Consequences of Misreporting*

Francesco Celentano† Jason Choi‡ Philip Coyle§

May 2026

Abstract

Detected misreporting triggers stock-price declines, reputational damage, and monetary sanctions, yet misreporting persists. We develop and estimate a dynamic heterogeneous-firm model in which managers inflate reported profits to lower the cost of external finance, at the risk of detection and the loss of access to manipulation. Estimated on U.S. firm-level data, each one-percentage-point increase in reported profitability lowers the per-unit cost of external finance by 1.8%; the directly punitive components of detection are quantitatively small. Almost the entire cost of being flagged operates through two channels: the firm cannot use misreporting to soften its financing wedge, and the wedge is priced off true rather than inflated profits. Prohibiting misreporting reduces shareholder value by 3.1%, of which two-thirds reflect the option value of the channel.

JEL Classification: G30, G32, G38, M41

Keywords: Firm dynamics, earnings misreporting, financial frictions, enforcement

*We are grateful to Martin Brown, Brent Bundick, Theodosios Dimopoulos, Amit Goyal, Julien Hugonnier, Oliver Levine, Nadya Malenko, Garen Markarian, Cyril Monnet, Erwan Morellec, Boris Nikolov, Antoinette Schoar, Norman Schürhoff, Rish Singhania and seminar participants at the University of Exeter, Kansas City Fed, University of Lausanne/EPFL, and Study Center Gerzensee. We thank Evgenii Maiorov for useful help on a previous version of this work, and Anastasia Zakolyukina for sharing her code. All remaining errors are our own

†University of Lausanne and Swiss Finance Institute. E-mail: francesco.celentano@unil.ch

‡University of Toronto. E-mail: jasonhj.choi@utoronto.ca

§Study Center Gerzensee. E-mail: philip.coyle@szgerzensee.ch

1 Introduction

Detected misreporting is a consequential event for the firm involved. Restatements driven by accounting fraud are accompanied by large declines in the firm's stock price, lasting reputational damage with customers, suppliers, and investors, lower sales in subsequent years, and elevated costs of external finance. The Securities and Exchange Commission (SEC) and the Department of Justice (DOJ) impose monetary sanctions and extended compliance monitoring on firms found to have misrepresented their financial statements. These events are not rare. About 1 percent of observations in our sample involve an intentional restatement, a figure that counts only the firms that were caught and so understates the true incidence of misreporting.

Given that misreporting persists despite costly consequences, this points to a private benefit large enough to, on average, outweigh the expected punishment. The benefit we focus on operates through the cost of external finance. Investors price loans, equity issuance, and covenant terms off the numbers a firm reports: stronger reported profitability buys tighter loan spreads, better issuance terms, and looser covenant restrictions. By inflating reported profits, a manager can relax these contractual frictions and raise external financing on cheaper terms. However, the reports are not free to produce. Building and maintaining a bias on the books absorbs real resources, including auditor side payments, fabricated supporting documentation, and internal coordination across multiple parties, and exposes the firm to the risk of detection. The trade-off depends on the magnitudes of the financing benefit, the resource cost of engineering bias, the probability of detection, and the severity of the punishment that follows.

In this paper we ask two questions. First, how do firms in the cross-section use the misreporting channel, and what happens to investment, financing, and firm value if access to the channel is removed? Second, holding the rest of the model fixed, which values of the detection and punishment parameters maximize cross-sectional firm value, and through which margins of firm behavior does each parameter bite? This is not a regulator's welfare problem: we do not model enforcement resource costs, investor welfare, or economy-wide losses from distorted reports. The object is deliberately narrower. Both questions require a framework that links the firm's reporting decision to its investment and financing decisions over time, and that disciplines the detection and punishment parameters jointly with the data.

We develop and structurally estimate a dynamic heterogeneous-firm model that maps this trade-off into firm-level investment, financing, and reporting choices. Firms differ in their productivity, capital stock, and accumulated liquid resources, and face a cost of external finance that investors price off the firm's reported numbers. Each period a firm chooses investment, savings, and, when unflagged, how aggressively to inflate its reported profits. Inflating reports cheapens external finance, but consumes real resources to coordinate the distortion across auditors, accountants, and internal records, and exposes the firm to detection.

If misreporting is uncovered, the firm enters a flagged state with three consequences. First, it pays a pecuniary penalty proportional to the size of the discovered distortion. Second, it bears a reputational surcharge on its cost of external finance, reflecting the additional cost investors impose on firms publicly identified as misreporters. Third, and most consequentially, increased regulatory monitoring makes further misreporting infeasible while the firm remains flagged, so the firm's cost of external finance is priced off true rather than inflated profits. The flagged state ends stochastically, after which the firm may once again misreport. The model contains four enforcement instruments: the detection probability, the persistence of the flagged state, the pecuniary penalty, and the reputational surcharge. Their relative bite is what the firm's optimal trade-off identifies in the data.

We bring the model to U.S. public firm data by simulated method of moments. The sample combines CRSP/Compustat for financing and investment over 2003–2019 with detected misreporting events from Audit Analytics. The estimation matches fourteen moments organized into four blocks: (i) production, (ii) external finance, (iii) investment and external financing in the immediate aftermath of being caught, and (iv) the frequency, magnitude, and duration of detection events themselves. The detection-event moments identify the monitoring technology, the issuance moments identify the cost of external finance, and the post-detection event-study moments identify the punishment intensity. The model fits the targeted moments closely.

The cost of external finance is highly sensitive to reported profits: each one-percentage-point increase in reported profitability lowers the per-unit cost of external finance by approximately 1.8%. The per-period detection probability of 2.1% implies that an actively misreporting firm is caught on average once every 48 years, and the flagged state lasts on average 3.1 years once entered. Taken on their own, the direct components of punishment are quantitatively modest. The implied pecuniary penalty averages roughly 1.2% of assets at the moment of detection (about $53 million for the mean firm), and the reputational surcharge raises a flagged firm's per-unit financing cost by only 0.2% above the baseline. How does this square with the large empirical costs of detection in the first paragraph? The main cost of being flagged is not the fine or the surcharge, but the loss of access to the misreporting channel itself: a flagged firm cannot inflate reports to soften its cost of external finance. Almost the entire post-detection penalty borne by caught firms operates through this loss rather than through fines or stigma.

How much do firms value the misreporting channel? Relative to baseline, outright prohibition reduces shareholder value by 3.1% and lowers the average investment rate by 7.5%. Firms can no longer use inflated reports to soften the cost of external finance, and real investment falls to the level consistent with truthful reporting. Investors are no longer misled by inflated reports, so the financing terms they offer are based on true rather than overstated profitability. The 3.1% firm-value loss splits into two components: the value loss for firms that are caught and lose access to the channel, and the option value of the misreporting channel itself for the broader cross-section. To separate them, consider an intermediate counterfactual in which firms can still choose to misreport but anyone caught is permanently flagged ($q_f = 1$, so the flagged state never ends). Firms whose state realization makes a one-shot spell worthwhile still take it; the unlucky ones who are caught lose access to the channel forever. Firm value falls by 1.2% relative to baseline in this case, isolating the value loss on the small group caught and permanently flagged. The remaining 1.9 percentage points to the full 3.1% loss is the option value: the value firms place ex ante on having the channel available at all, which outright prohibition removes. Variation in λ_π moves firm value far more than variation in any of the four enforcement instruments: halving it lowers firm value by 4.6% and doubling it raises value by 5.0%.

Which level of each enforcement parameter maximizes cross-sectional firm value in the model? For three of the four (the detection probability, the persistence of the flagged state, and the reputational surcharge), the firm-value-maximizing level is strictly positive; for the pecuniary penalty it is zero. These are not social-welfare optima. They are one-at-a-time comparative statics for average firm value, abstracting from the cost of implementing enforcement and from investor-side welfare. The interior maxima arise from a tension between two forces. At any given state, the firm's value function is weakly lower under stronger enforcement, so no individual firm prefers more enforcement to less. The cross-sectional average can still rise because stronger enforcement also shifts the stationary distribution: fewer firms carry accumulated bias on their books, more mass concentrates at $m_{-1} = 0$, and the average firm is in a state with a higher value function. Over a range of enforcement, this distributional shift outweighs the pointwise drop in firm values; beyond that range, the pointwise drop takes over. The pecuniary penalty is the exception: it takes cash out of the firm at every detection event, so its pointwise drag on value dominates the distributional shift across all positive values.

1.1 Literature Review

We contribute to three related strands of the literature. The first is the structural corporate finance literature that quantifies how financing frictions shape firms' intertemporal policies (Hennessy and Whited, 2007; Riddick and Whited, 2009; Bolton et al., 2011; Nikolov and Whited, 2014; Eisfeldt and Muir, 2016; Gao et al., 2021). A central insight of this work is that the wedge between internal and external funds is both economically meaningful and identifiable from joint patterns in financing, savings, and investment. Recent evidence links these frictions to accounting-based contracting terms and cash-flow covenants (Lian and Ma, 2021; Drechsel, 2023), providing the empirical foundation for a financing cost that depends on reported profits. Our closest paper in this tradition is Terry et al. (2023), who develop a dynamic model in which managers trade off misreporting against real investment distortions, emphasizing an information–investment channel. We adopt the same structural discipline but shift the mechanism: in our framework misreporting is a tool for managing the marginal cost of external funds rather than a substitute for real distortions, which lets us separately identify the reporting and financing margins.

The second strand is the quantitative accounting literature on earnings manipulation, detection, and punishment. Empirical work establishes that fraud detection has persistent consequences for firms and managers (Karpoff et al., 2008; Dyck et al., 2010), and most directly relevant to our mechanism, Graham et al. (2008) document that restatement raises loan spreads and tightens nonprice contract terms, providing an empirical anchor for the post-detection financing penalty in our model. On the structural and theoretical side, Benmelech et al. (2010) develop a dynamic rational expectations model in which stock-based compensation induces managers to conceal bad news and pursue suboptimal investment policies to sustain the pretense, while Povel et al. (2007) derive in an equilibrium framework that fraud peaks at the end of booms and is revealed in busts. Zakolyukina (2018) estimates intentional misreporting with detection and punishment as dynamic objects, and Bertomeu et al. (2022) develops a quantitative model of concealment dynamics. We build on this work but depart in what misreporting accomplishes: rather than treating the bias stock as a substitute for real activity, we model it as a direct lever over financing conditions, and we discipline the detection and punishment parameters using matched moments on restatement frequency and post-detection financing outcomes.

