

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

August 2014

ABSTRACT

We estimate the pervasiveness and the cost of corporate fraud. To identify the potential ‘iceberg’ of undetected fraud we take advantage of Arthur Andersen’s demise, which forces companies to change auditors and exposes preexisting frauds. This experiment suggests that only one quarter of frauds are detected in normal times, and leads us to infer that in the 1996-2004 period on average one out of seven large publicly-traded US firms was engaged in fraud. We obtain similar estimates by using an alternative approach. Firms that engage in fraud destroy on average one fifth of their value. These estimates set the average cost of fraud in large corporations to be \$380 billion a year.

* We thank Patricia Dechow, Michael Guttentag, Katherine Guthrie, Phil McCollough, Frank Partnoy, Joseph P. Weber, Michael Weisbach as well as 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, 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 is a consultant to the PCAOB.

More than ten years after the passage of Sarbanes-Oxley (SOX) regulation there is still a very active debate about its costs and benefits (Coates and Srinivasan, 2014). While the direct costs are (relatively) easy to quantify, its benefit – a reduction in corporate fraud – is much more difficult to assess. Not only is the reduction in the cost of corporate fraud due to SOX difficult to estimate, but so is the magnitude of the pre-existing problem, i.e., how pervasive and costly fraud was before SOX. If there were just a few rotten apples, large-scale intervention might have been a waste of energy and resources.

In this paper we try to tackle this question. The average cost of fraud around the passage of SOX represents an upper bound on the benefit of any regulation targeted at fraud reduction. Thus, if the average cost of fraud is small relative to the cost of SOX, we can easily sign the net benefit. If the cost of fraud is not small, then we can determine how big should be the reduction in the probability of fraud to justify the cost of regulation. An estimate of the cost of fraud would also help investors and boards tailor resources to mitigate the scope of the problem.

The main challenge in assessing the cost of fraud is that we only observe detected fraud: we do not know if the observed fraud is the whole iceberg, or just its visible tip. To resolve this problem and estimate the probability of corporate fraud we appeal to basic probability rules. The probability we observe a fraud is given by the joint probability that managers engage in a fraud and that they get caught, $\Pr(\text{engage, caught})$. Yet, what we are ultimately interested in is the unconditional probability of engaging in a fraud: $\Pr(\text{engage})$. This is the product of the detection likelihood $\Pr(\text{caught}/\text{engage})$ and the observed probability of engaging and getting caught. Thus, if we knew the detection likelihood we could easily calculate the probability of engaging in fraud. Our identification strategy exploits circumstances in which the likelihood of being caught increases to almost one. By comparing detection in this special circumstance to normal circumstances, we produce an estimate of the size of the iceberg below the water. From there, it is a short step to derive the unconditional likelihood of engaging in fraud.

Our primary test takes advantage of the natural experiment created by the 2002 demise of Arthur Andersen (AA) following the criminal indictment by the U.S. Department of Justice. That demise suddenly forced all firms that employed AA to seek another auditor. Any new auditor has the incentive to

reveal the mistakes of the previous one, before they become his own mistakes. This incentive was particularly acute, following the revelation of accounting fraud in Enron and WorldCom, when the market was putting accounting firms under the microscope. Thus, there was little reason for new auditors to cover up pre-existing fraud.

Consistent with this argument, we find that the incidence of fraud detection by auditors goes up significantly. The primary data for these tests comes from the Dyck, Morse and Zingales (2010) dataset (hereafter DMZ) that provides a sample of detected frauds and for each fraud identifies a whistleblower. We also take advantage of data on firms receiving a Securities and Exchange Commission Accounting, Auditing, and Enforcement Release (AAER), a widely used indicator of financial misrepresentation in the literature.

After the forced turnover following the demise of AA, the new auditors are roughly four times more likely to reveal that the firm had an ongoing accounting fraud than firms in the same period that did not have AA as their prior auditor. To be precise, the detection rate is 26.7%. This implies that the iceberg below the water is three times its visible tip. Since on average in the 1996-2004 period 4.0% of companies are detected to be engaging in fraud, it follows that the fraction of large publicly traded corporations engaging in fraud is 15% (i.e., $4.0 / 0.267$).

We validate these results by following an alternative approach. Conditional on a fraud being committed, the probability it is revealed is a positive function of the incentives and the opportunities for detection. For example, DMZ show that a fraud is more likely to be detected in a company followed by a large number of analysts. Thus, instead of focusing just on the changes in auditors' incentive to detect fraud (as for the AA case), we also look at changes in incentives and opportunities across all potential fraud detectors.

DMZ identify six sources of variations in the incentives/opportunities to detect fraud. By using these six sources of variation, we estimate what happens to the likelihood of detection when all incentives are simultaneously 'high.' By comparing the odds of detection in this 'high' detection state to the odds in

a ‘normal’ state, we derive an alternative estimate of the likelihood of detection. With this alternative method, we find a likelihood of detection similar to the one we find with the AA natural experiment.

Having determined the likelihood of engaging in a fraud, in Section 4 we offer an estimate of the net costs of fraud, detected and undetected. The costs we have in mind are the value destruction over and above poor fundamental performance which might have triggered the fraud in the first place. It is the value destruction due to the drop in reputation caused by the fraud, which will be capitalized in the firms’ enterprise value. As Karpoff and Lott (1993) and Karpoff, Lee and Martin (2008) show, these reputational losses can be quite large. In fact, we expect these losses to be larger in countries like the United States, where the governance system is relatively good and thus fraud is unexpected.

We construct a new measure of the cost of fraud, equal to the difference between the enterprise value after the fraud is revealed and what the enterprise value would have been in the absence of fraud. We construct this hypothetical value by making projections from the pre-fraud period, assuming the trajectory would have followed that of other firms in the same industry. Using this approach, we estimate that the median fraud firm loses 30 percent of its enterprise value when a fraud is detected.

Putting together our estimates of the pervasiveness of fraud with this estimate of the per-firm cost of fraud, we arrive to an expected cost of fraud equal to 3% of enterprise value of each firm. In 2004 the total enterprise value of U.S. firms with greater than \$750m in assets was \$21 trillion, this estimate implies a residual loss from detected and undetected fraud of \$630 billion. Since the average fraud lasts 1.67 years, the annual cost of fraud among large US corporations is \$380 billion.

Our paper builds on a rich literature that measures financial fraud, summarized in Karpoff, Koester, Lee and Martin (2013) (hereafter KKLM). Besides detected fraud, we focus on measuring undetected fraud, as in Wang (2013) and Zakolyukina (2013). In addition, we do not restrict ourselves to financial misrepresentations, but we also look at other types of fraud (e.g. lying about the future). Finally, we build on Karpoff and Lott (1993) and Karpoff, Lee and Martin (2008) to compute the total cost of this fraud.

The rest of the paper proceeds as follows. Section I describes the data and presents summary statistics on caught frauds. Section II provides our core estimate of undetected frauds, introducing our methodology and the quasi natural experiment provided by the demise of Arthur Anderson. Section III validate these estimates by applying a similar logic to a broader set of fraud detectors. Section IV provides costs estimates. In Section V we discuss the implications for regulation. Section VI concludes.

I. Data on Caught Fraud Incidence

Care is required to assemble data on fraud. KKLM provide a comprehensive overview of the strengths and weaknesses of the four most popular databases for capturing an important type of fraud: financial misrepresentations. They compare each of these databases against a hand-collected data set they compiled with the complete case histories of 1,099 examples. While they point to a number of possible biases in the conventional databases, the two most important ones are over-inclusiveness and over restrictiveness. Restatements include many non-material cases. For example, Hennes et al. (2008) categorize 73.6% of GAO restatements as unintentional misapplications of GAAP accounting. By contrast, focusing only on firms where there is a Securities and Exchange Commission Accounting, Auditing, and Enforcement Release (AAER), (e.g. Dechow, Sloan and Sweeney (1996), and Miller (2006)), first, eliminates frauds due to a lack of disclosure or misleading disclosure, and second, may be overly restrictive in that, because of a limited budget, the SEC cannot go after all frauds.¹

Instead, we focus our tests on the DMZ fraud sample, which are firms subject to securities class action lawsuits, as compiled by the Stanford Securities Class Action Clearinghouse (SCAC)². DMZ argue that this sample is close to the population of caught fraud for large (over \$750 million in assets) publicly-traded companies because of the incentive structure for law firms. Class action law firms have automated the mechanism of filing class action suits: every time there is a large drop in the share price of a large corporation, specialist attorneys start searching for a cause to file a suit. Since stock prices drop when a

¹ See AAER-1, 1982 SEC LEXIS 2565, May 17, 1982 and discussion in KKLM.

² The list of firms and the classification of whistleblowers is available from the websites of the authors.

serious corporate fraud is revealed, it is highly unlikely that SCAC would miss any sizeable corporate fraud, because a suit will be filed (Coffee, 1986). The evidence in KKLM is broadly supportive of the claim that SCAC data is inclusive, finding it had the lowest omission rate relative to their hand collected sample across the four databases they considered.

The biggest potential problem with using class action data is that it might over include frivolous cases. DMZ use a filtering process, summarized in the appendix of this paper, to remove this concern. The gist of the screening is to restrict attention: first, to cases after a 1995 change in the law 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 settlements amounts. We follow the legal literature's guidance as to what settlement amounts constitute nominal payments to make the suit go away.

It is worth noting that while we use the term fraud, these cases are better thought of as 'alleged frauds'. Security class action cases are almost always settled (to protect executives from personal liability), and settlements almost always involve no admittance of wrongdoing. This does not overly concern us since we are more interested in capturing extreme forms of costly opportunism and less interested in establishing intent. Nonetheless, in the rest of the paper we will use for simplicity the term fraud, and will not append the adjective "alleged".

In total, the sample includes 210 frauds detected in the 1996-2004 period.³ These frauds include all of the high profile frauds such as Enron, Worldcom, Adelphia and Healthsouth, as well as many others. The class action database provides start and end dates for the frauds.⁴ The frauds in the sample have an average duration (from the class action suit period) of approximately 1 year and 8 months (590 days). To gauge the pervasiveness of fraud, we also have to identify the possible population of firms that

³ We drop 6 frauds from the DMZ sample because they are not over the \$750 million threshold when the fraud started.

⁴ Because these dates can be, and often are, revised as suits progress, we use the most recent definition of the suit window from the legal filings. This definition of duration may be conservative in that the statute of limitations on class actions under Section 10(b) of the Exchange Act dictates that cases must be brought within one year after discovery of the alleged violation, and no more than three years after the violation occurred. This limit was loosened in 2002 as Sarbanes-Oxley legislation changed this to 2 years after discovery, and no more than 5 years after the violation occurred.

could have produced frauds. The relevant population for our purposes is the set of U.S. publicly-traded companies with \$750 million in assets. In Compustat, 1,748 companies on average per year meet this criterion.

With the information on start dates of frauds that are caught, duration of these caught frauds, and the population of firms “at risk” of fraud, we can calculate the percentage of caught frauds in the population of firms. Figure 1 illustrates this distribution. The grey bars represent the percentage of large U.S. publicly traded companies that start to commit fraud that year, while the black bars represent the percentage of all firms engaging in fraud that year. There seems to be a non-trivial level of fraud, with an average of 1.36 percent of firms starting to commit fraud each year and 3.4 percent of firms engaged in fraud that will eventually be caught.

It is important to estimate the incidence of fraud over a time period that involves booms and busts, as this period does, as fraud activity appears to be pro cyclical.⁵ Note the significant time series variation in these numbers, with the incidence of firms starting fraud peaking in 2000 (2.4 percent of large corporations), and the fraction of firms engaging in fraud peaking in 2001 (5.5 percent of large corporations).

Figure 2 corrects for the fact that there are some frauds taking place during our sample period that will be caught only after we ended our sample collection in 2004. To extrapolate these missing frauds we use the distribution of fraud duration for those cases which begin prior to the year 2000 to forecast how many cases are yet to be caught for frauds starting through 2003. Using the duration distribution, we then roll the distribution forward to forecast how many additional cases that began after 1999 will yet be caught. This correction raises our estimate of the overall incidence of firms starting fraud that will be caught to 1.43 percent per year and of the overall fraction of firms that in any year engage in frauds that will be caught to 4.0 percent of firms. This is the first key number for our calculations.

