Registered Report Proposal

PCAOB Monitoring and Auditor Effort: Evidence from Dynamic Model Estimation

Lisa Yao Liu and Shunsuke Matsuno*

February 2026

Abstract

We develop a dynamic structural model of auditor effort choice under PCAOB monitoring. Auditors face competing pressures: effort is costly, but low effort increases the probability that the PCAOB detects and reports deficiencies, which carry reputational and regulatory consequences. Because PCAOB monitoring intensity depends on deficiency history, auditors work harder not only to avoid current-period deficiencies but also to avoid triggering stricter future monitoring. We decompose auditor incentives into static deterrence (responding to current inspection risk) and dynamic deterrence (avoiding a poor regulatory record to prevent heightened future scrutiny), a distinction reduced-form methods cannot identify. Our proposed estimation recovers preference parameters measuring the cost of effort and the disutility from deficiencies. We design counterfactual analyses comparing dynamic enforcement to static alternatives and examining how the difference in monitoring intensity between auditors with clean versus poor regulatory records affects their behavior. The framework allows policymakers to evaluate alternative PCAOB inspection designs.

Keywords: Audit quality, PCAOB, regulation, enforcement, auditor effort
JEL Classification: G18, G38, K22, M42

*Lisa Yao Liu: Columbia Business School, Columbia University (yl4689@gsb.columbia.edu). Shunsuke Matsuno: Columbia Business School, Columbia University (SMatsuno26@gsb.columbia.edu). We appreciate helpful comments from two anonymous referees, Ehsan Azarmsa, Tim Baldenius, Phil Berger, Jeremy Bertomeu (discussant), Jon Glover, Changyong Kim, Joaquin Peris, Harsha Samaraweera (PCAOB), Nemit Shroff (editor), and participants at Columbia Business School and the 2024 PCAOB/TAR Registered Reports Conference. We thank Emmett Whitlock for his research assistance. We gratefully acknowledge financial support from Columbia Business School.

1 Introduction

Over the past two decades, public authorities have significantly increased their oversight of audit regulations. The Sarbanes-Oxley Act of 2002 (SOX) established the Public Company Accounting Oversight Board (PCAOB), a quasi-public agency that oversees and inspects audits of publicly traded companies in the United States. Heightened public oversight is usually justified by the presence of externalities and the challenges associated with self-regulation (e.g., the previous regime in which the American Institute of Certified Public Accountants (AICPA) conducted peer reviews). In the presence of such difficulties, mandatory inspections may enhance capital market welfare by improving audit quality. Despite a large body of evidence on the effects of PCAOB oversight (e.g., Lennox and Pittman, 2010; DeFond and Lennox, 2011; Lamoreaux, 2016; Aobdia and Shroff, 2017; DeFond and Lennox, 2017; Fung et al., 2017; Aobdia, 2018, 2020), debates continue over whether the benefits of monitoring justify their costs, and how inspection programs should be designed to improve deterrence.

A central question in this debate is how auditors respond to regulatory monitoring. Auditors operate in an environment with potentially conflicting pressures. On one hand, they derive revenue from client fees, and clients may prefer auditors who complete engagements efficiently without raising extensive concerns about accounting treatments. On the other hand, auditors face negative consequences when the PCAOB detects audit deficiencies, including reputational damage from public inspection reports, potential sanctions, and heightened future scrutiny. Understanding how auditors balance these competing pressures, and how regulatory design affects this balance, is essential for evaluating and improving PCAOB monitoring programs (see Hanlon and Shroff (2022) for a comprehensive discussion).

While prior studies provide useful evidence on the relation between PCAOB oversight and auditor behavior (e.g., DeFond and Lennox, 2011; Aobdia, 2018), important policy questions remain open. First, reduced-form estimates capture total effects but do not distinguish whether auditors respond primarily to *current-period* inspection risk or to the *future consequences* of a poor regulatory record. Second, reduced-form methods can evaluate the effects of observed policy variation but cannot predict auditor responses to counterfactual policies that have not been implemented. Our structural approach is designed to address these questions.

