt}} \right) m_{it}$. Using these measures, firm-level relative earnings is equal to the weighted mean of segment relative earnings, with the weight for each segment again equal to segment assets divided by total assets. The predicted relative earnings for each segment is simply $\widehat{\frac{e_{ijt}}{a_{ijt}}} - m_{it}$, where $\widehat{\frac{e_{ijt}}{a_{ijt}}}$ is as defined in Equations (10) or (13). Using these measures, we can let $s_{jt} - \widehat{s}_{jt}$ equal the difference between the real and predicted log weighted standard deviations of relative segment earnings and reestimate our baseline regression specification in Equation (12).

In Equation (12), $I_{jt}^{goodnews}$ is a dummy variable for whether overall firm news exceeds expectations. In our baseline specifications, $I_{jt}^{goodnews}$ indicates whether total firm earnings exceeds the same measure in the previous year. In tests focusing on relative segment earnings, we can instead let $I_{jt}^{goodnews}$ be an indicator for whether total firm earnings exceeds the industry mean (M_{ijt}). In supplementary tests, we find similar results if $I_{jt}^{goodnews}$ is an indicator for whether total firm earnings exceeds zero, the “break even” point.

Finally, we can measure firm news continuously as (1) the difference between total firm earnings and the same measure in the previous year, or (2) the difference between total firm earnings and the industry mean (M_{ijt}). Our theory does not predict that the consistency of segment earnings should increase *continuously* with firm performance. Rather, the theory predicts a jump in abnormal consistency when firm performance exceeds expectations. For example, the theory predicts that managers will increase consistency when firm news exceeds expectations, but not more so when firm news greatly exceeds expectations. However, the empirically-measured relationship between the consistency of segment earnings and firm performance may be smooth because we use noisy proxies for the expectations of those viewing the segment news disclosures.

Table 1
Summary Statistics

This table summarizes the data used in our baseline regression sample. Each observation represents a firm-year. Segment earnings equal segment EBIT divided by segment assets (the average of segment assets in the current and previous years). Segment sales are also scaled by assets. Firm earnings and sales are equal to the weighted means of segment earnings and sales, respectively, where the weights are equal to segment assets divided by total assets. All means and standard deviations are weighted and calculated using the segment data within each firm-year. *Good firm news* is an indicator for whether firm earnings in the current year exceeds the level in the previous year. *Good relative firm news* is an indicator for whether firm earnings exceeds the industry mean (calculated as in Section 4.2) in the same year. *Firm earnings* > 0 is an indicator for whether firm earnings is positive. Δ *Firm earnings* measures the continuous difference between firm earnings in the current and previous years. *Firm relative earnings* measures the continuous difference between firm earnings and industry mean earnings in the current year.

	Mean	Std. dev.	p25	p50	p75
Number of segments	2.575	0.936	2	2	3
Firm earnings (= mean earnings)	0.134	0.146	0.067	0.129	0.199
Std. dev. earnings	0.115	0.133	0.037	0.076	0.141
Log std. dev. earnings	-2.705	1.145	-3.309	-2.582	-1.962
Firm sales (= mean sales)	1.657	0.951	1.054	1.511	2.020
Std. dev. sales	0.545	0.573	0.184	0.371	0.701
Log std. dev. sales	-1.117	1.134	-1.694	-0.991	-0.356
Good firm news (dummy)	0.496				
Good relative firm news (dummy)	0.558				
Firm earnings > 0 (dummy)	0.895				
Δ Firm earnings (continuous)	0.013	0.097	-0.021	0.014	0.047
Firm relative earnings (continuous)	0.012	0.249	-0.052	0.008	0.073

Table 1 summarizes the data. Our baseline regression sample consists of 4,297 firms, corresponding to 23,276 firm-years observations. This final sample is derived from an intermediate sample of 60,085 segment-firm-year observations. For a firm-year observation to be included in the sample, we require that the firm reports the same set of segments in the previous year, which allows us to measure segment assets in the previous year as well as annual changes. The first year of a firm \times segment reporting format is excluded from the sample. We present summary statistics of the weighted means and standard deviations of segment earnings and sales. All measures of earnings and sales in this and future tables are scaled by assets unless otherwise noted.

Finally, we emphasize that throughout the empirical tests, we do not take a stand on whether investors, boards, or other receivers of firm earnings news are sophisticated or naïve about the distortion of consistency. The main prediction from the model does not require that receivers

rationally expect distortion, just that they use the consistency of earnings as a measure of the precision of the overall earnings signal and that managers react by manipulating consistency. If receivers do anticipate distortion, then as shown earlier, the same predictions apply.

4.3 Empirical results

Table 2
Consistency of Segment Earnings

The dependent variables in Columns 1 and 2 are the log standard deviation of segment earnings and sales, respectively. The dependent variables in Columns 3 and 4 are the abnormal log standard deviations of segment earnings, relative to predictions calculated using reported segment sales under a proportional costs assumption and industry-adjusted assumption, respectively. Control variables include the *good firm news* indicator, year fixed effects, and controls for the absolute means of segment earnings and sales. All means and standard deviations are weighted by segment assets divided by total assets. All earnings and sales measures are scaled by assets. Standard errors are clustered by firm. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

	SD Earnings	SD Sales	Abnormal SD Earnings	
	(1)	(2)	(3)	(4)
Good firm news	-0.0975*** (0.0150)	0.0146 (0.0138)	-0.111*** (0.0179)	-0.0884*** (0.0198)
Cost assumption			Prop	Ind adj
Control for mean	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
R ²	0.0643	0.129	0.179	0.0231
Obs	23276	23276	23276	23276

We begin by using regression analysis to test our model prediction that segment earnings are more consistent when firm news is good than when firm news is bad. Table 2 presents our main results. Column 1 regresses the log of the weighted standard deviation of segment earnings on an indicator for good firm news (whether firm earnings beat the same measure last year) as well as controls for year fixed effects and the absolute value of the mean of segment earnings and sales. We find support for the main model prediction that $\beta_1 < 0$. After controlling for general scale effects, the standard deviation of segment earnings is lower when firm news is good than when it is bad. In contrast, Column 2 shows that the standard deviation of segment sales does not vary significantly with the indicator for good firm news.

To more formally test whether the consistency of earnings appears to have been manipulated, Column 3 estimates our baseline specification as described in Equation (12), which compares the consistency of reported segment earnings to a benchmark consistency of segment earnings implied by segment-level sales data. This specification helps to account for other natural factors that may impact the variability of news across segments during good or bad times, assuming that these

factors similarly impact segment sales. We find that good firm news corresponds to an abnormal 11% decline in the standard deviation of segment earnings. In Column 4, we find similar results after using an industry-adjusted assumption to create a benchmark consistency of segment earnings from the reported sales data. Under an industry-adjusted assumption, good firm news corresponds to an abnormal 9 percent decline in the standard deviation of segment earnings. These results support a model in which management distort cost allocations so that good earnings news is consistent and bad earnings news is inconsistent.


