

Understanding the “Numbers Game”

Andrew Bird, Stephen A. Karolyi, and Thomas G. Ruchti

Tepper School of Business
Carnegie Mellon University

April 14, 2016

Abstract

We model the earnings management decision as the manager’s tradeoff between the costs and the capital market benefits of meeting earnings benchmarks. We estimate the benefits and realized distribution of earnings using a regression discontinuity design, and use these estimates as inputs to our model. Estimated model parameters yield the percentage of manipulating firms, magnitude of manipulation, and noise in manipulated earnings. These estimates also provide sufficient statistics for evaluating various proxies for “suspect” firms. Finally, we use the Sarbanes-Oxley Act as an experimental setting and show that it succeeded in reducing earnings management by 36%, through an increase in costs. This occurred despite an increase in benefits, as the market rationally became less skeptical of firms just meeting benchmarks.

We thank Brian Akins, Phil Berger, Brian Bushee, Alan Crane, Kevin Crotty, David De Angelis, Paul Fischer, Joseph Gerakos, Matthew Gustafson, Luzi Hail, Mirko Heinle, Burton Hollifield, Bob Holthausen, Peter Iliev, Andy Koch, Jason Kotter, Rick Lambert, Marios Panayides, K. Ramesh, Shiva Sivaramakrishnan, Chester Spatt, Chris Telmer, and Shawn Thomas for helpful comments and discussion and participants at the First Annual CMU-Pitt-PSU Finance Conference. We also thank the Tepper School of Business at Carnegie Mellon University for financial support.

1 Introduction

In his 1998 speech titled “The Numbers Game,” former Securities and Exchange Commission (SEC) chairman Arthur Levitt said, “I recently read of one major U.S. company that failed to meet its so-called ‘numbers’ by one penny, and lost six percent of its stock value in one day.” These “numbers” are earnings-per-share (EPS) and the game is played by equity analysts, corporate executives, and investors. The game is nuanced because the objectives, constraints, and often the actions of the players are uncertain or unobservable.

Earnings management, or the practice of intentionally distorting earnings-per-share to meet-or-beat benchmarks set by equity analysts, has become the focal point of an important and growing finance and accounting literature. Dichev et al. [2016] reports that the nearly 400 CFOs they survey believe that one-fifth of companies in a given fiscal quarter are intentionally managing earnings and that the distortion in earnings may be as large as 10% of the realized earnings. The surveyed CFOs further report that the top motivations for earnings management are “to influence stock price” and “the pressure to hit benchmarks”.

Two important observations emerge from the literature on earnings management. First, in the distribution of earnings surprise, the difference between realized earnings and analysts’ consensus EPS forecasts, we observe a larger mass of firms just above zero than just below zero (Brown and Caylor [2005], Dechow et al. [2003], Degeorge et al. [1999], Burgstahler and Dichev [1997], Hayn [1995]). Second, the stock market seems to reward firms that “just-meet-or-beat” their analysts’ consensus EPS forecasts (Payne and Thomas [2011], Keung et al. [2010], Bhojraj et al. [2009], Kasznik and McNichols [2002]). To allow for more general assumptions about the economic behavior of managers, analysts, and investors, we use a regression discontinuity approach to update, confirm, and generalize existing evidence on the benefits and presence of earnings management (McCrary [2008], Hahn et al. [2001]). With this empirical design, our evidence provides support for the presence of earnings manipulation and suggests that the stock market rewards “just-meet-or-beat” firms with approximately 1.5 percentage points higher cumulative market-adjusted returns

around their earnings announcements than “just-miss” firms.

Practitioners and academics have recognized the existence of short-term benefits of earnings management, and have even suggested explanations of this potentially suboptimal behavior (Zang [2011], Cohen et al. [2010], Caylor [2010], Roychowdhury [2006], Graham et al. [2005], Jensen [2005]). Given the large benefits that accrue to firms that meet-or-beat EPS benchmarks, on average, why is it that not all firms manipulate their earnings? There must be some unobservable frictions or costs associated with earnings management. Using our regression discontinuity estimates as inputs, we propose and implement a new structural approach based on the simulated method of moments to uncover the unobservable cost of earnings management.

In estimating the unobservable cost of earnings management, we make five contributions. First, we adopt modern econometric techniques to detect manipulation and measure the economic benefits of earnings management. Second, we formalize and implement a new methodology that quantifies the unobservable cost of earnings management, highlighting a new avenue for research on the key economic tradeoff that firms face in reporting earnings in the presence of conflicts between shareholders, management, and analysts. Third, we use our cost estimates to understand features of the marginal cost curve that give rise to the amount of earnings management that we observe and to understand how counterfactual cost curves would affect this equilibrium. Fourth, we derive implications for the commonly-used empirical proxies of suspect firms that may have engaged in earnings management and produce a new measure of the probability of earnings manipulation based on our structural estimates. Fifth, we use the Sarbanes-Oxley Act (SOX) as an experimental setting to understand how the regulation and enforcement of financial reporting changes the costs, benefits, and prevalence of earnings management. Our results suggest that SOX succeeded in curtailing the frequency and severity of manipulation by increasing the cost of earnings management, but also that investors rationally conditioned on this and began to reward meeting-or-beating earnings benchmarks even more.

Our structural model takes the perspective of managers who face a difficult decision. On one hand, they stand to gain significant short-term compensation via equity-based pay, which creates myopic incentives (Edmans et al. [2016], Matsunaga and Park [2001]). On the other hand, given the mass of firms just below analysts' consensus forecasts, some friction must exist that prevents them from accessing this boost in compensation. What constitutes this friction? Previous work has suggested that firms artificially meet-or-beat earnings benchmarks by (i) inducing biases in analyst forecasts (e.g. Cotter et al. [2006]), (ii) manipulating accounting information via accruals (e.g. Burgstahler and Dichev [1997]), revenue recognition (Caylor [2010]) or classification shifting (McVay [2006]), and (iii) altering real operating activities, including investment, R&D, discretionary expenses, or product characteristics (Ertan [2015], Cohen et al. [2010], Roychowdhury [2006]).

We endeavor to better understand this friction by estimating the marginal cost curve of earnings management that managers trade off against the observable and salient benefits. Unlike other papers that focus on proxies of analyst, accounting, or operating behavior, our estimation does not rely on observing cross-sectional variation in accounting or operating characteristics that may be the source of the friction. Because our estimation is agnostic about the source of the friction, it allows for all potential sources to be valued in the same scale, so unobservable opportunity costs that may not materialize until some future date, like the net present value of cutting R&D or investment, can be compared directly with presently realized costs.

We identify four parameters in the manipulation cost function, the marginal cost of earnings manipulation, the slope of the cost of earnings management, noise of manipulation, and heteroskedasticity. Our cost function therefore allows for variation in the cost of a single cent of manipulation, the increase in costs for additional earnings management, the degree to which earnings manipulation is noisy, and how that noise increases with increasing manipulation, respectively.

We assume that there is a latent, or unmanaged, distribution of firms in each earnings

surprise bin (for example, -1 cents to 0 cents earnings surprise, with respect to analyst consensus). Each firm in each bin draws a marginal cost curve, and chooses to manipulate short-term earnings to maximize utility, considering the benefit of just-meeting-or-beating analysts' consensus EPS forecast, as measured by our regression discontinuity estimate. This gives us manipulation propensities for each bin. Taking the empirical distribution of earnings surprise, as measured by regression discontinuity, we invert the optimal earnings management propensities for each bin to uncover the latent distribution of earnings surprise. We do this to fit cost parameters by trading off the latent distribution's smoothness and distance between to the empirical distribution.

Our main estimates show that the median marginal cost of earnings management is roughly 161 basis points per cent manipulated for manipulating firms,¹ and that the slope of the cost function is relatively convex, with an exponential parameter of 2.08. Further, we find that manipulation is uncertain, and that the variance of earnings manipulation is roughly 0.8 cent, and that the heteroskedasticity of earnings management increases nearly four-fold per additional cent of manipulation. These parameters lead to our finding that 2.62% of firms manipulate earnings over our sample. Conditional upon manipulating earnings, firms manipulate by 1.21 cents, on average, and 59.6% of these manipulating firms miss their manipulation targets. These estimates are precise primarily due to our large sample size and that we estimate a relatively small number of parameters.

Despite the apparent ability to manipulate earnings, a significant fraction, 6.1%, of firms just-miss their earnings benchmarks. Via counterfactual simulations, our structural estimates provide intuition for this surprising fact. The marginal cost and noise parameters of the cost function drive this behavior. The marginal cost parameter has straightforward consequences; as the marginal cost of manipulation increases, the optimal strategy of managers shifts toward avoiding manipulation. The noise parameter changes optimal manipulation in

¹Note that the median firm does not manipulate earnings. The marginal cost of the marginal manipulating firm is 104 basis points. This is less than the marginal benefit, because of the effect of noise—some firms that pay the cost to manipulate may not receive the benefit.

two ways. The optimal strategy for managers with low costs that expect to significantly miss earnings may be to manipulate up in the hopes of getting shocked into meeting-or-beating the benchmark. Reducing noise decreases the expected payoff of this strategy. Also, negative noise forces firms that expect to just-meet-or-beat to just-miss, so reducing noise increases the expected payoff of manipulating just above zero earnings surprise.

Analysts' forecasts may be biased due to managers issuing negative earnings guidance or strategic forecasting behavior to curry favor with managers (Kross et al. [2011], Chan et al. [2007], Burgstahler and Eames [2006], and Cotter et al. [2006]). Our estimation produces a latent distribution of unmanaged earnings surprise, which, if analysts' objectives are to minimize forecast error, should be symmetrically distributed around zero earnings surprise. Instead, 54.7% of our estimated latent distribution has a positive earnings surprise, which is similar to 57.0% of the empirical distribution. From this asymmetry, we can infer the presence of analyst bias that is not directly induced by the manager and quantify its contribution to the discontinuity in unmanaged earnings.

Workhorse empirical proxies for earnings management depend on the relative number of just-meet-or-beat firms and the number of just-miss firms. From our structural estimates, we uncover the proportion of firms in each cent bin that are manipulators, leading to a more nuanced and continuous proxy for suspect firms. This distribution of manipulation naturally produces a means to evaluate the commonly-used "suspect" bin empirical proxies for earnings management based on type 1 and type 2 errors.

Sarbanes-Oxley provides an important shift in the regulation and attention paid to accounting information. Whether or not this increased attention led to greater costs of earnings management remains unclear. We implement our estimation to the pre- (1999-2001) and post-SOX (2002-2004) periods, and compare estimates to uncover the effects of the regulation. We find that while the regression discontinuity estimates of the equity return benefits to just meeting or beating earnings increased between these periods, the marginal costs of manipulation, and particularly the incremental costs of earnings management, in-

creased even more. This had strong effects on the incentives to manage earnings—namely, the number of manipulating firms decreased by 36% following SOX.

The rest of the paper proceeds as follows: Section 2 formally describes and discusses our estimation of the benefits of earnings management, Section 3 presents our approach and estimation of the empirical distribution of earnings surprise, Section 4 describes our structural model, identification, and estimates, and Section 5 presents a series of consequences of our structural estimates, including counterfactual exercises and implications for future research on earnings management, and Section 6 concludes.

2 The Benefits of Earnings Management

An entire literature in finance and accounting has concerned itself with explaining phenomena related to short-term and long-term stock returns around earnings announcements (Bernard and Thomas [1989], Bernard and Thomas [1990], Frazzini and Lamont [2007], Barber et al. [2013], Foster et al. [1984], So and Wang [2014]). A growing component of this literature focuses on the role of short-term performance benchmarks, including analyst earnings-per-share (EPS) forecasts, lagged EPS, and zero earnings, in determining these stock returns (Athanasakou et al. [2011], Bhojraj et al. [2009], Bartov et al. [2002]). In particular, this literature has identified several means by which managers may use discretion in accounting or operations to generate a positive earnings surprise, which is defined as the difference between realized EPS and analysts' consensus EPS forecast (Edmans et al. [2016]). These papers, including survey evidence from Graham et al. [2005], suggest that managers' myopic incentives are to blame (Roychowdhury [2006], Baber et al. [1991], Jensen [2005]).