⁵ Galbraith (1961) was one of the first to conjecture such a relationship. This has been explored more recently in theoretical models in Povel, Singh and Winton (2007) and an empirical application in Wang, Winton and Yu (2010).

Again, we focus on the average over a period with cycles and boom as the data show significant time series variation with a much higher incidence of frauds starting prior to the demise of Enron and Arthur Andersen in 2001 and the passage of SOX in 2002.

II. Inferring Undetected Frauds and the Arthur Andersen Natural Experiment

The finding that there is an ongoing fraud in 4 percent of large traded firms suggests fraud is not as rare even in the United States. Yet, this result provides an incomplete picture, as it ignores the fact that some frauds are never caught. Without exploring the likelihood of undetected fraud, we do not know if these observed estimates are the whole iceberg, or just the tip of the iceberg. In this section we describe our strategy for identifying undetected frauds. We present some data to validate the assumptions in our identification strategy, and then test to see if undetected fraud is trivial or substantial, providing an estimate of the hidden portion of the iceberg.

II.1. Experiment Overview

Our identification strategy for inferring unobserved fraud relies on the Kolmogorov axiom of conditional probability. What we observe is the joint event of a firm engaging in fraud and being caught: $\Pr(\mathbf{F}, \mathbf{caught})$. (We will use the convention of bolding the variables we observe.) Our actual variable of interest is the probability of a firm engaging in fraud, regardless of whether it is caught or not: $\Pr(F)$. By the law of conditional probability, the unconditional probability of engaging in a fraud can be written as:

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

Thus, if we knew the denominator, the probability that a fraud is caught given that it is ongoing, $\Pr(\mathbf{caught} | F)$ (we will call this the detection likelihood), we could easily calculate $\Pr(F)$. In the rare events where $\Pr(\mathbf{caught} | F)$ is equal to one, then the unobserved $\Pr(F)$ would simply be equal to the observed $\Pr(\mathbf{F}, \mathbf{caught})$.

Our strategy exploits circumstances in which the likelihood of being caught by a particular fraud detector increases to close to one. If we assume that in these special situations $\Pr(\text{caught}/F)$ is exactly one, then to estimate the detection rate of fraud in normal times it is sufficient to compare the observed fraud rate in these circumstances to the observed fraud rate in normal circumstances. Note that if the true $\Pr(\text{caught}/F)$ is less than one, we end up overestimating the detection rate, underestimating the overall incidence of fraud. Thus, our estimates should be considered very conservative.

Our main experiment uses the sudden demise of Arthur Andersen (AA) as a situation in which the likelihood of being caught for certain types of fraud approaches one. This is an important event, since it involves all firms with AA as an auditor, roughly one fifth of all large publicly traded firms in 2001.

Since auditors can only detect frauds that have a financial component, we divide frauds between financial frauds (FF), which are catchable by an accountant, and non-financial frauds (NFF), which are not catchable by an accountant. For the moment we focus solely on detection likelihood for financial frauds.

II.2. Key Identification Assumptions

Although the experiment conditions are fairly intuitive, we want to be precise as to the assumptions we need. We begin with a baseline assumption about AA clients.

Assumption 1: $\Pr(FF / \bar{A}) = \Pr(FF / A)$

Assumption 1 says that, prior to 2002, financial fraud was equally likely in AA firms (A) and non-AA firms (\bar{A}).⁶

AA's indictment by the Department of Justice and its initial conviction for obstruction of justice may make this a surprising assumption. Yet a range of studies, largely in accounting, have come to the conclusion that there was no difference between AA and other auditors. In a matched sample, Agrawal & Chada (2005) find that the existence of AA as the auditor is not associated with firms having more

⁶ To keep notation simple, we do not include time subscripts, but for all the equations, the AA marker means that the firm was an Arthur Andersen client coming into the demise of 2001-2002.

restatements. Likewise, controlling for client size, region, time and industry, Eisenberg & Macey (2004) find that AA clients did not perform any better or worse than other firms. Furthermore, the initial conviction of AA was for obstruction of justice in a particular case (Enron), not for being a bad auditor, and it was unanimously overturned upon appeal.

Our sample of large U.S. corporations is different from the aforementioned studies, and we wanted to verify whether Assumption 1 holds for our sample. We test whether pre-indictment Arthur Andersen clients differ from non-AA clients using the earnings manipulation score “ProbM-Score” of Beneish (1999). We focus on the 1998-2001 period, immediately preceding the problems in 2001 and 2002. The components in the ProbM Score are drawn entirely from financial statements: days sales in receivables, gross margin, asset quality index, sales growth index, depreciation index, SGA index, leverage, and the ratio of accruals to assets. Beneish (1999) motivates how each of these subindices captures an aspect of manipulation. To construct the ProbM Score, we download the appropriate financial statement variables from Compustat, construct all the components directly following the data definitions in Beneish (1999), and use Beneish’s estimated coefficients to construct the ProbM Score. Appendix II details the equation calculation. ProbM is a score with no natural scale. The mean and standard deviation of ProbM are -2.325 and 1.357, respectively.

In Tables 1 and 2, we test whether AA clients and non-AA clients differ in their manipulation scores prior to AA’s demise. Panel A of Table 1 compares AA clients with two comparison groups: all non-AA clients that meet our size criteria; and, non-AA clients that are clients of the Big Five audit firms (perhaps a more appropriate reference group) and meet our size criteria. Panel A shows no significant differences between AA and non-AA clients in the probM score and in all of the 8 sub-components. Between AA and top-5 Non AA audited firms, there again are no differences in the probM score or any of its sub-components. The last row of Panel A also reports that AA clients are no more likely than other firms to start a fraud that is caught prior to 2001.

This test might not be altogether convincing, as the AA and non-AA firms could differ on other dimensions. To address this concern, we provide OLS and quantile (median) regressions of the probM

score on an AA dummy and a set of control variables for size, debt and profitability as well as 2-digit industry-crossed-with year fixed effects. Panel B of Table 1 sets the stage for this test providing summary statistics of the firm-level controls.

Table 2 reports our multivariate regression tests where the dependent variable is the ProbM score. We focus on the dummy for Arthur Anderson and find across all specifications that being an AA client does not significantly impact the probM score, and the economic magnitudes of the coefficients are small. In column 1, we use an OLS specification, include all firms, and the listed controls as well as year-industry (2 digit SIC) dummies to capture variation across industries and years. In column 2, we repeat this test in this case restricting our attention to the sample with a top-5 auditor. Columns 3 and 4 repeat the analysis using a median specification, with column 4 restricted to the sample with top-5 auditors. We conclude that, as was found in prior studies, AA clients are not statistically or economically different from other auditor's clients prior to 2002.

Our second identifying assumption is that, post-AA demise, the probability of detecting an ongoing financial fraud in a former AA client increases to one.

$$\textit{Assumption 2:} \quad \Pr(\textit{caught} \mid FF, A) = 1 \quad (2)$$

The time frame covers frauds starting prior to or during 2002, and catching done in 2002 or later.

The intuition for the assumption is as follows. In the fall of 2001, Enron's collapse triggered accusations about AA. AA was indicted in March 2002, and convicted in June, 2002. Over the period 2001-2002, all of AA's clients had to change their external auditor. Because new auditors do not want to face litigation risk or reputation risk for actions (or non-actions) taken by prior auditors, new auditors have a strong incentive to "clean house". These incentives were further strengthened by the circumstances prevailing at the time. Following the revelation of accounting fraud in Enron and WorldCom, the market was putting accounting firms under the microscope. Thus, there were few incentives to cover up. Furthermore, the disappearance of a large auditing firm and the sudden need of AA-audited firms to find

new auditors shifted the balance of power in favor of auditors. This conservative reporting bias among auditors of former AA clients is now well documented (Cahan and Zhang (2006), Krishnan (2007)).

Cleaning house implies that new auditors address any potentially misleading financial reporting, ranging from gross errors to overly aggressive financial reporting. An important advantage of this approach is that the line of causality is clear: the turnover leads to fraud revelation, rather than fraud leading to auditor's turnover.

Now we are in position to use the law of conditional likelihood, along with these two assumptions, to provide a formula that shows one can infer the detection likelihood in normal times solely from observables. The first step is to produce a version of equation (1) for both former AA clients and for firms that never had AA as an auditor. Rearranging this equation we construct the following ratio of financial frauds that were started and caught for AA firms and non AA ones⁷:

$$\frac{\Pr(\text{caught}|FF, \bar{A}) \Pr(FF|\bar{A})}{\Pr(\text{caught}|FF, A) \Pr(FF|A)} = \frac{\Pr(FF, \text{caught}|\bar{A})}{\Pr(FF, \text{caught}|A)}$$

Substituting in this equation Assumptions 1 and 2, it simplifies to

$$\hat{\Pr}(\text{caught} | FF, \bar{A}) = \frac{\Pr(FF, \text{caught} | \bar{A})}{\Pr(FF, \text{caught} | A)}. \quad (3)$$

Equation (3) makes it clear how in the period following AA's unexpected demise we can estimate the detection likelihood for financial fraud for non AA firms solely from two observables: the detection rate for financial frauds by auditors in non-AA firms; and, the detection rate by auditors for financial frauds in former AA firms. If we test and find that these two numbers are the same, then the detection likelihood is indistinguishable from 1, and we would conclude that there is no iceberg of undetected fraud. However, if the test shows that these two numbers are different, the inverse of the detection likelihood provides an estimate of the size of the iceberg.

⁷ Specifically, we rearrange the following: $\frac{\Pr(\text{caught}|FF, \bar{A}) \Pr(FF|\bar{A})}{\Pr(FF, \text{caught}|\bar{A})} = 1 = \frac{\Pr(\text{caught}|FF, A) \Pr(FF|A)}{\Pr(FF, \text{caught}|A)}$.

II.3 Tests to Determine the Detection Likelihood using the AA Experiment

We now provide tests of the detection likelihood using the AA experiment. So long as the AA demise produces a situation where the probability of detection approaches one for financial frauds, we can infer the detection likelihood for financial frauds by constructing the ratio of detection for non AA clients relative to detection for AA clients. For all tests, we take all large corporations that existed in 2002 and had an auditor identified in Compustat in either 2001 or 2002. We code the firm as an AA client if the auditor was AA in either 2002 or 2001. We put no restriction on survival post 2002 and do not include new firms entering post 2002. To capture the staggering of the changes in auditors and to allow the new auditors time to process all of the new client accounts, we measure a firm as having fraud if the fraud started prior to and including 2002, and the fraud was revealed in 2002 or later.

In the first test we use the DMZ data, and for our estimate of the frauds detectable by auditors we focus on those frauds in their sample where the whistleblower was an auditor, or the whistleblower is identified as the firm and the firm's financials received a qualified opinion from the auditor.⁸ In the second test, we also use the DMZ data and use as our estimate of detected financial frauds the frauds identified as financial misrepresentations, regardless of what whistleblower brought the fraud to light. We categorize a fraud as being a financial fraud if it restated financials after the revelation of the fraud. In the third test we use AAER data rather than the DMZ data. Firms that received an AAER are classified as having a financial fraud. This data is from the Center for Financial Reporting and Management, Berkeley.

Table 3 provides the results of these tests. In the first row of panel B we present the results for the DMZ auditor sample. We find that non AA firms (i.e. their auditor in 2001 or 2002 was not AA) had a 0.57 percent chance of having an auditor reveal a fraud that started prior to 2002 and was revealed in the 2002-2006 period. In contrast, we find that firms that had AA as their auditor in 2001 or 2002, and experienced a forced change in auditor, had a 2.17 percent chance of having an auditor reveal a fraud that

⁸ Results are quantitatively and qualitatively similar if we restrict our attention solely to observations where the auditor is the whistleblower.

started prior to 2002 and was revealed in the 2002-2006 period. We test if these numbers are different from each other, and find that they are at the 1% significance level. In panel C, we therefore construct the detection likelihood for financial frauds as the ratio of these numbers. It amount to 0.267, as we can see from the first column. In other words, our test suggests that in normal times only 27% of financial frauds detectable by an auditor are revealed, implying that the iceberg of undetected fraud is about three times the visible tip.