Figure 4

Real vs. predicted consistency of segment earnings

These graphs show how the abnormal standard deviation of segment earnings varies with overall firm news. The x-axis represents firm earnings in the current year minus firm earnings in the previous year. We measure the abnormal standard deviation of segment earnings as the difference between the log weighted standard deviation of segment earnings in the real data and the log weighted standard deviation of predicted segment earnings calculated from reported segment sales data. Predicted segment earnings are formed using a proportional costs assumption in Panel A and an industry-adjusted assumption in Panel B, as described in Section 4.2. We plot how the abnormal standard deviation of segment earnings varies with firm earnings, after controlling for fiscal year fixed effects. The curves represent local linear plots estimated using the standard rule-of-thumb bandwidth. Gray areas indicate 90 percent confidence intervals.

Panel A: Proportional costs assumption

Panel B: Industry-adjusted assumption



We can also visually explore the relationship between the consistency of segment earnings and firm-level news. Figure 4 shows that the abnormal standard deviation of segment earnings is lower when firm news is good than when firm news is bad. We measure firm-level news relative to expectations as the difference between firm earnings in the current year and the same measure in the previous year. We measure the abnormal standard deviation of segment earnings as the difference between the log standard deviation of segment earnings in the real data and the log standard deviation of predicted segment earnings calculated from reported segment sales data.

Focusing on the difference between the actual and predicted measures helps to account for potential unobserved factors that may cause the standard deviation of segment earnings differ between good and bad times. These factors should similar affect reported segment sales. However, segment sales are more difficult to manipulate because, unlike earnings, they are reported prior to the discretionary allocation of shared costs.

In Panels A and B, we use a proportional costs and industry-adjusted assumption, respectively, to predict how earnings would look using reported sales data. Our null hypothesis is that prediction error, i.e., the difference between the standard deviations of real and predicted earnings, should be uncorrelated with whether firm-level news exceeds expectations. Instead, we find that the abnormal standard deviation of earnings sharply declines around zero, when firm-level earnings exceeds the same measure in the previous year. This supports our model prediction that managers strategically allocate costs so that good news is consistent and bad news is inconsistent.

As noted previously, the model does not predict that the abnormal standard deviation of segment earnings should decline continuously with firm performance. Instead, the theory predicts a discrete drop in the abnormal standard deviation of earnings when firm performance exceeds expectations. However, because we proxy for expectations with error, the observed relationship may look more smooth. This is consistent with what we observe in both panels of Figure 4. The abnormal standard deviation of segment earnings declines approximately continuously with firm earnings. In addition, the slope is steepest near zero, when firm earnings in the current year is equal to earnings in the previous year, suggesting that the previous year's earnings represents a reasonable proxy for expectations.

Table 3
Consistency of Relative Segment Earnings

This table reestimates the results in Table 2 using performance measures relative to industry means. Relative earnings, sales, and abnormal relative earnings are as defined in Section 4.2. All other variables are as defined in Table 2. Standard errors are clustered by firm. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

	SD Relative Earnings	SD Relative Sales	Abnormal SD Relative Earnings	
	(1)	(2)	(3)	(4)
Good relative firm news	-0.282*** (0.0260)	-0.126*** (0.0236)	-0.180*** (0.0300)	-0.149*** (0.0324)
Cost assumption			Prop	Ind adj
Control for mean	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
R ²	0.0669	0.119	0.0728	0.0292
Obs	23276	23276	23275	23276

In our baseline analysis, we assume the receiver focuses on segment and firm earnings news in terms of the level of earnings (which is equal to ROA because we scale by assets). The receiver

of news may alternatively focus on performance relative other similar firms. In Table 3, we can extend our analysis to the case in which receivers focus on earnings relative to the industry mean. As described in Section 4.2, we can measure relative segment earnings as segment earnings minus the mean in the segment’s SIC2 industry. We measure firm level news as firm earnings minus the weighted mean performance of the firm’s associated segment industries. The *good relative firm news* indicator is equal to one if firm news exceeds the weighted industry mean. Using these relative performance measures, we again find evidence consistent with strategic distortion of segment news. Column 1 shows that the standard deviation of segment relative earnings is much lower when the firm is underperforming relative to its industry peers than when it is outperforming its peers. Column 2 shows that the standard deviation of segment sales is also significantly lower when the firm is outperforming, although the absolute magnitude of β_1 is smaller than that for segment earnings. This suggests that good times in terms of relative performance may naturally be associated with lower variance in segment news (or that the manager has already strategically manipulated the consistency of segment sales through targeted effort/resource allocation). However, Columns 3 and 4 show that the standard deviation of real segment earnings varies significantly more with overall firm news than predicted given the reported sales data. Under a proportional costs assumption in Column 3, good relative firm news corresponds to an abnormal 18 percent decline in the standard deviation of relative segment earnings. Under an industry-adjusted assumption in Column 4, good relative firm news corresponds to an abnormal 15 percent decline in the standard deviation of relative segment earnings.

Overall, the evidence is supportive of the model prediction that managers strategically allocate shared costs so that standard deviation of segment earning declines more with good firm news than implied by the reported sales data. Next, we compare the behavior of the consistency of segment EBIT within true multi-segment firms with the consistency of matched segments constructed using single-segment firm data. Unlike real multi-segment firms, the matched counterfactual firms mechanically cannot shift resources across segments to alter the consistency of reported earnings. However, if our results are driven by industry trends among connected segments during good or bad times, we expect to find similar results with industry-matched placebo segments.