The third strand is recent work documenting that misreporting has economically meaningful consequences for firms and the allocation of capital (Kedia and Philippon, 2009; Alawadhi et al., 2025). This work has focused primarily on measuring the consequences of detected misreporting. Our contribution is to embed those consequences in a structural framework that can trace how different detection and punishment parameters move firm behavior and cross-sectional firm value. To our knowledge, this is the first structural quantification of how average firm value responds to the mix of monetary sanctions and ongoing scrutiny, and of how those responses depend on the distinction between sanctions that remove the misreporting channel and those that transfer resources away from the firm.

2 Model

We develop a dynamic heterogeneous-firm model in which the risk of detection and the costs of regulatory scrutiny shape managers' incentives to misreport profits and accumulate capital. Section 2.1 describes the environment and the firm's problem; Section 2.2 examines the equilibrium outcomes.

2.1 Environment

Time is discrete and infinite, denoted by t. The economy is populated by a continuum of heterogeneous firms; to ease notation, we suppress firm subscripts throughout. Each firm has an idiosyncratic total factor productivity (TFP) z_t, a capital stock $k_t \geq 0$, liquid resources $\ell_t \geq 0$, an accumulated book-value bias $m_{t-1} \geq 0$, and a flag status $f_t \in \{0, 1\}$, where $f_t = 1$ denotes a *flagged* firm currently under regulatory scrutiny and $f_t = 0$ denotes an *unflagged* firm not currently under scrutiny. Flagged firms are prohibited from misreporting and face elevated financing costs; unflagged firms may choose to misreport profits.

Investors observe z_t, k_t, reported liquid resources $\hat{\ell}_t$, and reported profits $\hat{\pi}_t$, but do not observe the firm's book-value bias m_t, and detection of misreporting is probabilistic. Each period, firms choose investment, savings, and, when unflagged, the book-bias level, trading off the financing benefits of overstating profits against the risk of detection and the costs of adjustment.

2.1.1 Production technology

Each firm operates a decreasing-returns-to-scale production technology that generates period-t profits

$$\pi_t = z_t \, k_t^{\alpha} - \eta (m_t - m_{t-1})^2 \,, \tag{1}$$

where $\alpha \in (0, 1)$ governs returns to scale. The term $\eta(m_t - m_{t-1})^2$, with $\eta > 0$, is a quadratic adjustment cost on changes in the book-value bias. It represents the real resource costs the manager must incur to engineer a change in reported profits relative to the prior period, including side payments to accountants and auditors willing to sign off on doctored statements, the cost of fabricating supporting documentation, and the cost of coordinating across multiple parties to maintain the bias. The convex specification reflects that each successive unit of distortion requires progressively more elaborate concealment. Capital is accumulated subject to one-period time-to-build,

$$k_{t+1} = (1 - \delta)\, k_t + i_t, \tag{2}$$

where $\delta \in (0, 1)$ is the depreciation rate and i_t is gross investment. Idiosyncratic TFP follows an AR(1) in logs,

$$\ln z_{t+1} = (1 - \rho_z)\mu_z + \rho_z \ln z_t + \sigma_z \varepsilon_{t+1}, \qquad \varepsilon_{t+1} \overset{\text{iid}}{\sim} \mathcal{N}(0, 1), \tag{3}$$

where $\rho_z \in (0, 1)$ is the persistence of idiosyncratic TFP, μ_z its unconditional mean in logs, and $\sigma_z > 0$ the standard deviation of the innovations.

2.1.2 Misreporting

An unflagged firm ($f_t = 0$) may choose a book-value bias $m_t \geq 0$. Setting $m_t = 0$ corresponds to truthful reporting, while $m_t > 0$ corresponds to overstating profits and savings. The net distortion in reported earnings is the change in the bias stock, $\Delta m_t \equiv m_t - m_{t-1}$, and reported profits and reported liquid resources are

$$\hat{\pi}_t = \pi_t + m_t - m_{t-1}, \qquad \hat{\ell}_t \equiv \ell_t + m_{t-1}, \qquad \frac{\widehat{\ell}_{t+1}}{1 + r_\ell} \equiv \frac{\ell_{t+1}}{1 + r_\ell} + m_t. \tag{4}$$

The manager's reports overstate true savings by the current bias stock m_t, while the change Δm_t flows directly into reported earnings: building up the bias inflates current reported profits, while drawing it down depresses them.

The financing benefit of misreporting accrues through the period-on-period flow Δm_t via reported profits $\hat{\pi}_t$, but the spell itself is defined by the stock. As long as $m_t > 0$, the firm carries accumulated overstatement on its books and remains exposed to audit-based detection, even in a period in which it makes no incremental change to the bias. Ending a spell requires running the books back to zero by setting $m_t = 0$, which pays the adjustment cost ηm_{t-1}^2 and depresses reported profits in the wind-down period through $\Delta m_t = -m_{t-1} < 0$. A firm with accumulated m_{t-1} chooses each period whether to maintain the bias and bear the per-period detection hazard q_p, or to pay the wind-down cost and exit the spell undetected.

Misreporting carries the risk of detection. If an unflagged firm sets $m_t > 0$, it is caught at the end of the period with probability $q_p \in (0,1)$, which we interpret as the rate at which regulators (such as the SEC), auditors, or other external monitors uncover the bias. A caught firm is flagged in the next period ($f_{t+1} = 1$) and carries the bias stock m_t into the flagged state, where penalties are levied. If not caught, the firm remains unflagged and rolls m_t forward.

A flagged firm ($f_t = 1$) is prohibited from misreporting and must set $m_t = 0$. In addition to the elevated financing costs described below, a flagged firm pays a pecuniary penalty

$$p(m_{t-1}) = p_m \, m_{t-1}^2, \tag{5}$$

proportional to the square of the inherited bias stock and deducted from period-t profits. The parameter $p_m > 0$ represents the direct monetary costs the firm bears when it is caught cheating, including fines imposed by regulators, legal fees, settlement payments, and any other pecuniary sanctions tied to the magnitude of the discovered bias; the convex specification reflects that larger restatements typically attract proportionately larger enforcement responses. At the end of the period, a flagged firm exits scrutiny and returns to unflagged status with probability $1 - q_f$, and remains flagged with probability $q_f \in (0,1)$. The persistence parameter q_f controls the expected duration of regulatory scrutiny: a higher q_f means flagged firms remain under elevated oversight for longer. Because no new bias can accumulate while a firm is flagged, the bias stock entering the next period is zero regardless of the flag transition.

2.1.3 Financing

Firms finance operations using current profits, accumulated liquid resources, and costly external finance. Liquid resources carry over from one period to the next at a gross return $1 + r_\ell$, where $r_\ell < r$ is a fixed liquidity return that lies strictly below the firms' discount rate so that holding cash is costly in expectation. The wedge on external finance depends on the firm's reported performance and on its flag status, and acts as the friction that creates an endogenous role for reporting in the model. We interpret the wedge as the financing terms faced by the firm, not as money burned inside the firm. Lower reported-risk financing terms transfer value to the firm by reducing what it must give up to outside investors; when those terms are obtained through inflated reports, the corresponding gain to the firm comes at investors' expense. Let e_t denote the firm's net cash flow, with positive values corresponding to payouts and negative values to external issuance. For an **unflagged** firm, the external financing wedge is

$$\Lambda(\hat{\pi}_t, e_t, f_t{=}0) = \lambda_1 \exp\{-\lambda_\pi \hat{\pi}_t\} \, |e_t|, \tag{6}$$

where $\lambda_1 > 0$ is the baseline per-unit cost of external finance, reflecting standard frictions such as asymmetric information, compensation to outside investors for bearing risk, and issuance-related costs, and $\lambda_\pi > 0$ governs the reduction in financing costs from higher reported profits $\hat{\pi}_t$, reflecting the empirical regularity that firms with stronger reported performance face tighter loan spreads and better issuance terms. This last parameter is what gives misreporting its financing benefit: by raising $\hat{\pi}_t$, an unflagged firm can soften the per-unit wedge it faces when accessing external markets and retain value that investors would have captured under truthful reports. For a **flagged** firm, financing costs reflect the additional stigma of regulatory scrutiny:

$$\Lambda(\pi_t, e_t, f_t{=}1) = \big(\lambda_1 \exp\{-\lambda_\pi \pi_t\} + \lambda_f\big)|e_t|, \tag{7}$$

where $\lambda_f > 0$ is an additional financing surcharge reflecting the reputational cost of being publicly identified as a misreporter. Once a firm is flagged, capital providers know it has been caught cheating and demand a premium that persists for as long as the firm remains under scrutiny, raising its cost of external finance above what its fundamentals alone would warrant. Because flagged firms cannot manipulate their reports, their financing wedge is priced off true profits π_t rather than reported profits $\hat{\pi}_t$.

2.1.4 Timing

For each period t, the within-period timing is as follows:

1. **State realization.** The firm enters period t with state $(z_t, k_t, \ell_t, m_{t-1}, f_t)$. The firm-level TFP z_t, capital stock k_t, and flag status f_t are publicly observed, while the bias stock m_{t-1} is private.

2. **Firm decisions.** The firm chooses next-period capital $k_{t+1} \geq 0$ (equivalently, gross investment i_t via (2), which may be negative), next-period liquid resources $\ell_{t+1} \geq 0$, and, when unflagged, the book-value bias $m_t \geq 0$. Flagged firms are restricted to $m_t = 0$ and pay the penalty $p(m_{t-1})$. These choices determine true profits π_t via (1) and the public report vector $(\hat{\pi}_t, \hat{\ell}_t, \hat{\ell}_{t+1})$.

3. **Financing and payouts.** The net cash flow e_t is realized. If $e_t < 0$, the firm accesses external finance and pays the issuance cost $\Lambda(\cdot)$ given by (6) or (7), depending on its flag status.

4. **Detection and flagging.** If the firm set $m_t > 0$, it is caught with probability q_p. A caught firm is flagged next period ($f_{t+1} = 1$) and carries the bias stock m_t into the

flagged state, where it determines the subsequent penalty. If not caught, the bias stock m_t is carried forward and the firm remains unflagged. Flagged firms exit scrutiny with probability $1 - q_f$ ($f_{t+1} = 0$) or remain flagged with probability q_f ($f_{t+1} = 1$), with the bias reset to zero in either case.

Figure 1 summarizes the within-period sequence.

2.1.5 Firm problem

To write the firm's problem compactly, we adopt recursive notation: time subscripts are dropped, unprimed variables denote current-period values, primed variables denote next-period values, and the subscript -1 denotes lagged values. Firms take as given the external financing wedges (6)–(7), the law of motion for capital (2), and the stochastic process for idiosyncratic TFP (3).

