

Dyck, Alexander; Morse, Adair; Zingales, Luigi

Working Paper

How pervasive is corporate fraud?

New Working Paper Series, No. 327

Provided in Cooperation with:

George J. Stigler Center for the Study of the Economy and the State, The University of Chicago
Booth School of Business

Suggested Citation: Dyck, Alexander; Morse, Adair; Zingales, Luigi (2023) : How pervasive is corporate fraud?, New Working Paper Series, No. 327, University of Chicago Booth School of Business, Stigler Center for the Study of the Economy and the State, Chicago, IL

This Version is available at:

<https://hdl.handle.net/10419/268480>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



Stigler Center

for the Study of the Economy and the State

How pervasive is corporate fraud?

Alexander Dyck

University of Toronto

Adair Morse

University of California at Berkeley & NBER

Luigi Zingales

University of Chicago & NBER & CEPR

January 2023

New Working Paper Series No. #327

Abstract:

We provide a lower-bound estimate of the undetected share of corporate fraud. To identify the hidden part of the “iceberg,” we exploit Arthur Andersen’s demise, which triggered added scrutiny on Arthur Andersen’s former clients and thereby increased the detection likelihood of preexisting frauds. Our evidence suggests that in normal times only one-third of corporate frauds are detected. We estimate that on average 10% of large publicly traded firms are committing securities fraud every year, with a 95% confidence interval of 7%-14%. Combining fraud pervasiveness with existing estimates of the costs of detected and undetected fraud, we estimate that corporate fraud destroys 1.6% of equity value each year, equal to \$830 billion in 2021.

Stigler Center for the Study of the Economy and the State

University of Chicago Booth School of Business

5807 S Woodlawn Ave

Chicago, IL 60637



How pervasive is corporate fraud?

Alexander Dyck¹ · Adair Morse^{2,3} · Luigi Zingales^{3,4,5}

Accepted: 9 November 2022
© The Author(s) 2023

Abstract

We provide a lower-bound estimate of the undetected share of corporate fraud. To identify the hidden part of the “iceberg,” we exploit Arthur Andersen’s demise, which triggered added scrutiny on Arthur Andersen’s former clients and thereby increased the detection likelihood of preexisting frauds. Our evidence suggests that in normal times only one-third of corporate frauds are detected. We estimate that on average 10% of large publicly traded firms are committing securities fraud every year, with a 95% confidence interval of 7%–14%. Combining fraud pervasiveness with existing estimates of the costs of detected and undetected fraud, we estimate that corporate fraud destroys 1.6% of equity value each year, equal to \$830 billion in 2021.

Keywords Corporate governance · Corporate fraud · Detection likelihood · Cost–benefit analysis · Securities regulation · Arthur Andersen

JEL Classification G30 · G34 · K22 · M40

1 Introduction

Starting with Jensen and Meckling (1976), a very large literature documents the agency costs of public ownership. One of the less emphasized of these costs is fraud. To enrich themselves, managers who own a relatively small fraction of the stock might be willing to break the law, even when the cost of breaking the law from the company’s point of view exceeds its benefits. Is corporate fraud a major component

✉ Alexander Dyck
Alexander.Dyck@rotman.utoronto.ca

¹ University of Toronto, Toronto, ON, Canada

² University of California at Berkeley, Berkeley, CA, USA

³ NBER, Cambridge, MA, USA

⁴ University of Chicago, Chicago, IL, USA

⁵ CEPR, London, UK

of the agency costs of public ownership? Should regulation do something to reduce this cost?

Before we even try to answer these questions, we need to establish how pervasive corporate fraud is. Is the fraud we observe the whole iceberg or just its visible tip? The answer to this question requires an estimate of the ratio of the exposed tip to the submerged portion, also known as the detection likelihood. Thus far in the literature, approaches to estimating the detection likelihood (and by implication the hidden prevalence of fraud) have included (i) accounting misconduct prediction models (Beneish 1999; Dechow et al. 2011), (ii) corporate surveys (Dichev et al. 2013), and (iii) structural, partial-observability approaches (Wang 2013; Wang et al. 2010; Hahn et al. 2016; Zakolyukina 2018). In this paper, we introduce a fourth approach, which relies on a natural experiment: the 2001 demise of Arthur Andersen (AA).

The simple idea is that after the AA demise, former AA clients were subject to vastly increased scrutiny. They found themselves in the spotlight of the media, investment intermediaries, short-sellers, and their internal gatekeepers. In addition, they were forced to seek a different audit firm. Given the extreme cloud of suspicion that was covering AA clients immediately after the Enron scandal exploded (Chaney and Philipich 2002; Krishnamurthy et al. 2006), the new auditors, as well as all other fraud detectors, had strong incentives to uncover any fraud committed by former AA clients. Even if this increased scrutiny does not reveal all the existing fraud, the Kolmogorov axiom of conditional probability allows us to derive an upper-bound (and thereby conservative) estimate of the detection likelihood, which in turn provides us with a lower-bound estimate of the pervasiveness of corporate fraud.

As we explain in Section 2, we use the term “fraud” loosely, since what we measure is some form of misconduct or alleged fraud. To estimate the detection likelihood, we use several different measures. The first measure is auditor-detected securities fraud from Dyck et al. (2010) (hereafter “DMZ”). The second is SEC Accounting and Auditing Enforcement Releases (AAERs) from Dechow et al. (2011). The third is financial misreporting not due to simple clerical errors. The last one is all securities fraud under SEC section 10b-5.

We find that the detection likelihood varies depending on the measure used: 0.29 for auditor-detected securities fraud, 0.34 for accounting restatements, 0.47 for securities fraud at large, and 0.52 for AAERs. All of the estimates are statistically significant and robust to a host of potential omitted differences between the firms audited by Arthur Andersen and the rest (e.g., industry focus, regional variation, and importance of IPOs). Our estimates represent an upper bound of the detection likelihood, which implies that our estimates of unobserved fraud are conservatively low. The best estimate of the detection likelihood is 0.33, with a 95% confidence interval between 0.24 and 0.45.¹ Fraud is indeed like an iceberg with significant undetected fraud beneath the surface.

¹ As explained in Sect. 3, the detection likelihood is based on the ratio of two observables: the emergence of fraud in non-AA firms and the emergence of fraud in former AA firms. This experimental design has the most power for auditor-detected securities fraud and restatements, and our best estimate is the average of these measures.

With our estimates of the detection likelihood, we can calculate the pervasiveness of corporate fraud as a function of the definition of fraud we adopt. Accounting violations are widespread: in an average year, 41% of companies misrepresent their financial reports, even when we ignore simple clerical errors. Fortunately, securities fraud is less pervasive. In an average year, 10% of all large public corporations commit (alleged) securities fraud, with a 95% confidence interval between 7 and 14%.

Having estimated the pervasiveness of fraud, we can calculate the losses produced by fraud using cost estimates from the prior literature. For detected frauds, we use the decline in equity value at the revelation of the fraud as estimated by Karpoff et al. (2008) (henceforth “KLM”). For undetected frauds, we apply the average value lost due to cover-up, as estimated by Beneish et al. (2013). Combining these detected and undetected fraud losses – 25% and 11% of the equity value, respectively – to the proportion of frauds detected (0.33) and undetected (0.67) results in an average fraud cost of 16% of equity value. Putting together our estimate of the pervasiveness of fraud (10%) with the per-firm cost of fraud (16%), we arrive to an expected annual cost of fraud of 1.6% of the total equity value of US public firms (with a 95% interval between 1.2% and 2.2%). Hence, the annual cost of fraud among US corporations at the end of our sample period (2004) is \$254 billion or – bringing our cost calculation forward to 2021 – \$830 billion. If we compare the 2004 expected cost of fraud with the \$19.9 billion of annual SOX compliance cost (as estimated by Hochberg et al. (2009)), for the benefits of SOX to exceed its costs, SOX would have to reduce the probability of fraud initiation by 1 percentage point (equal to 10% of the baseline probability).

Our paper builds on a large literature seeking to understand the incidence of undetected corporate fraud or misconduct. Beneish (1999) and Dechow et al. (2011) estimate the likelihood of accounting misconduct through patterns in financial disclosure. Wang (2013) uses a bivariate probit to identify factors affecting the propensity to commit fraud and the vulnerability of fraud to detection. Hahn et al. (2016) integrate Wang’s bivariate probit with a flexible Bayesian prior. Finally, Zakolyukina (2018) uses a structural model of earnings manipulation. Instead of relying on statistical or economic models, our approach is based on a natural experiment, with all the pluses and minuses of natural experiments.

Our paper complements Krishnan et al. (2007), who document how after AA’s demise Big Four auditors were more likely to issue a going concern qualification when they audited large former AA clients than when they audited similar non-AA clients. They do not find a similar result for small firms, likely due to a greater rejection of small AA clients with dubious accounting by Big Four auditors. We find that when we look at all auditors (not just Big Four), the revelation of misconduct is more likely in small former AA clients than in similar-sized non-AA clients.

Finally, our paper also builds off the literature seeking to understand the costs associated with corporate fraud, both in direct reputation effects (Karpoff and Lott 1993; KLM) and in cover-up costs of avoidance (Beneish et al. 2013).

The rest of the paper proceeds as follows. Section 2 describes the data and presents summary statistics on caught frauds. Section 3 explains our methodology. Section 4 presents the detection likelihood and the pervasiveness of corporate fraud results. Section 5 provides estimates of the cost of such fraud, and Sect. 6 concludes.

2 Observed corporate fraud measures and incidence

We start by defining corporate fraud. As one court observed, “The law does not define fraud; it needs no definition; it is as old as falsehood and as versatile as human ingenuity.”² In a similar vein, the Fourth Circuit noted that “fraud is a broad term, which includes false representations, dishonesty and deceit.”³ Thus, it is not easy to define corporate fraud precisely. Securities law defines securities fraud as “to make any untrue statement of a material fact or to omit to state a material fact necessary in order to make the statements made, in the light of the circumstances under which they were made, not misleading.”⁴ Since our primary measure derives from securities class actions, we assume that definition and regard corporate fraud primarily as misrepresentation.

In spite of this ambiguity, the established legal practice requires the presence of three elements to label an act of misconduct fraud: misrepresentation, materiality, and intent. As we will explain below, all our measures of misconduct contain an element of misrepresentation. We do our best to set a minimum threshold to ensure materiality. Where we cannot deliver is on intent. Intent can only be proven in court, and all these alleged fraud cases are settled before they reach the final verdict, since directors’ and officers’ insurance does not indemnify executives if they are convicted in court (see, for example, the data and discussion in Black et al. (2006)). As a result, all the corporate fraud we discuss is *alleged* fraud that was settled out of court.

To capture alleged fraud we resort to four different measures. The first measure is financial misrepresentations exposed by auditors. To construct “auditor-detected fraud,” we start with SEC *Rule 10b-5* securities fraud from DMZ’s dataset of class actions suits. To ensure comprehensiveness, DMZ focus on large firms (those with over \$750 million in assets) so that a potential payoff from a suit would be sufficiently attractive to mobilize monitoring by legal actors.⁵ To remove frivolous suits, DMZ impose a series of filters, building on the existing legal literature.⁶ The final step is to restrict the DMZ sample to frauds that are revealed directly or indirectly by an auditor. Directly revealed by the auditors are frauds where the auditor in DMZ was the whistleblower. Indirectly revealed by the auditor are frauds where the auditor issued a qualified opinion and the whistleblower in DMZ is either the company itself or an analyst. For these latter cases, we reread the auditor’s qualifying opinion in the DMZ cases where the company itself or an analyst was the whistleblower and dropped two cases where the fraud could not be plausibly related to the reason why the audit opinion was qualified.