In this paper, we take market reactions and firm outcomes as given and estimate the unobservable cost function to learn about the tradeoff firms face in managing earnings and about their earnings management behavior. For this approach, we must accurately quantify the difference in short-term market reactions for firms that beat and miss their EPS benchmarks.

Here, our empirical approach diverges from the extant literature on earnings management. Empirically, papers that document the benefits of beating earnings benchmarks focus on the well-known difference in cumulative market-adjusted stock returns around earnings announcements between firms that just-miss and firms that just-meet-or-beat their analysts' consensus EPS forecast. Methodologically, these papers typically compare two subsamples of firms—those that just-miss have an earnings surprise between -1 and 0 cents or between -2 and 0 cents, and firms that just-meet-or-beat have an earnings surprise between 0 and 1 cents or between 0 and 2 cents.²

These tests of differences in means across two subsamples are useful and convincing estimators, but may not provide quantitatively accurate estimates of the benefits of beating earnings benchmarks. In particular, by focusing on two subsamples that make up approximately 14% of firm-year observations in a typical year, they eliminate almost all variation in earnings surprise, which means that they ignore potential trends in the conditional expectation function of market reaction given earnings surprise. For example, if market reactions have, on average, a positive linear relationship with earnings surprise, then the strategy of comparing just-miss firms with just-meet-or-beat firms will overstate the benefits of beating analysts' consensus EPS forecasts. Nonlinearities in the conditional expectation function of market reaction given earnings surprise yield even more nuanced empirical biases and inconsistencies. As shown in Figure 1, the true conditional expectation function of market reaction given earnings surprise is nonlinear and may even have different functional forms on either side of the zero earnings surprise.

The applied microeconometrics literature on regression discontinuity designs provides a solution for this problem. To estimate the difference in cumulative three day market-adjusted earnings announcement returns (CMAR) just-above and just-below the cutoff of zero earnings surprise, we implement two standard regression discontinuity estimators. These

²Keung et al. [2010] show that the market is becoming increasingly skeptical of firms in the [0,1) cent bin, with earnings response coefficients in that bin being much lower than those in adjacent bins. Consistent with Bhojraj et al. [2009], skepticism is warranted because earnings surprises in the [0,1) bin are minimally predictive of future earnings surprises.

estimators split firm-year observations into earnings surprise bins (as EPS is denoted in cents, we use conventional one cent bins), calculate the average CMAR for each bin, and estimate the conditional expectation function of CMAR given earnings surprise using either (i) global polynomial control functions over the full support of earnings surprise (Hahn et al. [2001]), or (ii) local polynomial/linear control functions within some bandwidth of the zero earnings surprise cutoff (Lee and Lemieux [2010]). Because the regression discontinuity estimators yield conditional expectation functions of CMAR given earnings surprise on each side of the zero earnings surprise cutoff, the difference in the two functions at the cutoff is an unbiased estimate of the “discontinuity” in CMAR, or benefits, that a firm receives by just-beating its analysts’ consensus EPS forecast relative to just-missing. Because the semiparametric control functions are flexible and can be estimated using the full distribution of firm-years, the estimators are not only statistically powerful and robust, but also provide unbiased estimates of the discontinuity in benefits at the zero earnings surprise cutoff (Hahn et al. [2001], Lee and Lemieux [2010]).

Using either global polynomial or local linear control functions, we estimate the following generalized regression discontinuity estimator:

$$CMAR_{it} = a + \mathcal{B} \cdot MBE_{it} + f^k(Surprise_{it}) + g^j(MBE \times Surprise_{it}) + e_{it} \quad (1)$$

Here, $CMAR_{it}$ is the cumulative three day market-adjusted earnings announcement returns for firm i in year t , $Surprise_{it}$ is the difference between firm i ’s realized EPS and its analysts’ consensus EPS forecast in year t , MBE_{it} is an indicator that equals one if firm i has a nonnegative earnings surprise in year t , $f^k(\cdot)$ and $g^j(\cdot)$ are order- k and order- j flexible polynomial functions of $Surprise_{it}$ on each side of the zero earnings surprise cutoff, and \mathcal{B} represents the discontinuity in capital market benefits of just-meeting-or-beating analysts’ consensus EPS forecast at the zero earnings surprise cutoff. We choose k and j using the Akaike Information Criterion and Bayesian Information Criterion (Hahn et al. [2001]), but

we report results where $k = j$ and range from 0 to 3. Like the extant literature (Bhojraj et al. [2009]), we measure earnings surprise in cents per share as analysts' EPS forecasts are widely reported and disseminated in cents and also because managers claim analysts' consensus EPS forecasts as the earnings target they manage toward most (Graham et al. [2005]). Data on analysts' EPS forecasts and firms' realized EPS come from I/B/E/S and Compustat, and daily stock returns data from which we calculate CMAR comes from CRSP.³ Our sample period spans 1994 to 2014.

We present estimates of \mathcal{B} in Table 2 for $k = j = 0, 1, 2, 3$. Panel A presents estimates using global polynomial control functions (Hahn et al. [2001]) and Panel B presents estimates using local linear control functions and bandwidth restrictions (Lee and Lemieux [2010]). In Panel A, these estimates range from 1.17% to 2.68%. Columns (2), (4), (6), and (8) include firm and year-quarter fixed effects to account for unobservable fixed cross-sectional differences and aggregate time-varying differences in earnings surprise and CMAR. These fixed effects estimates are consistent with estimates that use cross-sectional and time series variation in earnings surprise and CMAR. This suggests that the fundamental observation that investors reward firms for meeting or beating short-term performance benchmarks is not a fixed firm characteristic or a result of outlier years. Our preferred specification is in column (3) of Panel A, which suggests a discontinuity in the benefits of beating earnings benchmarks of 1.45%, because our model selection criteria (i.e., the AIC and BIC) choose $k = j = 2$.

The results in Panel B are consistent with those in Panel A, which means that our methodological choice of global polynomial control functions is not critical to identifying the discontinuity in benefits. Furthermore, it is encouraging that our estimates of \mathcal{B} are qualitatively similar across specifications within Panel A and consistent with differences-in-means estimates from the existing literature (Payne and Thomas [2011], Keung et al. [2010], Bhojraj et al. [2009], Kasznik and McNichols [2002]). Thus, our novel approach to identifying the discontinuity in the benefits of beating earnings benchmarks incorporates variation from

³We use raw forecast data unadjusted for stock splits to correct the ex-post performance bias from excessive rounding in the standard I/B/E/S database (see Diether et al. [2002]).

the broader earnings surprise distribution and relaxes restrictive assumptions about the relationship between CMAR and earnings surprise, but yields the same general conclusion that investors reward firms for meeting or beating short-term earnings benchmarks.

3 Detecting Manipulation

A large literature in finance and accounting has focused on identifying the presence and consequences of earnings management (see Healy and Wahlen [1999] for a review). Of course, because manipulation is not observed *ex ante* or, more importantly, *ex post*, identifying manipulation or the effects of manipulation relies on the use of proxies for manipulation, statistical tests of manipulation, or combinations of the two. These statistical tests of manipulation often follow those used to measure the benefits of beating earnings benchmarks: they compare the frequency of firms that just-miss to firms that just-meet-or-beat their analysts' consensus EPS forecast (Brown and Caylor [2005], Dechow et al. [2003], Degeorge et al. [1999], Burgstahler and Dichev [1997], Hayn [1995], Gilliam et al. [2015], Brown and Pinello [2007], Jacob and Jorgensen [2007]).

For these tests of differences in proportions to be valid tests of manipulation, one must assume that (i) the counterfactual frequencies of firms that just-miss and just-meet-or-beat are equal, (ii) firms have perfect control over earnings surprise, and (iii) the counterfactual earnings surprise for firms that just-miss and just-meet-or-beat are to switch to just-meet-of-beat and just-miss, respectively. These assumptions would be violated under mild conditions. For example, these assumptions would not hold if managers faced noise in their ability to manipulate earnings or if analysts exhibited strategic forecasting behavior. Our estimation relaxes these assumptions.

The challenge of identifying manipulation around financial reporting benchmarks is not unique to the earnings management literature (e.g., Berg et al. [2013], Howell [2015]). Again, we diverge from the extant literature on earnings management and turn to the applied

econometrics literature on regression discontinuity designs for methodological insight. In typical regression discontinuity designs, establishing quasi-random variation in treatment around the cutoff depends on rejecting that agents have precise control over which side of the cutoff on which they land. McCrary [2008] derives a formal test of manipulation around cutoffs to support or reject causal interpretation.

We aim to document and rigorously analyze the endogenous equilibrium behavior of managers facing both benefits and costs of earnings management. Therefore, rather than use the McCrary [2008] approach to reject precise control, we employ this regression discontinuity estimator to semiparametrically fit the distribution of earnings surprise while allowing for a discontinuity at the zero earnings surprise cutoff. This can be seen in Figure 2.

We estimate the following generalized regression discontinuity estimator:

$$Frequency_b = a + \Delta \cdot MBE_b + f^k(Surprise_b) + g^j(MBE \times Surprise_b) + e_b \quad (2)$$

Here, $Frequency_b$ is the proportion of firm-year observations in earnings surprise bin b , $Surprise_b$ is the earnings surprise in bin b , MBE_b is an indicator that equals one if bin b 's earnings surprise is positive, $f^k(\cdot)$ and $g^j(\cdot)$ are order- k and order- j flexible polynomial functions of $Surprise_b$, and Δ represents the discontinuity in frequencies at the zero earnings surprise cutoff. We choose k and j using the Akaike Information Criterion and Bayesian Information Criterion (Hahn et al. [2001]), but we report results where $k = j$ and range from 0 to 10.

We present estimates of Δ in Table 3. As in Table 2, Panel A presents global polynomial control function estimates with polynomials ranging in degree from 0 to 10, and Panel B presents local linear control function estimates with bandwidths ranging from 5 to 20 cents per share. In Panel A, these estimates range from 0.44% to 3.09%. Our preferred specification is in column (3), which estimates a discontinuity in the distribution of 2.63%, because our model selection criteria (i.e., the AIC and BIC) choose $k = j = 6$. This is

also the specification that we use in the next section to characterize the earnings surprise distribution in our cost estimation.

We are encouraged that these estimates produce an economically and statistically significant discontinuity in the earnings surprise distribution at zero cents per share, which, following McCrary [2008] is statistical evidence of earnings manipulation. In fact, even restrictive specifications that are unlikely to fit the earnings surprise distribution, like the ones with global linear control functions in column (2), produce a discontinuity. This evidence is broadly consistent with differences-in-proportions estimates from the existing literature and suggests that firms manipulate earnings to meet or beat short-term performance benchmarks, but our estimation allows for more general assumptions about the economic behavior of managers, analysts, and investors.

4 Costs of Managing Earnings

The goal of our structural model is to generate a cost function for managing earnings. In so doing, the model produces the implied manipulating firms that transform the observed distribution of earnings back to the underlying latent distribution of unmanaged earnings. We use, as inputs, the estimated earnings distribution as in Section 3, and \mathcal{B} , the benefit of meeting or beating earnings as estimated in Section 2, in basis points.

4.1 Model of Earnings Management

A manager⁴ receives an interim signal about her firm’s earnings per share, e , with respect to analyst forecasts. For example, earnings of $e = 0$ means that the firm just-meets the analysts’ consensus EPS forecast. The manager can then choose to manipulate earnings or not, trading off the capital market benefits of meeting or beating the analyst benchmark

⁴We abstract away from principal-agent problems, so that optimal behavior is determined by the benefits and costs to the firm. We therefore use the terms manager and firm interchangeably in the model discussion.

and the costs of earnings management, which could be real or accruals based.⁵ However, she knows that earnings manipulation is inherently uncertain, which we capture using noise. The result of earnings manipulation leads to the following report,

$$\text{Reported Earnings per Share} = \text{Interim Earnings per Share} + \text{Manipulation} + \text{Noise}$$

—or—

$$R = e + m + \varepsilon \tag{3}$$

where R is reported earnings per share, e is the interim signal, as above, m is the desired manipulation, and ε is noise, $\varepsilon \sim N(0, \sigma^2)$ where $\sigma^2 = (1 + \zeta(m - 1))\psi^2$, making noise a function of manipulation. Because the capital market observes reported earnings, the benefit to the manager is $\mathcal{B}(R)$, where $\mathcal{B}(\cdot)$ is the return for a stock that experiences a certain level of reported earnings, relative to the analyst benchmark.