Let $v(z, k, \ell, m_{-1}, f)$ denote the firm's value function. With gross investment $i = k' - (1 - \delta)k$, the firm's net cash flow is

$$e = \hat{\pi} - i - \frac{\widehat{\ell'}}{1 + r_\ell} + \widehat{\ell} \qquad \text{(unflagged)},$$

$$e = \pi - p(m_{-1}) - i - \frac{\ell'}{1 + r_\ell} + \ell \qquad \text{(flagged)},$$

where r_ℓ is the return on liquid savings. To compress notation, let $\mathcal{V}(m, f) \equiv \mathbb{E}_z[v(z', k', \ell', m, f)]$ denote the conditional expected continuation value, where the expectation is taken over the next-period idiosyncratic TFP shock. The firm's Bellman equation then takes the following form, with different continuation structures for unflagged and flagged firms.

Unflagged firms ($f = 0$)

$$\begin{aligned} v(z, k, \ell, m_{-1}, 0) = \max_{k', \ell', m \geq 0} \Bigg\{ & e - \mathbf{1}\{e < 0\}\, \Lambda(\hat{\pi}, e, 0) + \frac{\mathbf{1}\{m = 0\}}{1 + r}\, \mathcal{V}(0, 0) \\ & + \frac{\mathbf{1}\{m > 0\}}{1 + r} \big(q_p\, \mathcal{V}(m, 1) + (1 - q_p)\, \mathcal{V}(m, 0) \big) \Bigg\}. \end{aligned} \tag{8}$$

Flagged firms ($f = 1$)

$$\begin{aligned} v(z, k, \ell, m_{-1}, 1) = \max_{k', \ell' \geq 0} \Bigg\{ & e - \mathbf{1}\{e < 0\}\, \Lambda(\pi, e, 1) \\ & + \frac{1}{1 + r} \big(q_f\, \mathcal{V}(0, 1) + (1 - q_f)\, \mathcal{V}(0, 0) \big) \Bigg\}. \end{aligned} \tag{9}$$

A recursive competitive equilibrium consists of value functions $v(\cdot)$ and policy functions (k', ℓ', m) such that the policy functions solve the firm's Bellman equations (8)–(9), and the implied transitions of (z, k, ℓ, m_{-1}, f) are consistent with the idiosyncratic TFP process (3), the law of motion for capital (2), and the detection and flagging dynamics described above.

2.2 Results

We trace through the equilibrium outcomes of the model by examining the firm's policy functions, the lifecycle of a typical misreporting spell, and the implied cross-section of investor-side financing losses. Throughout this subsection, structural parameters are set to the estimates reported in Section 3.3 (Table 2). To keep the figures legible, we vary one state variable at a time in the policy-function plots and two at a time in the heatmaps, holding the remaining states at representative values from the simulation.

Figure 2 plots the next-period capital choice k' as a function of current capital k, the inherited bias stock m_{-1}, and idiosyncratic productivity z, separately for unflagged (solid blue) and flagged (dashed red) firms. The left panel shows that k' rises with current capital before levelling off near a long-run target consistent with one-period time-to-build and decreasing returns to scale, with a vertical gap between the unflagged and flagged curves that is largest at low and intermediate values of k. The middle panel shows that, conditional on being unflagged, k' is strictly decreasing in the inherited bias stock m_{-1}. The right panel shows that capital is increasing and concave in productivity, with flagged firms investing materially less at every level of z and a gap that widens with z.

The vertical gap between unflagged and flagged firms is the investment penalty associated with regulatory scrutiny. A flagged firm pays the surcharge λ_f and, more importantly, has lost the ability to use reported profits to soften its financing wedge, so it accumulates capital more slowly. The fact that the gap widens with productivity says that the value of the reporting channel is largest exactly for firms whose investment opportunities are most attractive. The negative slope in m_{-1} in the middle panel comes from three sources: a larger inherited bias raises the resource cost of maintaining the bias through $\eta(\Delta m)^2$, raises the expected pecuniary penalty $p_m m_{-1}^2$ should the firm be caught, and shortens the effective horizon over which the firm can continue manipulating reports without detection. All three draw resources away from physical investment.

Figure 3 plots the bias choice m for an unflagged firm. The left panel shows a hump-shaped relationship between m and current capital. The middle panel shows that m is strictly decreasing in the inherited bias stock. The right panel shows that m is increasing and concave in productivity.

The hump shape in capital reflects how external-finance needs scale with firm size. Very small firms operate at low scale with modest external-finance needs, so the marginal benefit of inflating reported profits is small. Mid-size firms invest aggressively, run negative net cash flows, and therefore benefit most from softening the financing wedge by overstating $\hat{\pi}$. Large firms have accumulated sufficient internal resources to fund operations without significant issuance, attenuating the misreporting incentive. The negative slope in the inherited bias stock comes from two mean-reverting forces: the convex adjustment cost $\eta(\Delta m)^2$ makes large period-on-period changes costly, and the quadratic detection penalty $p_m m^2$ raises the marginal cost of carrying a larger bias into the next period. These forces induce mean reversion in the bias process even absent detection. The positive slope in productivity reflects that more productive firms expand more, rely more heavily on external finance, and therefore extract more value from the multiplicative reduction $\exp\{-\lambda_\pi \hat{\pi}\}$ that an additional unit of bias produces.

The policy functions in Figures 2 and 3 jointly trace the lifecycle of a representative misreporting spell. A spell is most likely to initiate at a firm whose state (z, k, ℓ) maximizes the marginal benefit of inflating reports relative to the marginal cost of building bias: intermediate capital, low liquidity, and a positive productivity draw. During the build-up phase the firm accumulates m at a rate that is initially large but tapers as the inherited stock grows, both because $\eta(\Delta m)^2$ penalizes aggressive year-on-year increases and because the mean reversion of m in the inherited stock pulls the bias toward a state-dependent target level at which the marginal financing benefit just offsets the marginal expected punishment cost. Detection arrives stochastically; conditional on being caught, the firm carries its accumulated m into the flagged state, pays the pecuniary penalty $p_m m^2$, and is barred from further misreporting for the duration of regulatory scrutiny. While flagged, the firm's investment policy follows the dashed curve in Figure 2: it accumulates capital more slowly than an otherwise identical unflagged firm because its financing wedge is priced off true rather than reported profits and is therefore unmitigated by the manipulation lever. Upon exit from the flagged state, which takes an average of 3.1 years at the estimated q_f, the bias resets to zero and the firm re-enters the cross-section as an unflagged firm. Whether it re-initiates a new spell depends on the realized (z, k, ℓ) state at that moment.

Figure 4 maps $\Delta\lambda$, the per-firm gap between the financing cost a firm actually pays and the cost it would face if reported profits coincided with true profits. $\Delta\lambda$ is therefore a measure of how much each firm's misreporting reduces its private cost of external finance at a given state, holding the model's wedge schedule fixed. The three panels show $\Delta\lambda$ over the (ℓ, k), (k, z), and (ℓ, z) planes. The qualitative pattern is consistent across cuts: $\Delta\lambda$ is concentrated among firms with low capital, low liquidity, and high productivity, the firms whose combined need for external finance and capacity to inflate reported profits are largest. Well-capitalized or cash-rich firms have little need to issue and therefore little to gain from manipulation. Low-z firms generate low reported profits whether they cheat or not, so the multiplicative reduction $\exp\{-\lambda_\pi \hat{\pi}\}$ contributes a smaller dollar saving per unit of bias. The middle panel makes clear that $\Delta\lambda$ is steepest along the capital margin: even at moderate productivity, very small firms generate substantial gaps because their financing needs are large relative to their reported earnings base.

Figure 5 plots the corresponding misreporting policy m in the same three planes and confirms that the mapping between m and $\Delta\lambda$ is tight: regions of the state space where firms accumulate the most bias coincide with the regions where the private financing-cost reduction from misreporting is largest. Low-capital, low-liquidity, high-productivity firms are simultaneously the most active misreporters and those with the most to gain at the margin from inflating their reports.

These policy functions provide a mechanism for several patterns documented in Section 3.1. The hump shape of m in current capital matches the low frequency of detection events in the data (0.8% of firm-years): the smallest firms have modest external-finance needs and the largest firms can self-finance, so only a middle range of firms has strong incentives to inflate $\hat{\pi}$. The mean reversion of m in the inherited bias stock, combined with the convex adjustment cost $\eta(\Delta m)^2$, generates gradual bias build-up and decay rather than one-shot spikes, consistent with the average misreporting spell of 4.6 years observed in Audit Analytics. The concentration of $\Delta\lambda$ in the low-capital, low-liquidity, high-productivity corner of the state space identifies the firms whose financing needs are largest relative to their reported earnings base, and for whom the marginal value of inflating $\hat{\pi}$ is greatest.

3 Estimation

Section 3.1 describes the data. Section 3.2 reviews the externally calibrated parameters and lays out the simulated-method-of-moments strategy and identification logic for the twelve internally estimated parameters. Section 3.3 reports the estimates and assesses the model's fit.

3.1 Data

Our sample covers the period from 2003 to 2019 and is constructed from multiple data sources. Firms' financial and accounting data come from the annual CRSP/Compustat merged database. We keep only observations of firms incorporated in the U.S., listed on the

NYSE, Amex and Nasdaq, with non-missing and non-negative book assets and sales, and with reliable and primary links (LINKTYPE equal to LU or LC and LINKPRIM equal to P or C). We drop observations of regulated utilities (SIC 4900-4999), financial firms (SIC 6000-6999), and quasi-governmental and nonprofit firms (SIC 9000-9999). Data about the TFP were downloaded from Şelale Tüzel's website (İmrohoroğlu and Tüzel (2014)).[1]

Data about profits misreporting come from Audit Analytics. We follow a procedure similar to Terry et al. (2023) to distinguish intentional from involuntary misreporting, which is not of interest for us, and to compute the cumulative impact of a restatement on net income, which is our measure of the book-value bias. We apply an additional filter considering only misreporting cases in which the final book value bias is positive. Additionally, we require 7 years of data, for which 3 years must be consecutive to ensure sufficient time-series variation. We normalize all the nominal variables in 2019 U.S. dollars using the data available on the St. Louis Federal Reserve Economic Data (FRED) website.