We develop and estimate a dynamic structural model of auditor effort choice under PCAOB monitoring. We aim to inform the design of monitoring policies by explicitly modeling the dynamic incentives auditors face. The interaction between the PCAOB and auditors is not a one-shot game. When making effort choices, auditors consider both immediate consequences and future regulatory monitoring. The PCAOB's monitoring intensity depends on an auditor's deficiency history: specifically, whether the most recent inspection found deficiencies. Auditors with recent deficiency findings face heightened scrutiny, while those with clean records face less intensive monitoring. This state-dependent monitoring creates a dynamic tradeoff for the auditor. In the current period, higher effort reduces the probability of deficiency detection, lowering immediate regulatory costs. Looking forward, higher effort today increases the probability of having (or maintaining) a clean regulatory record, which reduces expected monitoring intensity and deficiency costs in future periods. This forward-looking component of auditor incentives, which we term *dynamic deterrence*, cannot be recovered from reduced-form evidence. By estimating a structural model, we can decompose auditor effort incentives into a static component (responding to current inspection risk) and a dynamic component (the value of a clean regulatory record).

Our model is based on the dynamic enforcement literature (e.g., Polinsky and Shavell, 1998, 2000). An auditor interacts with the regulator (the PCAOB) over time. In each period, the auditor exerts costly effort, which is unobservable to the regulator. The PCAOB inspects the auditor and may detect and report deficiencies. The probability of a reported deficiency depends on two factors: the auditor's effort choice (which affects underlying audit quality) and the PCAOB's monitoring intensity (which depends on the auditor's publicly observable state). We decompose the reported deficiency probability as the product of the intensity of the monitoring activities and audit quality.

The PCAOB commits ex ante to a stationary monitoring policy that depends only on the auditor's publicly observable regulatory state, determined by past inspection outcomes, and not on unobservable effort.[1] The auditor minimizes expected costs, trading off effort costs against the probability and consequences of deficiency findings, while considering how current choices affect future monitoring intensity.

We assume that the true audit effort is not observable or contractable; thus, the PCAOB cannot punish auditors for their level of effort. While certain audit inputs (such as hours worked) are observable, true audit effort, such as the rigor of testing, appropriateness of professional judgments, and sufficiency of evidence, is not directly observable. Moreover, even if audit quality could be measured, determining whether it was appropriate would need to consider engagement-specific factors such as client complexity and risk. The same level of audit quality may be adequate for a simple engagement but insufficient for a complex one. In our model, the PCAOB inspects whether the auditor's work is sufficient from the PCAOB's perspective, given the specific requirements of each engagement, and reports deficiencies if not. Because true auditor effort is unobservable, we decompose observed deficiency rates into regulatory monitoring intensity and underlying audit quality. This separation allows us to isolate auditor behavior from regulator policy. When we estimate the model, we construct a noisy proxy for effort from observable audit inputs, taking measurement error into account through a signal precision parameter.

An important feature of the financial reporting environment is that the clients are responsible for preparing financial statements, whereas auditors provide assurance. To focus on how auditors respond to the PCAOB's dynamic monitoring policy, our model abstracts from the demand side of the audit market. First, we treat client behavior as exogenous. If clients anticipate stricter auditor scrutiny, they may invest more in their own financial reporting quality; we abstract from this channel and treat client behavior as fixed. Second, we treat audit fees as exogenous. In practice, fees may adjust as auditors and clients bargain over how regulatory costs are shared (DeFond and Lennox, 2017), but explicitly modeling this would require us to specify auditor-client bargaining and audit market structure, which is beyond our current scope. These simplifications imply that our estimates

[1]This policy remains fixed over time; the regulator does not dynamically adjust monitoring intensity in response to individual auditor behavior. The assumption of ex-ante commitment is standard in the enforcement literature (e.g., Polinsky and Shavell, 1998; Blundell et al., 2020; Kang and Silveira, 2021). In our institutional setting, the PCAOB inspects under published auditing standards and inspection procedures, which can be interpreted as pre-committed monitoring policies.

capture the *direct* effect of PCAOB monitoring on auditor effort, holding client behavior and fees fixed. If client behavior responds positively to auditor effort, or if fee adjustments do not fully offset auditors' effort incentives, our estimates represent a lower bound on the total benefits of PCAOB monitoring.