We provide similar tests for the other two financial fraud samples: financial frauds in the DMZ sample (as captured by restatements), and firms receiving an AAER. We find that detection rates for the non AA sample are again significantly different than for the former AA sample. For financial frauds in the DMZ sample, these numbers are statistically different from each other at the 2% level, and for the AAER sample at the 1% significance level. We go on to construct the detection likelihood from these numbers, and find it is 0.451 for the financial misrepresentation sample, and 0.570 for the AAER sample. Since the detection likelihood hovering around a half, these tests imply that the part of the iceberg under the water is equal in size to the visible tip.

The lower economic magnitude using alternative samples is understandable given the nature of the AA shock. This shock most directly impacted auditor detectable financial frauds and had a less direct impact on financial misrepresentations more generally. Finally, for reference purposes we provide similar estimates using as our sample of frauds the complete DMZ fraud database, including financial and non-financial frauds. Again, the detection rates in the former AA and non-AA samples are statistically different from each other, and the implied detection likelihood is 0.54.

II.4 Estimates for Unconditional Probability of Engaging in Fraud

For investors, boards of directors and policy makers, the overall incidence of fraud is more important than whether it can be defined as a financial or a non-financial fraud. With this in mind, we also show how our data can shed light on fraud more generally, if one is willing to make one additional assumption. Since by construction the variable fraud is the combination of financial fraud and non-

financial fraud ($F = FF \cup NFF$), we can write the detection likelihood of any type of fraud for non AA firms as

$$\Pr(\text{caught}|F, \bar{A}) = (1 - \%FF) * \Pr(\text{caught}|NFF, \bar{A}) + \%FF * \Pr(\text{caught}|FF, \bar{A})$$

This equation shows that we can estimate this detection likelihood with three pieces of information: the detection likelihood for financial frauds for non AA clients (that we have already shown can be constructed from observables), the fraction of financial frauds among all frauds, and the detection likelihood for non-financial frauds. The DMZ database provides an estimate of the fraction of financial frauds amongst all frauds. In this database, a fraud is classified as a financial fraud if the firm subsequent to the fraud restates its financials. In the 1996-2004 period, the fraction of financial frauds is 64.7%. We do not have an empirical estimate for this third piece of information, so we will present two alternatives, what we think the most natural assumption, and a very conservative alternative assumption.

A natural assumption is that financial frauds are no more hidden than other types of fraud:

$$\textit{Assumption 3: } \Pr(\text{caught} | NFF, \bar{A}) = \Pr(\text{caught} | FF, \bar{A})$$

Under this assumption, equation (4), based entirely on observables, provides all of the required information

$$\hat{Pr}(\text{caught}|F, \bar{A}) = \hat{Pr}(\text{caught}|FF, \bar{A}) = \frac{\Pr(FF, \text{caught} | \bar{A})}{\Pr(FF, \text{caught} | A)} \quad (4)$$

We have no data to support or refute this assumption directly, but believe it to be conservative. Insiders need to produce financial statements, and likely misrepresentations can be identified by comparing this with past history and competitor information. Thus, given a fraud is committed, it is at least as likely that a financial fraud is detected as a non-financial fraud is.

We also consider the possibility that absolutely all non-financial frauds that start are caught.

$$\textit{Assumption 3A: } \Pr(\text{caught} | NFF, \bar{A}) = 1.$$

This is the most conservative alternative assumption. Under this assumption we obtain a lower bound estimate of $\Pr(\text{caught} | F, \bar{A})$ as

$$Pr(caught|F, \bar{A}) = (1 - \% FF) + \% FF * Pr(caught | FF, \bar{A}) \quad (5)$$

Our core estimates for detection likelihood for all frauds are based on assumption 3 that non-financial frauds are as detectable as financial frauds. The detection likelihoods remain the same as in the prior section at 0.267 for the DMZ auditor sample, 0.451 for the DMZ financial misrepresentation sample and 0.570 for the AAER sample. The iceberg is as big as 3 times the tip or roughly equal to the visible tip. Two factors make these core estimates conservative. First, they are conservative to the extent that the probability of revelation conditional on the natural experiment is less than one. Research suggests this is likely. The incentive to critically review past audit decisions and ‘clean house’ is dulled by the fact that for some firms the new auditor was the same individual as the old, other auditing firms hiring AA former auditors and they brought their clients with them (Blouin, Grein and Rountree (2005)). In addition, audit firms do not consider themselves as fraud detectors and explicitly state this in their assignments. Second, it is conservative to the extent that detection likelihood is lower for non-financial frauds.

If we make the most conservative assumption 3a, and unrealistically assume all non financial frauds are detected, not surprisingly detection likelihoods are reduced. We estimate detection likelihood in this most conservative case to be 0.526 for the DMZ auditor sample, 0.645 for the DMZ financial misrepresentation sample and 0.722 for the AAER sample. The iceberg below the surface varies from roughly the size of the visible tip, to 40% of the size of the visible tip.

Finally, we produce an estimate of the unconditional likelihood of fraud by taking the observed level of fraud that are started and caught described in the previous section and dividing it by the detection likelihood (the law of conditional probability). As reported in section I, the probability of committing a fraud of any type and getting caught in the 1996-2004 period is 4.0%. Our preferred estimate for the detection likelihood is 0.267 based on the DMZ auditor sample, and the assumption that non-financial frauds are as observable as financial frauds. This produces an estimate of firms that engage in frauds at a point in time of 15%, consisting of 4% of firms that engage in frauds that eventually will be detected and

11% of firms where frauds are taking place but are not detected. Using alternative samples provides smaller but still economically significant fraud incidence of 7-9%. Going further to use the most conservative assumption for non-financial frauds, the lowest estimate of fraud incidence is 5.5%.

III. Validation Methods to Identify Fraud

The AA experiment provides a natural experiment setting to identify hidden auditor-detectible frauds. To explore the validity of these findings, this section introduces an alternative approach to estimate fraud. The advantage of this alternative is that it is based on different shocks to detection than the forced auditor turnover in the AA situation. We will then compare our estimates with others existing in the literature.

III.1. Incentives and Opportunities Estimation: Approach

The AA example provided a clean natural experiment, but involved only one type of potential fraud detector. In this section, we employ an empirical design to uncover the hidden level of fraud by exploiting situations where incentives or opportunities for detections are high (H) across a village of fraud detectors. To illustrate the idea, consider one of the fraud detectors in DMZ: analysts. We know from Yu (2008) that firms with more analyst coverage engage in less earnings management. Thus, the number of analysts affect the probability accounting fraud is detected (to the extent earnings management equates to fraud). Our idea is to use this logic across five groups of known detectors identified by DMZ -- analysts, media, shortsellers, industry-specific regulators, employees with monetary incentives to whistleblowing -- to estimate the predicted level of detection in a setting in which firms face high incentives or opportunities for fraud detection for all of the village of detectors. In other words, we ask: what would be the detection probability if all the villagers were looking for fraud full force.

To estimate this setting of heightened detection, we can also employ a time series element; namely, *Sarbanes Oxley*. DMZ find higher incidence of fraud detection after *Sarbanes-Oxley*. Higher incentives for detection may result from the legislation itself or more simply from the considerable public scrutiny on fraud after the major scandals at the beginning of the new millennium.

More formally, we denote the potential for a firm i 's fraud to be detected in year t as D_{it} , where the firm is caught if D_{it} is positive:

$$\begin{aligned} D_{it} &= X_{it}^D \Gamma_D + v_{it} \\ caught_{it} &= 1 \text{ if } D_{it} > 0. \end{aligned} \tag{6}$$

D_{it} is a function of observables X_{it}^D . Our strategy is to include a set of indicators $[I_{it}^{HiIncentOpp}]$ in X_{it}^D that equal one when the circumstance leading to detection by each of fraud detectors is high. Following Wang (2013), we address the possibility that there are other observables that would also influence detection unrelated to detector incentives by including in X_{it}^D general firm and market characteristics $[x_{it}^D]$ that would lead to detection in all firms (e.g., size and stock performance). We denote these two groups:

$$X_{it}^D = \begin{bmatrix} I_{it}^{HiIncentOpp} & x_{it}^D \end{bmatrix}.$$

By comparing the ratio of detection in the ‘high’ setting indicated by H with that in the normal situation indicated by \bar{H} (i.e. when not all of the indicators are high), we produce an alternative estimate of detection likelihood. As in the prior section, it is instructive to make clear the assumptions we rely on to produce our estimates.

We start by assuming that the detection likelihood approaches 1 when all detection incentives and opportunities are high:

$$\textit{Assumption 4:} \quad \Pr(caught|F, H) = 1 \quad . \tag{7}$$

Assumption 4 is clearly conservative. More analysts, for example, provide greater opportunities for detection but this does not fundamentally change the limited incentives for analysts to bring frauds to light.

Initially, we also assume that the probability of engaging in a fraud is the same in both environments (the high detection and the low detection):

$$\textit{Assumption 5:} \quad \Pr(F|H) = \Pr(F|\bar{H}) \tag{8}$$

This assumption is unlikely to be satisfied. Potential fraudsters are aware of the presence of high incentives and opportunities for detection. Thus, they are less likely to commit fraud. Assumption 5, thus, will overestimate the detection rate, underestimating the unconditional probability of committing fraud.

With these two assumptions, we can produce an estimate of the likelihood of detection by comparing the observed detection rates when not all indicators are high and when all detection incentives and opportunities are high:

$$Pr(caught|F) = \frac{Pr(caught, F | \overline{H})}{Pr(caught, F | H)} \quad (9)$$

III.2 Incentives & Opportunities Estimation: Probit Results

Our first empirical estimate, based on a simple probit, holds assumption 5 to be true. Table 4 lists the variables we use in the initial probit estimations of what influences the likelihood of detection. Panel A defines the six indicator variables, showing how we split the sample into high and low analyst groups, high and low media groups, high and low shortability groups, industries with an industry regulator and not, industries with high monetary incentives for employees (i.e. Qui Tam Industries) and not, and post SOX and pre SOX.

Panel B lists key firm and market variables that we assume might also influence the likelihood of detection. These include firm-specific measures (company size, the stock return, ROA, leverage, R&D), measures of overall market conditions (VIX), and measures of competition (industry level HHI index, and a firms competitive position reflected in its market share). Disappointing accounting or stock market performance likely leads to lead to more detection. We include leverage under the assumption that high debt level raise risks for equity holders, sharpening the incentives of equity-based detectors (such as short sellers). We follow Wang (2013) and include R&D as a measure of the opacity of firm fundamentals, under the assumption that when financials are more opaque it is easier to engage in fraud and it is harder to detect it. We include the VIX under the assumption that more volatile market environments are more likely to spur detection. Finally, we introduce measures of competition. When markets are more competitive and rents are smaller, there is less scope to engage in self-dealing (e.g., Shleifer and Vishny

(1997)). As a measure of industry competition, we include a firm's industry market share and the Herfindahl-Hirschman Index (HHI) for its industry, which is defined as the sum of the squared market shares of publicly traded firms in the same sector.⁹

Panel C lists additional variants of the performance measures that we assume impact detection. Specifically, we include abnormal performance (stock and ROA). Note that by construction these are unanticipated movements in these control variables, thus they cannot affect the likelihood to start engaging in fraud before they materialize.

Table 5 presents our probit estimates of the determinants of detection. The reported coefficients show the marginal effects. The baseline specification in column (1) just includes the controls from panels B and C, but not the high incentive indicator variables of panel A. The control variables of Panel B are largely as expected. Fraud detection is higher in firms with higher leverage and lower accounting performance. The positive and significant coefficient on the VIX shows that detection is more likely when the stock price is more volatile. In this simple specification, inconsistent with Wang (2013), firms with high R&D have more detection. None of the other control variables, including the abnormal variables of panel C, are significant.

More important are the results in column (2) when we include the high incentive indicator variables. We find that 4 of the 6 indicators are positive and significant, and the remaining 2 insignificant. Specifically, we find that detection is significantly higher with high analysts, high shortability, and Qui Tam industries and post SOX. Control variables remain similar, with slightly lower magnitudes and the R&D variable loses its significance. The reported results are marginal effects.