We compute the external finance flows as in Eisfeldt and Muir (2016). Specifically, we compute cash flows to debt as the sum of reduction in long-term debt (Compustat item DLTR), changes in current debt (DLCCH), and net interest paid (XINT) minus long-term debt issuance (DLTIS). Cash flows to equity are computed as the sum of the purchase of common and preferred stocks (PRSTKC) and cash dividends (DV) minus sale of common and preferred stocks. The negative of this sum corresponds to our measure of external finance. As in Begenau and Salomao (2019), we use the adjustment of McKeon (2015) and consider only sales of common stock that are larger than 3% of the total market value to differentiate firm-initiated from employee-initiated stock issuances.

Table 1 reports descriptive statistics for the resulting panel of 40,311 firm-years from 3,083 distinct firms. The median sample firm has book assets of about $690 million and sales of $629 million in 2019 dollars (Panel A). Panel B shows that reported profitability is close to zero on average, the investment rate is roughly 13% of physical capital, and reported liquid resources average 37% of reported assets. Panel C summarizes the external-finance margin: firms have positive net external finance in 29% of firm-years, and conditional on issuing, the typical positive flow is 48% of reported assets; the equity and debt components contribute roughly symmetrically on average. Panel D documents that 0.8% of firm-years involve a detection event, with caught firms having accumulated a book-value bias averaging 5% of reported assets and an annual earnings inflation of 1%. Conditional on detection (Panel E), the average misreporting spell lasts 4.6 years and the average book-value bias at the time of detection is 5.7% of assets.

[1] The data were accessed at https://sites.google.com/usc.edu/selale-tuzel/home

3.2 Estimation strategy

We externally calibrate two parameters. The discount rate r is set to 1.5%, computed as the average annual T-bill rate over our sample period. The liquidity return rate r_ℓ is set to $0.8 \times r = 1.2\%$, following Gao et al. (2021), so that liquid resources carry over at a rate strictly below the discount rate. The top panel of Table 2 reports the values of all externally calibrated parameters.

We estimate the remaining twelve parameters by simulated method of moments: we choose the parameter vector that minimizes the distance between fourteen empirical moments and their model-simulated counterparts. Before computing moments in both data and simulation, we scale reported profits, new reported liquidity, external finance, and book bias by reported assets, defined as the sum of physical capital and reported initial liquidity. The moments fall into four blocks: production, financing, flag dynamics, and cheating. The discussion below highlights, for each parameter, the moment most informative at the margin; Appendix Table A.1 summarizes the mapping.

We discipline the idiosyncratic TFP process using the panel approach of Han and Phillips (2010). The first-order autocorrelation of log productivity identifies the persistence parameter ρ_z. The mean drift in the AR(1) regression identifies the unconditional log-mean μ_z. The variance of the residuals from that autocorrelation regression identifies the innovation standard deviation σ_z. The average level of reported profits identifies the returns-to-scale parameter α, since α governs the curvature of the production function and hence the typical level of operating profit a firm of a given size generates. The average investment rate identifies the depreciation rate δ: in the cross-section, gross investment must on average offset capital depreciation, so the mean investment rate pins δ down directly.

The intensity of the financing cost λ_1 is jointly identified by the frequency of positive external finance flows and the conditional average level of those flows in the full sample. A higher λ_1 raises the marginal cost of issuance, reducing both the share of firm-years with positive issuance and the typical issuance size, so these two moments together discipline both the extensive and intensive margins of external finance. The sensitivity of the financing cost to reported profits λ_π is identified by the correlation between the book-value bias m and positive external finance flows, computed over caught spells. During a cheating spell, firms manipulate reports to soften the financing wedge: a stronger λ_π means that incremental bias translates more strongly into cheaper external finance, producing a tighter positive co-movement between m and issuance.

We measure post-detection outcomes using firm-years within five years of a detection event, the horizon over which firms in the data plausibly remain under elevated regulatory scrutiny. The flagged financing surcharge λ_f is identified by the average investment rate of firms in this five-year post-detection window. A higher λ_f raises the cost of external finance while the firm remains flagged, lowering investment relative to peers; the depth of the post-detection investment drop therefore pins down the scrutiny premium. The persistence of the flagged state q_f is identified by both the frequency and the conditional average level of positive external finance flows in the same five-year post-detection window. A higher q_f keeps firms under elevated financing costs for longer, depressing the extensive and intensive margins of post-detection issuance simultaneously.

The detection probability q_p is identified by the fraction of firm-years involving a caught event: a higher q_p mechanically raises the share of firms that are detected each period. The punishment intensity p_m is identified by the average book-value bias at the time a firm is caught. A higher p_m raises the cost of carrying a large bias stock into the flagged state, leading firms to accumulate bias more conservatively and reducing the average bias observed at detection. The quadratic adjustment cost on the book-value bias η is identified by the average duration of a cheating spell conditional on being caught: a higher η makes rapid changes in m costly, inducing firms to build up the bias gradually and thereby lengthening the expected duration of a cheating spell.

3.3 Estimation Results

Table 2 reports the estimated parameter values. Most of the production parameters take values consistent with the prior literature: idiosyncratic TFP is moderately persistent ($\rho_z = 0.655$) with an annual innovation standard deviation of $\sigma_z = 0.253$; the depreciation rate of capital is $\delta = 0.071$; and returns to scale are $\alpha = 0.518$, well below one and indicating economically meaningful curvature in the production technology so that firm-size dispersion in the simulated panel is generated by productivity dispersion rather than scale economies. The detection probability $q_p = 0.021$ implies that, conditional on actively misreporting, a firm is caught at an average rate of roughly once every 48 years; combined with the data moment that 0.8% of firm-years involve a detection event, this implies a steady-state share of actively misreporting firms of roughly 38% of unflagged firm-years. The flag-state persistence $q_f = 0.682$ implies an expected duration of regulatory scrutiny of about 3.1 years. We focus the discussion below on the financing wedge parameters $(\lambda_1, \lambda_\pi, \lambda_f)$, the pecuniary penalty p_m, and the bias adjustment-cost intensity η.

Before discussing parameter values, we assess the model's fit to the fourteen targeted moments (Table 3). The production block fits tightly: the persistence, mean drift, and innovation variance of $\log z$ are within rounding of their data counterparts, and average reported profitability is 1.7% in the model against 1.5% in the data. The cheating block also fits well: the average misreporting spell is 4.4 years (data: 4.6), the average bias at detection is 4.8% of assets (data: 5.7%), and the detection frequency is 1.0% of firm-years (data: 0.8%). The financing block matches both the frequency (30.9% vs. 29.3%) and the conditional size (44.8% vs. 47.6%) of positive external-finance flows, as well as the bias-issuance correlation of 0.13 during caught spells (data: 0.12).

Three discrepancies are worth noting. The model overstates the average investment rate (15.4% vs. 12.9%), which we attribute to abstracting from non-financing investment frictions such as managerial inattention, capacity adjustment costs, and asymmetries in disinvestment. Post-detection external-finance activity is also somewhat too high: the frequency of positive external finance is 28.8% in the model versus 22.2% in the data. This gap suggests that detection in practice triggers consequences beyond the financing-wedge surcharge λ_f, such as reputational losses with customers and suppliers, executive turnover, and governance reforms, that a single reduced-form surcharge does not capture. Finally, the model slightly undershoots the average bias at detection (4.8% vs. 5.7%) while matching spell duration closely, consistent with the convex adjustment cost smoothing year-on-year bias build-ups more aggressively than is observed in the data. None of these gaps changes the qualitative direction of the targeted moments, but they bear on the magnitudes of the counterfactual exercises in Section 4.

The financing wedge $\lambda_1 \exp\{-\lambda_\pi \hat{\pi}\}$ is the per-unit cost faced by an unflagged firm with reported profits $\hat{\pi}$. The estimated baseline intensity is $\lambda_1 = 1.40$: a truthfully-reporting firm at zero reported profits gives up \$1.40 in financing terms for every dollar raised externally. We interpret this wedge as the price of external funds faced by the firm, not as a pure resource loss. Part of the wedge may reflect real issuance costs, but the central mechanism is redistributive: when reported profitability lowers the wedge, value shifts toward the firm and away from outside investors. The sensitivity parameter is $\lambda_\pi = 1.82$, so each one-percentage-point increase in reported profitability lowers the per-unit financing cost by approximately 1.80%. At a reported profitability of 5% (roughly the typical level among firms with positive flows), the per-unit cost falls by 8.7% relative to a firm reporting zero profits; at 10% it falls by 16.6%. This sensitivity is what makes misreporting economically valuable.

The estimated scrutiny surcharge is $\lambda_f = 0.0026$, only about 0.19% of the baseline wedge $\lambda_1 = 1.40$. The surcharge alone raises a flagged firm's per-unit financing cost from 1.40 to 1.4026, a negligible direct effect. The reason flagged firms face a sizeable financing disadvantage is not the surcharge itself but the loss of the misreporting channel: a flagged firm's wedge is priced off *true* profits rather than reported profits, so the multiplicative reduction $\exp\{-\lambda_\pi \hat{\pi}\}$ that an unflagged manager could engineer is unavailable. Quantitatively, almost all of the post-detection financing penalty is generated by the prohibition on misreporting, not by the regulatory stigma surcharge.

The estimated pecuniary penalty intensity is $p_m = 8.30 \times 10^{-4}$. The implied one-time fine upon detection is $p_m m^2$, which for the mean sample firm (reported book assets of approximately \$4,443 million) corresponds to roughly 1.2% of assets, or about \$53 million.[2] The fine is economically meaningful but remains within the range of legal and regulatory penalties observed in major financial misreporting cases.

Finally, the estimated bias adjustment-cost intensity is $\eta = 3.16 \times 10^{-3}$. Recall that $\eta(\Delta m)^2$ is the dissipative resource cost a manager incurs to engineer a change in reported profits by Δm, representing side payments to accountants and auditors, fabricated supporting documentation, and the internal coordination required to sustain the bias. The convex specification implies that aggressive year-on-year build-ups are disproportionately expensive: doubling the speed at which a firm grows its bias quadruples the per-period resource cost. The estimated magnitude is consistent with the average misreporting spell of 4.6 years observed in the data, since firms find it cheaper to accumulate bias gradually than in a single year.

4 Counterfactual experiments

We use the estimated model to study how the detection and punishment environment shapes firm behavior and value. Section 4.1 reports responses to one-at-a-time parameter changes, decomposing how each instrument moves the equilibrium. Section 4.2 then asks which values of those parameters maximize cross-sectional firm value in the model.