² Weiss v. United States, 122 F.2d 675, 681 (5th Cir. 1941), cert. denied, 314 U.S. 687 (1941).

³ See United States v. Grainger, 701 F.2d 308, 311 (4th Cir. 1983), cert. denied, 461 U.S. 947 (1983).

⁴ <https://ecfr.federalregister.gov/current/title-17/chapter-II/part-240/subpart-A/subject-group-ECFR271f57d9adc8b47/section-240.10b-5>.

⁵ The legal industry has set up an automatic process whereby every time a stock price experiences a large drop, specialist attorneys file suits and scour financial reports for misrepresentations. Large publicly traded companies offer lucrative possibilities for claims on a suit. For these, it is unlikely that the class action legal structure would miss an opportunity to detect a securities fraud (Coffee 1986).

⁶ As described in Appendix 1, DMZ restrict attention first to cases after a 1995 change in the law which forced courts to become more stringent about evidence to certify a case; second to cases that are not dismissed; and third to cases without low settlement amounts (below \$3 million).

The second measure is financial misrepresentation that led to an Accounting and Auditing Enforcement by the SEC, as contained in the AAER dataset created by Dechow et al. (2011). Since 1982, the SEC has issued Accounting and Auditing Enforcement Releases (AAERs) during or at the conclusion of an investigation against a company, an auditor, or an officer for alleged accounting or auditing misconduct. The benefit of this sample is that all AAERs are likely to be material violations, since an AAER release by the SEC is considered a severe step meant to attract attention. Dechow et al. (2011) refer to AAERs as misstatements, “although fraud is often implied by the SEC’s allegations,” because “firms and managers typically do not admit or deny guilt with respect to the SEC allegations.” We will refer to them simply as AAERs.

The main shortcoming of this measure is that the sample may be overly restrictive because the SEC does not have the budget to go after all frauds (e.g., Dechow et al. 1996; Miller 2006). Furthermore, as the SEC website notices, “[AAER] only highlights certain actions and is not meant to be a complete and exhaustive compilation of all of the actions that fall into this category.”⁷ For example, during the period 2001–2014, there were, on average, 54 AAERs per year, versus 220 securities class actions (Soltes 2019).

The third measure does not capture fraud – only restatements. To obtain such a measure, we start from the most widely employed data on restatements: the Audit Analytics database. We then use the Audit Analytics classification to separate clerical errors from intentional implementation of misstatements.

Finally, the fourth measure of fraud is the full set of SEC section 10b-5 securities fraud cases. While all 10b-5 cases arise from some material misstatements or omissions in providing material information, they also include non-auditing-related misrepresentations, like this example from DMZ: “The Department of Justice indicts Sotheby’s for price fixing after a long investigation of the art industry. Shareholders sue the firm for failing to disclose that a large portion of their revenues are being derived from the unsustainable high prices achieved through the pricing arrangement with Christie’s” (DMZ Appendix, 2010). In this example, no *accounting* violation has occurred, but omitting the economic implications of fixing prices represents an SEC rule 10b-5 violation. In DMZ, 35% of the 10b-5 violations are omissions of information concerning an underlying legal violation, as in the Sotheby’s example. Rather than using the original DMZ set of 10b-5 class actions, we use the updated dataset from the Securities Class Action Clearinghouse, as compiled by Kempf and Spalt (2022). We refer to this measure as *SCAC securities fraud*, again noting that these are alleged frauds.

Figure 1 plots the fraction of Compustat large US corporations engaged in (alleged) fraud or misconduct (henceforth, for simplicity, “fraud”) according to the various measures. The data are monthly from 1997 to 2005, and a company is classified as committing a fraud during a certain month if that month falls within the fraud period determined after detection. When the measure we use is accounting restatements, the fraction of firms involved in fraud is large (on average 13% per year, plotted with magnitudes on the right axis) and increasing over time (which might reflect the increased complexity of the accounting rules). All the other measures of misconduct produce a much lower incidence – between 1 and 4% of firms (plotted with magnitudes on the left axis) – with

⁷ <https://www.sec.gov/divisions/enforce/friactions.shtml>.




Fig. 1 Frequency of Observed Corporate Fraud. We plot the prevalence rates of corporate fraud, for cases of corporate fraud that are eventually caught. The magnitudes represent the percentage of all US large public corporations committing corporate fraud that are caught. Auditor-detected frauds are frauds in the DMZ sample of SEC 10(b) securities class actions which were triggered by an auditor, either by an auditor resignation or by the auditor issuing a qualified opinion and either the firm or analysts revealing the fraud. Restatements are from AuditAnalytics and refer to restatements triggered by accounting misapplication. AAERs are the SEC investigation releases used in Dechow et al. (2011). The sample SCAC securities fraud is from Kempf and Spalt (2022). Restatements are plotted with magnitudes on the right axis

a pronounced hump shape centered around the turn of the millennium. This pattern is consistent with Wang et al. (2010), who find that firms start fraud more in booms because monitoring is lower and the managerial reward from cheating is higher.

3 Methodology: inferring the detection likelihood using the Arthur Andersen experiment

3.1 Experiment overview

Figure 1 depicts the incidence of corporate frauds that are eventually caught. Our main agenda is to estimate what percentage of initiated frauds are caught. We refer to this percentage as the detection likelihood.

Our methodology relies on two key ingredients – the Kolmogorov axiom of conditional probability and the Arthur Andersen natural experiment. The pervasiveness of corporate fraud is the probability of a firm engaging in fraud (denoted F), regardless of whether it is caught or not: $Pr(F)$. What we observe instead is the joint event of a firm engaging in fraud and being caught: $\mathbf{Pr}(F, caught)$. (We will use the convention of bolding the variables we observe.) By the law of conditional probability, the unconditional probability of engaging in a fraud can be written as:

$$Pr(F) = \frac{\mathbf{Pr}(F, caught)}{Pr(caught|F)} \quad (1)$$

Thus, if we knew the denominator, $Pr(caught|F)$, we could calculate $Pr(F)$.

Our experiment derives an upper bound of $Pr(caught|F)$ and thus a lower bound of $Pr(F)$ by comparing the frequency of revealed fraud in “normal” times with the frequency of revealed fraud in a sample where scrutiny increased substantially. In our experiment, the sample of firms under enhanced scrutiny is the set of Arthur Andersen (AA) clients after the AA demise following the Enron scandal. On October 22, 2001, Enron acknowledged an SEC inquiry concerning possible conflicts of interest in various partnerships (Barton 2005). On November 23, 2001, the *New York Times* ran an article with the headline “From Sunbeam to Enron: Andersen’s Reputation Suffers.” Thus, we assume that the period of enhanced scrutiny started after November 30, 2001, a date to which we refer as “the watershed date.”⁸ Since roughly one-fifth of all large publicly traded firms had AA as their auditor in 2001, this event provides a natural experiment of increased scrutiny of firms. We look at the revelation of fraud that started being perpetrated before the watershed date. Thus, we can ignore any deterrence effect, since the enhanced scrutiny produced by Enron’s crisis was unexpected.

One source of the increased scrutiny is the new auditors (Chen and Zhou 2007). Kealey et al. (2007) find that the audit fees charged by AA successors to AA former clients when they switched after Enron are positively related to the length of AA tenure with those clients, suggesting that there is a perceived risk associated with AA clients that needs to be mitigated with additional monitoring. This enhanced scrutiny is also consistent with Nagy’s (2005) finding that smaller ex-AA clients have lower discretionary accruals after switching to a new auditor.

Nevertheless, new auditors are not the only source of enhanced scrutiny. As Chaney and Philipich (2002) and Krishnamurthy et al. (2006) show, after the Enron scandal exploded, all AA clients came under a cloud of suspicion. Investment intermediaries and internal gatekeepers of AA clients had stronger incentives to scrutinize their firms thoroughly to clear them of the shadow of suspicion. Likewise, short

⁸ As a robustness test, we shifted the beginning of the detection period one month backward or one month forward. The results are substantially unchanged.

sellers and the media had special incentives to focus their attention on AA clients in search for the next big scandal.

3.2 Identification Assumption 1: AA firms are not different in 1998–2000

Our first assumption is that in the period before AA's demise (1998–2000), fraud (F) was equally likely in AA firms and non-AA firms (\overline{AA})⁹:

$$\text{Assumption 1:} \quad Pr(F|\overline{AA}) = Pr(F|AA)$$

AA's indictment by the Department of Justice and its initial conviction for obstruction of justice may make this a surprising assumption. Yet, the initial conviction of AA was for obstruction of justice, not for being a bad auditor, and it was unanimously overturned upon appeal.

More convincingly, the accounting literature has concluded that there was no difference between AA and other auditors in terms of auditing rigor. In a matched sample, Agrawal and Chada (2005) find that the existence of AA as the auditor is not associated with firms having more restatements. Likewise, controlling for client size, region, time, and industry, Eisenberg and Macey (2004) find that AA clients did not perform any better or worse than other firms. Nevertheless, we retest this assumption for assurance and because our sample of large US corporations is different from the aforementioned studies.

Table 1 reports summary statistics comparing the characteristics and industry of AA clients and non-AA clients during the period 1998–2000. The firms in the set are large (>\$750 million in assets) US corporations with data in Compustat. We identify auditors using the Compustat "AU" item, mapping the auditor who signs a financial statement to the calendar month-year of the report. We use two sets of non-AA clients: all non-AA clients and the subset of all non-AA clients that are audited by another of the Big Five audit firms. Statistics are collapsed to one observation per firm over the time period. A firm appears in multiple columns only if it switches auditors. Because of auditor transitions and firms' delisting during the three-year period, the sample of AA firms is larger than in the subsequent tables, where we require a firm to be audited by AA in 2001.

Panel A reports that AA clients are statistically similar to other Big Five clients in assets, sales, and EBITDA, but have higher leverage (long term debt-to-assets). The leverage difference suggests a difference in industry composition, which is what we present in Panel B. For example, compared to Big Five auditors that are not AA, AA has at least 2 percentage points fewer clients in the industries of Banks & Insurance, Retail & Wholesale, and Computers, and 2 percentage points more clients in Utilities, Refining & Extractive, Communication & Transport, and Services & Healthcare. We return to these industry compositional differences in our empirical specifications.

⁹ To keep notation simple, we do not include time subscripts. We make the timing clear in the text.