By using (9), we can estimate the detection likelihood by comparing implied detection rates in the high state $\Pr(F, caught | H)$ to detection in the non-high state $\Pr(F, caught | \bar{H})$. $\Pr(F, caught | H)$ is

⁹ We construct both variables at the 2-digit SIC level using sales of publicly traded firms. We find similar results using a finer classification of 3 digit SIC. While there is a potential selection concern by not including private firms in constructing industry concentration, other papers such as Kadryzhanov and Rhodes-Kropf (2009) have found little difference in replacing these measures with census-based data measures of competition that include private in addition to publicly-traded firms.

8.44%, while $\Pr(F, caught | \bar{H})$ is 1.84%. These estimates produce a likelihood of detection equal to 0.217 (i.e. 1.84/8.44). Following the same procedure as in the AA experiment, the unconditional probability of engaging in fraud is 18% (0.04/0.217). As noted above, if assumption 5 is not satisfied, these results overstate the expected detection rates and underestimate the probability of engaging in fraud.

III.3 Incentives and Opportunities Estimation: Bivariate Probit Results

We now address empirically the real possibility that creating high incentives for detection will also influence the likelihood of engaging in fraud. Wang (2013) offers a solution, suggesting the possibility of applying the Poirier (1980) bivariate probit model for estimating dichotomous outcomes in partial observability settings to the case of corporate fraud. Our empirical methodology follows directly from Wang (2013).

Let E_{it} be the incentive for firm i to engage in fraud at time t . Fraud is committed if E_{it} is positive:

$$\begin{aligned} E_{it} &= X_{it}^E \Gamma_E + \mu_{it} \\ engage_{it} &= 1 \text{ if } E_{it} > 0. \end{aligned} \tag{10}$$

E_{it} is a function of observables X_{it}^E , which includes the high incentives and opportunities indicators and the variables in Table 3 panel B: $X_{it}^E = \begin{bmatrix} I_{it}^{HiIncentOpp} & X_{it}^E \end{bmatrix}$.

Identification in Poirier's model comes from two pieces. First, Poirier assumes that (μ_{it}, ν_{it}) are distributed bivariate standard normal. Second, identification depends on our ability to come up with variables which affect either firms' incentives to engage in fraud or detectors' ability to uncover fraud, but not both. Under the assumption that some of the X_{it}^E and some of the X_{it}^D are excluded from each other's set, the parameters in Γ_E and Γ_D can be identified using a bivariate probit model:

$$\Pr(engage_{it}, caught_{it} | X_{it}^E, X_{it}^D) = \Phi(X_{it}^E \Gamma_E, X_{it}^D \Gamma_D), \tag{11}$$

where $\Omega(\cdot, \cdot)$ denotes joint cumulative standard normal distribution over the two arguments.

We assume that the variables in panel C (abnormal ROA, abnormal stock return), affect the likelihood of detection, but not the likelihood of engaging in fraud as they would not be known at the time when a firm started to engage in fraud. In panel D we list two types of variables that we assume influence the likelihood of engaging in frauds but not detection. First, we include compensation-related measures as prior papers have found option grants to be an important predictor of fraud (e.g. Burns and Kedia (2006), Efendi, Srivastava and Swanson (2007)). We use Execucomp's valuation of all of the options held by top executives. We also measure the incentives provided by their most recent pay package to focus on future as opposed to current performance as captured in percentage of restricted stock grants divided by total compensation. Second, we include measures of market conditions under the assumption that the starting of frauds may be procyclical. Wang, Winton and Yu (2010) review theoretical bases for procyclicality in fraud and provide some evidence that positive sentiment (captured by IPO conditions) has positive, albeit non-linear, association with the likelihood to start committing fraud. To capture time trends in market conditions, we use the sentiment index of Baker and Wurgler (2006) and this index squared.

In Table 5 (columns 3.1 and 3.2), we report the results of the bivariate probit. Column 3.1 provides the coefficient estimates for the detection equation. Column 3.2 reports the coefficient estimates for the start equation. As one can see comparing column 3.1 and column 2, the results are very similar. Four of the six high incentive indicator variables are positive and significant (column 3.1), while the remaining two are insignificant. The identity of the significant monitors driving detection are the same as in the prior section, with significant impacts from high analysts, high shortability, qui tam industry and post SOX.

Column 3.2 presents the results for the start equation. The competition and market condition variables, which only appear in the start equation, have the expected signs, and the competition variables are significant. More importantly for our purposes, the high detection indicator variables in most cases have the opposite estimated effects on the likelihood of engaging in fraud. In 5 of the 6 indicators the sign flips between the start and detect equations. In two cases we find that the detection incentives have a

significant negative impact on starting. So, for example in the post-SOX period, we estimate both a lower likelihood of starting fraud, and a higher likelihood of detection.

With these estimates, we now recompute the detection likelihood. With a detection rate in the non-high state of 1.99% and a detection rate in the high state of 9.59%, from (9) we get a detection likelihood of 0.207 ($=1.99/9.59$). Thus, the estimate of the unconditional likelihood of fraud is 19% ($.04/0.207$).¹⁰ As predicted, relaxing assumption 5 does lower the detection likelihood, and raise the predicted unconditional estimate of fraud, but the change is not large economically. The results from both the probit specification and the bivariate probit specification are not far off from the preferred results based on the AA natural experiment, where we found a detection likelihood of 0.267 and an unconditional probability of engaging in fraud of 15.0%.

III.4 Related Literature

While obtained by using different samples and different definitions of what is fraudulent, the estimates that can be inferred from the existing literature are not very far from ours. For example, Wang, Winton and Yu (2010) examine frauds among the 3297 IPOs from 1995-2005. They classify as fraudulent firms that had an Accounting and Auditing Enforcement Release (AAER) and/or had been target of a securities class action that: i) was not dismissed; ii) exceeded the \$2mn threshold; iii) was related to financial reporting. Their main goal is to show that fraud is procyclical. Yet, their bivariate probit model produces predicted probabilities of engaging in fraud of 10-15%, very much in line with our estimates.¹¹

The literature on options backdating also provides another estimate of engaging in a particular type of fraud. Since prior to Lie (2005) none of this type of fraud was detected, this is an experiment in which detection rates go to 1 as researchers *ex post* look at the data to identify firms whose actions are

¹⁰ An attractive feature of this setting is we do not need to make an assumption about differences between the detectability of financial and non-financial fraud, as we needed to do in the last section (assumption 3 or 3a). By the law of conditional probability, all we need is the observable probability of engaging in fraud and being detected to generate the unconditional probability of engaging in fraud.

¹¹ We infer this from Figure 1, predicted probability of fraud, and summary statistics on the distribution of industry EPS growth available in the internet appendix.

most consistent with backdating. The DMZ sample includes no back dating cases, since the allegations came out after the end of our sample period.

Bebchuk, Grinstein and Peyer (2010) look in depth at the options backdating scandal first brought to attention by Lie (2005). They attempt to uncover the percentage of publicly-traded firms from 1996-2005 in which CEOs or directors were ‘lucky’ directors in that they received option grants on the lowest price day of the month, filtering out those that could have taken place simply due to luck. These lucky option grants increased the value of that grant by 20% and CEO pay in that year by 10%. By their estimate 12.4% of firms have such lucky CEOs and 7% of firms have lucky directors, with the percentage of “lucky grants” of 14.5% prior to SOX and 8.4% after. Their estimate of a 12.4% of incidence just for this type of governance problem is in line with our estimate from the AA experiment.

Research on earnings manipulation provides evidence suggesting the AA experiment results might be conservative. Dichev, Graham, Harvey and Rajgopal (2013), for example, survey 169 CFOs of public companies. A few of their questions focus on misrepresentations (they do not focus on fraud). Their survey suggests that 18.4% of firms manage earnings, meaning within-GAAP manipulation that misrepresents performance, and that for such firms, the manipulation is 10% of reported EPS, a substantial amount. Given the narrow scope of the questions, the authors suggest these answers provide a lower bound on actual earnings management.

Finally, Zakolyukina (2013) uses a structural model to explore detected and undetected manipulation. She finds that two thirds of CEOs manipulate their financials, with a detection likelihood of only 9%.

IV. How Expensive is Corporate Fraud?

To gauge whether this level of fraud is excessive it is not enough to estimate how frequent it is, we also need to determine how costly it is.

IV.1 Prior Estimates of the Cost of Detected Fraud

Prior research provides various measures of the cost of financial fraud. One method relies on an event study, which carefully measures the decline in equity (and sometimes debt) capitalization at the moment a fraud is revealed (e.g. Feroz, Park and Pastena (1991), Palmrose, Richardson and Scholz (2004), Gande and Lewis (2009)). This is a good measure under two conditions. First, it must be the first indication to the market of the fraud. If there was a prior partial leakage, then the event alone would be an underestimate of the actual costs, as Gande and Lewis (2009) and Karpoff et al. (2012) show. Second, and perhaps more challenging, the measured decline must be solely attributable to the fraud, rather than to the revelation of bad information about fundamentals which the fraud tried to cover up.

To address the second problem Karpoff, Lee and Martin (2008) compute an adjustment factor equal to the market value of the assets written off in the financial restatements. They compute the market value by using some industry multiples. Unfortunately, we cannot use this method because 35% of our frauds do not involve any restatement.

KLM estimate that the mean (median) misconduct loss is approximately 29% (28%) of equity value. Assuming that the value of debt is unaffected by the fraud and a 25% D/V (the median in our sample), a fraud is associated with destruction of 22% of an enterprise value.

IV.2. Method for Estimating the Cost of Detected Fraud

Rather than looking at value changes around specific event dates, we look at changes in value over the fraud period, using the beginning date and (revised) end date from security class action filings. This method does not require the knowledge of all the relevant event dates and thus it is less subject to errors of omission. To account for the value loss due to a deterioration in fundamentals we use industry multiples.

We express the firms' enterprise value after the fraud as a multiple. That is, we collect information both on post fraud enterprise value (value of equity plus book value of debt) and on post fraud performance (using EBITDA, Sales and Total Assets). We compare these 'actual' multiples with counterfactual multiples we construct. To build the counterfactual, we start from a firm's multiples in the pre-fraud period. Then, we estimate how multiples changed in a typical firm in the same industry

over the fraud period. Finally, we compute the counterfactual multiples by applying this typical multiple change in the industry to the firm's multiples in the pre-fraud period. By multiplying this counterfactual multiple by the post fraud *actual* EBITDA (or Sales or Total Assets) we obtain our 'non fraud implied enterprise value.'

The difference between the non-fraud implied enterprise value and the actual enterprise value after the fraud has been revealed gives us an estimate of the value loss from fraud. Appendix III provides more details of the calculations involved.

An example using the sales multiple illustrates the approach. Consider for concreteness a situation with a firm that pre fraud has an enterprise-value/sales multiple of 2. Suppose that this firm exaggerates sales during the fraud period and then, when it reports the correct sales numbers, its sales multiple becomes 1.5. Over the same period, the median firm in the same industry starts with a sales multiple of 3 and ends with a sales multiple of 2.7, thus experiencing an expected multiple change of 0.9 ($=2.7/3.0$). Our estimates focus on the difference between the post fraud multiple of 1.5, and the counterfactual multiple of 1.8 (i.e. the 2.0 pre-fraud multiple * industry multiple change of 0.9).

As the example shows, with a fraud firm having a pre-fraud multiple of 2x whereas the industry was 3x, our approach allows fraud firms to have different fundamentals than the median firm in the same industry. We just assume that during the fraud period the trend of the fraudulent firm's multiple absent fraud would have been the same as the trend of the multiple of the median firm in the industry.

For the post-fraud period we use the sales numbers after any restatements, so the actual multiple is not driven by fraudulent numbers. Finally, the contemporaneous changes in industry valuation are not incorrectly assumed to be a cost of fraud as they are deliberately excluded from the calculation.

IV.3. Costs of Detected and Undetected Fraud

Tables 6 and 7 present the results from this analysis. Table 6 Panel A provides summary statistics for market valuation measures (equity market, enterprise value) and key inputs into multiples (EBITDA, sales, assets) both pre and post fraud. Means always exceed medians showing positive skewness in size, and post fraud measures show the decline in valuation, while there is growth in assets over the fraud

period. Because the data have non-normal distributions we focus on medians. We use the book value of debt.¹²

The EV/EBITDA multiple is most easily comparable across firms, but we lose some data points when EBITDA is negative, so instead of focusing exclusively on this multiple we also look at sales and asset multiples. Our overall summaries will use the average across all three measures.