We take single-segment firms that have product lines that are comparable to segments in multi-segment firms, and assign them together to mimic multi-segment firms. Specifically, we match each segment-firm-year observation corresponding to a real multi-segment firm to a single segment firm in the same year and SIC2 industry that is the nearest neighbor in terms of Mahalanobis distance for the lagged EBIT, assets, and sales (all unscaled to also match on size). We then regress the log weighted standard deviation of the matched segment earnings, sales, or abnormal earnings on the firm’s performance measures, as in Equation (12). In Table 4, we find that in these artificial multi-segment firms, there is no negative relation between firm performance and the standard deviations

Table 4
Consistency of Matched Segment Earnings

This table reestimates the results in Table 2 using matched segment earnings and sales data. We match each segment-firm-year observation corresponding to a real multi-segment firm to a single segment firm in the same year and SIC2 industry that is the nearest neighbor in terms of the Mahalanobis distance for the lagged unscaled levels of EBIT, assets, and sales. Using the matched data, we calculated the log standard deviations and means of segment earnings, sales, and the abnormal standard deviation of earnings. All other variables are as defined in Table 2. Standard errors are clustered by firm. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

	SD Earnings	SD Sales	Abnormal SD Earnings	
	(1)	(2)	(3)	(4)
Good firm news	0.0232 (0.0184)	0.00911 (0.0166)	0.0176 (0.0222)	0.0287 (0.0244)
Cost assumption			Prop	Ind adj
Control for mean	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
R ²	0.0119	0.0825	0.211	0.0353
Obs	17192	17191	17192	17192

of segment earnings or sales. Further, the abnormal standard deviation of segment earnings is not significantly related to the *good firm news* indicator. This is further evidence that managers in real multi-segment firms use their control over cost allocations to distort reported segments earnings to convey consistent good news and inconsistent bad news.

4.4 Robustness

A potential alternative explanation for our empirical findings could be that the variation in shocks to segment-specific costs is larger when overall firm news is bad rather than good. For example, the firm may, for legal or business reasons, record a major write-down in a single segment, leading to low earnings for that segment, and a high standard deviation of earnings across segments. This negative cost shock to a specific segment may be so large that it substantially drags down the measure of overall firm earnings. If so, we would measure a larger standard deviation of earnings when overall firm news is bad rather than good. Our main benchmark using reported sales data would not account for this potential alternative explanation which focuses on a shock to segment costs rather than segment sales. However, we believe this story is unlikely to drive our results for two reasons. First, the story should lead to similar patterns in our sample of matched counterfactual firms, and we observe no such patterns in the data. Second, we continue to find a similar relation between the standard deviation of segment earnings and overall firm news even when the measure of firm news is not related to a negative cost shock.

In the first two columns of Table 5 Panel A, we proxy for firm-level news using only sales

data. We consider the firm news to be good when total sales this year exceeds the level in the previous year. We continue to find a negative relation between firm news and the abnormal standard deviation of segment earnings. This negative relation cannot be explained by a negative cost shock affecting a specific segment because the cost shock does not contribute to the sales-based measure of firm news. In Columns 3 and 4, we measure good firm news as whether total for earnings exceeded the level in the previous year, except we exclude the earnings of the segment that saw that greatest decline in earnings from the calculation. We again find a negative relation between firm news and the abnormal standard deviation of segment earnings, even though our measure of firm news is unrelated to a cost shock to the worst-performing segment.

In the second panel of Table 5, we explore the robustness of our results to alternative measures of good firm news more generally. So far in the analysis, we have proxied for whether overall firm news exceeds expectations using dummies for whether firm earnings exceeds the level in the previous year or the industry mean in the same year. In Columns 1 and 2, we define firm news to be good if earnings are positive, i.e., the firm is in the black rather than the red. We again find that the abnormal standard deviation of segment earnings is significantly lower when firm news is good, as measured by earnings being positive. Next, we move to two continuous measures of firm news. As discussed previously, the model predicts a discrete drop in the abnormal standard deviation of earnings when firm performance exceeds expectations. However, the empirically-measured relationship between the consistency of segment earnings and firm performance may be more smooth because we use noisy proxies for the expectations of those viewing the segment news disclosures. In Columns 3 through 6, we find that the abnormal standard deviation of segment earnings declines significantly with these continuous measures of firm news.

Next, we show that our results are robust to reasonable alternative sample cuts and control variables. In Columns 1 through 4 of Table 6 Panel A, we examine the data before and after the passage of SFAS No. 131 in 1997. SFAS No. 131 increased the prevalence of segment reporting among US public firms by requiring that firms disclose segment performance if the segments are evaluated internally as separate units, even if the segments operate in related lines of business. We find significant evidence of distortion of segment earnings as predicted by the model in the periods before and after the policy change.²⁵

In Columns 5 and 6, we limit the sample to firms-years in which all segments are similar in

²⁵Despite regulations mandating segment disclosure, firms may still have some leeway in terms of whether to disclose segment data and which segments to disclose. In relation to our model, it would be very interesting to explore how the choice of which segments to disclose depends on the manager's beliefs about whether overall firm news will be good or bad. In practice, empirical investigation of this question is challenging because firms without major restructuring episodes rarely switch reporting formats year to year and almost always continue to disclose a segment's performance if they disclosed it in the past. Isolated deviations from this norm, e.g., when Valeant Pharmaceuticals switched from reporting five to two segments in 2012, are noted with great suspicion in the financial press. Therefore, we focus our empirical tests on manipulation of consistency after the firm has committed to release performance measures for a set of segments.

terms of size. In our main analysis, we give more weight to larger segments because they contribute more to overall firm profits and may be more informative of managerial ability. By limiting the sample to firms with similarly-sized segments, we can further check that our results are not driven by relatively small and potentially anomalous segments. After excluding observations in which the size (as measured by assets in the previous year) of the largest segment exceeds the size of the smallest segments by more than a factor of two, we continue to find that the abnormal standard deviation of segment earnings is significantly smaller during good times than bad.

In Columns 1 and 2 of Panel B, we reestimate our baseline specification described by Equation (12), but exclude controls for the absolute values of mean segment earnings and sales within each firm-year. In regressions in which the dependent variable is the standard deviation of segment earnings or sales, these control variables help account for potential scale effects in which the standard deviation of larger numbers naturally tend to be larger. However, our baseline specification tests whether the difference between the actual and predicted standard deviations of earnings is lower when the firm is reporting good news overall. The predictions using segment sales should already account for potential scale effects, implying that we do not need to further control for scale. We find qualitatively similar estimates, slightly larger in absolute magnitude, if we exclude these scale controls.

In Columns 3 and 4, we show that our results are not caused by sales being bounded below by zero. Unlike earnings, sales cannot be negative. The zero lower bound for sales may mechanically limit the standard deviation of segment sales when overall firm news is bad. We find similar results after restricting the sample to observations in which the minimum level of segment sales in the previous firm-year is above the 25th percentile in the sample.

Finally in Columns 5 and 6, we explore how the abnormal consistency of segment earnings varies with firm news, after controlling for firm \times reporting format fixed effects (the set of reported segments remain constant within a reporting format). We do not control for these fixed effects in our baseline specifications because we wish to use both across-firm variation as well as within-firm variation over time. After controlling for firm \times reporting format fixed effects, we continue to find a strong and significant negative relationship between the abnormal standard deviation of segment earnings and the *good firm news* indicator.