4.1 Positive results

Table 4 reports the level of the baseline equilibrium and the percent change from baseline across nine counterfactual scenarios: halving or doubling the sensitivity of the financing wedge to reported profits ($\lambda_\pi/2$, $2\lambda_\pi$), halving or doubling the per-period detection probability ($q_p/2$, $2q_p$), removing or maximizing the persistence of the flagged state ($q_f = 0$, $q_f = 1$), removing the pecuniary penalty ($p_m = 0$), removing the reputational surcharge ($\lambda_f = 0$), and prohibiting misreporting entirely ("No Cheating"). The reported moments are firm value V, average true profits π, average investment rate i, average accumulated bias m, the realized financing wedge λ, the fraction of unflagged firms actively misreporting, and the average investor-side financing loss $\Delta\lambda$ conditional on a positive wedge.

[2]The dollar figure depends on the mapping between a unit of capital in the simulation and dollars in the data; we therefore emphasize the scaled penalty and use the dollar value only as an illustrative conversion.

The first two columns make clear that the value of the reporting channel is governed largely by λ_π. Halving the sensitivity ($\lambda_\pi/2$) cuts firm value by 4.6% even though the fraction of firms misreporting rises by 10% and the bias stock by 21%: when reported profits buy less of a financing discount, firms need more bias to obtain a given reduction in financing costs, and the firms that continue to cheat receive a smaller benefit per unit of manipulation. Doubling λ_π has the opposite effect: the fraction cheating falls by 79%, the bias stock falls by 82%, and firm value rises by 5%. The result is not that stronger reported-profit sensitivity mechanically encourages more cheating. A higher λ_π also makes truthful profitability more valuable in financing markets and raises the payoff from reporting strong true performance. In equilibrium, that direct financing benefit, together with the lower need to carry costly accumulated bias, dominates the stronger incentive to manipulate for most firms.

The detection-probability experiments ($q_p/2$, $2q_p$) show that scrutiny intensity reshuffles firms across the cheating margins but has a comparatively muted direct effect on V: halving q_p raises the fraction cheating by 20% but lowers V by only 0.45%, while doubling q_p produces nearly symmetric movements. The reason is that, conditional on cheating, the marginal benefit of an additional unit of bias is still driven by λ_π; q_p principally moves firms across the unflagged-flagged margin without changing the per-unit value of manipulation. These counterfactuals do not attach resource costs to raising detection, so the large values of q_p considered below should be read as model experiments rather than feasible enforcement prescriptions.

Flag-state persistence has a larger value impact. Setting $q_f = 0$, so detection is a one-period event with no carryover, raises the share of cheating firms by 42% and lowers firm value by 0.6% because firms internalize the punishment less. At $q_f = 1$, scrutiny is permanent: misreporting essentially disappears (the fraction cheating falls by 98%), but firm value also falls (by 1.2%) because the small fraction of firms unlucky enough to be caught are permanently barred from using the reporting channel.

The two "remove a punishment" experiments isolate the channels through which detection bites. Setting $p_m = 0$ has nearly no effect on V: most of the deterrent value of detection works through the loss of the manipulation lever, not through the direct fine. Setting $\lambda_f = 0$ also leaves V essentially unchanged but raises the share of cheating firms by 33%, reflecting the reputational surcharge's modest extensive-margin role.

Finally, the "No Cheating" column, which is equivalent to $p_m = \infty$, induces every firm to report truthfully. Firm value falls by 3.1% relative to baseline, the average investment rate drops by 7.5%, and average true profits rise by 40%. The 3.1% loss admits a decomposition against the $q_f = 1$ column. Under permanent scrutiny the option to misreport is still available, and the small fraction of firms whose state realization makes a one-shot spell worthwhile continue to use it; firm value nevertheless falls by 1.2% because firms caught in that experiment are permanently barred from using the channel. The residual 1.9 percentage points between the two columns is the option value of the misreporting channel: the contribution to average firm value from firms whose state realization (z, k, ℓ) makes the channel worth using when it is available, and which the "No Cheating" counterfactual effectively shuts off entirely. This is a statement about firm value. It does not net out investor-side financing losses, enforcement costs, or other social costs of misreporting.

4.2 Cross-sectional firm-value results

We ask which level of each enforcement parameter maximizes cross-sectional firm value in the estimated model. For each of the four parameters $(p_m, q_p, q_f, \lambda_f)$, we vary the parameter across a grid while holding the remaining structural parameters at their estimated values and re-solve the model. Figure 6 plots the equilibrium fraction of unflagged firms actively misreporting at each grid point; Figure 7 plots the corresponding equilibrium firm value, with the firm-value-maximizing point marked by the vertical red dashed line. These exercises vary one parameter at a time and do not solve a regulator's problem. In particular, they exclude the resource cost of enforcement, investor welfare, and any broader social cost of distorted reports.

Three of the four panels in Figure 7 display an interior maximum for cross-sectional firm value: V peaks at $q_p \approx 0.75$, $q_f \approx 0.85$, and $\lambda_f \approx 0.075$, and declines as the enforcement parameter is raised further. The fourth panel does not: the firm-value-maximizing pecuniary penalty is $p_m = 0$. The high value of q_p at the peak should be read in light of the narrow objective: the model captures how detection reshapes firm behavior, but not the resources required to generate such a detection probability.

The interior maxima are best read as properties of the cross-sectional value measure rather than as Pareto improvements available to individual firms. By construction, $V(\theta) = \mathbb{E}_{\mu(\theta)}[v(x;\theta)]$ averages the firm value function over the stationary distribution induced by the vector of enforcement parameters θ. At every state x, $v(x;\theta)$ is weakly decreasing in each enforcement instrument: a fixed firm always weakly prefers weaker enforcement because the firm's choice set expands. The non-monotonicity in V comes from the second moving piece. As q_p, q_f, or λ_f rises, the equilibrium policy functions shift toward lower m (Figure 6), and the stationary measure $\mu(\theta)$ shifts mass away from states with positive m_{-1} and toward $m_{-1} = 0$. Because v is decreasing in m_{-1}, the cross-section becomes weighted on states with higher firm value, raising the average over a range of θ. Eventually the pointwise decline in $v(x;\theta)$ dominates this distributional gain and the average falls; the peak of $V(\theta)$ marks the crossover. The natural reading of an interior maximum is the ex-ante value of being drawn at random from the stationary cross-section: under weak enforcement, the cross-section is heavily weighted on firms carrying accumulated bias, whose value functions are dragged down by costly bias maintenance and exposure to detection, even though each such firm chose voluntarily to enter that state.

The pecuniary penalty yields no interior maximum because, among the four instruments, it extracts the most cash from firms per unit of deterrence on the bias-accumulation margin. All four instruments work through the same ex-ante channel: they raise the expected cost of carrying bias into detection, leading firms to accumulate less m. Their incidence on firm cash flow differs. The penalty p_m is levied on the predetermined bias stock m_{-1}^2 at the moment of detection, so the firm has a single margin of response, through its choice of m in the prior unflagged period; once flagged with bias m_{-1}, it pays $p_m m_{-1}^2$ irrespective of any contemporaneous choice. The surcharge λ_f enters the per-unit cost of issuance in the flagged state, so the firm responds on two margins: ex ante through m, and ex post by scaling down $|e|$ once flagged. The transition-probability instruments q_p and q_f reshape continuation values without removing cash from the firm at any state. Cash extracted per unit of deterrence on the bias margin therefore ranks $p_m > \lambda_f > q_p, q_f = 0$, while the cross-sectional distributional shift toward $m_{-1} = 0$ is qualitatively similar across the four. The pointwise pecuniary drag on v from p_m is large enough to dominate the distributional gain at every positive value, so the maximum sits at zero.

5 Conclusion

We develop and estimate a dynamic heterogeneous-firm model in which managers may overstate reported profits to soften an external-financing wedge that depends on reported profitability, at the risk of detection and the loss of the manipulation lever for the duration of regulatory scrutiny. We discipline the model by fourteen targeted moments drawn from CRSP/Compustat and Audit Analytics, covering investment, external-finance flows, the frequency and magnitude of detection events, and the within-firm dynamics in the five years following a detection event.

The estimated sensitivity of the financing wedge to reported profits is large: each one-percentage-point increase in reported profitability cuts the per-unit cost of external finance by approximately 1.8%. The directly punitive components of detection, the pecuniary penalty and the reputational surcharge, are quantitatively small. As a consequence, almost the entire post-detection financing disadvantage of caught firms is generated by the prohibition on misreporting itself: a flagged firm's wedge is priced off true rather than reported profits, eliminating the multiplicative reduction in the per-unit financing cost that an unflagged manager could otherwise engineer. Cross-sectionally, investor-side financing losses $\Delta\lambda$ are concentrated among low-capital, low-liquidity, high-productivity firms, whose external-finance needs are largest relative to their reported earnings base.

The cross-sectional firm-value exercises identify interior maxima for three enforcement parameters: the detection probability, the persistence of the flagged state, and the reputational surcharge. The pecuniary penalty admits a corner solution at zero. These maxima are not social-welfare prescriptions. They hold the rest of the model fixed and omit enforcement costs, investor welfare, and other social costs of misreporting. The interior maxima reflect how each parameter moves the stationary distribution of firms. At every individual state, the firm's value function is weakly decreasing in each enforcement instrument; what changes with enforcement is the cross-sectional mass on states with positive accumulated bias. Stronger enforcement shifts that mass toward $m_{-1} = 0$, where value functions are higher because firms are not bearing the cost of carrying or unwinding accumulated bias, and the cross-sectional average V rises over a range of the instrument before the pointwise decline in v eventually takes over. Outright prohibition reduces firm value by 3.1% relative to baseline; about 1.9 percentage points of this loss reflect the option value of the channel for firms whose state realization makes it worth using, and the remaining 1.2 percentage points reflect the value loss for firms that continue to use the channel under permanent regulatory scrutiny and are caught. The pecuniary penalty does not yield an interior maximum because, among the four instruments, it extracts the most cash from firms per unit of deterrence: levied on the predetermined bias stock at detection, it leaves the firm no margin of response beyond its ex-ante choice of m, while the surcharge λ_f is partly absorbed by the firm's adjustment of issuance once flagged and the transition-probability instruments q_p and q_f extract no cash at any state.

Several extensions seem natural. One is to embed the firm-level decision in a setting where the financing-wedge schedule depends on the aggregate intensity of reporting bias, so that each firm's choice feeds back into the wedge faced by others. A second is to expand the objective beyond firm value by adding investor-side losses, enforcement resource costs, and other social costs of distorted reports. A third is to allow detection to affect real relationships with customers, suppliers, employees, or managers, not only financing terms. Finally, lenders could learn about firm-specific reporting credibility from public disclosures and restatements, making financing terms depend on both reported profitability and beliefs about manipulation.