Panel B allows us to compare actual multiples in fraud firms and counterfactual multiples. Post fraud, the median EBITDA multiple in fraud firms is 9.00. The counterfactual multiple for these firms is 12.13. This is based on a pre-fraud multiple of 12.17 that is multiplied by the change in industry multiple over the fraud period. There is little difference in the counterfactual before and after the fraud as on average the fraud period is less than 2 years. We get similar results if we use sales and assets multiples.

This drop in multiple levels regards detected fraud, what about undetected ones? Our main concern is the probability of detection is directly correlated with the magnitude of the fraud and hence with the magnitude of the value loss to the fraud. In this case, our estimate above would grossly overestimate the magnitude of the value loss from all fraud.

To address this concern we use the AA experiment. Since the forced turnover was unexpected, some fraud that “normally” would have never emerged did emerge. Thus, we can compare the costs of those with the costs of the one that did emerge. As Table 6 Panel C shows, the cost is indeed less for normally undetected fraud.

In Table 7 we provide the summary statistics for our calculation of the value loss due to fraud. We present both the dollar costs and the cost as a percentage of enterprise value. For the median fraud firm our estimate of the dollar cost of detected fraud ranges from \$1.2 to \$2.1 billion, depending on the

¹² If debt values are also affected by the fraud, our results will be conservative. To assess changes in debt, we collected monthly debt prices for all firms in our sample and computed the change in debt value around the announcement date for the fraud. We thank Sergei Davydenko for providing this information. Data was only available for a subset of our sample (90 firms). This data suggests that a measure of aggregate value destruction would be greater if we used such data, but the additional value destruction for the median firm that is our focus would only be minor. In the 90 observations the mean (median) decline in bond prices from the monthly bond price available before the fraud to the monthly bond price immediately after the fraud was 1.1% (5.5%) and the mean (median) bond price after was 0.979 (0.891).

multiple used. There is a significant variation around this median level: for example with sales multiple the loss varies from \$165 million (first quartile) to more than \$7.9 billion (third quartile). Expressing these losses as a percentage of enterprise value (computed the month prior to the announcement of the fraud), we find the median firm loses 25% of its enterprise value if we rely on the EBITDA multiple, 39% if we rely on the Sales multiple, and 36% if we rely on Assets multiple. Regardless of measure, the losses are substantial.

In Panel B we provide similar results for the normally undetected frauds that were exposed by the forced auditors' turnover. As expected, in this sample the value destruction is less, but the median firm still loses 19% of enterprise value based on the average across the three multiples.

IV.4 Net Costs of Fraud

The estimates we computed so far represent the cost of fraud borne by financial claim holders.¹³ Some of that cost, however, is not a social cost, since fraudster firms lose value also because they have to pay fines and other penalties, which are simply transfers to the government and the plaintiff law firms. To compute these transfers we rely on Karpoff, Lee and Martin's (2008), who estimate them at 3.7 percent of equity value. This is likely to be an over correction, because a large fraction of these fines is often covered by an insurance and thus does not burden the firm.

We present the results in Table 7 panel C. When we subtract the transfer estimate, we arrive at costs of detected fraud of 30% of enterprise value and costs of normally not-detected fraud of 17% of enterprise value. Assuming these estimates apply to fraud firms more generally, we project average fraud cost to be 20% of enterprise value ($= (0.04/0.15)*0.30+(0.11/0.15)*0.17$) suggesting that firms with fraud stand to lose roughly one fifth of their value. This estimate is almost identical to the previously reported estimate of value loss of 22% provided by Karpoff, Lee and Martin (2008), based on an alternative methodology.

¹³ This net cost is the cost born by financial claimholders and differs from the social cost in several respects. It ignores the cost imposed on employees and customers, while it also ignores that competitors can capture part of the value lost.

The size of the value loss with fraud is not so surprising. As noted in Karpoff and Lott (1993), corporations can gain in a variety of ways from being known as an honest actor that fulfills promises, delivers high product quality, etc.. These quasi rents from reputation can be lost when the market believes a fraud has been committed. Investors, consumers, suppliers and others will likely change the terms under which they interact with the firm.

In sections II and III we estimated that at any one time 15% of firms, or roughly 1 in 7 firms, are engaged in fraud, with the bulk of these frauds undetected. Combining this result on pervasiveness with the estimate of the net costs of fraud, we can estimate what the overall loss from fraud was before SOX. Note that this is a residual loss, after all the contractual and regulatory interventions to mitigate agency costs.

If the cost of fraud is 20% of enterprise value and the probability of engaging in a fraud in any given year is 15%, then the expected cost of fraud is 3% of enterprise value ($=.15*.20$). With US traded firms with more than \$750 million in assets collectively having an enterprise value of more than \$21 trillion, this implies a residual loss from detected and undetected fraud of \$630 billion. Since the average fraud lasts 1.67 years, these estimates suggest the annual cost of fraud among large US corporations is \$380 billion.

A reasonable objection to our estimates is that they overestimate the counterfactual by assuming that the fraud firm would have grown as the median firm in the industry. If firms are more likely to commit fraud when their growth slows down, then we should assume a less optimistic path of their multiples in the counterfactual. In panel D we use the 25th percentile level of growth instead of the median. When we do so, we get a loss of 11% of enterprise value for detected fraud, less than our previous estimate but still a substantial value.

V. Applications to Cost-Benefit Analysis of Regulation

The D.C. Circuit Court has struck down certain Securities and Exchange Commission rules based on its conclusion that the Commission failed to evaluate adequately a rule's economic impact (Business

Roundtable v. SEC, 647 F.3d 1144, 1150 (D.C. Cir. 2011)). While the same principles do not apply to Congress, there is an increasing trend to request a cost-benefit analysis for any kind of regulation. The possibility of such an analysis depends crucially upon the existence of reliable estimates of the potential benefits of regulation. Since one of the goals of regulation is to reduce fraud, our estimates of the cost and pervasiveness of fraud can be used for such an analysis. In what follows we sketch out two possible applications.

V.1 A Cost Benefit Analysis of SOX

Our results from Table 5 model 3 show that after the introduction of SOX fraud detection increased significantly and the number of frauds started decreased significantly. We can even go further and estimate the marginal effect of a post-SOX dummy on the likelihood of starting a fraud, albeit it is hard to disentangle how much of those changes are the result of the introduction of SOX and how much a temporary reaction to the discovery of some major scandals. If we attribute all the effect to SOX we obtain that the introduction of fraud reduced the probability of starting a fraud by 57%. With an annual cost of fraud of \$380 billion, the total expected benefit of SOX is \$215 billion a year.

What is the cost of compliance with SOX? We follow Hochberg, Sapienza and Vissing-Jorgensen (2009) in estimating compliance costs by exploiting survey data collected by Finance Executives International (FEI). They provide an annual survey on the costs associated with complying with SOX section 404, collecting an estimation of auditor attestation fees, as well as internal and external costs with compliance. We use their estimates from the March 2006 survey of 274 financial executives. They estimate an average cost at \$3.8 million per firm, with costs increasing with the annual sales revenue of the issuer.¹⁴ To calculate compliance costs in our sample of firms above \$750 million in capitalization, we multiply the compliance costs by firm sizes based on sales revenue, by the number of firms in our sample in the respective sales revenue categories, producing a total annual compliance cost of \$6.8 billion per year.

¹⁴ See Hochberg, Sapienza and Vissing-Jorgensen (2009), Table 11 on page 571 for the costs by size categories that we use in our estimates.

If we assume that all the effect of the post-SOX dummy is due to SOX, the net benefit of SOX is \$208 billion a year ($215 - 6.8$). Clearly, this is an overestimate, since, at least in part, the post SOX deterrence effect is due to a market reaction to the scandals. Yet, with our estimates we can conclude that for SOX to be justified on a cost-benefit analysis perspective it is sufficient that SOX reduced the probability of starting a fraud by 2%.

V.2 A Cost Benefit Analysis of Mandatory Auditor Rotation

A proposal contained in the initial version of SOX and then discarded was to mandate a periodic rotation of auditing firms, as it takes place in many countries. There are clear cost in mandatory rotation, but what are the benefits? Our above analysis provides us with some estimates.

We can think of the demise of Arthur Andersen as an unexpected mandatory rotation. Thus, we can compare the rate of detection of fraud post AA demise among former AA clients and among the rest. Notice that both the auditors of former AA firms and not former AA firms are subject to the same market pressure and the same degree of regulation post the AA demise and post SOX. The only difference between the two sets is that in first the auditor was changed, while in the second not.

From Table 3 Panel B we know that there is an increase in the detection of accounting fraud that goes from 0.6% to 2.2%. Thus, auditor turnover increases by 1.6 percentage points the detection of accounting fraud. We have seen that the real cost of fraud is 20% of the enterprise value and that the total enterprise value of this sample is \$21 trillion. Thus, auditor turnover has the potential of reducing the cost of fraud by \$67 billion ($0.016 * 0.2 * 21$ trillion). We say the potential because the cost is saved only if the fraud is prevented. We do not have here a good estimate of the deterrence effect: it can be larger (if even undetected fraud is prevented) or smaller (if even the certainty of detection does not dissuade people from committing fraud).

The General Accounting Office (GAO, 2003) estimates the initial year mandatory firm rotation could cost companies and audit firms between 43% and 128% more. The GAO (GAO, 2008) also puts the average audit fee for large companies at 0.08% of sales. Setting the GAO estimate at the midpoint (85.5%) and assuming a ratio of assets to sales equal 1, we can estimate that mandatory rotation would

cost our universe of firms above \$750 million in assets a total of \$14 billion ($0.0008 \times 0.855 \times 21$ trillion). Thus, a one off mandatory rotation to all firms above \$750 million in assets would have a benefit of \$67 billion and a cost of \$14 billion.

Any form of mandatory rotation, however, will not mandate one off or every year. Most likely, will be every 6 or 9 years. Thus, the annual cost of mandatory rotation would be divided by 6 or 9. It is less clear how the benefit will be affected. If the threat of rotation in 9 years is able to prevent a fraud today, then the annual benefit of rotation will remain constant, making mandatory rotation even more appealing. If, as it is likely, the benefit of preventing fraud decays if the rotation takes place far in the future, then the annual benefit might be smaller. The worst case scenario is that prevention works just the year before mandatory rotation. In that case the annual benefit will be divided by 6 or 9 as well. Yet, the conclusion is the same: the estimated benefits exceed the estimated costs.

VI. Conclusion

In this paper we try to quantify the pervasiveness of corporate fraud in the United States and to assess its costs. The major problem in any such a study is how to estimate the amount of undetected fraud. We follow different approaches to infer the unconditional probability that a fraud is committed whether or not it is subsequently caught. Regardless of the method used, we find that only 1 in 4 frauds is detected. Using these estimates we obtain that there is an ongoing fraud in about one out of 7 firms (15%).

Having established the incidence of fraud, we try to estimate how much frauds cost investors. We estimate that the cost of fraud to the median fraudulent firm is 22% of enterprise value, similar to prior studies. This estimate includes both the fraud that are normally detected and those that are not. With US traded firms with more than \$750 million in assets collectively having an enterprise of more than \$21 trillion, this implies a loss from detected and undetected fraud of \$630 billion. Since the average fraud lasts 1.67 years, the annual cost of fraud among large US corporations is \$380 billion.

Finally, we present two illustrations of how our estimates can be used to do cost-benefit analysis of regulation. First, we compare the benefits of the reduced incidence of fraud after the introduction of

SOX with the compliance cost associated with it. We arrive to an estimate of net annual benefit of SOX equal to \$208 billion. While this is might be an overestimate since not all the reduction in the incidence of fraud occurring after SOX is due to SOX, our estimates show that it would take a very minimal deterrence effect of SOX on fraud (2%) to justify its introduction. We also compare the benefits of mandatory auditors' rotation in terms of fraud discovery with the estimated additional costs it imposes on auditors and firms. We find that a one-off mandatory rotation would have a benefit of \$67 billion and a cost of \$14 billion.

These are just illustrative examples. More research is needed to settle these questions beyond any reasonable doubt. Regardless, these examples show the wide potential applications of our estimates.