The cost of earnings management is as follows:

$$c(m) = \beta m^\gamma \tag{4}$$

where m is the desired manipulation, $\beta \sim U[0, 2\eta]$ is the linear cost of manipulation, and γ is the exponential curvature of the cost of manipulation function.⁶ Therefore, given parameters $\theta = \{\eta, \gamma, \psi^2, \zeta\}$, the utility of the manager is,

$$u(e, m, \theta) = \int_{-\infty}^{\infty} \phi_{m,\theta}(\varepsilon) \mathcal{B}(e + m + \varepsilon) d\varepsilon - \beta m^\gamma \tag{5}$$

where $\phi_{m,\theta}(\varepsilon)$ is the normal pdf with mean 0 and variance $(1 + \zeta(m - 1))\psi^2$. Say T is

⁵These costs could be due to inducing bias in analyst forecasts, manipulating accounting information, and altering real operating activities. Manipulation decisions are potentially dynamic. Our model inherently incorporates the manager's subjective view of the future consequences of current behavior. Moreover, we focus on the analyst benchmark because rational analysts are likely to unravel the future implications of these decisions, such as accrual reversals.

⁶Because the analyst benchmark is measured in cents per share, the manipulation decision is also based on cents per share.

the largest earnings surprise number, the manager therefore chooses m such that utility is maximized,

$$m_e^* = \arg \max_{m \in [0, T-e]} u(e, m, \theta) \quad (6)$$

4.2 Simulated Method of Moments

To estimate model parameters, we apply the Simulated Method of Moments, following McFadden [1989] and Pakes and Pollard [1989].

There exists an empirical distribution of earnings surprise from $[-20, 20]$ cents per share, in one cent increments. Let $\boldsymbol{\pi}$ be the vector representing the number of firms with each per share earnings surprise. We index π_1 as the number of firms that have -20 cents in per share earnings surprise, and π_i the number of firms that have $-21 + i$ cents in per share earnings surprise, such that π_{41} is the number of firms that have 20 cents in per share earnings surprise.

We assume that the empirical distribution is the result of a latent distribution of unmanaged earnings surprise plus the effect of manipulation, or earnings management, on the part of each firm. While many firms go without managing earnings, some do so, which results in higher earnings surprise. The latent distribution of earnings surprise represents the number of firms that fall into each cent bin of earnings surprise before engaging in earnings management. Let \boldsymbol{x} be the model-implied latent distribution of earnings, indexed in a similar fashion, x_1, x_2, \dots, x_{41} , to $\boldsymbol{\pi}$. We identify costs by the proportion of firms, in each cent bin, that manipulate earnings, and tie this to the properties of the latent distribution of earnings surprise that we find.

Say we examine the -20 cents bin, which represents earnings surprise in the $[-20, -19)$ cents interval. Firms in this bin could have only one potential simulated latent earnings surprise, and that is -20 cents. A proportion of those firms chose not to manipulate their earnings numbers, and so landed in the -20 cents bin. We call the transition probability

from bin i to bin $j \geq i$, $p_{i,j}$. Therefore, we have $\pi_1 = p_{1,1}x_1$, or the number of -20 cents bin firms that did not manage earnings, and instead reported latent earnings, is equal to the transition probability from the -20 cents bin to the -20 cents bin multiplied by the latent number of -20 cents bin earners. Similarly, for the number of realized -19 cents bin firms, we have $\pi_2 = p_{1,2}x_1 + p_{2,2}x_2$, or equal to the number manipulating from latent -20 cents bin, and those that chose not to manipulate from -19 cents bin. We can see this below,

$$\begin{aligned}\pi_1 &= p_{1,1}x_1 \\ \pi_2 &= p_{1,2}x_1 + p_{2,2}x_2 \\ \pi_3 &= p_{1,3}x_1 + p_{2,3}x_2 + p_{3,3}x_3 \\ &\vdots\end{aligned}$$

In matrix form we get,

$$\begin{pmatrix} p_{1,1} & p_{2,1} & p_{3,1} & \cdots & p_{T,1} \\ p_{1,2} & p_{2,2} & p_{3,2} & \cdots & p_{T,2} \\ \vdots & & & \ddots & \\ p_{1,T} & p_{2,T} & p_{3,T} & \cdots & p_{T,T} \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_T \end{pmatrix} = \begin{pmatrix} \pi_1 \\ \pi_2 \\ \vdots \\ \pi_T \end{pmatrix}$$

or

$$\mathbf{P} \cdot \mathbf{x} = \boldsymbol{\pi} \tag{7}$$

—or—

(Manipulation, bin to bin) * (latent earnings) = empirical earnings

of course, it is important to note that $p_{i,j} = 0$, $\forall j < i$, that is, firms do not manipulate earnings down relative to analysts' consensus forecast, so our matrix of \mathbf{P} should be triangular,

$$\mathbf{P} = \begin{pmatrix} p_{1,1} & 0 & 0 & \dots & \dots & 0 \\ p_{1,2} & p_{2,2} & 0 & \dots & \dots & 0 \\ \vdots & & & & \ddots & \\ p_{1,T-1} & p_{2,T-1} & p_{3,T-1} & \dots & p_{T-1,T-1} & 0 \\ p_{1,T} & p_{2,T} & p_{3,T} & \dots & p_{T-1,T} & p_{T,T} \end{pmatrix}$$

To find the latent distribution, we invert the transition matrix and multiply by the empirical distribution, $\mathbf{x} = \mathbf{P}^{-1} \cdot \boldsymbol{\pi}$. We find the transition matrix \mathbf{P} by analyzing optimal firm behavior with respect to costs and benefits of earnings management. For each bin, b , we calculate the proportion of firms that land in each subsequent bin after managing earnings. That is, for b , we calculate $p_{b,b}, p_{b,b+1}, \dots, p_{b,T}$.

We calculate this proportion by simulating firm decisions to manage earnings. For each bin b , we simulate S firms. We take as given the parameters $\theta = \{\eta, \gamma, \psi^2, \zeta\}$. For each simulation, s , we draw a $\beta_{b,s} \sim U[0, 2\eta]$, the marginal cost of manipulation, and can calculate the utility of manipulation as in equation 5,

$$u_s(b, m, \theta) = \int_{-\infty}^{\infty} \phi_{m,\theta}(\varepsilon) \mathcal{B}(b + m + \varepsilon) d\varepsilon - \beta_{b,s} m^\gamma \quad (8)$$

Of course, bins are discrete, and therefore so is manipulation. We then assume that $\mathcal{B}(\cdot)$ is a vector, and that m is a whole number choice, making utility,

$$u_{s, Discrete}(b, m, \theta) = \sum_{-20}^{20} \phi_{m,\theta}(\varepsilon) \mathcal{B}(b + m + \varepsilon) d\varepsilon - \beta_{b,s} m^\gamma \quad (9)$$

where $\phi_{m,\theta}(\varepsilon)$ provides the discrete approximation of the continuous normal pdf with mean 0, and variance $(1 + \zeta(m - 1))\psi^2$. For various values of m , in whole cents from 0 to $T - b$,

we find the optimal earnings management,

$$m_{b,s}^*(\theta) = \arg \max_{m \in [0, T-b]} u(b, m, \theta) \quad (10)$$

From this we calculate the transition probability from each bin b to a weakly higher bin $j \geq b$,

$$p_{b,j,S}(\theta) = \begin{cases} \frac{\sum_{s=1}^S \mathbb{1}_{m_{b,s}^*(\theta)=j-b}}{S}, & \forall b, \text{ and } \forall j \geq b \\ 0, & \text{otherwise} \end{cases}$$

so that

$$\mathbf{P}_S(\theta) = \begin{pmatrix} p_{1,1,S}(\theta) & 0 & 0 & \dots & 0 \\ p_{1,2,S}(\theta) & p_{2,2,S}(\theta) & 0 & \dots & 0 \\ \vdots & & & \ddots & \\ p_{1,T,S}(\theta) & p_{2,T,S}(\theta) & \dots & \dots & p_{T,T,S}(\theta) \end{pmatrix}$$

similarly, using \mathbf{e} , the $T \times T$ matrix basis, such that $\mathbf{e}_{i,j}$ is the matrix with 0 in every position except (i, j) , which has value 1,

$$\mathbf{P}_S(\theta) = \frac{1}{S} \sum_{b=1}^T \sum_{s=1}^S \mathbf{e}_{b, m_{b,s}^*(\theta)} \quad (11)$$

Claim 1. *If \mathbf{K} is the true transition matrix from the latent distribution to the realized distribution of earnings surprise per share,*

$$\text{plim}_{S \rightarrow \infty} \mathbf{P}_S(\theta_0) = \mathbf{K} \quad (12)$$

Proof. For each bin, i , we know firms behave optimally. Therefore, for each $j \geq i$, we can calculate, for S , the number of firms that manipulate to j , and divide that by S to find

$p_{i,j,S}(\theta)$, for some θ . If θ_0 is the true set of parameters, then as $S \rightarrow \infty$, $p_{i,j,S}(\theta_0)$ should approach $K_{i,j}$ for all j , and then for all i , by the law of large numbers. ■

We assume that this is the unique set of parameters that generates this transition matrix, or

$$A 1. \theta_0 \equiv \arg \min_{\theta \in \Theta} (\mathbf{P}_S(\theta) - \mathbf{K})^2$$

We call the true distribution of earnings surprise $\boldsymbol{\ell}^* = \ell_1^*, \ell_2^*, \dots, \ell_T^*$, which represents earnings, before manipulation. We define the latent earnings surprise distribution

$$\begin{aligned} \mathbf{K} \cdot \boldsymbol{\ell}^* &= \boldsymbol{\pi} \\ \text{so,} \\ \boldsymbol{\ell}^* &= \mathbf{K}^{-1} \boldsymbol{\pi} \end{aligned} \tag{13}$$

as the smoothest curve (as measured by the differences between bins) that is close to the realized empirical earnings distribution, as follows,

$$A 2. \boldsymbol{\ell}^* \equiv \arg \min_{\boldsymbol{\ell} \in \mathbb{R}_{>0}^+} \left[\left(\begin{array}{c} \boldsymbol{\pi} \\ \boldsymbol{\ell}_{(2,\dots,T)} \end{array} \right) - \left(\begin{array}{c} \boldsymbol{\ell} \\ \boldsymbol{\ell}_{(1,\dots,T-1)} \end{array} \right) \right]' \Omega \left[\left(\begin{array}{c} \boldsymbol{\pi} \\ \boldsymbol{\ell}_{(2,\dots,T)} \end{array} \right) - \left(\begin{array}{c} \boldsymbol{\ell} \\ \boldsymbol{\ell}_{(1,\dots,T-1)} \end{array} \right) \right]$$

with the constraint that the mass under $\boldsymbol{\ell}^*$ is equal to the mass under the distribution $\boldsymbol{\pi}$.