Finally, we explore how the results vary using an alternative “average costs” assumption to predict how total costs within each firm-year would be associated with segments in the absence of strategic distortions. We use information on the average fraction of total costs assigned to each segment over time. For each segment, we calculate the average fraction of total costs that are allocated to the segment, based upon reported segment earnings and sales over the entire period in which a firm-reporting format exists in the data. For example, suppose a firm has three segments A, B, and C. On average, over a five year period, segment A is allocated 0.5 of total costs, segment

B is allocated 0.2 of total costs, and segment C is allocated 0.3 of total costs. We can then predict segment earnings in each year, assuming that, absent strategic cost allocations, segments A, B, and C would be assigned 0.5, 0.2, and 0.3, respectively, of total costs in each year. We then test whether the difference between the standard deviation of real segment earnings and these predicted earnings varies with whether the firm is releasing good news overall. In other words, we test whether cost allocations differ from the within-firm mean over time in a manner predicted by the model.

Table 7 shows the results using this average costs assumption. Because our predictions use information on average cost allocations, our regression coefficients are identified from within-firm variation in firm-level news over time. To focus on this within-firm variation, we show specifications that include firm-reporting format fixed effects in Columns 2 and 4. In addition, we also present results using continuous measures of overall firm performance in Columns 3 and 4 to allow for greater variation in firm-level news over time. In all specifications, we continue to find that the abnormal standard deviation of segment earnings is lower when firm-level news is good. These results support our theory of consistent good news and inconsistent bad news and show that our findings are robust to a variety of different cost allocation benchmarks.

Table 5**Alternative definitions of good firm news**

This table explores alternative measures of firm-level news. Columns 1 and 2 of Panel A measure good firm news using an indicator for whether firm sales exceed the level in the previous year. Columns 3 and 4 measure good firm news using an indicator for whether total segment earnings, calculated excluding the segment with the greatest drop in earnings, exceeded the level (estimated using the same subset of segments) in the previous year. In Columns 1 and 2 of Panel B, we measure good firm news using an indicator for whether firm earnings exceeds zero. Columns 3 and 4 use a continuous measure of firm news equal to the difference between firm earnings in the current and previous year. Columns 5 and 6 use a continuous measure of firm news equal to the difference between firm earnings in the current year and the weighted industry mean, calculated as described in Section 4.2. All other variables are as defined in Tables 2 and 3. Standard errors are clustered by firm. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Panel A						
	Abnormal SD Earnings					
	(1)	(2)	(3)	(4)		
Good firm sales news	-0.171*** (0.0251)	-0.128*** (0.0285)				
Good firm news (excluding worst-performing segment)			-0.0914*** (0.0198)	-0.0714*** (0.0221)		
Cost assumption	Prop	Ind adj	Prop	Ind adj		
Control for mean	Yes	Yes	Yes	Yes		
Year FE	Yes	Yes	Yes	Yes		
R ²	0.179	0.0234	0.178	0.0226		
Obs	23276	23276	23276	23276		

Panel B						
Abnormal SD:	Earnings				Relative Earnings	
	(1)	(2)	(3)	(4)	(5)	(6)
Firm earnings > 0	-0.694*** (0.0398)	-0.546*** (0.0454)				
ΔFirm earnings (cont)			-0.495*** (0.110)	-0.293** (0.114)		
Firm relative earnings (cont)					-0.669*** (0.0955)	-0.497*** (0.100)
Cost assumption	Prop	Ind adj	Prop	Ind adj	Prop	Ind adj
Control for mean	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
R ²	0.197	0.0332	0.179	0.0226	0.0739	0.0294
Obs	23276	23276	23276	23276	23275	23276

Table 6
Robustness

This table reestimates Columns 3 and 4 of Table 2 using alternative sample restrictions and control variables. In Panel A, Columns 1 and 2 are restricted to fiscal years ending on or before 1997 (the year when SFAS No. 131 passed) while Columns 3 and 4 are restricted for fiscal years ending after 1997. Columns 5 and 6 are restricted to firm-year observations in which the assets (measured in the previous year) of the largest segment did not exceed the assets of the smallest segment by more than a factor of 2. In Panel B, Columns 1 and 2, we exclude control variables for the absolute mean of segment earnings and sales. In Columns 3 and 4, we restrict the sample to observations in which the minimum level of segment sales in the previous firm-year is above the 25th percentile in the sample in each year. In Columns 5 and 6, we add in control variables for fixed effects for each firm x segment reporting format (a segment reporting format is a consecutive period in which the firm reports the same set of segments). All other variables are as defined in Table 2. Standard errors are clustered by firm. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Panel A						
Abnormal SD Earnings	Year ≤ 1997		Year > 1997		Similarly Sized Segments	
	(1)	(2)	(3)	(4)	(5)	(6)
Good firm news	-0.132*** (0.0230)	-0.0931*** (0.0248)	-0.0760*** (0.0286)	-0.0887*** (0.0333)	-0.138*** (0.0273)	-0.0972*** (0.0311)
Cost assumption	Prop	Ind adj	Prop	Ind adj	Prop	Ind adj
Control for mean	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
R ²	0.164	0.0189	0.203	0.0324	0.159	0.0221
Obs	14275	14275	9001	9001	10818	10818
Panel B						
Abnormal SD Earnings	No Mean Controls		Exclude Low Sales		Firm-Reporting Format FE	
	(1)	(2)	(3)	(4)	(5)	(6)
Good firm news	-0.253*** (0.0193)	-0.134*** (0.0196)	-0.107*** (0.0207)	-0.100*** (0.0229)	-0.0902*** (0.0217)	-0.0559** (0.0244)
Cost assumption	Prop	Ind adj	Prop	Ind adj	Prop	Ind adj
Control for mean	No	No	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
R ²	0.0134	0.00974	0.166	0.0251	0.673	0.614
Obs	23276	23276	17454	17454	23276	23276

Table 7**Predicted earnings using an average costs assumption**

This table presents results using an average costs assumption to estimate the predicted segment earnings benchmark. For each segment, we calculate the average fraction of total costs that are allocated to the segment based upon reported segment earnings and sales data over the entire period in which a firm-reporting format exists in the data. We predict that, absent strategic distortions, each segment would be associated with this average fraction of total costs for each segment-year in the data. We then estimate the abnormal standard deviation of segment earnings as the difference between the the standard deviation of real segment earnings and these predicted earnings. To use within-firm variation over time, we introduce firm-reporting format fixed effects in Columns 2 and 4. To allow for greater within-firm variation in firm-level news, we use a continuous measure of overall firm performance in Columns 3 and 4. All other variables are as defined in Table 2. Standard errors are clustered by firm. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