Appendix I: Dyck, Morse, and Zingales (2010) Filters to Eliminate Frivolous Fraud

First, they 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 U.S. 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 U.S. 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.¹⁵ 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 previous studies suggested to divide frivolous suits from meritorious ones.¹⁶ Fifth, they exclude those security frauds that Stanford classifies as non-standard, including mutual funds, analyst, and IPO allocation frauds.¹⁷ 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 settled to avoid the negative publicity.¹⁸

¹⁵ 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.

¹⁶ Grundfest (1995), Choi (2007) and Choi, Nelson, and Pritchard (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 \$2.5 - \$1.5 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.

¹⁷ Stanford Class Action Database distinguishes these suits for the reason that 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.

¹⁸ 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 when the firm or its executives issued a positive outlook statement for which they are later sued.

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

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

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

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

For a complete description of the motivation for each item as an indicator of potential for manipulation and for the compustat codes leading to the calculation of the indices, please see the paper referenced above. We followed their compustat definitions exactly to construct the ProbM Score yearly for the large corporations in our sample. According to Beneish (1999), a score greater than -2.22 indicates a strong likelihood of a firm being a manipulator

Appendix III – Process for Calculating Implied Value Loss not Attributable to Changes in Fundamentals

We follow Berger and Ofek’s (1995) multiples approach with modification to exploit firm-specific information. Assume that a fraud begins right after time s and ends before time t . The pre-fraud enterprise multiple, specific to firm i , which resides in industry j , is:

$$m_{ijs} = \frac{\text{Long Term Debt}_{is} + \text{Market Equity}_{is}}{Y_{is}},$$

where we consider several valuation bases, $Y \in \{EBITDA, \text{revenue}, \text{fixed assets}\}$. Likewise, we define a pre-fraud industry multiple, M_{js} , as the revenue-weighted average multiple for SIC 2-digit industries, indexed by j . We exclude the fraud firm in this calculation. We do the same procedure at time t , the year ending *after* the fraud revelation date to get M_{jt} . We use the change in the industry multiple as the benchmark for how the firm’s multiple would have evolved over the time period if it was just impacted by factors affecting the industry; i.e.:

$$\hat{m}_{ijt} = m_{ijs} \frac{M_{jt}}{M_{js}}.$$

The idea is to compare the fraud firm’s value of debt and equity at time t with the debt and equity which would be projected by the firm’s pre-fraud multiple adjusted to a growth or decline rate in its industry benchmark multiples. The estimated “but-for” or counterfactual valuation is thus the EBITDA, sales, or fixed assets implied enterprise value at time t , calculated as:

$$\text{Counterfactual Enterprise Value} = \hat{m}_{ijt} Y_{it},$$

$$\text{for } Y_{it} \in \{\text{revenue}_{it}, \text{fixed assets}_{it}, \text{EBITDA}_{it}\}.$$

The next step is to compare the counterfactual with the actual enterprise value post fraud revelation to produce a dollar loss per firm arising from the fraud. To ensure comparability across firms we also express this dollar loss relative to the pre-fraud enterprise value to define the fraud loss as a percentage of enterprise value.

$$\text{Cost Caught Fraud}_{it} = \text{Counterfactual Enterprise Value}_{it} - (\text{Long Term Debt}_{it} + \text{Equity}_{it}).$$

References

- Agrawal and Chada, (2005), "Corporate Governance and Accounting Scandals." *Journal of Law and Economics*.
- Baker, M. and J. Wurgler, (2006), "Investor Sentiment and the Cross-Section of Stock Returns," *Journal of Finance* vol. 61, p.1645-1680.
- Bebchuk, Lucien, Olav Grinstein and Urs Peyer, (2010). "Lucky CEOs and Lucky Directors," *Journal of Finance*, (65): 2363-2401.
- Beneish, M.D., (1999), "The Detection of Earnings Manipulation, *Financial Analysts Journal*, 24-36.
- Berger, Philip G. and Eli Ofek. (1995). "Diversification's Effect on Firm Value." *Journal of Financial Economics*, vol. 37(1), 39-65.
- Blouin, J., B. Grein, B. Rountree, (2005). "An analysis of forced auditor rotation: the case of former Arthur Andersen clients." *Accounting Review*.
- Burns, N. and S. Kedia. (2006). The impact of performance-based compensation on misreporting. *Journal of Financial Economics* 79: 35-67.
- Cahan and Zhang, (2006). "After Enron: Auditor Conservatism and Ex-Andersen Clients," *The Accounting Review*, (81): 49-82
- Choi, Stephen 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, Stephen J., Karen 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, IV, John C. and Srinivasan, Suraj, "SOX after Ten Years: A Multidisciplinary Review," *Accounting Horizons*, 2014.
- Coffee, John 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): 669-727.
- Coffee, John C. (2006), *Gatekeepers: The Professions and Corporate Governance*, Oxford University Press: New York.
- Dechow, P. M., R. G. 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.
- Dichev, Ilija, Joh Graham, Campbell Harvey and Shiva Rajgopal, (2013). "Earnings Quality: Evidence from the Field," working paper,
- Dyck, Alexander, Adair Morse and Luigi Zingales. (2010). "Who Blows the Whistle on Corporate Fraud?" *Journal of Finance*. 65 (6): 2213-2254.

- Dyck, Alexander, Natalya Volchkova, and Luigi Zingales (2008), "The Corporate Governance Role of the Media: Evidence from Russia," *Journal of Finance*, 63 (3): 1093-1136.
- Efendi, Jap, Anup Srivastava, and Edward Swanson, (2007). "Why Do Corporate Managers Misstate Financial Statements? The Role of in-the-money Options and Other Incentives," *Journal of Financial Economics*, 85(3): 667-708.
- Eisenberg, Macey, (2004). "Was Arthur Andersen Different? An Empirical Examination of Major Accounting Firm Audits of Large Clients" *Journal of Empirical Legal Studies*,
- Fama, Eugene (1980), "Agency Problems and the Theory of the Firm," *The Journal of Political Economy*, Vol. 88, No. 2, pp. 288-307
- Fama, Eugene and Michael Jensen (1983), "Agency Problems and Residual Claims," *Journal of Law and Economics*, Vol. 26, No. 2, pp. 327-349.
- Feroz, E., K. Parek and V. Pastena, (1991). "The Financial and Market effects of the SEC's Accounting and Auditing Enforcement Releases," *Journal of Accounting Research*, 29: 107-148.
- Galbraith, J.K, (1961). *The Great Crash 1929*, Pelican.
- General Accounting Office, 2003, "Public Accounting Firms. Required Study on the Potential Effects of Mandatory Audit Firm Rotation," United States General Accounting Office, 2003.
- General Accounting Office, "Audits of Public Companies: Continues Concentration in Audit Market for Large Public Companies Does Not Call for Immediate Action'," United States General Accounting Office, 2008.
- Gande, Amar and Craig Lewis, (2009). "Shareholder-Initiated Class Action Lawsuits: Shareholder Wealth Effects and Industry Spillovers," *Journal of Financial and Quantitative Analysis*, 44 (4): 823-850.
- Gilson, Ron and Renier Kraakman, (1984), "The Mechanisms of Market Efficiency," *Virginia Law Review*, Vol. 70. 549- 619.
- Grundfest, Joseph A. (1995). "Why Disimply?" *Harvard Law Review*, (108): 740-741.
- Hennes, K, Leone, A., B. Miller, (2008). "The Importance of Distinguishing Errors from Irregularities in Restatement Research: The Case of Restatements and CEO/CFO Turnover," *The Accounting Review*, 83(6): 1487-1519.
- Hochberg, Yael V., Paola Sapienza and Annette Vissing-Jorgensen. 2009. A Lobbying Approach to Evaluating the Sarbanes-Oxley. *Journal of Accounting Research*. 47(2): 519-583.
- Jensen, Michael and William Meckling, (1976). "The Theory of the Firm: Managerial Behavior, Agency Cost, and Ownership Structure." *Journal of Financial Economics* 3, no. 4 (October): 305-60

- Karpoff, J.M. and J.R. Lott, Jr., (1993). "The Reputational Penalty Firms Bear from Committing Criminal Fraud." *Journal of Law and Economics*, (36): 757-802.
- Karpoff, Jonathan M., D. Scott Lee, and Gerald S. Martin, (2008). "The Cost to Firms of Cooking the Books," *Journal of Financial and Quantitative Analysis* 43 (3): p. 581–612.
- Karpoff, Jonathan, Allison Koester, D. Scott Lee and Gerald Martin, (2013). "An Analysis of Database Challenges in Financial Misconduct Research," University of Washington, working paper.
- Krishnan, Gopal V., (2007). "Did earnings conservatism increase for former Andersen clients?" *Journal of Accounting, Auditing and Finance*. 22(2): 141-163.
- Lie, Eric (2005), "On the Timing of CEO Stock Option Awards," *Management Science*, 51(5): 802-812.
- Miller, Gregory S.(2006). "The Press as a Watchdog for Accounting Fraud." *Journal of Accounting Research* 44, no. 5 (December): 1001-1033.
- Palmrose, Z.-V.; V. Richardson; and S. Scholz., (2004). "Determinants of Market Reactions to Restatement Announcements." *Journal of Accounting and Economics* 37 : 59–89.
- Palmrose, Z.-V., and S. W. Scholz. (2004). The circumstances and legal consequences of non-GAAP reporting: Evidence from restatements. *Contemporary Accounting Research*. 21 (1): 139-180.
- Povel, Paul, Rajdeep Singh and Andrew Winton, 2007. "Booms Busts and Fraud," *Review of Financial Studies*, (20): 1219-1254.
- Roe, Mark, (2002), "Corporate Law's Limits," *Journal of Legal Studies*, Vol 31.
- Wang, Tracy Yue, (2013). "Corporate Securities Fraud: Insights from a New Empirical Framework", *Journal of Law, Economics and Organization*, 29 (3): 535-568.
- Wang, Tracy Yue, Andrew Winton and Xiaoyun Yu, (2010). "Corporate Fraud and Business Conditions: Evidence from IPOs," *Journal of Finance*,
- Yu, F. (2008). "Analyst Coverage and Earnings Management," *Journal of Financial Economics* (88): 245–71.
- Zakolyukina, Anastasia, (2013). "Measuring Intentional Manipulation: A Structural Approach," University of Chicago, working paper.

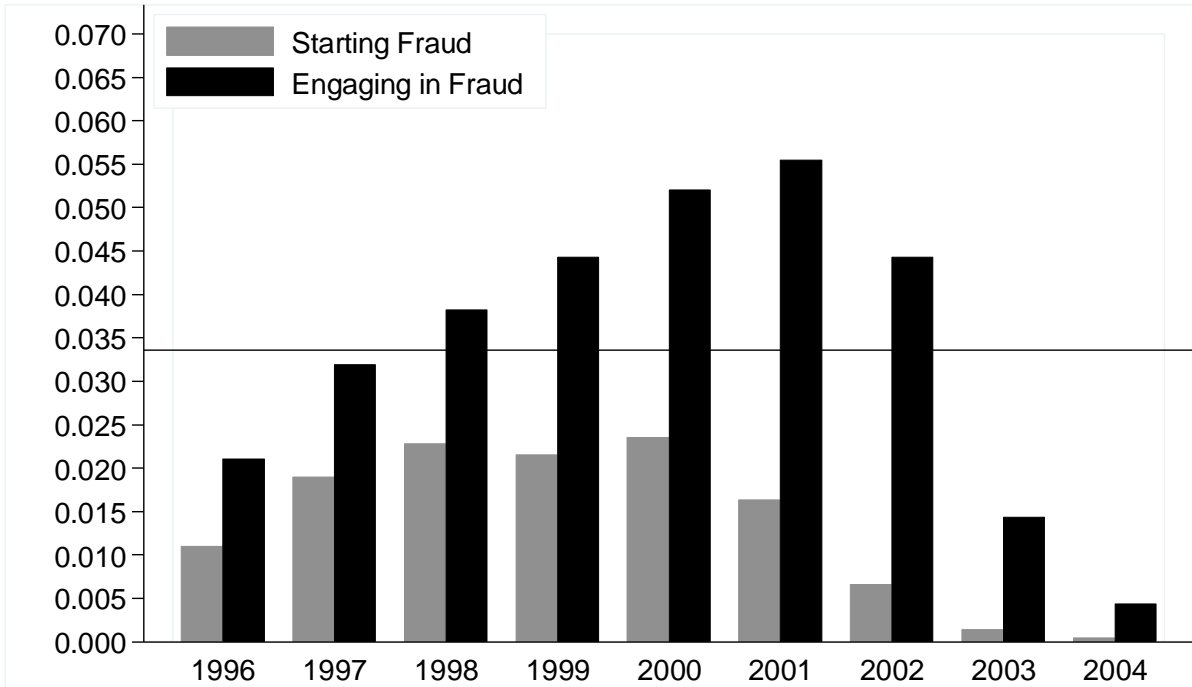


Figure 1: Percentage of Large Corporations Starting and Engaging in Fraud
 The grey and black bars respectively report the percent of firms starting and engaging in fraud. The reference line at 0.034 is the overall mean percent of firms engaging in fraud.