Further, let

$$\mathbf{x}_S(\theta) \equiv \mathbf{P}_S(\theta)^{-1} \cdot \boldsymbol{\pi} \tag{14}$$

Our SMM estimator finds the true parameters when the simulated distribution of latent earnings coincides with the true distribution of earnings surprise. From equations 11 and

14, using these T moments, we propose our second stage estimator,

$$\hat{\theta}_{T, 2\text{-stage}}^S \equiv \arg \min_{\theta \in \Theta} [\boldsymbol{\ell}^* - \boldsymbol{x}_S(\theta)]' \Omega [\boldsymbol{\ell}^* - \boldsymbol{x}_S(\theta)] \quad (15)$$

Proposition 2. $\text{plim}_{S \rightarrow \infty} \hat{\theta}_{T, 2\text{-stage}}^S = \theta_0$

Proof. Simply,

$$\begin{aligned} 0 &= [\boldsymbol{\ell}^* - \boldsymbol{\ell}^*]' \Omega [\boldsymbol{\ell}^* - \boldsymbol{\ell}^*] \\ &= [\boldsymbol{\ell}^* - \mathbf{K}^{-1} \cdot \boldsymbol{\pi}]' \Omega [\boldsymbol{\ell}^* - \mathbf{K}^{-1} \cdot \boldsymbol{\pi}] \end{aligned} \quad (\text{equation 13})$$

And from Claim 1, we know that

$$\begin{aligned} &[\boldsymbol{\ell}^* - \mathbf{K}^{-1} \cdot \boldsymbol{\pi}]' \Omega [\boldsymbol{\ell}^* - \mathbf{K}^{-1} \cdot \boldsymbol{\pi}] = \\ &\text{plim}_{S \rightarrow \infty} [\boldsymbol{\ell}^* - \mathbf{P}_S(\theta_0)^{-1} \cdot \boldsymbol{\pi}]' \Omega [\boldsymbol{\ell}^* - \mathbf{P}_S(\theta_0)^{-1} \cdot \boldsymbol{\pi}] \end{aligned} \quad (16)$$

Which, with Assumption 1, implies that

$$\begin{aligned} &\text{plim}_{S \rightarrow \infty} [\boldsymbol{\ell}^* - \mathbf{P}_S(\theta)^{-1} \cdot \boldsymbol{\pi}]' \Omega [\boldsymbol{\ell}^* - \mathbf{P}_S(\theta)^{-1} \cdot \boldsymbol{\pi}] \\ &= \text{plim}_{S \rightarrow \infty} [\boldsymbol{\ell}^* - \boldsymbol{x}_S(\theta)]' \Omega [\boldsymbol{\ell}^* - \boldsymbol{x}_S(\theta)] \end{aligned} \quad (\text{equation 14})$$

$$= \text{plim}_{S \rightarrow \infty} \hat{\theta}_{T, 2\text{-stage}}^S \quad (\text{equation 15})$$

is minimized uniquely at $\theta = \theta_0$. ■

We translate this two-step estimator, which first estimates $\boldsymbol{\ell}^*$, then estimates $\hat{\theta}_{T, 2S}^S$ using $\boldsymbol{\ell}^*$, into a one-step estimator. From Assumption 2 and equation 15, we have that,

Proposition 3.

$$\hat{\theta}_T^S \equiv \arg \min_{\theta \in \Theta} \left[\left(\begin{array}{c} \boldsymbol{\pi} \\ \mathbf{x}_S(\theta)_{(2,\dots,T)} \end{array} \right) - \left(\begin{array}{c} \mathbf{x}_S(\theta) \\ \mathbf{x}_S(\theta)_{(1,\dots,T-1)} \end{array} \right) \right]' \Omega \quad (17)$$

$$\left[\left(\begin{array}{c} \boldsymbol{\pi} \\ \mathbf{x}_S(\theta)_{(2,\dots,T)} \end{array} \right) - \left(\begin{array}{c} \mathbf{x}_S(\theta) \\ \mathbf{x}_S(\theta)_{(1,\dots,T-1)} \end{array} \right) \right]$$

is equivalent to $\hat{\theta}_{T, 2S}^S$

Proof. The intuition is clear given the definitions of $\boldsymbol{\ell}^*$ and $\hat{\theta}_{T, 2S}^S$. $\boldsymbol{\ell}^*$ is defined as the distribution of earnings surprise that balances being similar to the empirical distribution, $\boldsymbol{\pi}$ with having the smallest differences between bins (i.e., smoothness). Defining $\hat{\theta}_{T, 2S}^S$, above, as the model-implied latent distribution closest to $\boldsymbol{\ell}^*$ is the same as applying the same criterion to $\mathbf{x}_S(\theta)$. So, $\hat{\theta}_T^S$ is defined as the model-implied latent distribution of earnings surprises that is closest to the empirical distribution while also having the smallest differences between bins. ■

4.3 Estimation

We identify parameters, $\theta_0 = \{\eta_0, \gamma_0, \psi_0^2, \zeta_0\}$ using the feasible one-step estimator, $\hat{\theta}_T^S$ as defined in Equation 17. Using a grid search over potential parameters, we employ partial derivatives to find a global minimum. For each θ , we run 1500 simulations for every cent bin to find the transition matrix from the latent distribution to the empirical distribution.

η is identified using the discontinuity benefit and the gross manipulation that occurs, mainly from the [-1,0) cent to [0,1) cent bins.

γ is the convexity of the cost curve, and is particularly important for the incentives of firms in the [-2,-1) cent bin and lower to manipulate. The higher this is, the smaller the proportion of firms that manipulate more than one cent to just meet or beat expected

earnings.

ψ^2 is the noisiness with which firms manipulate earnings. That is, a firm that manipulates has a chance to over- or under-shoot their target. This smoothes the process of earnings manipulation and leads to some stochasticity in manipulation. For example, this means [-3,-2) cent bin firms may manipulate up to the [-1,0) bin, for the chance to meet the analyst benchmark if they are lucky. It also leads some [-1,0) cent bin firms to manipulate beyond the benchmark to protect themselves against a negative shock.

ζ is the increasing noisiness with which firms manipulate earnings. That is, a firm in the [-5,-4) cent bin is more likely to manipulate with the goal of on average reaching the [-1,0) cent bin if that gives the firm a chance of just meeting or beating earnings. This matters for these firms that are farther out, on either side, of just meeting or beating. On the positive side, a [-2,-1) cent bin firm will be more likely to land in the [2,3) cent bin when on average aiming for just meeting or beating earnings. This parameter helps explain some of the spillover of manipulation above the [0,1) cent bin.

4.4 Results

Table 5 presents estimates of the model. The parameters we estimate are η , the marginal cost; γ , the exponential slope of the cost curve; ψ^2 , the noisiness of earnings manipulation; and ζ , the heteroskedasticity of earnings manipulation noise. We find that marginal costs, η , are 161 basis points. We find that the slope of the cost function is reasonably convex, with costs increasing with exponent 2.08. However, we find that earnings management is not certain, and that manipulation has a 0.82 cent variance (outcomes can deviate from planned), increasing nearly fourfold (heteroskedasticity of 3.71) for each additional cent of manipulation. Figure 3 shows the empirical and latent distributions, as estimated by the model.

The marginal benefit of just-meeting-or-beating analysts' consensus forecast to firms in the [-1,0) cent bin is 145 basis points, which is less than the marginal cost for the median

firm. Also, noise in manipulation means that some firms in the $[-1,0)$ bin that choose to manipulate up will miss their target, meaning the marginal benefit of manipulating is even lower. We calculate that the marginal cost of the marginal manipulator is 104 basis points, meaning that the median firm in the $[-1,0)$ bin does not manipulate earnings. Further, the average marginal cost of manipulation for manipulating firms, is 37 basis points. This is less than half of the marginal cost of the marginal manipulator, which reflects the fact that some manipulators are manipulating up by more than one bin.

Our model takes these estimates and can produce the proportion of firms that manipulate with respect to analysts' consensus forecast. Namely, we find that 2.62% of firms in our sample manipulate earnings. We also find the extent of manipulation, and show that manipulating firms do so by 1.21 cents on average. Because of stochasticity in manipulation, 59.6% of firms do not hit their intended targets, missing either above or below. Figure 4 illustrates the effect of the equilibrium manipulation strategy on the proportion of firms in each surprise bin. Panel A shows the fraction of manipulators in each earnings surprise bin, categorized by the amount of manipulation. While most firms do not manipulate at all, a nontrivial fraction of firms manipulate from one up to five cents. In Panel B, we normalize the mass of firms in each bin to one.

In the remaining rows of Table 5, we show that our results are not only robust to model specification, but also sensitive to salient features of the empirical distribution. We use bins from $[-20,-19)$ to $[20,21)$ cents in our main estimates, but most of the manipulating firms are in bins closer to the discontinuity. As such, we show that reducing the number of bins we use to fit our model does not meaningfully change our results. When we estimate our model on data from the $[-15,-14)$ to $[15,16)$ cent bins, or further reducing to the $[-10,-9)$ to $[10,11)$ bins, or the $[-5,-4)$ to $[5,6)$ bins, our results are similar. This is particularly true for important parameters η , which varies from 154 to 167 basis points (vs. a baseline of 161) and ψ^2 , which varies from 0.76 to 0.79 (vs. a baseline of 0.82). In addition, $\%manip$ is relatively unchanged, varying from 2.56% to 2.84% (vs. a baseline of 2.62%), as is the $\%noise$, varying

from 55.7% to 60% (vs. a baseline of 59.6%).

Our estimation equally weights similarity to the empirical distribution and smoothness as qualities of the latent unmanaged earnings surprise distribution. We therefore reestimate our model overweighting empirical sameness by 10% and then run it again instead overweighting smoothness by 10%.⁷ We find that our results do not significantly change, and cost parameters and manipulation move in the direction one would expect. Namely, costs go up (170 vs. 161 basis points) and manipulation goes down (2.45% vs. 2.62% of firms manipulate) when empirical sameness is emphasized, whereas costs go down (146 basis points) and manipulation goes up (3.11% of firms manipulate) when smoothness is emphasized. Because ζ is possibly the least important parameter in our model, we estimate the model without ζ (setting $\zeta = 0$). Our estimates change the most in this robustness check. While ζ does not have large effects on equilibrium behavior, we can see from these results that it is still an important parameter in fitting this model to data.

We investigate whether our approach is sensitive to features of the empirical distribution by mechanically changing the empirical distribution of earnings. First, we reduce the peak found in the [0,1) cent bin, halving its size relative to the [1,2) cent bin. Using this *Reduce peak* empirical distribution of earnings, the implied marginal costs increase by 27%, while the slope increases by only 5%. If we remove the peak entirely, making the [0,1) cent bin equal to the [1,2) cent bin, then the slope remains unchanged, but marginal costs increase by 70% relative to the baseline. Further, reducing the peak cuts down on the implied number of manipulators, down to 1.91% and 1.36% of firms, respectively. These alternative estimates demonstrate that our results are sensitive to earnings surprise discontinuities, and that removing the discontinuity yields much lower levels of manipulation, via higher cost parameters.

In untabulated results, we employ the full distribution of capital markets benefits for different levels of earnings relative to the benchmark. As shown in Figure 1, the main

⁷This can be done through weighting of the matrix Ω in our estimator, effectively increasing the importance of one or the other of the set of empirical sameness or smoothness moments.

difference is the existence of positive marginal benefits of moving rightward from just meeting the benchmark. When we estimate the model, we obtain similar parameter estimates for the cost function, though because of the increased benefits, this cost function yields somewhat higher levels of manipulation.

4.5 Identification

In estimating models of this type, it is important that we demonstrate the role the parameters in our model play in generating the *empirical* earnings surprise distribution from the latent earnings surprise distribution. Table 6 shows how the manipulation strategy varies while altering each parameter, in turn, by reducing it to 10% of our estimated level. For example, we estimate a marginal cost of 161 basis points, but here ask what equilibrium earnings management behavior would we observe if firms instead faced marginal costs of 16 basis points. By reducing the marginal cost parameter, we see that the number of manipulating firms increases to 12.6%, roughly five times our estimated 2.62%. We also find that mean manipulation increases to 1.92 cents, which is significantly larger than our estimate of 1.21 cents. Similarly, decreasing the slope of the cost curve increases the number of manipulating firms to 9.45% and mean manipulation to 10.16 cents, which is more than eight times larger than our estimates. There is also an increase to 78.4% (vs. 59.6% estimated) of firms missing their target, either above or below, due to heteroskedasticity. If we decrease the noise parameter, we see a large reduction in missing manipulation targets, and only a modest increase in the number of manipulators (3.22%) and mean manipulation (1.23 cents). Altering the heteroskedasticity parameter has negligible effects, pointing to the fact that heteroskedasticity, while likely present, is not a crucial component in modeling noise in earnings management.