Table 1 Statistics: Arthur Andersen Clients (AA) and Non-AA Clients, 1998–2000

Panel A: Firm Characteristics									
	AA Firms			All Non-AA			Big 5 Non-AA		
	Observations = 483			Observations = 1,792			Observations = 1,426		
	Mean	Median		Mean	p-value	Median	z-value	Mean	p-value
Assets	6,998	1,991		9,802	0.120	1,982	0.920	9,088	0.148
Sales	3,645	1,471		4,297	0.212	1,384	0.784	4,431	0.144
EBITDA	568.6	196		668.9	0.318	182	0.174	677	0.233
Leverage	0.34	0.311		0.279	0.000***	0.256	0.000***	0.279	0.000***
Panel B: Industry Distribution									
	AA Firms			All Non-AA			Big 5 Non-AA		
	Observations = 483			Observations = 1,792			Observations = 1,426		
	Percent of Distribution (*100)			Percent of Distribution (*100)			Percent of Distribution (*100)		
Agriculture	0.41			0.17			0.14		
Banks & Insurance	12.42			21.09			19.57		
Chemicals	3.11			4.58			4.70		
Communication & Transport	16.15			10.44			9.40		
Computers	4.35			10.04			9.68		
Durable Manufacturing	10.35			12.39			13.32		
Food & Tobacco	1.24			2.85			2.59		
Lumber, Furniture, Printing	4.14			4.30			4.42		
Mining & Construction	2.07			1.56			1.68		
Pharmaceuticals	1.66			2.90			3.16		
Refining Extractive	3.93			1.79			1.82		
Retail & Wholesale	7.25			9.49			9.33		
Services & Healthcare	11.39			7.37			7.36		
Textile & Apparel	1.04			1.06			0.98		
Utilities	20.50			9.99			11.85		

Table 1 (continued)

Pearson's Chi-Square test of Distribution Equivalence of non-AA Samples to AA Sample	
Statistic	105.2
P-value	0.000
	85.8
	0.000

The sample is all Compustat firms with \$750 million in assets in the 1998–2000 period. The designation of auditor is the calendar time designation in financial statements, not the experiment designation starting in Table 3. All statistics are collapsed to a single observation per firm, except where a firm switches auditors as per the column designation. In those cases, we allow the firm to appear in multiple columns. Because of this time series collapsing, we have more AA firms here than in the experiment tables. In Panel A, the p-values are from tests for the mean comparisons, and the z-values are from ranksum tests for the medians. Leverage is long-term debt divided by assets. In Panel B, we present the industry distribution of firms and present Pearson distribution equivalence tests

Table 2 Did Arthur Andersen clients exhibit excess corporate fraud in the pre-period of 1998–2000?

Panel A: Univariate tests. Probability of manipulation scores (ProbM Score) are calculated using Compustat data following Beneish (1999), winsorized at 0.025 for extremes outside Beneish's range. Fraud scores (FScores) are from Dechow et al. (2011); auditor-caught frauds are from Dyck et al. (2010); AAERS are from the CFRM at Berkeley-Haas; accounting restatements are from AuditAnalytics; and SCAC are from Kempf and Spalt (2022). The designation of auditor is the calendar time designation in financial statements. Panel A reports cross sectional t-test of differences in mean fraud rates, comparing 1998–2000 large corporations with AA as auditor to all others, or to those with another Big Five auditor												
	AA Firms			All Non-AA			Big 5 Non-AA			p-value		
	Mean	Obs	Obs	Mean	Obs	Obs	Mean	Obs	Obs			p-value
ProbM Score (Beneish)	-1.711	435	435	-1.548	1,603	1,603	-1.601	1,269	1,269	0.014		0.096
ProbM > -2.2 threshold	0.585	435	435	0.624	1,603	1,603	0.613	1,269	1,269	0.078		0.217
FScore (Dechow et al. 2011)	3.217	420	420	4.378	1,559	1,559	3.664	1,246	1,246	0.301		0.650
FScore > 1.4 threshold	0.323	420	420	0.413	1,559	1,559	0.391	1,246	1,246	0.000		0.005
Auditor-Caught Fraud	0.0186	483	483	0.0201	1,792	1,792	0.021	1,426	1,426	0.839		0.747
AAERS	0.0352	483	483	0.0458	1,792	1,792	0.0463	1,426	1,426	0.313		0.302
Restatements	0.164	483	483	0.137	1,792	1,792	0.139	1,426	1,426	0.134		0.183
SCAC Securities Fraud	0.0704	483	483	0.0854	1,792	1,792	0.0799	1,426	1,426	0.287		0.498

Panels B & C: Multivariate Tests. The dependent variables are accounting scores predicting fraud (Panel B) or realized corporate fraud (panel C). Columns (3), (6), (9), and (12) used indicators for the score being above the relevant threshold of concern as identified in the original literature. Industry fixed effects, when included, are those of Table 1. Estimation is either by ordinary least squares (OLS) or by logit, as designated. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively. Standard errors in brackets are clustered at the firm level												
Panel B: Unobservable Measures of Fraud Dependent Variables												
Dependent var.:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Estimation:	Accounting Manipulation (ProbM Score)			Fraud Score (FScore)			Fraud Score (FScore)			Logit FScore > 1.4		
	OLS ProbM	OLS ProbM	Logit ProbM > -2.2	OLS ProbM	OLS ProbM	Logit ProbM > -2.2	OLS FScore	OLS FScore	Logit FScore > 1.4	OLS FScore	OLS FScore	Logit FScore > 1.4
Arthur Andersen	-0.142*** [0.0514]	-0.0224 [0.0509]	0.11 [0.0960]	-0.104* [0.0530]	0.00648 [0.0518]	0.158 [0.0990]	-0.384 [0.516]	0.179 [0.546]	0.0229 [0.120]	-0.155 [0.515]	0.318 [0.554]	0.0683 [0.124]
Log Assets	-0.106*** [0.0164]	-0.0938*** [0.0157]	-0.100*** [0.0341]	-0.112*** [0.0176]	-0.106*** [0.0166]	-0.113*** [0.0371]	-0.184 [0.277]	0.00543 [0.266]	0.0411 [0.0388]	-0.0758 [0.296]	0.117 [0.283]	0.0663 [0.0428]
Sales / Assets	0.352*** [0.0411]	0.115** [0.0498]	0.366*** [0.108]	0.342*** [0.0427]	0.121** [0.0512]	0.364*** [0.117]	-1.025*** [0.252]	-1.608*** [0.373]	-0.0978 [0.0806]	-0.772*** [0.246]	-1.252*** [0.356]	-0.0859 [0.0879]

Table 2 (continued)

EBITDA / Sales	-2.495*** [0.332]	-2.849*** [0.390]	-3.155*** [0.623]	-2.394*** [0.360]	-2.604*** [0.426]	-3.345*** [0.700]	-23.28*** [4.199]	-22.19*** [4.779]	-2.909*** [0.565]	-19.25*** [4.143]	-17.70*** [4.566]	-3.280*** [0.650]
LT Debt/Assets	-0.189 [0.141]	-0.107 [0.145]	0.00651 [0.188]	-0.234 [0.156]	-0.145 [0.157]	-0.203 [0.209]	-7.047*** [1.486]	-5.010*** [1.208]	-0.345* [0.209]	-5.521*** [1.481]	-3.355*** [1.147]	-0.432* [0.236]
Sample	All Clients											
Year FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Industry FE	N	Y	Y	N	Y	Y	N	Y	Y	N	Y	Y
Industry*Year FE	N	Y	N	N	Y	N	N	Y	N	N	Y	N
Observations	4,963	4,963	4,963	4,167	4,167	4,167	4,855	4,855	4,855	4,121	4,121	4,121
R2 / Pseudo R2	0.082	0.18	0.123	0.085	0.193	0.125	0.031	0.066	0.172	0.025	0.061	0.183
Panel C: Observable Measures of Fraud Dependent Variables												
Dependent var.:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Non-AA sample:	Auditor-Caught Securities Fraud											
Estimation:	-	-	Big 5	Big 5	-	-	Big 5	Big 5	-	-	Big 5	Big 5
Arthur Andersen	OLS	OLS	OLS	OLS	OLS	OLS	OLS	OLS	OLS	OLS	OLS	OLS
	-0.0375 [0.395]	0.158 [0.426]	-0.0865 [0.397]	0.0369 [0.430]	-0.143 [0.290]	0.0227 [0.306]	-0.145 [0.297]	-0.0213 [0.313]	0.119 [0.149]	0.189 [0.156]	0.127 [0.153]	0.181 [0.160]
Log Assets	0.547*** [0.106]	0.619*** [0.122]	0.540*** [0.117]	0.594*** [0.130]	0.471*** [0.0884]	0.562*** [0.0910]	0.460*** [0.0965]	0.536*** [0.0967]	-0.104* [0.0541]	-0.0548 [0.0550]	-0.0818 [0.0584]	-0.0428 [0.0585]
Sales / Assets	0.297* [0.152]	0.0184 [0.212]	0.282* [0.167]	0.0406 [0.215]	0.419*** [0.103]	0.309** [0.131]	0.425*** [0.110]	0.330** [0.138]	0.109 [0.0769]	-0.0824 [0.106]	0.0979 [0.0875]	-0.0836 [0.117]
EBITDA / Sales	-0.883 [2.324]	-1.898 [1.431]	-1.494 [2.286]	-2.163 [1.472]	-0.409 [1.043]	-0.925 [0.910]	-0.653 [1.152]	-0.99 [1.014]	-1.208* [0.643]	-1.359* [0.711]	-1.214* [0.713]	-1.193 [0.784]
LT Debt/Assets	0.715 [0.728]	0.873 [0.775]	0.305 [0.860]	0.368 [0.846]	0.48 [0.489]	0.879* [0.477]	0.259 [0.558]	0.612 [0.510]	-0.0784 [0.262]	0.0165 [0.281]	0.0233 [0.293]	0.0821 [0.310]

Table 2 (continued)

Fixed Effects:												
Industry	N	Y	N	Y	N	Y	N	Y	N	Y	N	Y
Year	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Industry*Year	N	N	N	N	N	N	N	N	N	N	N	N
Observations	5,465	4,993	4,641	4,228	5,465	5,105	4,641	4,331	5,465	5,465	4,641	4,641
R-squared	0.0586	0.103	0.0597	0.117	0.062	0.109	0.0602	0.0989	0.0133	0.0402	0.0134	0.0388
Dependent var.: SCAC Securities Fraud 10b-5												
Non-AA sample:												
Estimation:												
Arthur Andersen												
Log Assets												
Sales / Assets												
EBITDA / Sales												
LT Debt/Assets												
Fixed Effects:												
Industry	N	Y	N	Y	N	Y	N	Y	N	Y	N	Y
Year	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Industry*Year	N	Y	N	Y	N	Y	N	Y	N	Y	N	Y
Observations	5,465	5,237	4,641	4,152	5,465	5,105	4,641	4,331	5,465	5,465	4,641	4,641
R-squared	0.0493	0.114	0.0434	0.11	0.062	0.109	0.0602	0.0989	0.0133	0.0402	0.0134	0.0388

Were AA clients more fraudulent? In Table 2, we consider this question using eight measures of fraud, collapsed to one observation per firm for the 1998–2000 period before AA’s demise (allowing, as before, a firm to have two observations if it switched to or from AA in this period). For the discrete variables of caught fraud – *auditor-caught frauds*, *AAERs*, *restatements*, and *SCAC securities fraud* – we define a fraud event if the event ever occurs in the three-year period. Thus, to compare the detected fraud in this table to detected frauds in Fig. 1 (and to those in the subsequent tables), we need to divide the reported number by three. In addition to these measures of fraud, we include measures of fraud likelihood based on a continuous measure from the accounting literature. We include both the probability of accounting manipulation score, *ProbM* (Beneish 1999),¹⁰ and the probability of fraud score, *FScore* (Dechow et al. 2011).¹¹ Following Beneish (1999) and Dechow et al. (2011), we also consider a discrete version of the *ProbM* and *FScore* in the form of a threshold dummy variable, with the authors considering a firm above the threshold as indicating a strong likelihood of the firm’s being a manipulator.

Panel A of Table 2 reports that for all eight measures of fraud, AA clients are never more likely to be committing fraud. In fact, based on the measures from the accounting literature, AA clients show lower scores, indicating that they are less likely to be committing fraud. These results do not account for the above-mentioned leverage and industry differences.