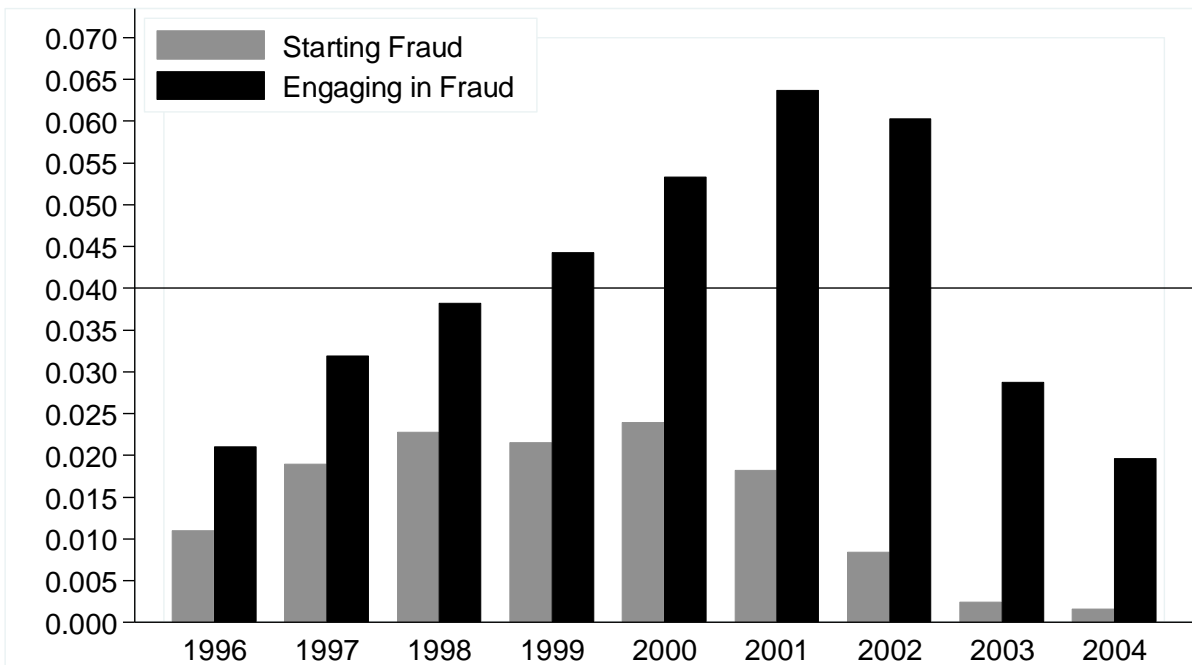


Figure 2: Adjusted Percentage of Large Corporations Starting and Engaging in Fraud
 The grey and black bars respectively report the percent of firms starting and engaging in fraud. The reference line at 0.04 is the overall mean percent of firms engaging in fraud. The figure adjusts for the empirical distribution of frauds which in expectation will be caught.

Table 1: Were AA Firms Different from non-AA Firms: Univariate Stats

The sample is all Compustat firms with \$750 million in assets and compares in the 1998-2001 period, The probM score (Beneish, 1999) is the probability of manipulation based on the accounting components listed (see the Appendix for a further description). ProbM is a score with no natural scale. The mean and standard deviation of ProbM are -2.325 and 1.357, respectively. The Fraud Started variable is from Dyck, Morse and Zingales (2010), the year a fraud started that later resulted in a non-frivolous class action suit, also 1998-2001. In panel A the p-values are for ttests that the non-AA clients differ from AA clients on each of the statistics. In all comparisons we compare AA firms to non-AA firms and to only other Big 5 firms. Panel B provides summary statistics for the control variables used in Table 2, which offers multivariate tests of the Panel A univariate results.

Panel A: Univariate Tests for Difference of AA Firms in Manipulation and Fraud

	AA Firms		All Non-AA			Big 5 Non-AA		
	Mean	Obs.	Mean	Obs.	P-value	Mean	Obs.	P-value
ProbM Score	-2.328	1,008	-2.324	5,135	0.929	-2.333	3,835	0.925
Days Sales Receivables	1.049	1,089	1.717	5,393	0.629	1.944	4,058	0.574
Gross Margin Index	1.037	1,105	0.995	5,514	0.690	0.967	4,162	0.560
Asset Quality Index	1.273	1,109	2.571	5,528	0.146	2.631	4,170	0.132
Sales Growth Index	1.227	1,105	1.242	5,516	0.692	1.245	4,163	0.677
Depreciation Index	35.61	1,109	190.5	5,521	0.433	171.8	4,164	0.513
SG&A Expense	1.223	1,024	1.440	5,248	0.717	1.387	3,931	0.783
Accruals/Assets	-0.062	1,135	-0.058	5,723	0.492	-0.057	4,320	0.344
Leverage Index	1.069	1,109	1.084	5,527	0.751	1.092	4,169	0.683
Fraud Started	0.0202	1,141	0.0188	5,749	0.757	0.0184	4,343	0.701

Panel B: Summary Statistics of Control Variables for pre-2000 AA and non-AA Firms

	AA Firms			All Non-AA			Big 5 Non-AA		
	Mean	Median	Obs.	Mean	Median	Obs.	Mean	Median	Obs.
Assets	9,549	2,659	1,135	14,591	2,528	5,732	13,612	2,805	4,329
LT Debt	3,069	843	1,135	2,532	542	5,730	2,265	623	4,327
Leverage	0.322	0.301	1,135	0.238	0.210	5,730	0.248	0.226	4,327
Sales	4,301	1,838	1,135	5,197	1,387	5,724	5,473	1,607	4,321
EBITDA	704.1	248	1,135	865.2	201	5,724	900	234	4,321

Table 2: Were AA Firms Different from non-AA Firms: Multivariate Tests

The sample is all Compustat firms with \$750 million in assets and compares in the 1998-2001 period. The dependent variable is the prob-M score (Beneish, 1999), the probability of manipulation based on the accounting components listed (see the Appendix for a further description). ProbM is a score with no natural scale. The mean and standard deviation of ProbM are -2.325 and 1.357, respectively. We test whether AA firms differ from non AA firms in their probM score using OLS and quantile (median) regressions, including controls for other firm attributes including year crossed with industry fixed effects (2 digit SIC), assets, and the following variables each scaled by assets: sales, EBITDA and long term Debt. The even columns restrict the non-AA sample to only Big Five clients. ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively. Robust (for OLS) standard errors are in brackets.

Estimation	Dependent Variable: ProbM Score of Beneish (1999)			
	(1)	(2)	(3)	(4)
	OLS	OLS	Medians	Medians
Arthur Andersen	0.042 [0.0508]	0.060 [0.0513]	0.024 [0.0208]	0.032 [0.0231]
Log Assets	-0.0058 [0.0168]	-0.0149 [0.0192]	0.0018 [0.00599]	0.0029 [0.00745]
Sales / Assets	-0.243*** [0.0387]	-0.241*** [0.0422]	-0.0712*** [0.0162]	-0.0693*** [0.0195]
EBITDA / Sales	-1.400*** [0.428]	-1.116** [0.453]	-0.575*** [0.107]	-0.683*** [0.129]
LT Debt / Assets	-0.520*** [0.123]	-0.489*** [0.133]	-0.131*** [0.0431]	-0.197*** [0.0534]
Observations	6,143	4,843	6,143	4,843
R-squared (pseudo)	0.104	0.11	0.066	0.072
Industry * Year F.E.	Yes	Yes	Yes	Yes
Limit Sample to Big 5	No	Yes	No	Yes

Table 3 – Main Results: Pervasiveness of Fraud Based on AA Natural Experiment

The table presents the calculations described step-by-step in Section II of the text. Panel A re-writes the text equation definitions and defines the fraud measures for convenience. Panel B presents the frequency of revealed fraud in various samples. Panel C reports the calculations to go from the likelihood of detection to the unconditional probability of fraud. Data for the calculations are from Compustat, DMZ(2010), and, for AAERs, the Center for Financial Reporting and Management, Haas School of Business, Berkeley. The sample is U.S. publicly traded corporations with more than \$750 million in assets. A firm is identified as being an Arthur Andersen (AA) client if it was audited by AA in 2001 or 2002. In DMZ, a fraud is identified as a financial fraud if the firm involved restated its financials after the revelation of the fraud. In the AAER sample, a firm is identified as having a financial misrepresentation if it received an AAER from the SEC. The detection likelihoods, defined in the text, represent the ratio of the AA to non-AA probabilities for each fraud measure. The difference between Core and Most Conservative estimates concerns whether non financial frauds have similar detection likelihoods as financial frauds (the core assumption) versus the non-financial frauds having a detection likelihood of 1 (perfect detection), the most conservative assumption. For calculations of detection likelihoods for all frauds and unconditional probability of all frauds, the percentage of financial frauds amongst all frauds is set at 64.7%, the observed fraction of financial frauds over the 1996-2004 period in the DMZ data. For the unconditional probability of fraud calculations in panel C, we use DMZ's estimate of the probability of committing fraud and getting caught of 4% presented in Figure 2.

Panel A: Definitions from Text Equations & Measures

Equation Terms

Pr (F or FF, caught |A) : Probability of a AA client in 2001 committing fraud which started prior
 Pr (F or FF, caught | \bar{A}) : Probability of a non-AA firm committing fraud which started prior to

Measures

Sample of Fraud (F) F: Class actions : The DMZ sample of SCAC class action frauds, filtered to remove frivolous fraud. This sample includes misrepresentation frauds (64.7% of observations), as well as failed or false disclosures, regulatory incompliance, and self dealing.

Samples of Financial Fraud (FF) FF1: Auditor-Identified Fraud: The subset of class actions in which the auditor revealed the fraud or the firm revealed the fraud and the auditor did not issue a perfectly unqualified audit opinion prior to revelation.
 FF2: Misrepresentations: The subset of class actions with misrepresentation (64.7% of cases)
 FF3: AAERs: Firms for which the SEC issued an AAER.

Panel B: Test Statistics: Observed Probabilities of Fraud

	for A firms Pr (F or FF, caught A) =	for \bar{A} firms Pr (F or FF, caught \bar{A}) =	p-value for difference ttest
FF1 Sample (Auditor-Identified Fraud)	0.0217	0.0057	0.0058***
FF2 Sample (Misrepresentations)	0.0397	0.0179	0.0186**
FF3 Sample (AAERs)	0.0902	0.0514	0.0096***
F Sample (Class Actions)	0.0578	0.0312	0.0254**
Observations	277	1,730	

Panel C: Fraud Detection Likelihood and Main Estimates of Fraud Probability

	Detection Likelihoods			Conditional probability	Unconditional Probability of Fraud	
	$Pr(caught FF)$	$Pr(caught F)$		$Pr(F, caught \bar{A})$	$P(F) = \frac{Pr(F, caught \bar{A})}{Pr(caught F)}$	
		Core Estimates	Most Conservative Estimates	Estimated (Figure 2) for 1996-2004	Core Estimates	Most Conservative Estimates
FF1 Sample (Auditor-Identified Fraud)	0.267	0.267	0.526	0.040	0.150	0.076
FF2 Sample (Misrepresentations)	0.451	0.451	0.645	0.040	0.089	0.062
FF3 Sample (AAERs)	0.570	0.570	0.722	0.040	0.070	0.055
F Sample (Class Actions)	--	0.540	0.540	0.040	0.074	0.074

Table 4: Variables in Probit/Biprobit Estimations of Fraud Detection

The table presents the variables, their description, and the sources of the data used in Table 4. Variables in panels A-C are used in the probit estimates of Table 4. Variables in panels A-D are used in the partial observability bivariate probit in Table 4. Panel A defines the high incentives and opportunities indicator variables. Panel B defines the primary control variables that affect detection and engaging in fraud. Panel C defines variables used in the detecting equation only. Panel D defines variables only used in the engaging equation.