Figure 5 provides further intuition for how the parameters are estimated by showing how the empirical distribution would differ according to the piecemeal parameter reductions as mentioned above, and as shown in Table 7. The curves shown in red represent the empirical

distribution and the curves shown in blue are the counterfactual empirical distributions. Specifically, we see that the marginal cost parameter, η , is the most important for identifying manipulation close to the discontinuity, as reducing the marginal costs induces larger jumps out of the $[-2,-1)$ and $[-1,0)$ cent bins. This effect is straightforward to interpret.

However, the slope parameter, γ , is associated with the degree to which firms from farther away from the discontinuity manipulate. This can be seen by the greater participation in manipulation by firms from -20 cents to -5 cents. If we reduce the noise parameter, ψ^2 , to 10% of its estimated value, a divot appears in the $[-1,0)$ cent bin, implying that noise corrects for the discontinuous nature of manipulation, smoothing both the benefits of manipulation as well as the outcomes of manipulation itself. This parameter reflects the uncertainty inherent in manipulation. Consistent with the results presented in Table 7, we see that the heteroskedasticity parameter, ζ , has insignificant effects on manipulation. While it allows for greater manipulation from farther away from the discontinuity and manipulation beyond the discontinuity, this type of distributional smoothing is only marginally important, and we, therefore, conclude that heteroskedasticity is not a crucial component of our model.

5 Consequences

An understanding of the costs of earnings management has many implications for how we interpret firm behavior and for the validity of measures employed in the literature to identify earnings management. In this section, we discuss a number of applications of the results described above.

5.1 Why do so few firms manage earnings?

Given the significant benefit to meeting earnings benchmarks and the results of surveying managers themselves (Dichev et al. [2016]), our estimated fraction of firms managing earnings is perhaps surprisingly low. Of particular interest are firms in the $[-1,0)$ earnings surprise

bin, since these are the ones with the most potential to take advantage of the discontinuity in capital market responses to earnings announcements. The counterfactual simulations discussed above can help in understanding why so many firms remain below, and yet very close to, the benchmark.

Two of the parameters in the cost function appear to drive this behavior—the marginal cost parameter, η , and the noise parameter, ψ^2 . In the latent distribution, 6.1% of firms are in the $[-1,0)$ cent bin. Cutting η to 10% of its estimated level, as in Table 6, reduces this fraction by 3.1%. This reduction is a straightforward consequence of decreasing the marginal cost of manipulation. If ψ^2 is simultaneously cut to 10% of its estimated level, only 0.9% of firms remain in the negative one bin. The effect of reducing noise has two key effects. The optimal strategy for managers with relatively low costs in bins farther away from the benchmark may be to manipulate up to negative one in the hopes of getting shocked into meeting the benchmark⁸—reducing noise decreases the expected payoff of this strategy. The second effect is mechanical: negative noise drops firms from meeting the benchmark (the most likely target) into the negative one bin.

The results of the counterfactuals show that both the marginal cost effect and the noise effect are economically significant explanations for the relative lack of manipulation out of the negative one bin. Determining the exact contribution of the two forces is complicated by the fact that there is an interaction between the two. As marginal costs decrease, more firms target just-meeting the benchmark, but the mechanical effect of noise means that a higher noise parameter shocks more of these firms back down into the negative one cent bin. However, this interaction accounts for only nine percent of the overall effect of the counterfactual.

⁸Although managers are risk neutral (so that noise is not necessarily harmful), there is an asymmetry to the effect of noise on the manager’s objective function. From the $[0,1)$ cent bin, which is the most likely target for managers, a negative shock hurts the manager much more than she is helped by a positive shock, because of the nature of the discontinuity in benefits. From the $[-1,0)$ bin, the asymmetry is reversed, with the expected benefit of noise becoming positive.

5.2 Identifying “suspicious” firms

Our model estimates reveal the incidence (and amount) of manipulation in each earnings surprise bin. This is of interest given the prevalent use of proxy measures of manipulation based on whether a firm ends up in the $[0,1)$ cent bin or not (e.g., Bhojraj et al. [2009], Cohen et al. [2008], Roychowdhury [2006]). Table 7 quantifies how well a variety of definitions of “suspicious” firms fare in accurately identifying manipulation. The first row shows the percentage of firms in each surprise bin that manipulated to get there and the second row captures the percentage of all manipulators from the sample that end up in that surprise bin. The third and fourth rows replicate the first and second, respectively, while weighting the manipulating firm observations according to the mean level of manipulation in each surprise bin (e.g. the firm counts twice if it manipulated up two cents).

The first column of Table 7 describes the firms captured by the typically used $[0,1)$ “suspicious” bin. Of the firms in this bin, 11% manipulated to get there, which is perhaps surprisingly low, but in line with our overall estimates of the extent of manipulation in aggregate. This subset covers 53.1% of all manipulators in the sample. On these two metrics, this bin performs the best at identifying manipulation, restricting to single bin measures. The next best bin to add to the measure would be $[1,2)$, which reduces accuracy to 9.9% but then accounts for 80.9% of the total manipulators. Given the salience of the zero cutoff in capital market response to earnings surprise, it is interesting that the bin with the third most manipulators is the $[-1,0)$ bin; as described above, this is because the possibility of a positive shock makes manipulating into this bin the optimal strategy for some firms with an intermediate level of costs. If the priority were to account for nearly all of the manipulators, the $[-1,3)$ range may be desirable in that 97.4% of manipulators fall into this range and the accuracy falls only to 7.3%. Relative to the $[0,1)$ bin, this reflects an 83% increase in the fraction of all manipulators included, at the cost of reducing accuracy by 34%.

Turning to the third and fourth rows of Table 7, we find a similar pattern of results, though with a clear tendency toward improvement from including more bins. This results

from bins farther away from the benchmark having higher expected manipulation, conditional on manipulating at all, relative to bins closer to the benchmark. For example, in this case, the $[0,2)$ range performs better than $[0,1)$ on both margins, increasing the weighted fraction of manipulators and significantly increasing the fraction of total manipulation covered. Extending to $[-1,3)$ yields an even larger increase in coverage and reduces accuracy by only 13% relative to the $[0,2)$ range.

Overall, Table 7 shows how different definitions of “suspicious” bins trade off Type 1 and Type 2 error and can help guide this measurement choice in a variety of contexts. Further, it may be preferred for some applications to use a continuous measure of the probability of manipulation by bin, which naturally places more weight on bins close to the benchmark and less on those farther away, rather than making a discrete choice of which bins to include in the measure and which to exclude.

5.3 Do analysts bias their forecasts?

So far, we have studied the management of earnings towards a particular benchmark, which we assume to be known at the time the managers make their manipulation decisions. However, it is well known in the literature (Kross et al. [2011], Chan et al. [2007], Burgstahler and Eames [2006], and Cotter et al. [2006]) that the benchmark itself is subject to manipulation. This can be accomplished by managers through negative earnings guidance, which indirectly affects the analyst’s choice of benchmark, or may directly reflect the strategy of analysts, who would like to curry favor with managers. Indeed, Burgstahler and Eames [2006] show evidence that firms avoid negative earnings surprise through downward management of analysts’ forecasts, incremental to their management of the underlying earnings number.

One of the outputs of our estimation process, as described above, is the unmanaged distribution of earnings. If the benchmark is chosen by analysts with the objective of minimizing forecast error, one would expect this distribution to be symmetric, with an equal probability of positive and negative errors, especially in the neighborhood of the benchmark.

Interestingly, this is not the case—the asymmetry is reduced in the latent distribution relative to the empirical distribution but is not eliminated.⁹ Specifically, we find that 57.0% of observations are on the positive side of our window in the empirical distribution; this falls to 54.7% in the latent distribution. We can infer from this asymmetry the presence of analyst bias, and can quantify the contribution of this bias to the discontinuity in earnings relative to that of short-term earnings management. The remaining bias is striking because analysts themselves should be able to understand the potential for manipulation and so might be expected to remove the effects of bias from their forecasts, leading to a distribution symmetric around this forecast. We therefore interpret the remaining asymmetry as the analysts’ own preference to bias forecasts, in the absence of active managerial intervention, consistent with findings in the literature (Brown et al. [2015], Richardson et al. [2004]).

5.4 Sarbanes-Oxley and the cost of earnings management

After a string of accounting scandals in the late 1990s and early 2000s, including Enron and Worldcom, the Sarbanes-Oxley Act (SOX) was enacted in 2002. In line with the SEC’s stated objective of protecting investors, SOX initiated broad changes in corporate disclosures, with the goal of improving the reliability of financial information disclosed in company reports. SOX consisted of three main regulatory changes. First, it established the Public Company Accounting Oversight Board (PCAOB) to set rules on auditing, quality control, and independence. Second, it provided for greater auditor independence by requiring that independent directors sit on audit committees and that firms rotate audit partners at least every five years. Third, it produced a code of ethics for senior financial officers and required chief executive and financial officers to certify and take personal responsibility for public financial statements and to forfeit bonus compensation upon accounting restatements. Along with the other eleven sections of the bill, these provisions were intended to provide enforceable

⁹Brown [2001] shows that there was a shift in the distribution of earnings surprise rightward from 1984-1999. Our data, which mostly come from years after this period, confirm the observation that the median earnings surprise is a small positive number.

rules and penalty-based incentives for managers to engage in reliable financial reporting behavior.

Our methodology can be used to uncover the effects of such regulation on earnings management by exploiting time series variation in the empirical earnings surprise distribution and the benefits of meeting-or-beating earnings benchmarks. SOX provides an ideal setting for such analysis because the effects of SOX have been of great interest to practitioners and in the academic literature. Cohen et al. [2008] and Bartov and Cohen [2009] find that earnings management declined after the introduction of SOX, and attribute this decline particularly to reduced accruals-based earnings management (as well as a decline in downwards expectation management using negative managerial guidance). Cohen et al. [2008] note that the three years prior to SOX were characterized by an abnormal increase in earnings management relative to earlier in the 1990s.

To investigate the effects of SOX using our methodology, we look at subsamples of the three years before SOX (1999-2001) and the three years after SOX (2002-2004).¹⁰ Table 8 presents estimates from these two subsamples. Consistent with the work discussed above, our diagnostic test statistic of the extent of manipulation, the discontinuity in the earnings distribution (Δ), falls post-SOX. The benefit of meeting the benchmark (\mathcal{B}) increases after SOX by 80 basis points. Together, these two findings suggest that SOX had its intended effect of improving the quality of financial reporting and that capital markets responded with larger rewards for meeting the benchmark. This is likely due to an increased willingness by markets to attribute the meeting of earnings benchmarks to true performance rather than manipulation.

Next, we use our structural model to estimate the consequences of these changes for the costs of earnings management as well as the extent of manipulation. As expected, we find significant increases in costs. Table 8 shows that both the marginal cost parameter and the slope parameter increase significantly. This suggests that both the cost of the first cent of

¹⁰Our results are not sensitive to these specific definitions of the pre- and post-SOX periods. For example, dropping either 2001 or 2002 yields similar estimates.

managed earnings and the cost of incremental manipulation both increased, consistent with the goals of the reform.¹¹ Earnings management also became much less noisy after SOX—we conjecture that this is due to a switch away from riskier, less certain strategies, towards more conservative use of discretion. In aggregate, manipulation fell from 4.05% of firms to 2.61%, a 36% reduction, evidence that SOX had its desired effect on financial reporting and earnings management.

6 Conclusion

In this paper, we employ a novel empirical approach to estimate the capital market benefits to managing earnings as well as the underlying distribution of unmanaged earnings. We use these as inputs to a structural model to estimate the costs of earnings management. This cost function provides numerous insights about firm behavior and can be used, for example, to assess the validity of identifying manipulation using a “suspicious” earnings surprise bin and also to estimate the effect of SOX on the financial reporting process. In future research, we plan to study more broadly how costs change over time and whether they vary according to the benchmark and to the characteristics of the firm and the incentives of its managers.

¹¹There are a number of mechanisms through which the costs could have increased. For example, auditors likely increased their scrutiny of reported earnings both because of increased regulation on their behavior, and as a rational response to the increased risk to their survival made evident by the collapse of Arthur Andersen.