References

Alawadhi, A., Karpoff, J. M., Koski, J. L., and Martin, G. S. (2025). The prevalence and costs of undetected financial misrepresentation. SSRN working paper.

Begenau, J. and Salomao, J. (2019). Firm financing over the business cycle. *The Review of Financial Studies*, 32(4):1235–1274.

Benmelech, E., Kandel, E., and Veronesi, P. (2010). Stock-based compensation and CEO (Dis)incentives. *Quarterly Journal of Economics*, 125(4):1769–1820.

Bertomeu, J., Marinovic, I., Terry, S. J., and Varas, F. (2022). The dynamics of concealment. *Journal of Financial Economics*, 143(1):227–246.

Bolton, P., Chen, H., and Wang, N. (2011). A unified theory of Tobin's q, corporate investment, financing, and risk management. *Journal of Finance*, 66(5):1545–1578.

Drechsel, T. (2023). Earnings-based borrowing constraints and macroeconomic fluctuations. *American Economic Journal: Macroeconomics*, 15(2):1–34.

Dyck, A., Morse, A., and Zingales, L. (2010). Who blows the whistle on corporate fraud? *The Journal of Finance*, 65(6):2213–2253.

Eisfeldt, A. L. and Muir, T. (2016). Aggregate external financing and savings waves. *Journal of Monetary Economics*, 84:116–133.

Gao, X., Whited, T. M., and Zhang, N. (2021). Corporate money demand. *The Review of Financial Studies*, 34(4):1834–1866.

Graham, J. R., Li, S., and Qiu, J. (2008). Corporate misreporting and bank loan contracting. *Journal of Financial Economics*, 89(1):44–61.

Han, C. and Phillips, P. C. B. (2010). Gmm estimation for dynamic panels with fixed effects and strong instruments at unity. *Econometric Theory*, 26(1):119–151.

Hennessy, C. A. and Whited, T. M. (2007). How costly is external financing? evidence from a structural estimation. *The Journal of Finance*, 62(4):1705–1745.

İmrohoroğlu, A. and Tüzel, S. (2014). Firm-level productivity, risk, and return. *Management Science*, 60(8):2073–2090.

Karpoff, J. M., Lee, D. S., and Martin, G. S. (2008). The cost to firms of cooking the books. *Journal of Financial and Quantitative Analysis*, 43(3):581–611.

Kedia, S. and Philippon, T. (2009). The economics of fraudulent accounting. *The Review of Financial Studies*, 22(6):2169–2199.

Lian, C. and Ma, Y. (2021). Anatomy of corporate borrowing constraints. *The Quarterly Journal of Economics*, 136(1):229–291.

McKeon, S. B. (2015). Employee option exercise and equity issuance motives. Working paper.

Nikolov, B. and Whited, T. M. (2014). Agency conflicts and cash: Estimates from a dynamic model. *Journal of Finance*, 69(5):1883–1921.

Povel, P., Singh, R., and Winton, A. (2007). Booms, busts, and fraud. *Review of Financial Studies*, 20(4):1219–1254.

Riddick, L. A. and Whited, T. M. (2009). The corporate propensity to save. *The Journal of Finance*, 64(4):1729–1766.

Terry, S. J., Whited, T. M., and Zakolyukina, A. A. (2023). Information versus investment. *The Review of Financial Studies*, 36(3):1148–1191. Advance Access published 2022-08-09.

Zakolyukina, A. A. (2018). How common are intentional gaap violations? estimates from a dynamic model. *Journal of Accounting Research*, 56(1):5–44.

Tables

Table 1: Summary statistics

	N	Mean	SD	p25	p50	p75
Panel A: Firm size						
Book assets (millions)	40,311	4442.975	12427.619	161.003	689.723	2751.401
Sales (millions)	40,311	3882.096	10903.203	128.506	628.516	2486.162
Gross PP&E (millions)	40,130	2045.397	6200.442	37.499	216.497	1097.000
Panel B: Firm characteristics						
Reported profitability ($\hat{\pi}$/assets)	39,053	0.015	0.371	-0.040	0.063	0.168
Investment rate (i/k)	40,070	0.129	0.110	0.058	0.097	0.163
Reported liquidity ($\hat{\ell}$/assets)	39,052	0.372	0.389	0.079	0.249	0.555
Panel C: External finance						
Net external finance / assets	39,053	-0.000	0.436	-0.149	-0.043	0.012
Net equity issuance / assets	39,053	-0.005	0.305	-0.080	-0.010	0.000
Net debt issuance / assets	39,053	0.002	0.280	-0.064	-0.006	0.013
Pr(positive external finance)	40,311	0.293	0.455	0.000	0.000	1.000
Positive external finance / assets	10,548	0.476	0.777	0.058	0.191	0.552
Panel D: Misreporting (firm-year)						
Pr(caught misreporting)	40,311	0.008	0.091	0.000	0.000	0.000
Book-value bias (m/assets)	1,128	0.051	0.102	0.006	0.018	0.063
Earnings bias (Δm/assets)	1,128	0.013	0.035	0.000	0.002	0.011
Panel E: Detection events						
Book-value bias at detection	333	0.057	0.141	0.006	0.018	0.060
Years cheating (when caught)	334	4.626	1.360	4.000	5.000	5.000
Observations	40311					

Notes: The sample covers U.S. public firms in CRSP/Compustat from 2003 to 2019, excluding regulated utilities (SIC 4900–4999), financial firms (SIC 6000–6999), and quasi-governmental and nonprofit firms (SIC 9000–9999). Firms are required to have at least seven years of data, of which three must be consecutive. Panel A levels are reported in 2019 millions of U.S. dollars; firm age is measured as years since first appearance in the panel. In Panel B, reported profitability, reported liquidity, and the book-value and earnings biases are scaled by reported assets, defined as physical capital plus reported lagged liquidity; investment is scaled by physical capital. In Panel C, net external finance equals minus the sum of cash flows to debt and cash flows to equity (Eisfeldt and Muir, 2016); net equity issuance and net debt issuance correspond to the two components, both scaled by reported assets. All continuous variables are winsorized by year at the 1st and 99th percentiles. Panel E is computed over firm-years in which a misreporting firm is caught.

Table 2: Estimated and externally calibrated parameters

Parameter	Description	Value
Panel A: Externally calibrated		
r	Discount rate	0.015
r_ℓ	Liquidity return	0.012
Panel B: Internally estimated — Production		
ρ_z	Persistence of log TFP	0.655
μ_z	Unconditional log-mean of TFP	−0.294
σ_z	Innovation std. of log TFP	0.253
α	Returns to scale	0.518
δ	Capital depreciation rate	0.071
Panel C: Internally estimated — Financing		
λ_1	Baseline per-unit financing wedge	1.400
λ_π	Reported-profit sensitivity of wedge	1.816
λ_f	Flagged reputational surcharge	2.64×10^{-3}
Panel D: Internally estimated — Flag and cheating		
q_f	Persistence of flagged state	0.682
q_p	Per-period detection probability	0.021
p_m	Pecuniary penalty intensity	8.30×10^{-4}
η	Bias adjustment-cost intensity	3.16×10^{-3}

Notes: Panel A reports externally calibrated parameters: the discount rate r is set to the sample-period average annual T-bill rate, and the liquidity return r_ℓ is set to $0.8 \times r$ following Gao et al. (2021). Panels B–D report the twelve internally estimated parameters, obtained by simulated method of moments using the fourteen targeted moments listed in Table 3. Identification logic is discussed in Section 3.2.

Table 3: Targeted moments: data and model

Moment	Data	Model
Panel A: Production		
Autocorrelation of $\ln z$	0.659	0.655
Mean drift of $\ln z$	-0.106	-0.101
Var. residuals from AR(1) of $\ln z$	0.066	0.069
Average reported profitability	0.015	0.017
Average investment rate	0.129	0.154
Panel B: Financing		
Frequency of EF > 0	0.293	0.309
Mean EF $\mid$ EF > 0	0.476	0.448
Corr(m, EF > 0 $\mid$ caught spells)	0.121	0.130
Panel C: Flag dynamics		
Average investment rate $\mid$ post-detection	0.118	0.131
Frequency of EF > 0 $\mid$ post-detection	0.222	0.288
Mean EF $\mid$ EF > 0 $\mid$ post-detection	0.333	0.372
Panel D: Cheating		
Fraction of firms caught cheating	0.008	0.010
Average bias m at detection	0.057	0.048
Average years of misreporting spell	4.626	4.376

Notes: The table lists the fourteen moments used to estimate the twelve parameters in Panels B–D of Table 2. Data moments are computed over the 2003–2019 CRSP/Compustat sample described in Section 3.1; model moments are the corresponding simulated values at the parameter estimates of Table 2. The three Panel A moments on $\ln z$ (autocorrelation, mean drift, and residual variance) are computed using the panel approach of Han and Phillips (2010), which accommodates intercept heterogeneity across firms. "EF" denotes net external finance (Equation (6)); profitability, the bias stock, and EF are scaled by reported assets $\hat{a} \equiv k + \hat{\ell}_{-1}$, while investment is scaled by the physical capital stock. "Post-detection" restricts to firm-years within five years of a detection event; "caught spells" restricts to firm-years inside an eventually-detected misreporting episode. Appendix Table A.1 summarizes the moment-to-parameter mapping discussed in Section 3.2.

Table 4: Model moments under counterfactual scenarios

	Baseline	% Change from Baseline								
		$\lambda_\pi/2$	$2\lambda_\pi$	$q_p/2$	$2q_p$	q_f=0	q_f=1	p_m=0	λ_f=0	No Cheating
V	6.85	−4.63	5.04	−0.45	0.44	−0.62	−1.19	−0.10	−0.50	−3.10
π	0.016,6	181.03	−70.99	−3.00	2.82	−3.00	164.33	0.21	−4.32	39.56
i	0.154,2	−17.18	0.56	0.19	−0.20	0.30	−16.32	−0.02	0.34	−7.48
m	0.023,3	21.08	−81.63	18.21	−19.15	29.73	−96.79	4.42	18.48	−100.00
λ	0.000,2	175.68	−81.78	−5.42	9.90	−10.94	383.02	−0.38	−9.03	68.89
Frac. Cheating	0.485,1	10.17	−78.61	19.92	−29.54	41.90	−98.26	2.49	33.48	−100.00
$\Delta\lambda \mid \lambda > 0$	0.639,6	25.73	−70.41	8.81	−12.48	16.82	−95.37	1.36	11.87	−100.00

Notes: The table reports the level of each moment in the baseline equilibrium (column 1) and the percent change relative to baseline under nine counterfactual scenarios. The counterfactuals vary one parameter at a time while holding the remaining structural parameters at their estimated values in Table 2: $\lambda_\pi/2$ and $2\lambda_\pi$ halve or double the financing-wedge sensitivity to reported profits; $q_p/2$ and $2q_p$ halve or double the per-period detection probability; $q_f = 0$ removes flag-state persistence and $q_f = 1$ makes scrutiny permanent; $p_m = 0$ removes the pecuniary penalty; $\lambda_f = 0$ removes the reputational surcharge on flagged firms; "No Cheating" restricts all firms to $m_t = 0$. Reported moments: V is average firm value; π, i, m, and λ are averages of true profits, the investment rate, the bias stock, and the realized financing wedge; "Frac. Cheating" is the share of unflagged firms with $m_t > 0$; and $\Delta\lambda \mid \lambda > 0$ is the average investor-side financing loss conditional on a positive wedge.