	Abnormal SD Earnings (Average Costs Assumption)			
	(1)	(2)	(3)	(4)
Good firm news	-0.0859*** (0.0145)	-0.0725*** (0.0206)		
Δ Firm earnings (cont)			-0.664*** (0.0857)	-0.665*** (0.135)
Control for mean	Yes	Yes	Yes	Yes
Firm-reporting format FE	No	Yes	No	Yes
Year FE	Yes	Yes	Yes	Yes
R ²	0.0883	0.512	0.0900	0.514
Obs	23247	23247	23247	23247

5 Conclusion

These results show that selective news distortion to affect the consistency and hence persuasiveness of news can be an important strategy in sender-receiver environments. Theoretically, we show when a more precise signal is more persuasive in that it moves the posterior distribution and posterior estimate more strongly away from the prior, and when more consistent news implies that the mean of the news is a more precise signal. We then show that the most persuasive strategy for a sender is to increase or decrease consistency in good or bad times by focusing on distorting the least or most favorable news respectively. These incentives for selective distortion lead to the testable implication that reported news is more consistent when it is favorable than unfavorable.

We test the model on firms that report earnings data for multiple segments. We show that earnings are more consistent when the firm is doing well than when the firm is doing poorly. This same pattern does not arise in the consistency of reported sales across firms. Similarly the same pattern does not arise in a sample of stand-alone firms matched to segments in multi-segment firms. Therefore, the evidence supports the interpretation that multi-segment firms shift the allocation of costs across segments to manipulate the consistency of reported earnings.

6 Appendix

Proof of Property 1. For given q the likelihood ratio $g(\bar{x} - q|s')/g(\bar{x} - q|s)$ is increasing in \bar{x} if $g(\bar{x} - q|s)$ is log-supermodular in (\bar{x}, s) . Since log-supermodularity is preserved by integration, from (2) this holds if

$$\frac{1}{(\sigma_\varepsilon \sqrt{2\pi})^n} e^{-\frac{ns^2 + n(\bar{x}-q)^2}{2\sigma_\varepsilon^2}} \quad (14)$$

is log-supermodular in $(\bar{x}, s, \sigma_\varepsilon)$ (Lehmann (1955) - see discussion of Lemma 2 in Athey (2002)). This holds if all the cross-partials are non-negative (Topkis, 1978). Checking from (2), the cross-partials of the log of (14) are

$$\begin{aligned} \frac{d}{d\bar{x}} \frac{d}{ds} &= 0 \\ \frac{d}{d\sigma_\varepsilon} \frac{d}{ds} &= \frac{2ns}{\sigma_\varepsilon^3} \\ \frac{d}{d\sigma_\varepsilon} \frac{d}{d\bar{x}} &= \frac{2n(\bar{x} - q)}{\sigma_\varepsilon^3} \end{aligned} \quad (15)$$

which are all non-negative for $\bar{x} \geq q$. Hence $g(\bar{x} - q|s')/g(\bar{x} - q|s)$ is increasing in \bar{x} for $\bar{x} \geq q$ and, by symmetry, $g(\bar{x} - q|s')/g(\bar{x} - q|s)$ is decreasing in \bar{x} for $\bar{x} \leq q$. ■

Proof of Property 2. Since μ is unrestricted, assume WLOG that $y = 0$. Then, using the symmetry of g ,

$$\begin{aligned} E[q|y = 0, \rho] &= \frac{\int qg(-q|\rho)f(q)dq}{\int g(-q|\rho)f(q)dq} \\ &= \frac{\int_{-\infty}^0 qg(-q|\rho)f(q)dq + \int_0^\infty qg(-q|\rho)f(q)dq}{\int_{-\infty}^0 g(-q|\rho)f(q)dq + \int_0^\infty g(-q|\rho)f(q)dq} \\ &= \frac{\int_0^\infty qg(q|\rho)(f(q) - f(-q))dq}{\int_0^\infty g(q|\rho)(f(q) + f(-q))dq} \\ &= \frac{\int_0^\infty qg(q|\rho)(f(q) - f(-q)) \frac{f(q)+f(-q)}{f(q)+f(-q)} dq}{\int_0^\infty g(q|\rho)(f(q) + f(-q))dq} \\ &= \int_0^\infty v(q)dZ(q) \end{aligned} \quad (16)$$

where

$$v(q) = q \frac{f(q) - f(-q)}{f(q) + f(-q)} \quad (17)$$

and

$$Z(q|\rho) = \frac{\int_0^q g(\bar{q}|\rho)(f(\bar{q}) + f(-\bar{q}))d\bar{q}}{\int_0^\infty g(\bar{q}|\rho)(f(\bar{q}) + f(-\bar{q}))d\bar{q}}. \quad (18)$$

Integrating (16) by parts,

$$E[q|y = 0, \rho] - E[q|y = 0, \rho'] = \int_0^\infty v'(q) (Z(q|\rho') - Z(q|\rho)) dq. \quad (19)$$

so if $Z(q|\rho') \succ_{FOSD} Z(q|\rho)$ then $E[q|y=0, \rho] - E[q|y=0, \rho']$ depends on the sign of $v'(q)$.

Regarding FOSD,²⁶ by assumption $g(q|\rho') \succ_{UV} g(q|\rho)$, so for $\rho' > \rho$ and $q > 0$, $g(q|\rho') \succ_{MLR} g(q|\rho)$, i.e., $g(q|\rho')/g(q|\rho)$ is increasing. So for any $a > b$,

$$\frac{g(a|\rho') (f(a) + f(-a))}{g(a|\rho) (f(a) + f(-a))} > \frac{g(b|\rho') (f(b) + f(-b))}{g(b|\rho) (f(b) + f(-b))} \quad (20)$$

or

$$\frac{g(a|\rho') (f(a) + f(-a))}{g(b|\rho') (f(b) + f(-b))} > \frac{g(a|\rho) (f(a) + f(-a))}{g(b|\rho) (f(b) + f(-b))}. \quad (21)$$

So integrating over q for $q \geq a$ implies

$$\frac{\int_a^\infty g(q|\rho) (f(q) + f(-q)) dq}{g(b|\rho) (f(b) + f(-b))} > \frac{\int_a^\infty g(q|\rho') (f(q) + f(-q)) dq}{g(b|\rho') (f(b) + f(-b))} \quad (22)$$

or

$$\frac{g(b|\rho') (f(b) + f(-b))}{\int_a^\infty g(q|\rho') (f(q) + f(-q)) dq} < \frac{g(b|\rho) (f(b) + f(-b))}{\int_a^\infty g(q|\rho) (f(q) + f(-q)) dq}. \quad (23)$$

And integrating over q for $0 \leq q \leq a$ implies

$$\frac{\int_0^a g(q|\rho) (f(q) + f(-q)) dq}{\int_a^\infty g(q|\rho) (f(q) + f(-q)) dq} < \frac{\int_0^a g(q|\rho') (f(q) + f(-q)) dq}{\int_a^\infty g(q|\rho') (f(q) + f(-q)) dq} \quad (24)$$

or

$$\frac{Z(a|\rho)}{1 - Z(a|\rho)} > \frac{Z(a|\rho')}{1 - Z(a|\rho')} \quad (25)$$

and hence, substituting $q = a$, $Z(q|\rho) > Z(q|\rho')$ for all $q > 0$.