References

- Vasiliki Athanasakou, Norman C. Strong, and Martin Walker. The market reward for achieving analyst earnings expectations: does managing expectations or earnings matter? *Journal of Business Finance & Accounting*, 38(1-2):58–94, 2011.
- William R. Baber, Patricia M. Fairfield, and James A. Haggard. The effect of concern about reported income on discretionary spending decisions: The case of research and development. *The Accounting Review*, 66(4):818–829, 1991.
- Brad M. Barber, Emmanuel T. De George, Reuven Lehavy, and Brett Trueman. The earnings announcement premium around the globe. *Journal of Financial Economics*, 108(1):118–138, 2013.
- Eli Bartov and Daniel A. Cohen. The “Numbers Game” in the pre- and post-Sarbanes-Oxley eras. *Journal of Accounting, Auditing & Finance*, 24(4):505–534, 2009.
- Eli Bartov, Dan Givoly, and Carla Hayn. The rewards to meeting or beating earnings expectations. *Journal of Accounting and Economics*, 33(2):173–204, 2002.
- Tobias Berg, Manju Puri, and Jorg Rocholl. Loan officer incentives and the limits of hard information. *Working paper*, 2013.
- Victor L. Bernard and Jacob K. Thomas. Post-earnings-announcement drift: delayed price response or risk premium? *Journal of Accounting Research*, 27:1–36, 1989.
- Victor L. Bernard and Jacob K. Thomas. Evidence that stock prices do not fully reflect the implications of current earnings for future earnings. *Journal of Accounting and Economics*, 13(4):305–340, 1990.
- Sanjeev Bhojraj, Paul Hribar, Marc Picconi, and John McInnis. Making sense of cents: An examination of firms that marginally miss or beat analyst forecasts. *The Journal of Finance*, 64(5):2361–2388, 2009.
- Lawrence D. Brown. A temporal analysis of earnings surprises: Profits versus losses. *Journal of Accounting Research*, 39(2):221–241, 2001.
- Lawrence D. Brown and Marcus L. Caylor. A temporal analysis of quarterly earnings thresholds: Propensities and valuation consequences. *The Accounting Review*, 80(2):423–440, 2005.
- Lawrence D. Brown and Arianna Spina Pinello. To what extent does the financial reporting process curb earnings surprise games? *Journal of Accounting Research*, 45(5):947–981, 2007.
- Lawrence D Brown, Andrew C Call, Michael B Clement, and Nathan Y Sharp. Inside the black box of sell-side financial analysts. *Journal of Accounting Research*, 53(1):1–47, 2015.
- David Burgstahler and Ilia Dichev. Earnings management to avoid earnings decreases and losses. *Journal of Accounting and Economics*, 24(1):99–126, 1997.

- David Burgstahler and Michael Eames. Management of earnings and analysts' forecasts to achieve zero and small positive earnings surprises. *Journal of Business Finance & Accounting*, 33(5-6):633–652, 2006.
- Marcus L. Caylor. Strategic revenue recognition to achieve earnings benchmarks. *Journal of Accounting and Public Policy*, 29(1):82–95, 2010.
- Louis K.C. Chan, Jason Karceski, and Josef Lakonishok. Analysts' conflicts of interest and biases in earnings forecasts. *Journal of Financial and Quantitative Analysis*, 42(04):893–913, 2007.
- Daniel Cohen, Raj Mashruwala, and Tzachi Zach. The use of advertising activities to meet earnings benchmarks: Evidence from monthly data. *Review of Accounting Studies*, 15(4):808–832, 2010.
- Daniel A. Cohen, Aiysha Dey, and Thomas Z. Lys. Real and accrual-based earnings management in the pre-and post-Sarbanes-Oxley periods. *The Accounting Review*, 83(3):757–787, 2008.
- Julie Cotter, Irem Tuna, and Peter D. Wysocki. Expectations Management and Beatable Targets: How Do Analysts React to Explicit Earnings Guidance? *Contemporary Accounting Research*, 23(3):593–624, 2006.
- Patricia M. Dechow, Scott A. Richardson, and Irem Tuna. Why are earnings kinky? An examination of the earnings management explanation. *Review of Accounting Studies*, 8(2-3):355–384, 2003.
- Francois Degeorge, Jayendu Patel, and Richard Zeckhauser. Earnings management to exceed thresholds. *The Journal of Business*, 72(1):1–33, 1999.
- Ilia Dichev, John Graham, Campbell R. Harvey, and Shiva Rajgopal. The misrepresentation of earnings. *Financial Analysts Journal*, 72(1):22–35, 2016.
- Karl B. Diether, Christopher J. Malloy, and Anna Scherbina. Differences of opinion and the cross section of stock returns. *The Journal of Finance*, 57(5):2113–2141, 2002.
- Alex Edmans, Vivian W. Fang, and Katharina A. Lewellen. Equity vesting and managerial myopia. *Working paper*, 2016.
- Aytekin Ertan. Real earnings management in the financial industry. *Working paper*, 2015.
- George Foster, Chris Olsen, and Terry Shevlin. Earnings releases, anomalies, and the behavior of security returns. *The Accounting Review*, 59(4):574–603, 1984.
- Andrea Frazzini and Owen Lamont. The earnings announcement premium and trading volume. *Working paper*, 2007.
- Thomas A. Gilliam, Frank Hefflin, and Jeffrey S. Paterson. Evidence that the zero-earnings discontinuity has disappeared. *Journal of Accounting and Economics*, 60(1):117–132, 2015.

- John R. Graham, Campbell R. Harvey, and Shiva Rajgopal. The economic implications of corporate financial reporting. *Journal of Accounting and Economics*, 40(1):3–73, 2005.
- Jinyong Hahn, Petra Todd, and Wilbert Van der Klaauw. Identification and estimation of treatment effects with a regression-discontinuity design. *Econometrica*, 69(1):201–209, 2001.
- Carla Hayn. The information content of losses. *Journal of Accounting and Economics*, 20(2):125–153, 1995.
- Paul M. Healy and James M. Wahlen. A review of the earnings management literature and its implications for standard setting. *Accounting Horizons*, 13(4):365–383, 1999.
- Sabrina T. Howell. Financing constraints as barriers to innovation: Evidence from R&D grants to energy startups. *Working paper*, 2015.
- John Jacob and Bjorn N. Jorgensen. Earnings management and accounting income aggregation. *Journal of Accounting and Economics*, 43(2):369–390, 2007.
- Michael C. Jensen. Agency costs of overvalued equity. *Financial Management*, 34(1):5–19, 2005.
- Ron Kasznik and Maureen F. McNichols. Does meeting earnings expectations matter? Evidence from analyst forecast revisions and share prices. *Journal of Accounting Research*, 40(3):727–759, 2002.
- Edmund Keung, Zhi-Xing Lin, and Michael Shih. Does the stock market see a zero or small positive earnings surprise as a red flag? *Journal of Accounting Research*, 48(1):91–121, 2010.
- William J. Kross, Byung T. Ro, and Inho Suk. Consistency in meeting or beating earnings expectations and management earnings forecasts. *Journal of Accounting and Economics*, 51(1):37–57, 2011.
- David S. Lee and Thomas Lemieux. Regression Discontinuity Designs in Economics. *Journal of Economic Literature*, 48:281–355, 2010.
- Steven R. Matsunaga and Chul W. Park. The effect of missing a quarterly earnings benchmark on the CEO’s annual bonus. *The Accounting Review*, 76(3):313–332, 2001.
- Justin McCrary. Manipulation of the running variable in the regression discontinuity design: A density test. *Journal of Econometrics*, 142(2):698–714, 2008.
- Daniel McFadden. A method of simulated moments for estimation of discrete response models without numerical integration. *Econometrica*, 57(5):995–1026, 1989.
- Sarah E. McVay. Earnings management using classification shifting: An examination of core earnings and special items. *The Accounting Review*, 81(3):501–531, 2006.

- Ariel Pakes and David Pollard. Simulation and the asymptotics of optimization estimators. *Econometrica*, 57(5):1027–1057, 1989.
- Jeff L. Payne and Wayne B. Thomas. The torpedo effect myth or reality? *Journal of Accounting, Auditing & Finance*, 26(2):255–278, 2011.
- Scott Richardson, Siew Hong Teoh, and Peter D. Wysocki. The Walk-down to Beatable Analyst Forecasts: The Role of Equity Issuance and Insider Trading Incentives. *Contemporary Accounting Research*, 21(4):885–924, 2004.
- Sugata Roychowdhury. Earnings management through real activities manipulation. *Journal of Accounting and Economics*, 42(3):335–370, 2006.
- Eric C. So and Sean Wang. News-driven return reversals: Liquidity provision ahead of earnings announcements. *Journal of Financial Economics*, 114(1):20–35, 2014.
- Amy Y. Zang. Evidence on the trade-off between real activities manipulation and accrual-based earnings management. *The Accounting Review*, 87(2):675–703, 2011.

Figure 1: Benefits of Earnings Surprise

This figure presents a regression discontinuity plot of cumulative earnings announcement returns around EPS surprise. The running variable is *EPS Surprise*, which is defined as realized earnings-per-share (EPS) minus analysts' consensus EPS forecast, and the dependent variable is $CMAR^{3day}$, which is defined as three day cumulative market-adjusted returns. The scatterplot presents dots that correspond to the average $CMAR^{3day}$ within each *EPS Surprise* bin and two fitted polynomials in grey, which represent the best-fitting (both second degree) polynomials on either side of the EPS surprise cutoff of zero cents. The intersection of the fitted polynomials at $EPS Surprise = 0$ represents our preferred discontinuity estimate of the benefits of meeting-or-beating analysts' consensus EPS forecast.

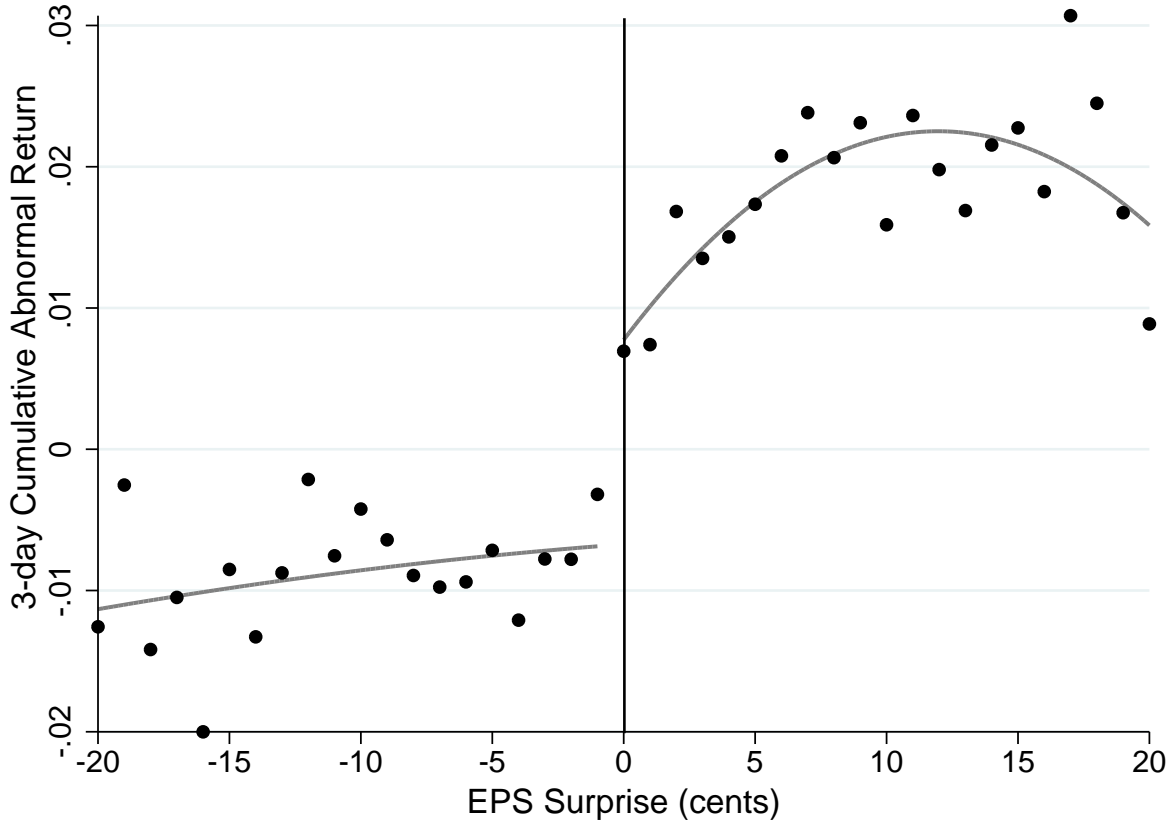


Figure 2: Fitted EPS Surprise Frequencies

This figure presents a regression discontinuity plot that corresponds to a McCrary [2008] test of manipulation of earnings-per-share (EPS) around analysts' consensus earnings forecasts. The running variable is *EPS Surprise*, which is defined as actual earnings minus analysts' consensus earnings forecast, and the dependent variable is the frequency of firm-quarter observations in each earnings surprise bin. The scatterplot presents dots that correspond to the frequency of observations within each *EPS Surprise* bin and two fitted polynomials in grey, which represent the best-fitting (both sixth degree) polynomials on either side of the EPS surprise cutoff of zero cents. The intersection of the fitted polynomials at *EPS Surprise* = 0 represents our preferred diagnostic test statistic of the prevalence of earnings management around the zero earnings surprise cutoff.