In Panel B, Table 2, we report estimates from multivariate estimations of fraud likelihood using accounting literature measures of *probM*, the indicator *probM* > -2.2, *FScore*, and the indicator *FScore* > 1.4. For each estimation, our variable of interest is an AA dummy, and we include a set of controls for size (log assets and sales/assets), debt (LT Debt/Assets), and profitability (EBITDA/Sales). When we estimate the threshold models (columns 3, 6, 9, and 12), we estimate via logit instead of OLS. All columns include year fixed effects. Columns 2, 3, 5, 6, 8, 9, 11, and 12 include industry fixed effects (where the industries are those of Table 1). Columns 2, 5, 8, and 11 include industry*year fixed effects. (We do not include this further interaction for the logit, as the estimation does not consistently converge uniquely.) In columns 4–6 and 10–12, we restrict the non-AA firms to Big Five accounting firm clients only.

Across all specifications, the coefficient on the AA dummy is never positive and significant. In columns 1 and 3, the AA dummy is negative and significant (i.e., AA firms are associated with *less* probability of manipulation). Being an AA client does not significantly affect the positive likelihood of manipulation or fraud in the pre-demise period.

In Panel C, we repeat these tests using, as dependent variables, *auditor-caught fraud*, *AAERs*, *restatements*, and *SCAC securities fraud*. In all even-numbered

¹⁰ Appendix 2 details the calculations for the ProbM score. ProbM is a score with no natural scale. The mean and standard deviation of ProbM in our sample are -2.325 and 1.357, respectively. According to Beneish (1999), a score greater than -2.22 indicates a strong likelihood of a firm being a manipulator.

¹¹ We thank Weili Ge for the FScore data. We use their FScore measure from model 2 of the paper, which includes financial statement variables and market data. FScore is a score variable, scaled to imply that a score of 1.00 indicates that the firm has no more probability of AAER fraud than the unconditional probability. In our sample, the mean and standard deviation of Fscore are 1.785 and 10.57, respectively.

specifications, we control for industry crossed with year fixed effects, except for auditor-detected fraud (whose small sample forces us to include only industry fixed effects). As in Panel B, the coefficient on the AA dummy is never positive and significant.

Krishnan and Visvanathan (2009) find that earnings management was more pronounced in the AA Houston office than in the Houston offices of other Big Five auditors. Since Enron was headquartered in Houston, this is not that surprising. Nevertheless, to address the concern that in certain regions AA-audited firms may have been more likely to commit fraud, in Appendix Table A1, we restrict our sample to large corporations located in the state of Texas. For each AA client, we find a match among other Big Five clients within the two-digit SIC code, based on a propensity score to be an AA client. We generate the propensity score based on assets, sales, EBITDA, and leverage within industry. In Table A1 (columns 1 and 2), the dependent variable is the ProbM score. An AA dummy variable (equal to one for an AA client) is not significant either in the Texas sample (Panel A) or in the United States sample (Panel B). Across the remaining columns, we repeat the same tests using other measures of uncaught (*FScore*) and caught fraud (*auditor-caught fraud*, *AAERs*, *restatements*, and *SCAC securities fraud*). The results are unchanged. We find no evidence that AA clients are statistically or economically different from other auditors' clients, consistent with Assumption 1.

3.3 Identification Assumption 2

Our methodology is based on the assumption that after the watershed date, AA clients were subject to an enhanced level of scrutiny not only by the new auditors but also by gatekeepers, short sellers, investment intermediaries, and the media. This assumption implies that $Pr(\text{caught}|F, AA) > Pr(\text{caught}|F, \overline{AA})$. While we cannot test this inequality directly, the hypothesis is consistent with the decline in the stock price of AA clients at the announcement of AA's problems (Chaney and Philipich 2002) and indictment (Krishnamurthy et al. 2006). Furthermore, as we will show in Table 3, this inequality is true in the data for all our measures of fraud.

To obtain a point estimate on the amount of undetected fraud, we further assume:

Assumption 2: $Pr(\text{caught}|F, AA) = 1$.

When it comes to purely financial fraud or financial restatements, Assumption 2 may be plausible since this was the fraud discovered in Enron (and later in World-Com) and since this is the type of fraud every gatekeeper was fearful of missing. However, when we use, as a measure of fraud, all SCAC cases (including for instance price fixing), the level of alertness was less. The reason is that some of the detectors with enhanced incentives (like the new auditors) have no role in uncovering these types of fraud. Thus, for some types of fraud, it is more likely that $Pr(\text{caught}|F, AA) > Pr(\text{caught}|F, \overline{AA})$, but $Pr(\text{caught}|F, AA) < 1$. When this is

the case – as we explain below – our estimate will yield an upper bound of the detection likelihood and a lower bound on the amount of undetected fraud.

In some cases – Kohlbeck et al. (2009) estimate that the number is 25% – AA clients switched audit firms but not audit partners. The former AA engagement partner moved to another audit firm, bringing the account. This continued relationship might reduce the new auditor's willingness to "clean the dirty laundry"; yet it is hard to imagine that such relationships will completely eliminate the power of the experiment. Speaking about a firm that went from AA to Grant Thornton, the CEO of Grant Thornton testified at trial that "we converted it to a Grant Thornton audit approach and Grant Thornton audit-quality controls, and we had other people review the engagement."¹² Thus, even when the audit partner remains the same, the new audit firm performs an extra screening.

A more general concern is that the increased attention on corporate fraud that followed the Enron and WorldCom scandals might have prompted the SEC (or the other auditors not affected by the turnover) to become more active in detecting fraud. In the limit, if this enhanced scrutiny exposed all fraud in all firms, there will be no difference between the amount of fraud revealed in AA clients and the amount of fraud revealed in clients of other audit firms, invalidating our experiment. Yet, as long as some enhanced scrutiny affects all firms but AA firms are affected more, our methodology will work, but will underestimate undetected fraud. Thus, our results should be interpreted as a lower bound on the pervasiveness of fraud.

Finally, the AA demise was completely unexpected, so it could not have altered the ex ante incentives to commit fraud. In this sense, it is different from a mandatory turnover, which is anticipated.

3.4 Conditional probabilities within the AA experiment

With Assumptions 1 and 2 in hand, we can use the law of conditional probability to derive an estimator for the detection likelihood of fraud. If we write down two versions of Eq. (1) (one for AA and one for non-AA), bringing the denominators to the other side and dividing one by the other, we have a relationship that puts only observable detections of fraud on the right-hand side:

$$\frac{\Pr(\text{caught}|F, \overline{AA}) \cdot \Pr(F|\overline{AA})}{\Pr(\text{caught}|F, AA) \cdot \Pr(F|AA)} = \frac{\Pr(F, \text{caught}|\overline{AA})}{\Pr(F, \text{caught}|AA)}. \quad (2)$$

Substituting Assumptions 1 and 2 into the left-hand side of Eq. (2) implies:

$$\Pr(\text{caught}|F, \overline{AA}) = \frac{\Pr(F, \text{caught}|\overline{AA})}{\Pr(F, \text{caught}|AA)}. \quad (3)$$

¹² <https://www.wsj.com/articles/SB112951490246670395>.

Since the normal level of scrutiny is the one experienced by non-AA clients, $Pr(caught|F, \overline{AA}) = Pr(caught|F)$. Thus, Eq. (3) provides an estimator for the detection likelihood of fraud that is based solely on two observables: the emergence of fraud in non-AA firms and the emergence of fraud in former AA firms.

Note that to the extent that not all fraud emerges in former AA clients (Assumption 2 becomes $Pr(caught|F, AA) < 1$), then our estimate $Pr(F)$ of the pervasiveness of fraud represents a lower bound of the real amount of fraud. To see this, consider that

$$Pr(F) = \frac{\Pr(F, caught)}{Pr(caught|F)} > \frac{\Pr(F, caught)}{Pr(caught|F)}, \quad (4)$$

where the inequality derives from (3) and $Pr(caught|F, AA) < 1$.

Note also that we derive the detection likelihood by comparing fraud detection in two groups (former AA clients and non-AA clients) at the same time. Thus, this estimate should not be affected by fluctuations in the probability of committing fraud, as long as these fluctuations are similar in the two groups. Yet, these fluctuations will matter for the level of fraud pervasiveness at any point in time, since this is likely to vary over the business cycle as shown by Wang et al. (2010).

3.5 Empirical methodology

Having explained why the detection likelihood given in Eq. (3) is the number we want to estimate, we now describe how we will estimate it. For each company, we observe whether fraud is discovered during the so-called detection period, both for companies that were AA clients and for companies that were not. This experiment can be thought of as a sum of Bernoulli trials, which can be modeled using a Poisson probability distribution, where the mean λ is the proportion of firms discovered as fraudulent per unit of time. The maximum likelihood estimate of λ in a Poisson distribution is simply the frequency of fraud during the detection period. To facilitate a system estimation (see below), we will estimate this mean λ with a Poisson regression, where $E(Y|X) = \lambda = e^{\theta x}$.

In this framework, the object of our interest (i.e., the detection likelihood $Pr(caught|F, \overline{AA}) = \frac{\Pr(F, caught|AA)}{\Pr(F, caught|AA)}$) is simply the relative risk (RR) of a firm being caught committing a fraud if it is not an AA former client versus if it is an AA former client. Thus, the detection likelihood is simply the ratio of the frequency of fraud in the two samples. Given the paucity of fraud observations, we compute the exact p-values of this ratio in finite samples. Following Agresti (1992), we do so by enumerating all possible draws in the sample and then computing the frequency of outcomes that are at least as different from the null hypothesis as the one observed.

4 Results

The timing of the experiment is as follows. We define a company as having AA as an auditor if AA signed a financial report anytime in the calendar year 2001, irrespective of the firm's fiscal year. All companies without AA as auditor during this period are non-AA clients. We consider that a fraud was revealed during the detection period if the fraud started before the watershed date of November 30, 2001, and came to light between the watershed date and the end of 2003 (the detection period). Note that our estimates are obtained from the ratio of the fraud caught in AA and non-AA firms. Thus, as long as the amount of fraud committed in AA and non-AA clients responds to cyclical fluctuations in a similar fashion, our estimates are not affected by business cycle fluctuations in the *aggregate* level of fraud, an important phenomenon documented by Wang et al. (2010).

The count of observations is shown at the bottom of Table 3, Panel A. The natural experiment includes 353 AA clients in the pre-period who survive at least one year in the post-period, and 2,404 non-AA client firms. The number of detected fraud events depends on the fraud measure: 168 restatements, 63 SCACs, 59 AAERs, and 21 auditor-detected fraud.

4.1 Detection likelihood results

Table 3, Panel A reports the main detection likelihood estimates across the four fraud measures – *auditor-detected fraud*, *AAERs*, *restatements*, and *SCAC securities fraud*. For each of these measures, the estimated coefficient of the Poisson regression we report is the coefficient of a dummy variable equal to one if a company was a former AA client. Thus, if former AA clients have an average frequency of fraud of $\lambda_{AA} = e^{\theta_{AA}x}$ and non-AA clients have an average frequency of fraud of $\lambda_N = e^{\theta_Nx}$, the estimated coefficient we report is $\theta_{AA} - \theta_N$. Since the detection likelihood ratio is nothing but $\frac{\lambda_N}{\lambda_{AA}} = \frac{e^{\theta_N}}{e^{\theta_{AA}}} = e^{\theta_N - \theta_{AA}}$, we obtain the detection likelihood ratio as $\frac{1}{e^{\theta_{AA} - \theta_N}}$, or the inverse of e raised to the power of the estimated coefficient. We report both the asymptotic and the exact p-values for the hypothesis that the detection likelihood is equal to 1.