Regarding the sign of $v'(q)$,

$$v' = \frac{f(q) - f(-q)}{f(q) + f(-q)} + q \frac{(f(q) + f(-q)) (f'(q) - f'(-q)) - (f(q) - f(-q)) (f'(q) + f'(-q))}{(f(q) + f(-q))^2} \quad (26)$$

If $\mu > 0$ then by symmetry and quasiconcavity $f(q) > f(-q)$ for $q > 0$, so $v' > 0$ if

$$(f(q) + f(-q)) (f'(q) - f'(-q)) - (f(q) - f(-q)) (f'(q) + f'(-q)) > 0 \quad (27)$$

or

$$f(-q)f'(q) - f(q)f'(-q) > 0, \quad (28)$$

or noting that $f(-q) = f(q + 2\mu)$, if

$$\frac{d}{dq} \ln f(q) > \frac{d}{dq} \ln f(q + 2\mu) \quad (29)$$

which holds for $\mu > 0$ by logconcavity of f . If instead $\mu < 0$ then $f(q) < f(-q)$ for $q > 0$, and following similar steps, $v' < 0$.

Therefore, since $y = 0$ was WLOG, (19) implies for $\rho' > \rho$ that $E[q|y, \rho] < E[q|y, \rho']$ if $y < \mu$ and $E[q|y, \rho] > E[q|y, \rho']$ if $y > \mu$. ■

²⁶The following approach to showing FOSD of the posteriors based on MLR dominance of the news is based on that in Milgrom (1981).

Property 3 If $g(y - q|\rho)$ is a symmetric quasiconcave density where $g(y - q|\rho') \succ_{UV} g(y - q|\rho)$ for $\rho' > \rho$ and f is a symmetric logconcave density, then (i) $F(q|y, \rho) \leq F(q|y, \rho')$ for all $q < y$ if $y \geq \mu$ and (ii) $F(q|y, \rho) \geq F(q|y, \rho')$ for all $q > y$ if $y \leq \mu$.

Proof. (i) Assume WLOG that $\mu = 0$. Since $g(y - q) = g(q - y)$ and also $f(q) = f(-q)$ by $\mu = 0$,

$$\begin{aligned} F(y|y, \rho) &= \frac{\int_{-\infty}^y g(q - y|\rho)f(q)dq}{\int_{-\infty}^y g(q - y|\rho)f(q)dq + \int_y^{\infty} g(q - y|\rho)f(q)dq} \\ &= \frac{\int_{-\infty}^y g(y - q|\rho)f(q)dq}{\int_{-\infty}^y g(q - y|\rho)(f(q) + f(q - 2y))dq} \\ &= \int_{-\infty}^y v(q)dZ(q - y|\rho), \end{aligned} \quad (30)$$

where now

$$v(q) = \frac{f(q)}{f(q) + f(q - 2y)} \quad (31)$$

and now

$$Z(q|y, \rho) = \frac{\int_{-\infty}^q g(\bar{q} - y|\rho)(f(\bar{q}) + f(\bar{q} - 2y))d\bar{q}}{\int_{-\infty}^y g(\bar{q} - y|\rho)(f(\bar{q}) + f(\bar{q} - 2y))d\bar{q}}. \quad (32)$$

Integrating (30) by parts,

$$F(y|y, \rho) - F(y|y, \rho') = \int_{-\infty}^y v'(q) (Z(q|y, \rho') - Z(q|y, \rho)) dq \quad (33)$$

so if $Z(q|y, \rho') \succ_{FOSD} Z(q|y, \rho)$ then $F(y|y, \rho) - F(y|y, \rho')$ depends on the sign of $v'(q)$.

Regarding FOSD, by UV dominance, for $q < q' < y$, $g(q' - y|\rho)/g(q' - y|\rho') > g(q - y|\rho)/g(q - y|\rho')$ or $g(q' - y|\rho)/g(q - y|\rho) > g(q' - y|\rho')/g(q - y|\rho')$. So, following the same steps as (20)-(25), for any $a < y$,

$$\frac{\int_{-\infty}^a g(q - y|\rho)(f(q) + f(q - 2y))dq}{\int_a^y g(q - y|\rho)(f(q) + f(q - 2y))dq} < \frac{\int_{-\infty}^a g(q - y|\rho')(f(q) + f(q - 2y))dq}{\int_a^y g(q - y|\rho')(f(q) + f(q - 2y))dq} \quad (34)$$

or, for any $q < y$,

$$Z(q|y, \rho) < Z(q|y, \rho'). \quad (35)$$

Regarding the sign of $v'(q)$, it is the same as that of

$$\frac{d}{dq} \ln f(q) - \frac{d}{dq} \ln f(q - 2y) \quad (36)$$

which, by logconcavity of f , is positive for $y > 0$, and hence, since $\mu = 0$ is WLOG, for $y > \mu$.

Therefore, for $\rho' > \rho$, from (33) $F(y|y, \rho) < F(y|y, \rho')$ for $q < y$.

(ii) The case of $y < \mu$ follows by symmetry. ■

Property 4 If $g(y - q|\rho)$ is a symmetric quasiconcave density where $g(y - q|\rho') \succ_{UV} g(y - q|\rho)$ for $\rho' > \rho$ and f is a symmetric logconcave density, then (i) $\int_{-\infty}^a F(q|y, \rho) da \leq \int_{-\infty}^a F(q|y, \rho') da$ for all a if $y \geq \mu$ and (ii) $\int_{-\infty}^a F(q|y, \rho) da \geq \int_{-\infty}^a F(q|y, \rho') da$ for all a if $y \leq \mu$.