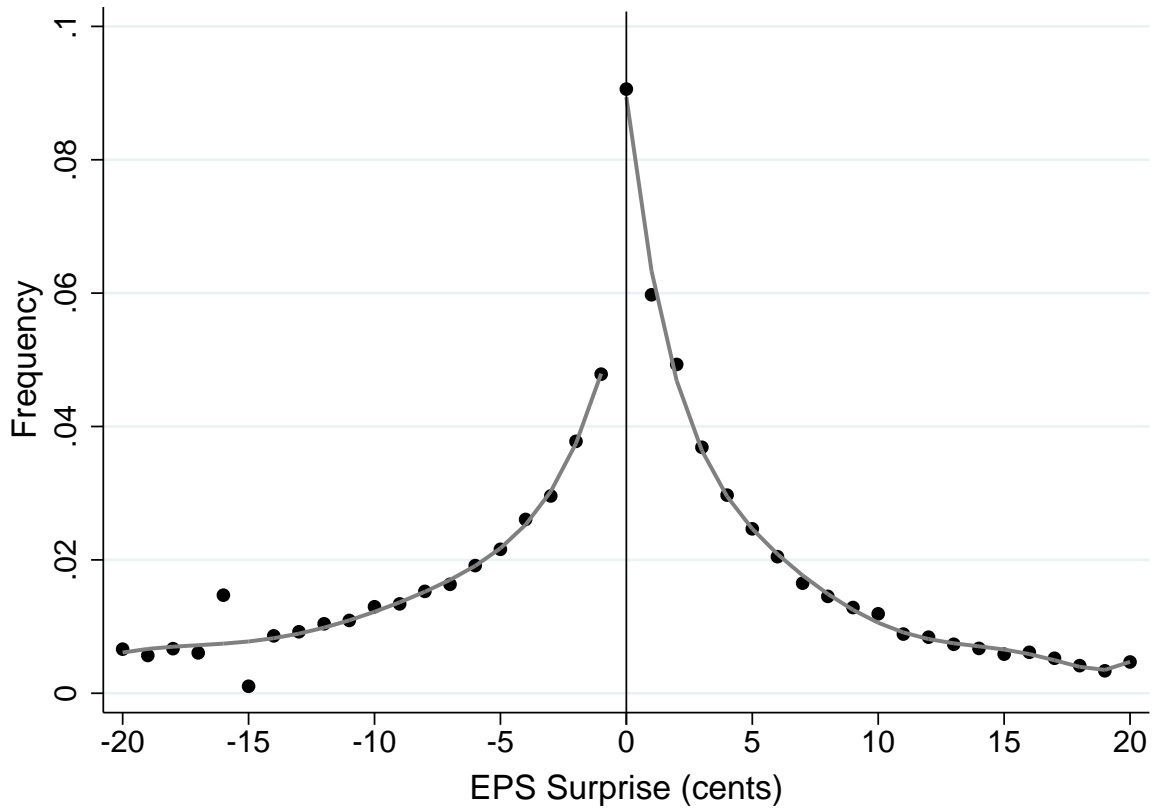
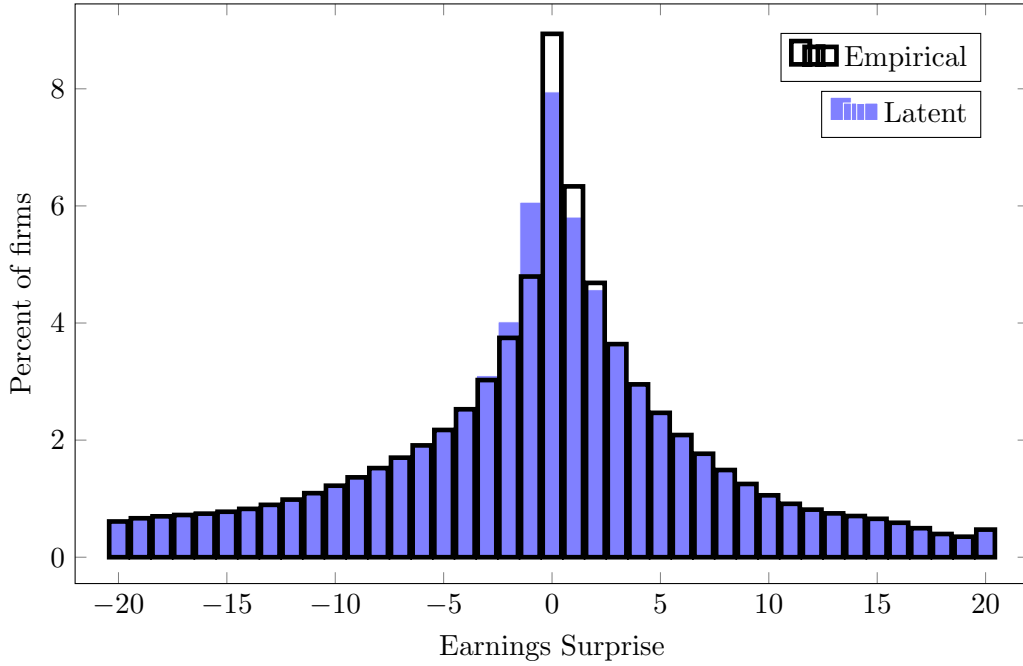


Figure 3: Latent and Empirical Distributions

This figure presents the empirical and latent distributions of earnings surprise, which is defined as the realized earnings-per-share (EPS) minus analysts' consensus EPS forecast. The empirical distribution is estimated using a regression discontinuity approach (McCrary [2008]) which allows for a discontinuity in the frequencies of observations around the zero earnings surprise cutoff. The latent distribution is an output of our structural estimation that accounts for managers' optimal manipulation decisions and trades off distance to the empirical distribution with smoothness. Panel A presents these distributions from -20 to 20 cents of earnings surprise, whereas Panel B presents them from -5 to 5 cents of earnings surprise.

(a) -20 to 20 cents



(b) -5 to 5 cent bins


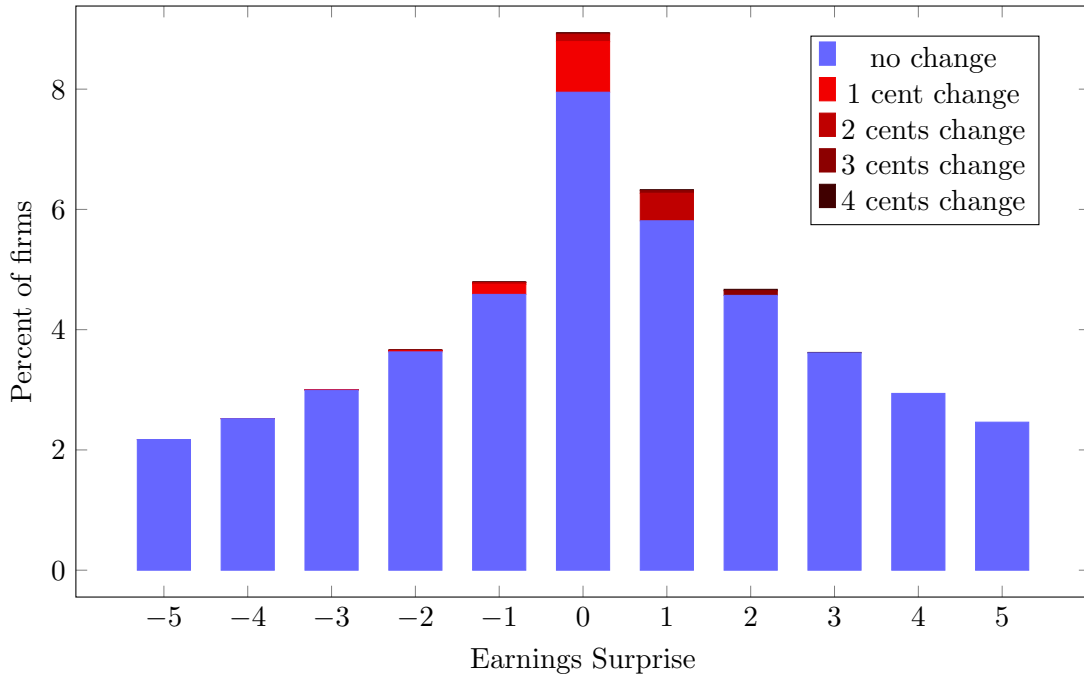


Figure 4: Cent Manipulation

This figure presents the empirical distribution of earnings surprise, which is defined as the realized earnings-per-share (EPS) minus analysts' consensus EPS forecast, between -5 and 5 cents of earnings surprise. The figures highlight the fraction of manipulating firms in each earnings surprise bin and the mean amount of manipulation that the manipulating firms choose in shades of red. Panel A presents manipulation across the empirical distribution manipulation and Panel B normalizes the empirical distribution to highlight the prevalence and amount of manipulation in each surprise bin.