As we can see in Panel A of Table 3, the hypothesis that the detection likelihood is bigger or equal to 1 is rejected at least at the 10% level in all four measures and in three out of four cases at the 5% level or better. These findings verify the assumption that after the demise of AA more frauds are caught in AA firms than in non-AA firms ($Pr(\text{caught}|F|AA) > Pr(\text{caught}|F|\overline{AA})$).

More specifically, when we look at *auditor-detected fraud*, the estimated detection likelihood is 0.29, which is different from 1 at the 3% level. For *AAERs*, the detection likelihood estimate is 0.52, different from 1 at the 7% level. The detection likelihood estimate for accounting violations is 0.34, different from one at the 0.1% level. For the *SCAC securities fraud* measure, the detection likelihood is 0.47, different from 1 at the 2% level.¹³ In sum, despite the different sources

¹³ To show that our results do not depend upon IPO-induced securities class actions, we also report SCAC securities fraud purged of three IPO frauds. Using this measure, the detection likelihood estimate is 0.53, which is very close to the 0.47 reported in the text.

Table 3 Detection likelihood estimates**Panel A: Exact Poisson Estimation of Poisson-Binomial Draws**

	Severe Financial Reporting Violations		Accounting Violations	Securities Fraud
	Auditor-Detected	AAERs	Restatements	SCAC
Former Arthur Andersen Clients				
Coefficient	1.225	0.655	1.088	0.755
Detection Likelihood	0.294	0.520	0.337	0.470
Asymptotic p-value	0.008	0.037	0.000	0.011
Exact p-value	0.026	0.067	0.000	0.024
Observations (# of clients)	2,757	2,757	2,757	2,757
AA clients	353	353	353	353
Non-AA clients	2,404	2,404	2,404	2,404
Total Number of Frauds	21	59	168	63
Pseudo R-squared	0.0067	0.0067	0.0278	0.0094

Panel B: Simultaneous, Maximum Likelihood Estimation of System of Poisson-Binomial Equations, Constrained to Have Equal Coefficients across Equations

	Equation Dependent Variables: Auditor-Detected Fraud, AAERs, Restatements, SCAC		Equation Dependent Variables: Auditor-Detected Fraud, Restatements
Former Arthur Andersen Clients			
Coefficient	0.957		1.104
Detection Likelihood	0.384		0.332
Confidence Interval	0.300	0.493	0.243
Asymptotic p-value	0.000		0.000
Observations (# of clients)	2757 in each of 4 equations		2757 in each of 2 equations

Presented are main Poisson-binomial estimates of the detection likelihood (odds-ratio) using the AA experiment. The sample is US publicly traded corporations with more than \$750 million in assets. A firm is identified as being a former Arthur Andersen (AA) client if it was audited by AA in the year 2001. In all cases, the frauds considered are those starting prior to the watershed date of December 1, 2001, which are caught after the watershed date and before the end of 2003. Auditor-detected frauds are frauds in the Dyck et al. (2010) sample of class action frauds which were detected by an auditor either by an auditor resignation or by the auditor issuing a qualified opinion and either the firm or analysts revealing the fraud. Restatements are from AuditAnalytics and refer to restatements triggered by accounting misapplication. AAERs are the SEC investigation releases used in Dechow et al. (2011). The sample SCAC securities fraud is from Kempf and Spalt (2022). Panel A reports Poisson-binomial coefficient estimates of the asymptotic and exact p-values for the AA coefficient. The detection likelihood is calculated directly in Poisson estimation as $odds\ ratio = 1/\exp(coefficient)$. Panel B reports maximum likelihood estimates of the system of all Poisson equations estimated simultaneously, with the coefficient on AA constrained to be the same. The first column of Panel B includes all four fraud definitions, and the second column includes only auditor-detected fraud and restatements, per the discussion in the text that the experiment is most precise for these definitions. Also included in Panel B are the confidence intervals for the detection likelihood estimates

and the different definitions of fraud, there is a clear result: a substantial amount of corporate fraud remains undetected. With detection likelihood between 29 and 52%, there is indeed an iceberg of undetected fraud that ranges between 48 and 71% of total fraud.

If we want to move beyond this simple result and estimate more precisely the amount of undetected fraud, we need to pool together all the possible observations toward the estimation of a single detection likelihood. This is what we do in Table 3, Panel B. The first column reports the estimates obtained by estimating simultaneously the four Poisson regressions, with the restriction that the coefficient should be the same in all four. When we pool the four measures of fraud, the estimated detection likelihood is 0.38, with a 95% confidence interval between 0.30 and 0.49. Thus, we can say with 95% confidence that between 51 and 70% of all fraud is undetected.

These estimates of undetected fraud are downward biased because Assumption 2 (that after the Enron scandal all the fraud in former AA firms was revealed) is unlikely to be satisfied. This is true for all the measures, but particularly so for SCACs and AAERs. The SCAC securities fraud measure includes frauds that are not misrepresentation of accounting information but rather failures to disclose material information. It is unlikely that the additional scrutiny triggered by the AA demise would expose all these failure-to-disclose cases. In such cases, the experiment design implies that 0.47 detection likelihood is only an upper bound (biasing downward the proportion of undetected fraud). Similarly, the detection of an AAER-type of fraud depends on the willingness of the SEC to bring an enforcement action. In any economic analysis of crime and punishment (see Becker 1968), the SEC's incentives to bring a case against a defunct firm like AA are small. Thus, the 0.52 AAER estimate is also likely an upper bound.

By contrast, Assumption 2 is more likely to hold for misconducts that auditors are more likely to catch, such as accounting restatements and auditor-detected fraud. Thus, in column 2 of Table 3, Panel B, we report the detection likelihood estimate obtained by pooling together only the two fraud measures for which Assumption 2 is more likely to hold. The estimated detection likelihood is 0.33, with a 95% confidence interval of 0.24–0.45. Thus, only one-third of corporate fraud is detected, and the total amount of corporate fraud is three times the corporate fraud that is observed. In the rest of the paper we will treat this as our best point estimate, with a 95% confidence interval that the total amount of fraud is between 2.2 and 4.1 times the fraud observed.

4.2 Detection likelihood robustness

4.2.1 Small firms

Krishnan et al. (2007) document that Big Four auditors were much more likely to issue a going concern qualification in their audits of large former AA clients than in their audits of large non-former AA clients (specifically, 5.6% versus 2.3%). In our experiment, this ratio would imply a detection likelihood of 0.41, quite consistent

with our results. Yet, Krishnan et al. (2007) find the opposite for small firms: former AA clients received less than half of the going concern opinions of non-AA clients (6.2% to 14.5%). As the authors recognize, this difference in the results for large and small firms is driven by sample selection. Krishnan et al (2007) restrict their attention to AA firms that are subsequently audited by Big Four auditors. It is likely that Big Four auditors accept all the large former AA clients but only the best of the small ones, leaving the rest to less prestigious auditors. In this paper we can test this conjecture by looking at all former AA clients, not just those subsequently audited by a Big Four auditor. To do so, however, we need to expand the sample to small firms.

Table 4, Panel A reproduces our detection likelihood estimation setup for smaller public corporations – those that were never over \$750 million in assets in the pre-period years. We run the experiment in the same setup as Table 3. Since DMZ does not collect auditor-detected data for small firms, the first measure of misconduct is not available for this sample.

We find, in Panel A, that for all three measures the detection likelihood is less than 1 and that for two out of three measures the detection likelihood is significantly less than 1 at the 1% level. This analysis, suggested by an anonymous referee, provides an out-of-sample validation of our methodology to use the AA-forced turnover to identify misconduct that generally remains undetected.

For the AAER and restatement measures, the detection likelihood estimates obtained using small firms in Table 4 are very similar to the ones obtained using large firms and reported in Table 3: 0.64 versus 0.52 for AAERs and 0.42 versus 0.34 for restatements. By contrast, for securities fraud cases, the detection likelihood falls drastically from 0.47 (large firms) to 0.29 (smaller firms). Most likely, this drop is due to the fact that the probability of a class action suit among small non-AA clients is only 0.5%, versus the 2% observed in large firms. This is hardly surprising. The main reason why DMZ restrict their sample to large companies is that the discovery of a fraud will always lead to a class action suit in large companies, where there is plenty of money to pay the lawyers, but not necessarily in small companies. Thus, the lower detection likelihood reflects a real difference between large and small companies: more fraud goes undetected in small companies.

4.2.2 Placebo

It is possible that AA clients were in businesses intrinsically more prone to corporate fraud or its detection. One way to help to rule out this possibility is to observe these firms in a different time period and check that they are not behaving differently. To this purpose, we reproduce the same experiment comparing former AA and non-AA clients (minus a few firms that do not survive) in the two years after the detection period (i.e., 2004–2006) and measure fraud detection in these years. Since the enhanced scrutiny following the demise of AA cannot last forever, we regard this exercise as a placebo test. This placebo period coincides with the beginning of the implementation of SOX. A large literature has tried to establish what the effects of the introduction of SOX are (see Coates and Srinivasan (2014)). Our test, however, is unaffected by any impact of SOX, since it compares the detection rate in former AA clients and former non-AA clients at the same time.

Table 4 Detection Likelihood Robustness: Small Corporations & Placebo**Panel A: Small Corporations**

	AAERs	Restatements	SCAC
Former Arthur Andersen Clients			
Coefficient	0.438	0.867	1.241
Detection Likelihood	0.645	0.420	0.289
p-value	0.322	0.000	0.000
Observations (# of clients)	7,553	7,553	7,553
AA clients	772	772	772
Non-AA clients	6,781	6,781	6,781
Total Number of Frauds	40	394	46
Pseudo R-squared	0.0018	0.0134	0.0134

Panel B: Placebo of Detection Time Period = 2004–2006

	AAERs	Restatements	SCAC
Former Arthur Andersen Clients			
Coefficient	0.021	-0.117	0.460
Detection Likelihood	0.979	1.124	0.631
Asymptotic p-value	0.961	0.429	0.190
Exact p-value	1.000	–	0.267
Observations (# of clients)	2,445	2,445	2,445
AA clients	314	314	314
Non-AA clients	2,131	2,131	2,131
Total Number of Frauds	46	450	53
Pseudo R-squared	0.0000	0.0003	0.0030

Presented are Poisson-binomial estimates of the detection likelihood (odds-ratio) using the AA experiment. Panel A differs from Table 3 in that instead of covering US publicly traded corporations with more than \$750 million in assets, we report estimates from firms with less than \$750 million in assets. Panel B differs in that the detection period covers frauds detected between 2004 and 2006, as opposed to being caught after the watershed date and before the end of 2003 (for Panel A and Table 3). The sample is US publicly traded corporations with more than \$750 million in assets. A firm is identified as being a former Arthur Andersen (AA) client if it was audited by AA in the year 2001. In all cases, the frauds considered are those starting prior to the watershed date of December 1, 2001, which are caught after the watershed date and before the end of 2003. Auditor-detected frauds are frauds in the Dyck et al. (2010) sample of class action frauds which were detected by an auditor either by an auditor resignation or by the auditor issuing a qualified opinion and either the firm or analysts revealing the fraud. Restatements are from AuditAnalytics and refer to restatements triggered by accounting misapplication. AAERs are the SEC investigation releases used in Dechow et al. (2011). The sample SCAC securities fraud is from Kempf and Spalt (2022). Both panels report Poisson-binomial coefficient estimates of the asymptotic p-values for the AA coefficient. Panel B also reports exact p-values, but the small sample is too large for computation of these statistics in Panel A and for restatements in Panel B. The detection likelihood is calculated directly in Poisson estimation as $odds\ ratio = 1/exp(coefficent)$

Table 4, Panel B reports the results. The percentage of firms caught committing fraud is very similar between former AA clients and former non-AA clients (detection likelihood estimates for all three measures are not different from 1), suggesting that there is not a natural proclivity of former AA clients to commit more fraud. In

addition, the similarity between the fraud revealed in AA and non-AA clients suggests that the enhanced disclosure of fraud during the treatment period is not just an acceleration of the discovery that would have taken place regardless, but a net increase in discovery.