Proof. (i) The likelihood ratio of the posterior densities is

$$\begin{aligned} \frac{f(q|y, \rho)}{f(q|y, \rho')} &= \frac{f(q)g(q - y|\rho)}{\int f(q)g(q - y|\rho) dq} / \frac{f(q)g(q - y|\rho')}{\int f(q)g(q - y|\rho') dq} \\ &= \frac{g(q - y|\rho)}{g(q - y|\rho')} / \frac{\int f(q)g(q - y|\rho) dq}{\int f(q)g(q - y|\rho') dq}, \end{aligned} \quad (37)$$

so

$$\frac{d}{dq} \frac{f(q|y, \rho)}{f(q|y, \rho')} \propto \frac{d}{dq} \frac{g(q - y|\rho)}{g(q - y|\rho')}, \quad (38)$$

and hence $g(q - y|\rho') \succ_{UV} g(q - y|\rho)$ implies $f(q|y, \rho') \succ_{UV} f(q|y, \rho)$. This implies that $f(q|y, \rho')$ starts above $f(q|y, \rho)$, and crosses twice (Whitt, 1985), so $F(q|y, \rho')$ crosses $F(q|y, \rho)$ once from above. By Property 1 above $E[q|y, \rho] > E[q|y, \rho']$ or, integrating by parts,

$$\int_{-\infty}^{\infty} F(q|y, \rho) dq < \int_{-\infty}^{\infty} F(q|y, \rho') dq. \quad (39)$$

Since $F(q|y, \rho')$ crosses $F(q|y, \rho)$ once from above,²⁷ this implies for all a that

$$\int_{-\infty}^a F(q|y, \rho) dq < \int_{-\infty}^a F(q|y, \rho') dq. \quad (40)$$

(ii) Follows by the same logic as (i). ■

Proof of Proposition 3. The naive receiver case follows from the discussion in the text. For the sophisticated receiver case let \tilde{x}^* be the conjectured equilibrium strategy in the statement of the proposition. First suppose $\bar{x} \leq \mu$, in which case $\tilde{x}^* = (x_1 - d/2, x_2, \dots, x_{n-1}, x_n + d/2)$. If news such that $\tilde{x}_n - \tilde{x}_{n-1} \geq d/2$ and $\tilde{x}_2 - \tilde{x}_1 \geq d/2$ is observed the news is on the path and the receiver believes accordingly that $x = (\tilde{x}_1 + d/2, \tilde{x}_2, \dots, \tilde{x}_{n-1}, \tilde{x}_n - d/2)$. Hence any deviation to $\tilde{x}' \neq \tilde{x}^*$ that is on the path implies the sender is either not distorting in the most s increasing direction or is not using the entire d distortion budget to increase s , so $s(\tilde{x}^{*-1}(\tilde{x}')) < s(\tilde{x}^{*-1}(\tilde{x}^*)) = s(x)$. News such that $\tilde{x}_n - \tilde{x}_{n-1} < d/2$ or $\tilde{x}_2 - \tilde{x}_1 < d/2$ is off the equilibrium path. Let $X(\tilde{x})$ be the feasible set of x that satisfy (6) for an observed \tilde{x} that is off the equilibrium path. Let \underline{x} be the $x \in X(\tilde{x})$ that implies the lowest s , which is unique by (5). Suppose the receiver puts all weight on \underline{x} . Since \tilde{x}^* is feasible and implies at least a high s as \underline{x} by definition of \underline{x} , no type benefits by deviating. So \tilde{x}^* is a PBE strategy.

To show that \tilde{x}^* survives D1, we need to show that beliefs putting all weight on \underline{x} are not eliminated by the refinement. We adopt standard D1 notation for best responses by a receiver to

²⁷This approach to showing SOSD of the posteriors based on the means and a crossing from above is Karlin and Novikov's "cut criterion".

our case of values of s inferred by the receiver. For subsets I of $X(\tilde{x})$ let $S(I, \tilde{x})$ be the set of s generated by beliefs concentrated on I : $s(I, \tilde{x}) = \cup_{\{\mu: \mu(I)=1\}} s(\mu, \tilde{x})$. Define D_x as the set of s that cause type x to deviate, $D_x = \{s \in S(X(\tilde{x}), \tilde{x}) : U^*(x) < U(\bar{x}, s)\}$ and D_x° as the set of s that make x indifferent, $D_x^\circ = \{s \in S(X(\tilde{x}), \tilde{x}) : U^*(x) = U(\bar{x}, s)\}$. The D1 refinement requires that for any x if there exists x' such $D_x \cup D_x^\circ \subseteq D_{x'}$ that zero weight be put on type x in the receiver's beliefs. For any $x \neq \underline{x}$, $U^*(x) > U^*(\underline{x})$ so by continuity there exists a feasible $x' \neq \underline{x}$ such that $U^*(x) > U^*(x')$, so type x' is strictly willing to deviate if type x is weakly willing to deviate. Therefore $D_x \cup D_x^\circ \subseteq D_{x'}$ so type x must have zero weight. For type $x = \underline{x}$, no such x' can exist. Hence D1 not only permits beliefs putting all weight on \underline{x} , but such beliefs are the only beliefs surviving D1.

Now suppose $\bar{x} \geq \mu$, in which case any report is an equilibrium report for some realization of x , so the only question is whether any types benefit by deviating from \tilde{x}^* given equilibrium beliefs. If a report is observed such that any extreme news is pooled, $\tilde{x}_1 = \tilde{x}_2 = \dots = \tilde{x}_j$ and $\tilde{x}_k = \dots = \tilde{x}_{n-1} = \tilde{x}_n$ for some $j > 1$ and/or $k < n$, then the report is not invertible as equilibrium strategies imply a range of possible x and a corresponding distribution of receiver beliefs over s . Otherwise, even if some non-extreme news is identical, the report is inferred as fully separating and hence invertible.

First consider when x is such that \tilde{x}^* is separating. It is not feasible to deviate to an equilibrium report with pooling since by construction the sender pools the most extreme news when feasible. And, by the same argument as above, deviating to another invertible report leads the receiver to believe that s is higher than from the equilibrium strategy. So no type whose equilibrium strategy is separating benefits by deviating.