(a) Manipulation by Bin



(b) Normalized Manipulation by Bin

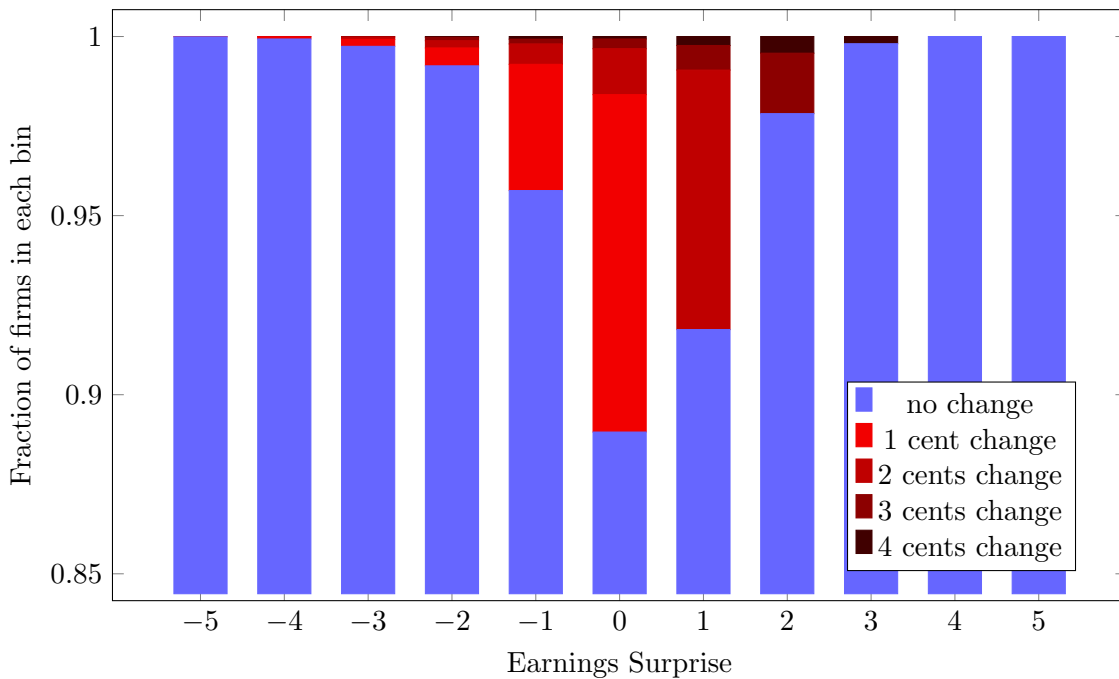

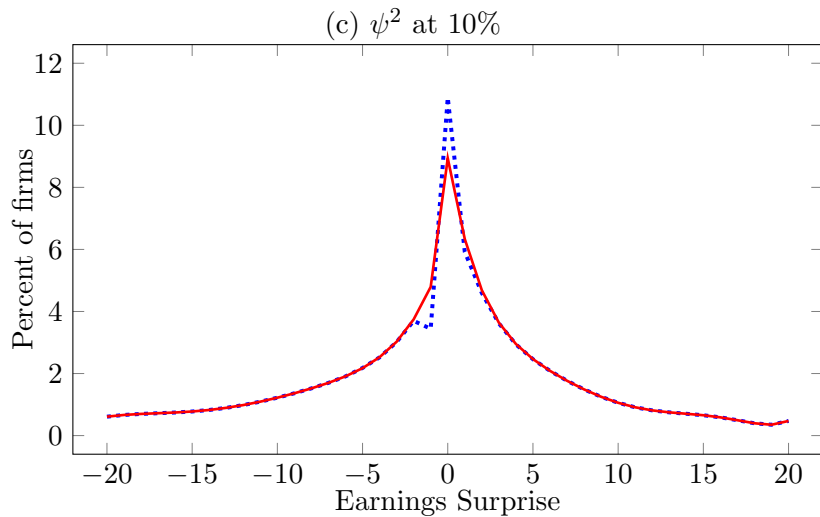




Figure 5: Counterfactual Parameters

This figure contrasts the empirical distribution of earnings surprise, which is defined as the realized earnings-per-share (EPS) minus analysts' consensus EPS forecast, between -20 and 20 cents of earnings surprise with counterfactual distributions of earnings surprise that arise by cutting each estimated structural parameter to 10% of its estimated value and then recalculating optimal manipulation. The dashed blue lines represent counterfactual empirical distributions and the solid red lines represent the empirical distribution. These counterfactuals provide intuition for the role of each cost function parameter in shaping earnings manipulation behavior. Panel A, B, C, and D present the counterfactuals of setting η , γ , ψ^2 , and ζ to 10% of estimated values, respectively.



— Empirical
- - - Counterfactual

Table 1: Summary Statistics

This table presents summary statistics (i.e., mean, standard deviation, 25th percentile, median, and 75th percentile) of key characteristics used in our regression discontinuity tests to follow.

	Mean	SD	25%	Median	75%
<i>EPS Surprise</i>	-0.236	7.521	-4	0	3
<i>EPS Surprise</i> ≥ 0	58.40%	-	-	-	-
<i>CMAR</i> ^{3 day}	0.552	8.327	-2.797	0.457	3.811

Table 2: The Benefits of Beating Earnings Benchmarks

This table presents regression discontinuity tests of three day cumulative earnings announcement returns around earnings benchmarks. The running variable is *EPS Surprise*, which is defined as realized earnings-per-share (EPS) minus analysts' consensus EPS forecast, and the dependent variable is $CMAR^{3day}$, which is defined as three day cumulative market-adjusted returns. We estimate the discontinuity in $CMAR^{3day}$ at the cutoff $EPS Surprise = 0$. Panel A uses global polynomial controls. Panel B uses local linear controls with bandwidth restrictions. Robust standard errors are clustered at the firm level and are reported in parentheses. ***, **, * represent statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A. Benefits test - Global Polynomial Controls

Dependent variable: $CMAR^{3day}$								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$EPS Surprise \geq 0$	2.494*** (0.083)	2.678*** (0.099)	1.815*** (0.120)	2.006* (0.144)	1.447** (0.168)	1.514*** (0.198)	1.173*** (0.228)	1.300*** (0.271)
<i>Fixed effects :</i>								
<i>Firm</i>		X		X		X		X
<i>Year – quarter</i>		X		X		X		X
<i>Polynomial degree</i>	0	0	1	1	2	2	3	3
R^2	0.0210	0.2235	0.0233	0.2254	0.0241	0.2260	0.0242	0.2261
<i>Observations</i>	103,819							
<i>Bins</i>	41							

Panel B. Benefits test - Local Linear Controls

Dependent variable: $CMAR^{3\text{ day}}$

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$EPS\ Surprise \geq 0$	1.172*** (0.200)	1.324*** (0.244)	1.620*** (0.148)	1.733*** (0.179)	1.773*** (0.132)	1.903*** (0.158)	1.815*** (0.120)	2.006* (0.144)
<i>Fixed effects :</i>								
<i>Firm</i>		X		X		X		X
<i>Year – quarter</i>		X		X		X		X
<i>Bandwidth</i>	5	5	10	10	15	15	20	20
R^2	0.0168	0.2548	0.0209	0.2353	0.0217	0.2292	0.0233	0.2254
<i>Observations</i>	67,659	67,659	86,707	86,707	96,138	96,138	103,819	103,819
<i>Bins</i>	41							

Table 3: Distributional Evidence of Manipulation Around Earnings Benchmarks

This table presents McCrary [2008] tests of manipulation of earnings-per-share (EPS) around analysts' consensus earnings forecasts. The running variable is *EPS Surprise*, which is defined as actual earnings minus analysts' consensus earnings forecast, and the dependent variable is the frequency of firm-quarter observations in each earnings surprise bin. We estimate the discontinuity in π_i at the cutoff $EPS Surprise = 0$. Panel A uses global polynomial controls. Panel B uses local linear controls with bandwidth restrictions. Robust standard errors are presented in parentheses. ***, **, * represent statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A. Manipulation test - EPS Surprise - Global Polynomial Controls

Dependent variable: π_i											
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
<i>EPS Surprise</i> ≥ 0	0.438 (0.553)	1.614* (0.939)	2.416*** (0.884)	2.775*** (0.697)	2.818*** (0.474)	2.625*** (0.371)	2.631*** (0.555)	3.093*** (0.610)	1.789* (1.080)	2.901 (2.435)	2.728** (1.055)
<i>Polynomial degree</i>	0	1	2	3	4	5	6	7	8	9	10*
R^2	0.4324	0.6987	0.9186	0.9720	0.9854	0.9886	0.9899	0.9902	0.9914	0.9917	0.9913
<i>Observations</i>	103,819										
<i>Bins</i>	41										

* $k = 10$ and $j = 8$ due to degrees of freedom restriction.

Panel B. Manipulation test - EPS Surprise - Local Linear Controls

Dependent variable: π_i

	(1)	(2)	(3)	(4)
<i>EPS Surprise</i> ≥ 0	2.752** (0.896)	2.322*** (0.992)	1.831* (0.964)	1.614* (0.939)
<i>Bandwidth</i>	5	10	15	20
<i>R</i> ²	0.9248	0.8378	0.7793	0.6987
<i>Observations</i>	103,819			
<i>Bins</i>	41			

Table 4: Variable Definitions

This table presents definitions of variables and parameters used in estimation of the model.

Notation	Definition
B	Estimated capital market benefit of just-meeting-or-beating the earnings benchmark (in equity market returns)
Δ	Estimated discontinuity in the earnings distribution (in % of firms)
η	Firm marginal cost to manage earnings (in equity market returns)
γ	Exponential slope parameter for cost to manage earnings (> 1 implies increasing costs of earnings management)
ψ^2	Stochastic noise of earnings management (in cents)
ζ	Heteroskedasticity of earnings management noise (increases noise with greater manipulation)

Table 5: Structural estimates

This table presents our structural estimates as well as statistics regarding the percentage of manipulators, the expected amount of manipulation by manipulators and the fraction of manipulators shocked out of their target bin by noise.

	η	γ	ψ^2	ζ	$\%manip$	\overline{manip}	$\%noise$
Baseline	0.0161 (0.0004)	2.0817 (0.0517)	0.8201 (0.0236)	3.7105 (0.0941)	2.62%	1.21	59.6%
<i>robustness</i>							
-15 to 15 cents	0.0157	2.2398	0.7761	3.8515	2.71%	1.15	55.7%
-10 to 10 cents	0.0154	1.9381	0.7942	3.9199	2.84%	1.27	60.0%
-5 to 5 cents	0.0167	2.0950	0.7599	4.1904	2.56%	1.19	57.3%
overweight empirical	0.0170	2.2847	0.7992	2.6322	2.45%	1.15	57.3%
overweight smoothness	0.0146	1.8871	0.7901	2.5761	3.11%	1.31	58.9%
$\zeta = 0$	0.0211	2.0907	0.2983	0	2.29%	1.29	26.4%
<i>alternate distributions</i>							
Reduce peak	0.0205	2.1867	0.8178	2.2439	1.91%	1.17	73.5%
Kill peak	0.0273	2.0523	1.1699	0.9738	1.36%	1.20	64.3%

Table 6: Counterfactuals

This table shows parameter estimates and statistics based on counterfactual simulations. These simulations involve, one at a time, reducing the value of each parameter to 10% of its estimated value. The first row reproduces the baseline estimates for comparison purposes.

	η	γ	ψ^2	ζ	$\%manip$	\overline{manip}	$\%noise$
Baseline	0.0161	2.0817	0.8201	3.7105	2.62%	1.21	59.6%
η at 10%	<i>0.0016</i>	2.0817	0.8201	3.7105	12.6%	1.92	64.3%
γ at 10%	0.0161	<i>0.2082</i>	0.8201	3.7105	9.45%	10.16	78.4%
ψ^2 at 10%	0.0161	2.0817	<i>0.0820</i>	3.7105	3.22%	1.23	6.9%
ζ at 10%	0.0161	2.0817	0.8201	<i>0.3711</i>	2.56%	1.29	57.8%

Table 7: Identifying manipulators by earnings surprise bins

This table illustrates the performance of various definitions of “suspect” bins on identifying manipulation. The second set of percentages weights manipulation by the amount of manipulation, i.e. a firm manipulating two cents gets twice the weight of a firm only manipulating one cent.

	$[0, 1)$	$[0, 2)$	$[0, 3)$	$[-1, 1)$	$[-1, 2)$	$[-1, 3)$
% Manipulating	11.0%	9.9%	8.1%	8.7%	8.5%	7.3%
% of Manipulators	53.1%	80.9%	86.4%	64.2%	92.0%	97.4%
% Manipulating (weighted)	12.8%	14.2%	12.4%	10.2%	12.1%	11.1%
% of Manipulators (weighted)	39.6%	77.3%	88.3%	48.3%	86.0%	97.0%

Table 8: SOX and the cost of earnings management

This table shows how the costs of earnings management changed around SOX. The pre-SOX period is 1999-2001 and the post-SOX period is 2002-2004. Δ is the estimated discontinuity in the earnings distribution, estimated as in Table 3. \mathcal{B} is the benefit to meeting the analyst benchmark, estimated as in Table 2. η , γ , ψ^2 and ζ are the estimated parameters of the cost function. The final three rows describe the estimated equilibrium manipulation in each period.

	Pre-SOX	Post-SOX	Diff.	<i>p-value</i>
Δ	4.196	3.012	-1.184	0.008
\mathcal{B}	1.229	2.027	0.798	0.034
η	0.0097	0.0278	0.0181	<0.0001
γ	1.942	2.808	0.865	<0.0001
ψ^2	0.646	0.202	-0.444	<0.0001
ζ	3.040	1.316	-1.724	<0.0001
$\overline{\%manip}$	4.05%	2.61%	-1.44%	-
\overline{manip}	1.29	1.06	-0.23	-
$\overline{\%noise}$	54.2%	16.0%	-38.1%	-