We estimate the model using a standard two-step approach in estimating dynamic discrete choice models. In the first step, we non-parametrically estimate deficiency rates and state transition probabilities. Unlike standard settings in estimating dynamic discrete choice models, true effort is unobservable in our setting. Thus, we construct a noisy effort signal from audit inputs (such as hours) and recover choice probabilities by adjusting for signal imprecision.

In the second step, we estimate the structural preference parameters: the cost of effort (η) and the disutility from deficiency findings (θ_R). We estimate these parameters by matching the model's predicted effort choices to their empirical counterparts. Identification comes from variation in monitoring intensity across regulatory states: auditors with prior deficiencies face heightened scrutiny, and their differential effort response reveals how they trade off current effort costs against both immediate and future consequences of poor performance. This forward-looking behavior is central to our analysis of dynamic enforcement. The ratio θ_R/η quantifies the revealed-preference "price" auditors place on avoiding PCAOB findings.

Our paper contributes to the literature on PCAOB monitoring and auditor behavior. First, we provide new economic insights into the mechanisms through which PCAOB monitoring affects auditor effort. Prior studies document that auditor behavior changes around PCAOB inspections (e.g., DeFond and Lennox, 2011), but these studies cannot distinguish between static deterrence, where auditors work harder when facing current inspection risk, and dynamic deterrence, where auditors work harder today to avoid heightened scrutiny in the future. By decomposing deterrence into static and dynamic components, we can assess whether auditors respond primarily to current inspection risk or to the forward-looking value of a clean regulatory record.

Second, we conduct quantitative evaluations of counterfactual monitoring policies. Prior research evaluates the effects of observed PCAOB policies (Lennox and Pittman, 2010; DeFond and Lennox, 2011; Aobdia, 2018) but cannot assess hypothetical alternatives. Our structural estimates allow us to simulate auditor behavior under counterfactual enforcement policies. We consider static monitoring (uniform scrutiny regardless of past performance) and policies that vary the degree of heightened scrutiny imposed on auditors with or without prior deficiencies. Quantifying the dynamic deterrence effect and analyzing these counterfactuals would provide insights directly relevant to current policy debates about PCAOB inspection design. Such insights are difficult to obtain from reduced-form methods that rely on observed policy variation.

Third, we contribute to the growing literature on structural estimation in accounting (Bertomeu et al., 2023; Breuer et al., 2024). In the audit setting, Cheynel and Zhou (2024) develop a dynamic structural model of auditor tenure and misstatement detection, and Cheynel et al. (2025) estimate the value of auditing services for private firms using conditional choice probabilities. We complement their work by focusing on the PCAOB's regulatory monitoring regime and its deterrent effect on auditor effort in public company audits. Our work also relates to the theoretical and empirical literature on regulatory enforcement (Becker, 1968; Polinsky and Shavell, 1998; Blundell et al., 2020; Christensen et al., 2020; Kang and Silveira, 2021), applying insights from this literature to audit oversight.

2 Institutional Background

2.1 PCAOB Overview

The PCAOB is a quasi-public agency established by the SOX Act, funded largely by SEC-registered public companies ("issuers"), and overseen by the Securities and Exchange Commission (SEC). The PCAOB registers, oversees, and inspects all audit firms employed by these issuers. This new regime represents a major shift from self-regulation to public oversight of the audit profession.[2]

The SOX Act and the PCAOB's rules stipulate that for audit firms with over 100 issuers, the PCAOB conducts an annual inspection. For audit firms with 100 or fewer issuers, the PCAOB conducts (at least) one inspection every three years. The PCAOB has substantial enforcement authority and a wide array of sanctions.