4.2.3 Industry and geography

In Sect. 3.2, we showed that AA had more clients in Communications & Transport, Refining & Extractive, Services & Healthcare, and Utilities, and fewer clients in Banks & Insurance, Retail & Wholesale, and Computers. If, for unspecified reasons, corporate fraud was more prevalent among sectors in which AA was overrepresented just prior to the detection period and not afterward, then our detection likelihood estimate could be biased. To address this concern, we report the detection likelihood estimates in various subsamples that remove industries where AA was either overrepresented or underrepresented. (Appendix Table A2, Panels B and C report the industry and regional distribution of fraud for all the fraud measures.) As Table 5, Panel A shows, the detection likelihoods remain substantially unchanged.

The same concern could arise from regional variation in AA clients. For this reason, in Table 5, Panel B, we repeat the same exercise excluding some regions or some large states. Once again, the detection likelihood results appear stable.

4.3 Pervasiveness of corporate fraud results

We now can use the detection likelihood estimates in Eq. (1), $Pr(F) = \frac{Pr(F, caught)}{Pr(caught|F)}$, to estimate the overall pervasiveness of corporate fraud. The numerator in Eq. (1) is the observable incidence of fraud that is caught. The denominator in Eq. (1) is the detection likelihood. Table 6 reports observed caught frequencies in Panel A, detection likelihood estimates from Table 3 in Panel B, a baseline estimate of the pervasiveness of fraud across the measures of misconduct and alleged fraud in Panel C, and a best estimate of the pervasiveness of fraud across the measures of misconduct and alleged fraud using the best estimate detection likelihood in Panel D. Since AA firms during the detection period are assumed to have a probability of detection equal to one, we exclude them from Panel A, which computes the frequency of caught frauds under normal circumstances.

Focusing first on Panel A, recall that Fig. 1 showed that observed incidences of misconduct vary widely depending on the definition of corporate fraud and the time period of reference. Since fraud may be cyclical (Wang et al. 2010), we do not want to rely on a specific point in time, instead preferring to estimate pervasiveness over a full cycle of boom and bust years. The start in January 1998 and the end point in December 2005 are each almost exactly halfway through the respective expansion periods; thus the period covers one full business cycle from mid-point to mid-point.¹⁴

¹⁴ According to the NBER Business Cycle Dating Committee, the two periods of expansion are March 1991-March 2001 and November 2001-December 2007, while the one recession is March 2001-November 2001.

Table 5 Detection Likelihood Robustness: Industry and Region Subsampling**Panel A: Robustness in Excluding Industries**

	Counts:AA	Detection Likelihood			
	[Non-AA]	Auditor Detected	AAERs	Restatements	SCAC
Full Sample (Table 3)	353 [2404]	0.294**	0.520*	0.337***	0.470**
Excluding Less Represented Industries					
Banks & Insurance	313 [1592]	0.294**	0.608	0.355***	0.544*
Computers	340 [2236]	0.329**	0.593	0.320***	0.468**
Retail & Wholesale	332 [2231]	0.276***	0.458**	0.313***	0.456***
Excluding More Represented Industries					
Refining & Extractive	338 [2360]	0.286***	0.496**	0.327***	0.481**
Communication & Transport	306 [2230]	0.594	0.656	0.357***	0.524**
Utilities	276 [2214]	0.229***	0.536*	0.374***	0.499**
Services & Health	313 [2261]	0.237***	0.426***	0.334***	0.425***

Panel B: Robustness in Excluding Regions and States

	Counts:AA	Detection Likelihood			
	[Non-AA]	Auditor Detected	AAERs	Restatements	SCAC
Full Sample (Table 3)	353 [2404]	0.294**	0.520*	0.337***	0.470**
Excluding Regions:					
Southeast	267 [1928]	0.277**	0.615	0.353***	0.396***
Northeast	290 [1822]	0.398	0.385***	0.311***	0.525*
Midwest	258 [1918]	0.247***	0.523*	0.399***	0.553
Mountain	333 [2325]	0.266***	0.496**	0.349***	0.430***
Southwest	307 [2132]	0.247***	0.537*	0.301***	0.422***
West	323 [2067]	0.365**	0.568*	0.316***	0.529**
Excluding Large States:					
Texas	320 [2172]	0.253***	0.549*	0.323***	0.442***
California	328 [2137]	0.358**	0.558*	0.310***	0.519**
New York	335 [2186]	0.306**	0.460**	0.338***	0.507**

The table presents robustness detection likelihood estimates, where each detection likelihood results from its own Poisson-binomial estimation of each fraud definition, and the odds ratio is the transformation of the coefficient on former Arthur Andersen client, as in Table 3. The only variation is that in Panels A and B we filter out certain industries (Panel A) and geographies (Panel B), as noted in the row labels. The overall sample is US publicly traded corporations with more than \$750 million in assets. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% level from Poisson p-values, testing that the AA coefficient is significantly different from zero or, equivalently, that the odds ratio is different from unity.

Panel A reports that auditor-detected securities frauds hover around 1%, with a peak of 1.1% in 2000–2001. AAER investigations average 2.6%, with a peak of 3.5% in 2000–2001 and a trough of 2.0% in 1998–99. Accounting violations measured by non-clerical restatements average 13.5% during the entire sample period, with a peak of 18.3% in 2002–2003 and a trough of 7.2% in 1998–1999. The broader SCAC securities fraud averages 3.4%, with a peak of 4.8% in 2000–2001 and a trough of 2.3% in 2004–2005.

Table 6 How Pervasive Is Corporate Fraud?

Panel A: Annual Ongoing Misconduct Eventually Caught as % of Firms					
	<i>All</i>	<i>1998–1999</i>	<i>2000–2001</i>	<i>2002–2003</i>	<i>2004–2005m1</i>
Auditor-Detected Fraud	0.81%	1.06%	1.09%	0.62%	–
		[fraud = 44; fraud-years = 113]			
AAERs	2.64%	1.96%	3.54%	2.69%	1.98%
		[fraud = 122; fraud-years = 369]			
Restatements (filtered)	13.46%	7.21%	12.01%	18.29%	17.99%
		[fraud = 676; fraud-years = 1905]			
SCAC Securities Fraud	3.39%	2.90%	4.83%	2.72%	2.31%
	[fraud = 241; fraud-years = 464]				
Panel B: Detection Likelihood Estimates from Table 3					
	<i>Detection Likelihood</i>				
Auditor-Detected Fraud	0.294				
AAERs	0.520				
Restatements	0.337				
SCAC Securities Fraud	0.470				
Best Estimate (average of auditor-detected and restatements)	0.332				
Panel C: Baseline Estimates Pervasiveness of Corporate Fraud					
	<i>Fraud Pervasiveness = Observed Fraud / Detection Likelihood</i>				
	<i>All</i>	<i>1998–1999</i>	<i>2000–2001</i>	<i>2002–2003</i>	<i>2004–2005m1</i>
Auditor-Detected Fraud	2.8%	3.6%	3.7%	2.1%	–
AAERs	5.1%	3.8%	6.8%	5.2%	3.8%
Restatements	40.0%	21.4%	35.7%	54.3%	53.4%
SCAC Securities Fraud	7.2%	6.2%	10.3%	5.8%	4.9%
Panel D: Best Estimate of Pervasiveness of Corporate Fraud					
	<i>Fraud Pervasiveness = Observed Fraud / Detection Likelihood Best Estimate</i>				
	<i>All</i>	<i>1998–1999</i>	<i>2000–2001</i>	<i>2002–2003</i>	<i>2004–2005m1</i>
Auditor-Detected Fraud	2.45%	3.19%	3.30%	1.88%	–
AAERs	7.97%	5.92%	10.69%	8.11%	5.98%
Restatements	40.59%	21.74%	36.22%	55.15%	54.24%
SCAC Securities Fraud	10.21%	8.73%	14.56%	8.20%	6.98%

The table presents estimates of corporate fraud pervasiveness in publicly traded US firms with more than \$750 million in assets, along with the inputs to generate these estimates. Panel A presents the average likelihood a firm is engaging in corporate fraud sometime in monthly data during the date range specified by columns and by the corporate fraud measure in the rows. (The corporate fraud observed percentages in Panel A are all corporate frauds, not just those aligning with the AA experiment timing.) Since AA firms during the detection period are assumed to have a probability of detection equal to one, they are excluded from this calculation. Panel B presents the detection likelihoods from Table 3. Panel C uses the Panel A and B estimates to calculate the pervasiveness of corporate fraud as the observed fraud divided by the detection likelihood. Panel D presents our best estimate, using the Panel B of Table 3 estimate, for the detection likelihood.

Using these observed misconduct averages in Panel A and detection likelihood estimates from Table 3 reproduced in Panel B, Table 6, Panel C reports estimates of the pervasiveness of fraud. A concern with Panel C is that, as we discussed with regard to Table 3, some of the detection likelihood estimates are clearly more biased than others. The biases are always conservative, but they do not help with the precision of our fraud pervasiveness estimates. In particular, our experimental design of using AA's demise has the most power to identify the whole iceberg of hidden fraud (Assumption 2 holding as an equality) for auditor-detected securities fraud and restatements, as opposed to AAERs and the broader SCAC securities fraud cases, since no matter how rigorous the monitoring of the new auditors is, they will find it difficult to identify cases of price fixing such as Sotheby's. Thus, the best estimate detection likelihood is based on the pooled detection likelihood of auditor-detected and restatement measures, or 0.33, as highlighted before. We use this detection likelihood to calculate the best estimate of corporate misconduct across all of our fraud measures in Panel D.

Our findings as to the pervasiveness of corporate fraud depend on the measure of misconduct we use. We find that in any year averaged across the business cycle, 2.5% of large corporations are committing severe financial misreporting that auditors can detect. Auditor-detected securities fraud is a subcategory of SCAC alleged securities fraud; thus, it is not surprising that it has a low frequency. Such a measure is interesting for the detection likelihood estimation, given that it maps well to our AA demise design, but of more interest to speak to fraud at large are the SCAC securities frauds.

We find that, during an average year over the business cycle, 10% of large corporations are committing a misrepresentation, an information omission, or another misconduct that can lead to an alleged securities fraud claim settled for at least \$3 million (with a 95% confidence interval between 7 and 14%). This result from the SCAC data is our main estimate of the pervasiveness of corporate fraud, since SCAC cases are indeed (alleged) securities fraud. By using the AAER measure, we arrive at similar estimates: 8% of fraud pervasiveness, with a 95% confidence interval range of 6%-11%. This magnitude is similar to the SCAC estimate, even though AAERs have lower observed frequencies because their existence requires the SEC to act. (Recall that the SEC failed to act on Madoff despite six substantive complaints.¹⁵)

Accounting violations, less severe than alleged securities fraud, are more prevalent, with an average annual pervasiveness of 41% (95% confidence interval between 30 and 55%). We do not want to conclude from this estimate that each year 41% of large corporations commit a severe misreporting. To reach this conclusion we would need some more substantive filters to eliminate inconsequential misreporting. Nevertheless, this estimate does not bode well for the US auditing system. In spite of all the regulation, roughly half of the US financial statements suffer from misreporting more serious than pure clerical errors.