<i>Panel A: High Incentives and Opportunities Indicator Variables</i>		
Variable	Description	Source
High Analyst Coverage Indicator	A dummy variable that takes the value of 1 for companies with higher than the median value of analyst coverage in companies with more than \$750 million in assets.	I/B/E/S
High Media Coverage Indicator	A dummy variable that takes the value of 1 if the firm has higher than the median value of media coverage in companies with more than \$750 million in assets. We manually collect media coverage by searching the Wall Street Journal print edition and recording the number of media hits for the year 1995.	Factiva
High Shortability Indicator	A dummy variable that takes the value 1 for companies with a greater than median level of institutional shareholding in the prior year.	Compact-D
Regulated Firm Indicator	A dummy variable that take the value of 1 if the firm is in the following categories: financials, transportation equipment, transportation, communications, electric, gas and sanitary services, drug, drug, proprietaries and druggists sundries, petroleum and petroleum products wholesalers pharmaceuticals, healthcare providers, and healthcare related firms in business services.	Industries identified in Winston (1998) and others
Qui-Tam Industry Indicator	A dummy variable that takes the value of 1 if the industry is one in which qui tam lawsuits are possible. Included are healthcare and defense contractor industries.	Civil Division, Department of Justice
Post Sox Indicator	A dummy variable that takes the value of 1 if the time period is post-SOX.	Legislation date
<i>Panel B: Primary Control Variables in Both Detecting and Engaging Equations</i>		
Company Size	Log of total book assets	Compustat
Stock Return	Total return on stock	CRSP
R&D	Log of R&D expenditures / total assets	Compustat
Leverage	Long term Debt/ total assets	Compustat
ROA (lag)	Operational income after depreciation/Total Assets	Compustat
Stock Return (lag)	$(P_t - P_{t-1} + D) / P_t$	CRSP

VIX	Implied volatility of S&P 500 index options	CRSP
Industry HHI	The level of industry concentration, based on the sum of the squared market shares using sales (HHI) for publicly traded firms, using 2-digit SIC industries	Compustat
Firm Market Share	The firm market share using sales and 2-digit SIC industry definition.	Compustat
<i>Panel C: Variables in Detecting Equation Only</i>		
Abnormal ROA	Residual from regression with i denoting company; j industry; and t time: $ROA_{ijt} = \alpha_0 + \alpha_1 ROA_{ijt-1} + \alpha_2 \overline{ROA}_{jt} + \epsilon_{ROA,ijt}$, where \overline{ROA}_{jt} denotes the industry average. This estimation removes serial correlation and the industry effect.	Compustat
Abnormal Stock Return	Residual from CAPM regression: $r_{it} = r_{ft} + \beta_i(r_{mt} - r_{ft}) + \epsilon_{r,it}$. r_{mt} , r_{it} , and r_{ft} denote the market return, the firm return, and the risk free rate, all in quarter t .	CRSP
<i>Panel D: Variables in Engaging Equation Only</i>		
Lop Options Held	The sum of the in-the-money exercisable options for all executives.	Execucomp
Incentive Pay percentage (lag)	The average of the ratio of restricted stock grants divided by total compensation across executives for a firm-year.	Execucomp
Sentiment index	Annual data on sentiment used and described in Baker-Wurgler (2006). Sentiment index is first principal component of six standardized sentiment proxies (value-weighted dividend premium, IPO volume, First day returns, closed-end fund discount, equity share in new issues, log of NYSE turnover)	Jeffrey Wurgler website
Sentiment index squared	The square of the sentiment index.	Jeffrey Wurgler website

Table 5: Probit and Partial Observability Bi-Variate Probit Estimates of Fraud Detection

The dependent variable in models (1) and (2) is an indicator variable whether a firm is caught engaging in fraud in that year. In these models, we use a probit estimate and report the marginal effects of changes in the independent variables. The independent variables for (1) are defined in Table 3 panels B and C, and for (2) in panels A-C. In model (3) we use the partial observability bivariate probit model of Poirier and the dependent variable in the caught equation is the same as in the probit, and the dependent variable in the start equation is a variable that takes the value 1 in the year in which the firm first engages in the fraud. The first column of model 3, (3.1), is the estimate for the caught equation, the second column (3.2) provides the marginal effects from the caught equation, and the third column (3.3) provides the estimates on the independent variables for the start equation. The variables used in the caught equation are defined in Table 4, panels A-C, and the variables used in the start equation are defined in Table 4, panel A,B, and D. Standard errors are in brackets. ***, **, and ** denote significance at the 1%, 5%, and 10% confidence intervals, respectively.

Caught Equation	(1)	(2)	(3.1)	(3.2)	(3.3)
	Probit (Caught=1) MFX	Probit (Caught=1) MFX	BiProbit (Caught, Start) Caught	MFX Con- ditional: P(Caught Start)	BiProbit (Caught, Start) Start
<i>Abnormal Leverage (lag)</i>	0.0036 [0.0042]	0.0060 [0.0046]	0.0905* [0.0519]		
<i>Abnormal ROA (lag)</i>	0.0000 [0.0002]	-0.0001 [0.0002]	0.0015 [0.00262]		
<i>Abnormal Stock Return (lag)</i>	-0.0069 [0.0108]	-0.0145 [0.0099]	-0.0517 [0.0911]		
<i>Abnormal VIX</i>	-0.0012*** [0.0004]	0.0000 [0.0005]	0.0040 [0.00514]		
<i>Log Options Held</i>					-0.0347*** [0.0123]
<i>Incentive Pay % (lag)</i>					0.286** [0.124]
Hi Analysts (lag)		0.0087*** [0.0026]	0.176** [0.0840]	0.0271	-0.135** [0.0653]
Hi Media (lag)		-0.0043 [0.0029]	0.0690 [0.102]	0.0008	-0.0688 [0.113]
Hi Shortability (lag)		0.0055** [0.0026]	0.185** [0.0722]	0.0142	-0.139** [0.0593]
Qui Tam Industry		0.0215** [0.0084]	3.014*** [0.832]	0.0002	-2.969*** [0.803]
Regulated		-0.0040 [0.0034]	-0.1440 [0.481]	-0.0081	0.1730 [0.499]
PostSOX		0.0165*** [0.0043]	1.521*** [0.513]	0.0006	-1.594*** [0.526]
Log Assets (lag)	0.0024** [0.0009]	0.0014 [0.0010]	-0.0796 [0.0636]	0.0275	0.129* [0.0713]
Log R&D (lag)	0.0014*** [0.0005]	0.0003 [0.0006]	-0.0839 [0.0703]	0.0108	0.0889 [0.0703]
Leverage (lag)	0.0219*** [0.0073]	0.0147** [0.0072]	-0.0837 [0.221]	0.0289	-0.0560 [0.215]
ROA (lag)	-0.0276* [0.0159]	-0.0241* [0.0146]	12.83*** [3.449]	-0.0125	-12.68*** [3.435]
Stock Return (lag)	-0.0027 [0.0027]	-0.0010 [0.0024]	1.277*** [0.358]	-0.0131	-1.176*** [0.347]
Vix	0.0029*** [0.0005]	0.0021*** [0.0004]	0.0005 [0.0667]	-0.0131	0.0051 [0.0688]
Observations	8114	7671	5973		
LR/Wald Chi-Square	83.2	133.4	42.3		
Sum of Marginal Effects		4.39		3.48	

Table 6: Summary Statistics for Cost of Fraud Estimates

The table presents the statistics setting up the counterfactual exercise to estimate the cost of corporate fraud provided in Table 8. The statistics are reported only for fraud firms of DMZ's original sample of 216 firms which have statistics for pre and post periods. The pre and post columns represent the same set of firms, with the counts depending on availability of financial items. Panel C restricts the sample to the set of firms who were Arthur Andersen clients in 2000 or 2001 and subsequently were revealed to have started fraud during Arthur Andersen's watch. Panel A reports the valuation of equity, long-term debt and enterprise value as well as the financial statement line items which enter the multiples analysis. The numbers are in millions of USD. The firm counterfactual in panels B and C is based on the pre-fraud firm value that is projected to experience the same multiple change as a typical firm in the industry, defining industry as the same2 digit SIC. In panels B and C the reported numbers are medians.

Panel A: Statistics

	Pre-Fraud				Post-Fraud			
	Median	Mean	StDev	Frequency	Median	Mean	StDev	Frequency
Market Capitalization	4,976	15,820	30,287	196	2,214	9,925	19,472	196
Long Term Debt	854	2,861	7,574	196	884	6,110	32,357	196
Enterprise Value	6,811	18,682	32,380	196	5,632	19,481	46,303	196
EBITDA	460	1,264	1,960	179	458	1,713	3,313	179
Sales	2,519	6,927	9,194	196	3,378	8,411	11,598	196
Assets	3,500	14,047	33,386	196	4,205	23,233	77,445	196

Panel B: Firm Value Pre-Fraud and Firm Value Post Fraud for Fraud Firm and Counterfactual Firm

	Pre-Fraud	Post-Fraud	
		Fraud Firm	Counterfactual Firm
EBITDA Multiple	12.27	9.00	12.13
Sales Multiple	2.36	1.38	2.20
Assets Multiple	1.42	0.77	1.25

Panel C: Multiples for 'Normally Not-Detected' Frauds using Arthur Andersen Natural Experiment

	Pre-Fraud	Post-Fraud	
	(Median Multiple)	Fraud Firm	Counterfactual Firm
EBITDA Multiple	10.60	8.12	10.04
Sales Multiple	2.20	1.22	2.21
Assets Multiple	1.02	0.62	0.85

Table 7: Cost of Fraud

The table presents the statistics setting up the counterfactual exercise to estimate the cost of corporate fraud provided in Table 8. The statistics are reported only for fraud firms of DMZ's original sample of 216 firms which have statistics for pre and post periods. The pre and post columns represent the same set of firms, with the counts depending on availability of financial items. Panel C restricts the sample to the set of firms who were Arthur Andersen clients in 2000 or 2001 and subsequently were revealed to have started fraud during Arthur Andersen's watch. Panel A reports the valuation of equity, long-term debt and enterprise value as well as the financial statement line items which enter the multiples analysis. The numbers are in millions of USD. The firm counterfactual in panels B and C is based on the pre-fraud firm value that is projected to experience the same multiple change as a typical firm in the industry, defining industry as the same2 digit SIC. In panels B and C the reported numbers are medians.

Panel A: Cost of Detected Fraud in Dollars and as a Percentage of Enterprise Value

<i>Based on:</i>	Distribution of Costs in \$ Million		
	25 th	Median	75 th
EBITDA Multiple	-116	1,175	6,079
Sales Multiple	165	2,063	7,916
Assets Multiple	274	1,556	6,054
<i>Average</i>		1598	

<i>Based on:</i>	Distribution of Costs as % of Enterprise Value		
	25 th	Median	75 th
EBITDA Multiple	-0.066	0.246	0.627
Sales Multiple	0.063	0.390	1.133
Assets Multiple	0.098	0.361	0.900
<i>Average</i>		0.332	

Panel B: Cost of 'Normally Not-Detected Fraud' as a Percentage of Enterprise Value

Sample of Firms with Arthur Andersen as Auditor in 2000 or 2001

<i>Based on:</i>	Median
	Distribution of Costs as % EV
EBITDA Multiple	0.083
Sales Multiple	0.253
Assets Multiple	0.232
<i>Average</i>	0.189

Panel C: Cost of Fraud as a Percentage of Enterprise Value Excluding Legal Costs

Sample of Firms with Arthur Andersen as Auditor in 2000 or 2001

Median Detected Fraud (average of all Multiples)	Median 'Normally Not-Detected' Fraud (average of all Multiples)
30%	17%

Panel D: Alternative Cost of Detected Fraud Based on 25th Percentile Growth Rate

<i>Based on:</i>	25th percentile Growth rate
	EBITDA Multiple
Sales Multiple	0.178
Assets Multiple	0.128
<i>Average</i>	0.113