Figures

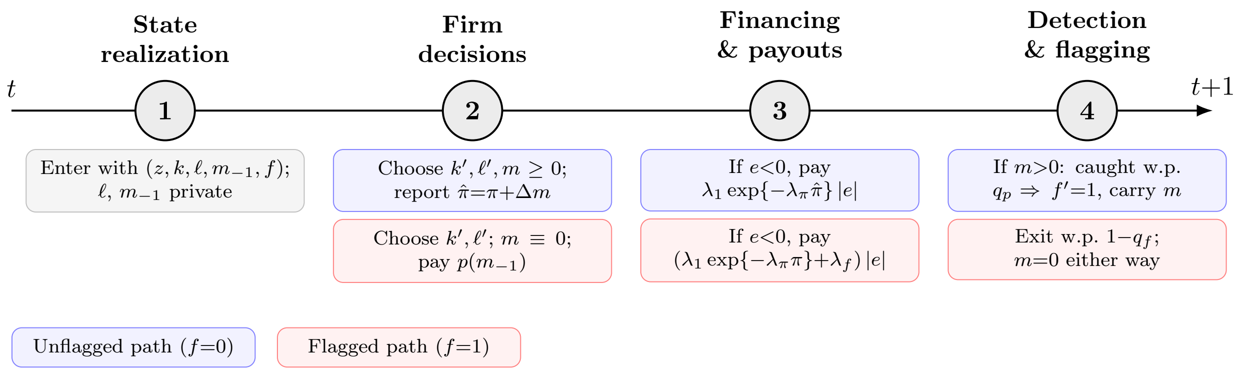

Figure 1: Within-period timing

Notes: The four circles trace the within-period sequence described in Section 2.1.4. Step 1 is common to all firms; in steps 2–4 the firm's actions and the transitions it faces depend on its flag status, so each step splits into an unflagged path (blue) and a flagged path (red). Hatted quantities denote reported values; unhatted quantities denote true values.

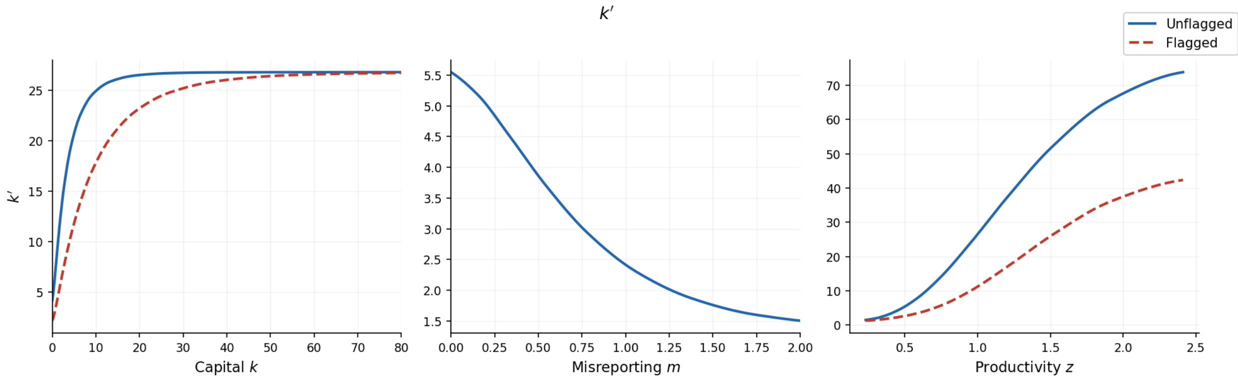

Figure 2: Capital policy k' across one-dimensional state variations

Notes: Each panel plots the next-period capital choice k' as one state variable is varied while the remaining states are held at their unconditional means. The solid blue (dashed red) line denotes unflagged (flagged) firms. The middle panel is plotted only for unflagged firms because $m \equiv 0$ for flagged firms. Structural parameters are at the estimates in Table 2.

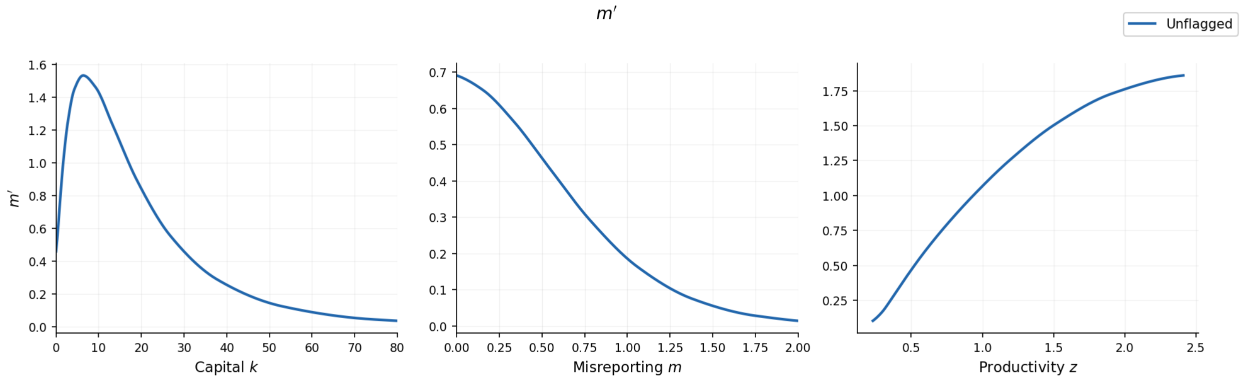

Figure 3: Misreporting policy m across one-dimensional state variations

Notes: Each panel plots the bias choice m for an unflagged firm as one state variable is varied while the remaining states are held at their unconditional means. Flagged firms are constrained to $m = 0$ and are therefore omitted. Structural parameters are at the estimates in Table 2.

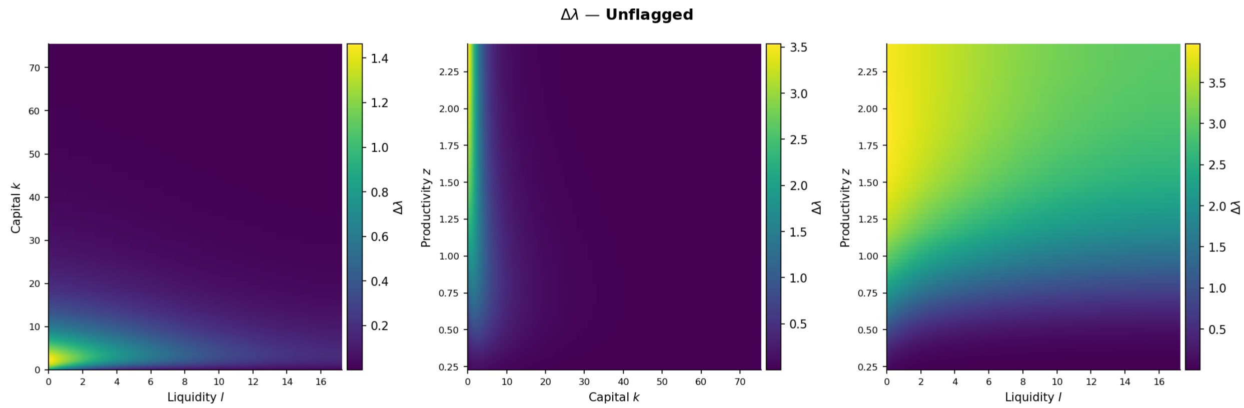

Figure 4: Investor financing-cost loss $\Delta\lambda$ across pairs of state variables, unflagged firms

Notes: Each panel plots a heatmap of $\Delta\lambda$, defined as the difference between the external-finance cost the firm actually pays and the cost the firm would have paid if reported profits coincided with true profits, over one pair of state variables holding the remaining state at its unconditional mean. Brighter colors denote larger investor-side financing losses. Panels show, from left to right, the (ℓ, k), (k, z), and (ℓ, z) planes. Structural parameters are at the estimates in Table 2.

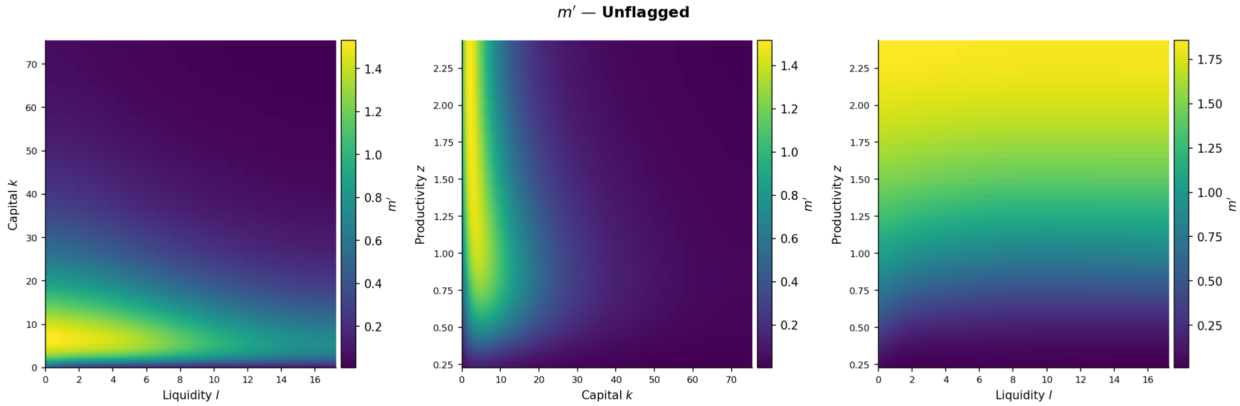

Figure 5: Misreporting policy m across pairs of state variables, unflagged firms

Notes: Each panel plots a heatmap of the bias choice m for an unflagged firm over one pair of state variables, with the remaining state held at its unconditional mean. Brighter colors denote larger bias accumulation. Panels show, from left to right, the (ℓ, k), (k, z), and (ℓ, z) planes. Structural parameters are at the estimates in Table 2.