¹⁵ <https://www.nytimes.com/2009/09/03/business/03madoff.html>.

4.3.1 Comparison with pervasiveness estimates in the literature

Taking the estimate of 10% as our main estimate of the pervasiveness of corporate fraud, we now ask how our estimates lines up with the literature. As the summary reported below shows, our estimate is at the low end of the pervasiveness of corporate fraud found in the literature. This is not surprising, since if Assumption 2 is violated, our estimate represents a lower bound.

We begin with evidence concerning the appetite for corporate misdoing for personal gain. Prior to Lie (2005), the existence of the options backdating practice had not been understood, and thus firms wanting to commit fraud in this manner could do so with little detection threat. Bebchuk et al. (2010) look back over the pre-2005 period and identify the percentage of publicly traded firms from 1996 to 2005 in which CEOs or directors were “lucky” in that they received option grants on the day of the month when the stock price was the lowest. By their estimate, 12% of firms had such lucky CEOs, suggesting that the appetite for fraudulent behavior was present in at least 12% of firms.

That at least 12% of firms are committing some fraud is supported also by survey evidence and accounting prediction models. Dichev et al. (2013) survey 169 CFOs of public companies and find that 18% of firms manage earnings to misrepresent performance. Beneish et al. (2013) build a model, based on Beneish (1999), that out of sample was able to predict successfully 71% of the most important accounting violations. They apply this model to estimate the pervasiveness of accounting manipulation and find that 18% of firms are fraudulent each year during the period 1997–2005.

How do the structural approaches to modelling the hidden iceberg of corporate misconduct line up with the evidence reported thus far? Wang et al. (2010) examine financial fraud among the 3,297 IPOs from 1995 to 2005. While their main goal is to show that fraud is procyclical, their bivariate probit model produces predicted probabilities of engaging in fraud of 10%–15%, very much in line with our estimates.¹⁶ Similarly, in their flexible Bayesian priors approach to partial observability, Hahn et al. (2016) estimate a pervasiveness of SEC-investigated accounting misconduct of 15%. Zakolyukina (2018) uses a structural model to explore detected and undetected GAAP manipulation (more in line with our accounting violation measure) and reports that 73% of CEOs manipulate their financials, with a detection likelihood of only 0.06.

5 How expensive is corporate fraud?

5.1 Expected cost of fraud

The finding that 10% of large corporations are committing fraud does not necessarily imply that the economic cost of fraud is large. To draw this conclusion, we need

¹⁶ We infer this from Fig. 1, predicted probability of fraud, and summary statistics on the distribution of industry EPS growth available in the [internet appendix](#).

to estimate the cost of these corporate frauds. We do not innovate here, but rely on existing estimates of the cost of fraud. As in the prior literature, we focus on the cost of fraud borne by equity holders of firms involved in fraud, ignoring other potential losses (e.g., loss in debt value, loss of employee pensions and jobs, loss of supplier payables, loss of customer prepayments, loss of taxes owed, professional fees and administrative expenses for firms in bankruptcy) and spillover costs borne by competitors (e.g., Bower and Gilson 2003). In this sense, the equity cost estimates provide a lower bound on costs from a social welfare perspective.¹⁷ We build our calculations off two prior papers.

Using an event-study methodology, KLM estimate the reputational loss of detected fraud at 25% of the equity value of a fraudulent firm. This cost, which is almost entirely due to a loss in reputation, represents the present value of the decline in expected cash flows as firms' investors, suppliers, and customers change the terms on which they interact with a fraudulent firm.

KLM's cost estimates are for detected fraud and do not necessarily apply to undetected fraud. Even when a fraud is not yet in the public domain, the firm incurs costs for two reasons. First, the fraud is unlikely to remain a secret for customers and employees, who will seek business relationships or employment elsewhere, demand a premium to remain, or take advantage of the fraud themselves. For example, Bernie Madoff's employees like Frank Di Pasquale were lavishly paid to ensure their silence. In addition, they stole money for themselves.¹⁸ Second, the biggest cost is often the cover-up. For 20 years, the Japanese company Olympus was able to hide a \$730 million financial loss from 1990 with a series of bad acquisitions and accounting tricks. The bad acquisitions alone cost \$300 million.¹⁹

It is hard to put a number on these costs. Yet, if we assume that, at least in the medium term, the stock market is strong-form efficient, the abnormal low returns of companies that are likely to have committed a fraud but were never exposed as fraudulent can provide an estimate of these hidden costs. Beneish et al. (2013) perform this exercise. They compare the annual buy-and-hold return of firms with a high probM score with that of firms with a low probM score. After controlling for a four-factor model, they estimate an annual 10.9% difference in returns. We take this underperformance as an estimate of the costs of undetected fraud.²⁰

We are now in the position to compute the total cost of fraud. Our estimates suggest that 10% of firms are committing fraud. If the detected fraud (representing 33%

¹⁷ These spillovers are even more difficult to estimate than the costs we consider. We ignore them not because they are unimportant, but because we do not need them to reach our conclusion on the benefit of financial regulation purely from an investors' perspective.

¹⁸ <https://www.justice.gov/sites/default/files/usao-sdny/legacy/2012/04/16/20090811dipascaliinformatiionsigned.pdf>.

¹⁹ <https://www.nytimes.com/2011/12/09/business/deep-roots-of-fraud-at-olympus.html>.

²⁰ We use Beneish et al.'s figures as the cost of undetected fraud, conditional on committing fraud. If Beneish et al.'s numbers are both selection and treatment, they represent the unconditional effect of fraud, not the conditional one. In such a case, our method would underestimate the reputational cost of fraud. Beneish et al.'s figures will be an overestimate if in the year after being classified as a high probM firm, that firm reduces its fraudulent financial reporting, as then the estimate would include what KLM call a "readjustment effect."

of cases) costs the firm 25% of market value and the undetected fraud (representing 67% of cases) costs 10.9%, then the cost of an average fraud is 15.6% of firms' market capitalization. Thus, the cost of fraud is 1.6% of a firm's equity value per year (i.e., fraud pervasiveness (10%) multiplied by the loss of market capitalization for an average firm in case of fraud (15.6%)). The 95% confidence interval for cost of fraud is between 1.2% and 2.2% of the firm's equity value.

In 2004, the total capitalization of the US equity market was \$16 trillion.²¹ Since, on average, 10% of firms are engaged in fraud, the annual fraud cost is \$254 billion a year (with a 95% confidence interval between \$189 billion and \$347 billion). If we repeat the calculation at the end of 2021, the expected annual cost of fraud is \$830 billion (with a 95% confidence interval between \$5609 billion and \$1,133 billion).

5.2 An application to cost–benefit analysis

When a firm is 100% owned by one individual, the cost of fraud is fully internalized by the owner. In publicly traded companies, where equity is dispersed and insiders often own only a tiny fraction of the outstanding equity, this is not the case. These agency costs are one of the justifications for the introduction of regulations like Sarbanes–Oxley (SOX). As an illustration of the wide applicability of our estimates, we sketch how it can be used for a cost–benefit analysis of SOX. Note that reducing agency costs is only one of the benefits of SOX. We are completely silent on other potential benefits, like reducing the lemon discount due to the asymmetry of information in the financial accounts.

Hochberg et al. (2009) exploit survey data collected by Finance Executives International (FEI) to arrive at an estimate of \$3.8 million of compliance costs per firm, with costs increasing in the issuer's size.²² To calculate the total compliance costs, we multiply these average compliance costs by the number of publicly traded firms in 2004 (5,226) to obtain an annual compliance cost of SOX of \$19.9 billion per year. Note that faced with a regulatory mandate to increase monitoring in some ways, firms may cut down some of the monitoring they were previously doing in other ways (Karpoff and Lott 1993). Thus, \$19.9 billion represents an overestimate of SOX compliance costs.

The potential benefit of SOX is equal to the net reduction in the probability of fraud times the cost of fraud when it occurs. If instead of 10% of firms engaged in fraud, a policy change resulted in only 9% of firms engaged in fraud, the cost of fraud would drop by \$25 billion (1% times 15.6%, which is the average cost of fraud, times \$16 trillion, which is the 2004 total equity market capitalization). Thus, for the benefits of SOX regulation to exceed the costs, it would suffice if SOX were able to reduce the probability of fraud by one percentage point (a 10% reduction). Establishing whether this was indeed the case is beyond the scope of our paper. The purpose here is simply to illustrate the usefulness of an estimate of the pervasiveness of fraud, which enables us to derive an estimate of the total cost of fraud. With this estimate, we can easily assess what is the minimum level of effectiveness (in terms

²¹ US equity market capitalization statistics for 2004 and 2021 are from <https://statistics.world-exchanges.org/>.

²² See Hochberg et al. (2009), Table 11 on page 571, for the costs by size categories that we use in our estimates.

of reduction in the probability of starting a fraud) that any system of controls needs to achieve to justify its cost.

6 Conclusion

In this paper, we use the natural experiment provided by the sudden demise of a major auditing firm, Arthur Andersen, to infer the fraction of corporate fraud that goes undetected. This detection likelihood is essential to quantify the pervasiveness of corporate fraud in the United States and to assess the costs that this fraud imposes on investors. We find that two out of three corporate frauds go undetected, implying that, pre Sox, 41% of large public firms were misreporting their financial accounts in a material way and 10% of the firms were committing securities fraud, imposing an annual cost of \$254 billion on investors.

These figures project a dismal picture of the effectiveness of financial auditing pre SOX. Whether SOX reduced or eliminated the problem, our paper is unable to answer. Yet, the magnitude of the problem suggests that some action was warranted. Based on our estimates, we can also infer that if the new regulation reduced by 10% the probability of starting a fraud, its cost would be fully justified.

Appendix 1: Dyck et al. (2010) Filters to Eliminate Frivolous Fraud

First, Dyck et al. (2010) restrict attention to alleged frauds in the period of 1996–2004, specifically excluding the period prior to passage of the Private Securities Litigation Reform Act of 1995 (PSLRA) that was motivated by a desire to reduce frivolous suits and among other things, made discovery rights contingent on evidence. Second, they restrict attention to large US publicly traded firms, which have sufficient assets and insurance to motivate law firms to initiate lawsuits and do not carry the complications of cross-border jurisdictional concerns. In particular, they restrict attention to US firms with at least \$750 million in assets in the year prior to the end of the class period (as firms may reduce dramatically in size surrounding the revelation of fraud). Third, they exclude all cases where the judicial review process leads to their dismissal.²³ Fourth, for those class actions that have settled, they only include those firms where the settlement is at least \$3 million, a level of payment that previous studies suggested to divide frivolous suits from meritorious ones.²⁴ Fifth, they exclude those security frauds that Stanford classifies as

²³ They retain cases where the reason for dropping the suit is bankruptcy, for in this instance the cases could still have had merit, but as a result of the bankruptcy status, plaintiff lawyers no longer have a strong incentive to pursue them.