Now consider when x is such that \tilde{x}^* involves some pooling of extreme news, so equilibrium beliefs imply a distribution over possible values of s . Since U is monotonically decreasing in s , a distribution that is first-order stochastic dominated is better for the sender. For any report \tilde{x} the least favorable case in the support of receiver beliefs is that $x = (\tilde{x}_1 - d/2, \tilde{x}_2, \dots, \tilde{x}_{n-1}, \tilde{x}_n + d/2)$, which is the same as separating. Therefore deviating to any report that is for a separating type never benefits the sender. So the remaining case is deviating to another pooling report. Clearly the sender will never benefit from deviating to a higher \tilde{x}_i if $\tilde{x}_i > \bar{x}$ or a lower \tilde{x}_i if $\tilde{x}_i < \bar{x}$, or ever benefit from using less than the full distortion budget on each side, so deviations on either side of the mean by a total of $d/2$ can be considered separately. Consider deviations from the upper pooling region. Deviating to another pooling report implies reducing some report or reports that are pooled in equilibrium to below other pooling reports, and hence reducing the pooled reports less. The effect on s from reducing the pooled reports is least for the case in the support of receiver equilibrium beliefs where $x = (\tilde{x}_1 - d/(2j), \dots, \tilde{x}_j - d/(2j), \tilde{x}_{j+1}, \dots, \tilde{x}_{k-1}, \tilde{x}_k + d/(2(n+1-k)), \dots, \tilde{x}_n + d/(2(n+1-k)))$. But even in this case, by (5) the effect on s is larger than from reducing any reports below the pool,

so such a deviation hurts the sender. By the same argument, deviations from the lower pooling region must also hurt the sender, so no type benefits from deviating from \tilde{x}^* . ■

References

- [1] Athey, Susan (2002). “Monotone comparative statics under uncertainty,” *Quarterly Journal of Economics* 117(1), 187–223.
- [2] Babbage, Charles (1830). *Reflections on the Decline of Science in England, and on Some of Its Causes*. London.
- [3] Berger, Philip G. and Rebecca Hann (2003). “The impact of SFAS No. 131 on information and monitoring,” *Journal of Accounting Research* 41(2), 163–223.
- [4] Berger, Philip G. and Rebecca N. Hann (2007). “Segment profitability and the proprietary and agency costs of disclosure,” *The Accounting Review* 82(4), 869–906.
- [5] Chakraborty, Archishman and Rick Harbaugh (2010). “Persuasion by cheap talk,” *American Economic Review* 100(5), 2361–2382.
- [6] Chambers, Christopher P. and Paul J. Healy (2012). “Updating toward the signal,” *Economic Theory* 50, 765–786.
- [7] Chevalier, Judith and Glenn Ellison (1999). “Career concerns of mutual fund managers,” *Quarterly Journal of Economics* 114(2), 389–432.
- [8] Cho, In-Koo and David M. Kreps (1987). “Signaling games and stable equilibria,” *Quarterly Journal of Economics* 102(2), 179–221.
- [9] Daniel, Kent, Mark Grinblatt, Sheridan Titman, and Russ Wermers (1997), “Measuring mutual fund performance with characteristic-based benchmarks,” *Journal of Finance* 52(3), 1035–1058.
- [10] Dawid, A.P. (1973). “Posterior expectations for large observations,” *Biometrika* 60(3), 664–667.
- [11] Epstein, Marc J., and Krishna G. Palepu (1999). “What financial analysts want,” *Strategic Finance* 80(10), 48–52.
- [12] Finucan, H.M. (1973). “The relative position of the two regression curves,” *Journal of the Royal Statistical Society. Series B (Methodological)* 35(2), 316–322.

- [13] Gentzkow, Matthew, and Jesse M. Shapiro (2006). “Media bias and reputation,” *Journal of Political Economy* 114(2), 280–316.
- [14] Givoly, Dan, Carla Hayn, and Julia D’Souza (1999). “Measurement errors and information content of segment reporting,” *Review of Accounting Studies* 4, 15–43.
- [15] Hautsch, Nikolaus, Dieter Hess, and Cristoph Müller (2012). “Price adjustment to news with uncertain precision,” *Journal of International Money and Finance* 31(2), 337–355.
- [16] Holmstrom, Bengt (1982). “Managerial incentive problems – a dynamic perspective,” festschrift for Lars Wahlback, republished in *Review of Economic Studies*, 66(1),169–182.
- [17] Holmstrom, Bengt and Paul Milgrom (1991). “Multitask principal-agent analyses: incentive contracts, asset ownership, and job design,” *Journal of Law, Economics, & Organization* 7, 24–52.
- [18] Jacob, John (1996). “Taxes and transfer pricing: income shifting and the volume of intrafirm transfers,” *Journal of Accounting Research* 34(2) 301–312.
- [19] Kamenica, Emir and Matthew Gentzkow (2011). “Bayesian persuasion,” *American Economic Review* 101(6), 2590–2616.
- [20] Kirschenheiter, Michael and Nahum D. Melumad (2002). “Can ‘big bath’ and earnings smoothing co-exist as equilibrium financial reporting strategies?” *Journal of Accounting Research* 40(3), 761–796.
- [21] Lehmann, E. L. (1955). “Ordered families of distributions,” *Annals of Mathematical Statistics* 26(3), 399–419.
- [22] Meyer, Jack (1987). “Two-moment decision models and expected utility maximization,” *American Economic Review* 77(3), 421–430.
- [23] Milgrom, Paul R. (1981). “Good news and bad news: Representation theorems and applications,” *Bell Journal of Economics* 12(2), 380–391.
- [24] O’Hagan, A. (1979). “On outlier rejection phenomena in Bayes inference,” *Journal of the Royal Statistical Society. Series B (Methodological)* 41(3), 358–367.
- [25] Oreskes, Naomi and Erik M. Conway (2010). *Merchants of Doubt: How a Handful of Scientists Obscured the Truth on Issues from Tobacco Smoke to Global Warming*. Bloomsbury Press, New York.

- [26] Prendergast, Canice, and Lars Stole (1996). “Impetuous youngsters and jaded old-timers: Acquiring a reputation for learning.” *Journal of Political Economy* 104(6), 1105–1134.
- [27] Scharfstein, David S. and Jeremy C. Stein (2000). “The dark side of internal capital markets: Divisional rent-seeking and inefficient investment,” *Journal of Finance* 55(6), 2537–2564.
- [28] Stein, Jeremy C. (1989). “Efficient capital markets, inefficient firms: A model of myopic corporate behavior,” *Quarterly Journal of Economics* 104(4), 655–669.
- [29] Stone, Daniel F. (2015). “A few bad apples. Communication in the presence of strategic ideologues,” working paper.
- [30] Student (1908). “The probable error of a mean,” *Biometrika* 6(1), 1–25.
- [31] Subramanyam, K. R. (1996). “Uncertain precision and price relations to information,” *The Accounting Review* 71(2), 207–220.
- [32] Topkis, Donald (1978). “Minimizing a submodular function on a lattice,” *Operations Research* 26(2), 305–321.
- [33] Whitt, Ward (1985). “Uniform conditional variability ordering of probability distributions,” *Journal of Applied Probability* 22(3), 619–633.
- [34] You, Haifeng (2014). “Valuation-driven profit transfer among corporate segments,” *Review of Accounting Studies* 19(2), 805–838.