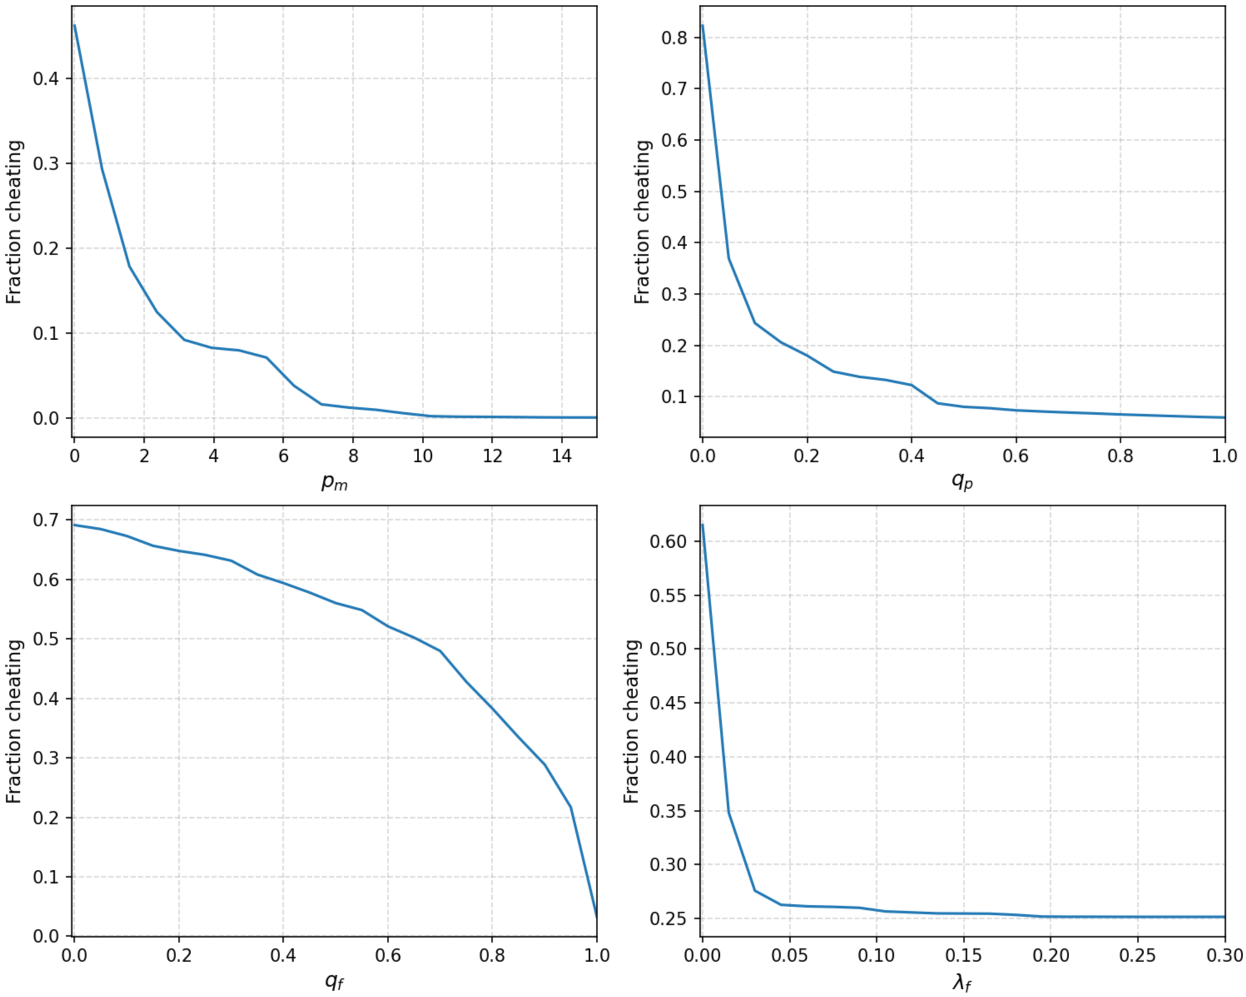

Figure 6: Equilibrium fraction of unflagged firms misreporting as enforcement parameters vary

Notes: Each panel varies one of the four enforcement parameters $(p_m, q_p, q_f, \lambda_f)$ across a grid while holding the remaining structural parameters at their estimated values in Table 2, and plots the equilibrium share of unflagged firms with $m_t > 0$ at each grid point.

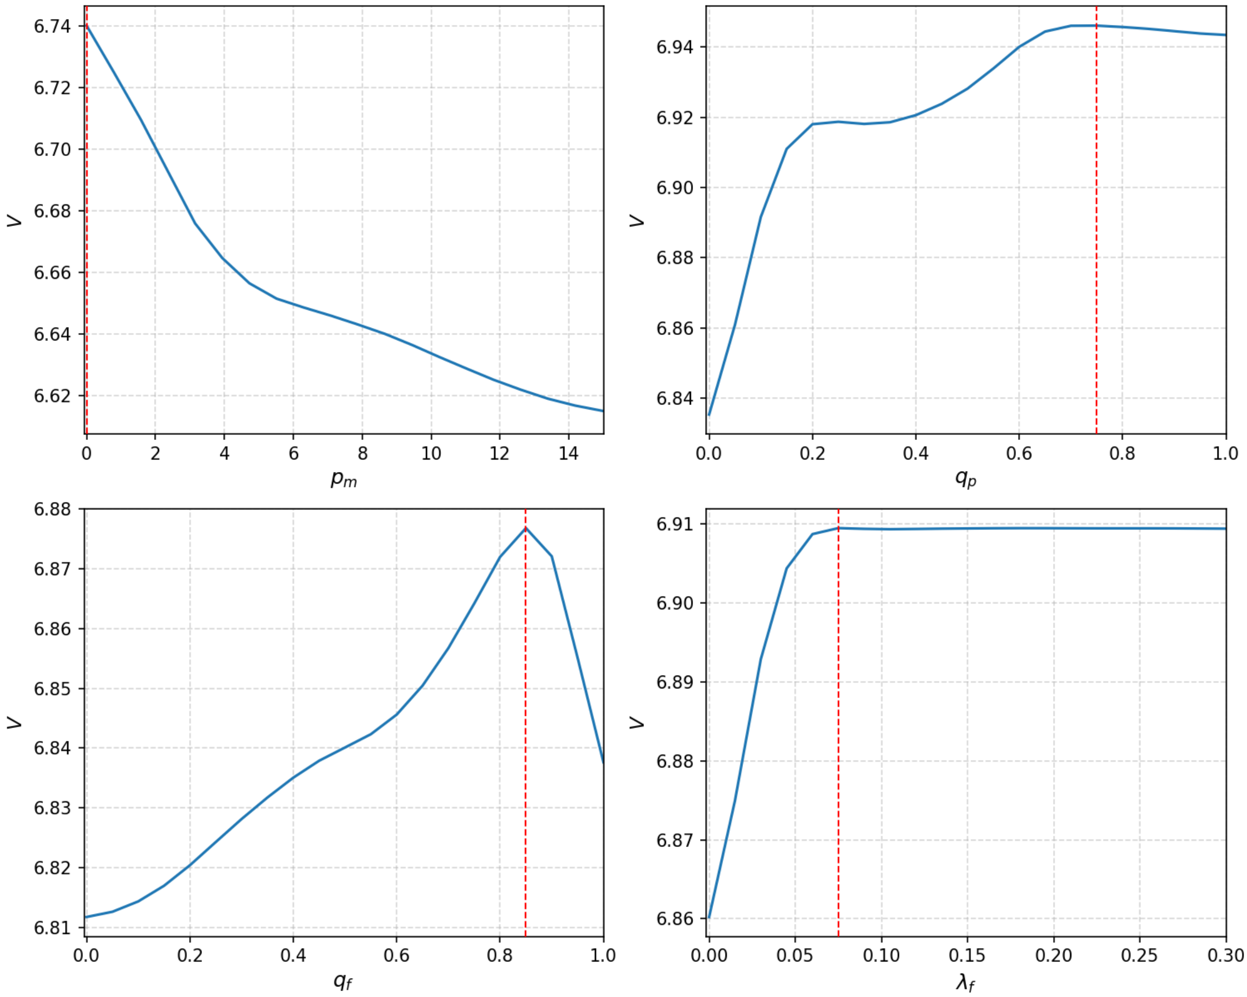

Figure 7: Equilibrium firm value as enforcement parameters vary

Notes: Each panel varies one of the four enforcement parameters $(p_m, q_p, q_f, \lambda_f)$ across a grid while holding the remaining structural parameters at their estimated values in Table 2, and plots the resulting equilibrium average firm value V at each grid point. The vertical red dashed line marks the firm-value-maximizing grid point in each panel.

Appendix

Table A.1: Identification summary

Parameter	Description	Identifying moment(s)
Production		
ρ_z	Persistence of log TFP	Autocorrelation of $\ln z$
μ_z	Unconditional log-mean of TFP	Mean drift of $\ln z$
σ_z	Innovation std. of log TFP	Variance of residuals from AR(1) of $\ln z$
α	Returns to scale	Average reported profitability
δ	Capital depreciation rate	Average investment rate
Financing		
λ_1	Baseline per-unit financing wedge	Frequency of EF > 0 and mean EF $\mid$ EF > 0
λ_π	Reported-profit sensitivity of wedge	Corr. of m and EF > 0 during caught spells
Flag dynamics		
λ_f	Flagged reputational surcharge	Average investment rate post-detection
q_f	Persistence of flagged state	Frequency and mean of EF $>$ 0 post-detection
Cheating		
q_p	Per-period detection probability	Fraction of firm-years caught
p_m	Pecuniary penalty intensity	Average bias m at detection
η	Bias adjustment-cost intensity	Average duration of misreporting spell

Notes: The table summarizes which moment is most informative about each parameter at the margin. All twelve parameters jointly affect every simulated moment in equilibrium; see Section 3.2 for discussion of the identifying logic. "EF" denotes net external finance; "post-detection" restricts to firm-years within five years of a detection event; "caught spells" restricts to firm-years inside an eventually-detected misreporting episode.