²⁴ Grundfest (1995), Choi (2007) and Choi et al. (2009) suggest a dollar value for settlement as an indicator of whether a suit is frivolous or has merit. Grundfest establishes a regularity that suits which settle below a \$1.5–\$2.5 million threshold are on average frivolous. The range on average reflects the cost to the law firm for its effort in filing. A firm settling for less than \$1.5 million is most almost certainly just paying lawyers' fees to avoid negative court exposure. To be sure, we employ \$3 million as our cutoff.

nonstandard, including mutual funds, analyst, and IPO allocation frauds.²⁵ The final filter removes a handful of firms that settle for amounts of \$3 million or greater but where the fraud, upon their reading, seems to have been settled to avoid the negative publicity.²⁶

Appendix 2: Calculation of Beneish's Probability of Manipulation Score (ProbM Score)

The components in the ProbM Score include days' sales in receivables, gross margin, asset quality index, sales growth index, depreciation index, SGA index, leverage, and the ratio of accruals to assets. (Please refer to Beneish (1999) for motivation of how each of these subindices captures an aspect of manipulation.) To construct the ProbM Score, we use Compustat data to construct the variable components following Beneish (1999) and apply his estimated coefficients.

The probability of manipulation, ProbM Score, of Beneish (1999) is calculated as follows:

$$ProbM = -4.84 + 0.92*DSR + 0.528*GMI + 0.404*AQI + 0.892*SGI + 0.115*DEPI + 0.172*SGAI + 4.679*ACCRUALS - 0.327*LEVI$$

The variable codes are defined as follows:

DSR	Days Sales in Receivables
GMI	Gross Margin Index
AQI	Asset Quality Index
SGI	Sales Growth Index
DEPI	Depreciation Index
SGAI	Sales, General and Administrative expenses Index
ACCRUALS	Total Accruals to total assets
LEVI	Leverage Index

²⁵ Stanford Class Action Database distinguishes these suits because all have in common that the host firm did not engage in wrongdoing. IPO allocation cases focus on distribution of shares by underwriters. Mutual fund cases focus on timing and late trading by funds, not by the firm in question. Analyst cases focus on false provision of favorable coverage.

²⁶ The rule they apply is to remove cases in which the firm's poor ex post realization could not have been known to the firm at the time the firm or its executives issued a positive outlook statement for which they are later sued.

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s1142-022-09738-5>.

Acknowledgements We thank Patricia Dechow, Michael Guttentag, Katherine Guthrie, Phil McColough, Frank Partnoy, Richard Sloan, Joseph P. Weber, Michael Weisbach, an anonymous referee, and participants at the NBER Crime Working Group, the University of Illinois Symposium on Auditing Research, the European Finance Association, the University of Chicago, USC, the Rotman School at the University of Toronto, Berkeley-Haas Accounting, Berkeley-Haas Finance, and Queens University. Alexander Dyck thanks the Connaught Fund of the University of Toronto, and Adair Morse and Luigi Zingales thank the Center for Research on Security Prices, the Stigler Center, and the Initiative on Global Financial Markets at the University of Chicago, for financial support. Luigi Zingales has worked as a consultant for the PCAOB.

Data Availability The article uses four different measures of fraud. Financial misrepresentations exposed by auditors are identified by Dyck et al. (2010) derived from data collected from the Stanford Securities Class Action Clearinghouse [<https://securities.stanford.edu/>]. A summary of cases and whistleblowers identified by Dyck et al. (2010) is available here [https://docs.google.com/spreadsheets/d/1e-1fNFIjJC4yHvNLCE1rFy-JZFvbFPHW/edit?usp=share_link&ouid=112993024848888180836&rtpof=true&sd=true]. SEC section 10b-5 securities fraud cases are identified by Kempf and Spalt (2022), and derived from data collected from the Stanford Securities Class Action Clearinghouse [<https://securities.stanford.edu/>]. Restatements are based on data from the Audit Analytics database. Financial Misrepresentations that led to an Accounting and Auditing Enforcement Release are identified by Dechow et al. (2011) based on original data from the SEC website [<http://www.sec.gov/divisions/enforce/friactions.shtml>] and their updated dataset is available here [<https://sites.google.com/usc.edu/aaerdataset/home>].

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Agresti, A. 1992. A survey of exact inference for contingency tables. *Statistical Science* 7 (1): 131–153.
- Agrawal, A., and S. Chada. 2005. Corporate Governance and accounting scandals. *Journal of Law and Economics* 48 (2): 371–406.
- Barton, Jan. 2005. Who cares about auditor reputation?*. *Contemporary Accounting Research* 22 (3): 549–586. <https://doi.org/10.1506/C27U-23K8-E1VL-20R0>.
- Bebchuk, L., Y. Grinstein, and U. Peyer. 2010. Lucky CEOs and lucky directors. *Journal of Finance* 65 (6): 2363–2401.
- Becker, G.S. 1968. Crime and punishment: An economic approach. *Journal of Political Economy* 76 (2): 169–217.
- Beneish, M.D. 1999. The detection of earnings manipulation. *Financial Analysts Journal* 55 (5): 24–36.
- Beneish, M.D., C.M.C. Lee, and D. Nichols. 2013. Earnings manipulation and expected returns. *Financial Analysts Journal* 69 (2): 57–82.
- Black, B., B.R. Cheffins, and M. Klausner. 2006. Outside director liability. *Stanford Law Review* 58 (4): 1055–1159.
- Bower, J.L., and S.C. Gilson. 2003. The social cost of fraud and bankruptcy. *Harvard Business Review* 81 (12): 20–22.
- Chaney, Paul K., and Kirk L. Philipich. 2002. Shredded reputation: The cost of audit failure. *Journal of Accounting Research* 40 (4): 1221–1245. <https://doi.org/10.1111/1475-679X.00087>.

- Chen, K.Y., and J. Zhou. 2007. Audit committee, board characteristics and auditor switch decisions by Andersen's clients. *Contemporary Accounting Research* 24 (4): 1085–1117.
- Choi, S.J. 2007. Do the merits matter less after the private securities litigation reform act? *Journal of Law, Economics and Organization* 23 (3): 598–626.
- Choi, S.J., K.K. Nelson, and A.C. Pritchard. 2009. The screening effect of the securities litigation reform act. *Journal of Empirical Legal Studies* 6 (1): 35–68.
- Coates, J.C., and S. Srinivasan. 2014. SOX after ten years: A multidisciplinary review. *Accounting Horizons* 28 (3): 627–671.
- Coffee, J.C. 1986. Understanding the Plaintiff's Attorney: The implications of economic theory for private enforcement of law through class and derivative actions. *Columbia Law Review* 86 (4): 669–727.
- Dechow, P.M., R. Sloan, and A. Sweeney. 1996. Causes and consequences of earnings manipulation: An analysis of firms subject to enforcement actions by the SEC. *Contemporary Accounting Research* 13 (1): 1–36.
- Dechow, P.M., G.E. Weili, C.R. Larson, and R.G. Sloan. 2011. Predicting material accounting misstatements. *Contemporary Accounting Research* 28 (1): 17–82.
- Dichev, I., J. Graham, C. Harvey, and S. Rajgopal. 2013. Earnings quality: Evidence from the field. *Journal of Accounting and Economics* 56: 1–33.
- Dyck, A., N. Volchkova, and L. Zingales. 2008. The corporate governance role of the media: Evidence from Russia. *Journal of Finance* 63 (3): 1093–1136.
- Dyck, A., A. Morse, and L. Zingales. 2010. Who blows the whistle on corporate fraud? *Journal of Finance* 65 (6): 2213–2254.
- Eisenberg, T., and J.R. Macey. 2004. Was Arthur Andersen different? An empirical examination of major accounting firm audits of large clients. *Journal of Empirical Legal Studies* 1 (2): 263–300.
- Grundfest, J.A. 1995. Why disimply? *Harvard Law Review* 108: 727–747.
- Hahn, P.R., J.S. Murray, and I. Manolopoulou. 2016. Bayesian partial identification approach to inferring the prevalence of accounting misconduct. *Journal of the American Statistical Association* 111 (513): 14–26.
- Hochberg, Y.V., P. Sapienza, and A. Vissing-Jorgensen. 2009. A lobbying approach to evaluating the Sarbanes-Oxley. *Journal of Accounting Research* 47 (2): 519–583.
- Jensen, M., and W. Meckling. 1976. The theory of the firm: managerial behavior, agency cost, and ownership structure. *Journal of Financial Economics* 3 (4): 305–360.
- Karpoff, J.M., and J.R. Lott. 1993. The reputational penalty firms bear from committing criminal fraud. *Journal of Law and Economics* 36 (2): 757–802.
- Karpoff, J.M., D.S. Lee, and G.S. Martin. 2008. The cost to firms of cooking the books. *Journal of Financial and Quantitative Analysis* 43 (3): 581–612.
- Kealey, B.T., H.Y. Lee, and M.T. Stein. 2007. The Association between audit-firm tenure and audit fees paid to successor auditors: evidence from Arthur Andersen. *Auditing: A Journal of Practice & Theory* 26 (2): 95–116.
- Kempf, E., and O. Spalt. 2022. Attracting the sharks: Corporate innovation and securities class action lawsuits. *Management Science*, forthcoming.
- Kohlbeck, Mark, Brian W. Mayhew, Pamela Murphy, and Michael S. Wilkins. 2008. Competition for Andersen's clients. *Contemporary Accounting Research* 25 (4): 1099–1136. <https://doi.org/10.1506/car.25.4.6>.
- Krishnamurthy, S., J. Zhou, and N. Zho. 2006. Auditor reputation, auditor independence and the stock market impact of Andersen's indictment on its client firms. *Contemporary Accounting Research* 23 (2): 465–490.
- Krishnan, G.V. 2007. Did earnings conservatism increase for former Andersen clients? *Journal of Accounting, Auditing and Finance* 22 (2): 141–163.
- Krishnan, J., K. Raghunandan, and J.S. Yang. 2007. Were former Andersen clients treated more leniently than other clients? Evidence from going-concern modified audit opinions. *Accounting Horizons* 21 (4): 423–435.
- Krishnan, Gopal, and Gnanakumar Visvanathan. 2009. Do auditors price audit committee's expertise? The case of accounting versus nonaccounting financial experts. *Journal of Accounting, Auditing & Finance* 24 (1): 115–144. <https://doi.org/10.1177/0148558X0902400107>.
- Lie, E. 2005. On the timing of CEO stock option awards. *Management Science* 51 (5): 802–812.
- Miller, G.S. 2006. The press as a watchdog for accounting fraud. *Journal of Accounting Research* 44 (5): 1001–1033.

- Nagy, A.L. 2005. Mandatory audit firm turnover, financial reporting quality, and client bargaining power: The case of Arthur Andersen. *Accounting Horizons* 19 (2): 51–68.
- Soltes, E.F. 2019. The frequency of corporate misconduct: Public enforcement versus private reality. *Journal of Financial Crime*. 26 (4): 923–937.
- Wang, T.Y. 2013. Corporate securities fraud: Insights from a new empirical framework. *Journal of Law, Economics and Organization* 29 (3): 535–568.
- Wang, T.Y., A. Winton, and X. Yu. 2010. Corporate fraud and business conditions: Evidence from IPOs. *Journal of Finance* 65 (6): 2255–2292.
- Zakolyukina, A. 2018. How common are intentional GAAP Violations? Estimates from dynamic model. *Journal of Accounting Research* 56 (1): 5–44.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.