[2] The discussion in this section is largely adapted from the PCAOB website and the CAQ (Center for Audit Quality) document (CAQ, 2021). Additional institutional details are provided in Appendix A.

2.2 Inspection Process and Deficiency Reporting

The PCAOB selects engagements to inspect using a combination of risk-based methods and random selection. Risk factors include the nature of the company, such as the industry and the market capitalization, audit issues likely to be encountered, and whether the company has significant operations in emerging markets. The audit firm has no influence on the PCAOB's selection of audits inspected.

During inspections, PCAOB staff review audit work papers and interview relevant personnel. If, during their fieldwork, inspectors identify potential deficiencies in one or more engagements, the PCAOB gives the auditor the opportunity to respond. Deficiencies that are not satisfactorily addressed are included in the public inspection report as Part I findings, which assert that the engagement team's work did not sufficiently support the audit opinion.

2.3 Implications for Modeling

Several features of the PCAOB inspection regime are important for our model. First, the PCAOB's risk-based selection implies that auditors with prior deficiencies face heightened scrutiny, creating state-dependent monitoring intensity. Second, the public disclosure of Part I findings creates reputational consequences that affect auditor incentives. Third, the dynamic nature of the inspection process, where past performance influences future monitoring, motivates our focus on auditors' forward-looking behavior. Fourth, the PCAOB reviews the auditor's work and the circumstances of each engagement to assess compliance with applicable standards. This motivates our modeling assumption that the regulator does not observe the auditor's true effort but instead observes only a noisy signal of that effort.

3 Related Literature

A large body of research examines the effects of PCAOB inspections on audit quality using reduced-form methodologies (one notable exception of Hanlon and Shroff (2022), a survey paper that offers valuable insights into the inspection process and its economic impacts). Studies document that inspections have positive effects on audit quality for both domestic and foreign auditors. Aobdia (2018) finds that audit firms and clients respond to the PCAOB engagement inspection process. Their response often aligns with the standards set by the Part I Findings, and their remediation efforts following Part II Findings are aimed at improving the financial reporting quality of the issuer. Aobdia (2020) documents a negative correlation between quality control deficiencies and overall audit quality. Audit firms are incentivized to avoid deficiencies to mitigate litigation risks (Christensen et al., 2021) and avert negative perceptions (e.g., Aobdia and Shroff, 2017; Shroff, 2020).

Inspections prompt lower-quality auditors to exit the market (DeFond and Lennox, 2011), without significantly impacting the remaining audit firms' market shares (Lennox and Pittman, 2010). Moreover, inspections incentivize audit firms to address deficiencies (DeFond and Lennox, 2017). For foreign auditors, inspections increase audit quality and expand market share (Lamoreaux, 2016; Fung et al., 2017; Aobdia and Shroff, 2017).

Research also examines the geographic deterrent effect beyond directly inspected engagements. Lamoreaux et al. (2022) find that large audit firm offices improve audit quality following enforcement actions against other offices within their firm, while small firm offices respond to enforcement against local competitors. Boone et al. (2015) document that PCAOB disciplinary actions against Deloitte reduced the firm's ability to retain clients and slowed audit fee growth.

A related literature examines how audit fees and client behavior respond to regulatory oversight (DeFond and Lennox, 2017). We abstract from demand-side responses to focus on auditors' effort incentives, treating fees and client behavior as exogenous, as discussed in the introduction.

We build on this reduced-form evidence by developing a structural model that enables three types of analyses that are otherwise difficult. First, we decompose the channels through which monitoring operates, distinguishing whether auditors respond to current-period inspection risk or to future consequences of a poor regulatory record. Second, we recover preference parameters with direct economic interpretation: the cost of effort and disutility from deficiencies, whose ratio quantifies the revealed-preference "price" auditors place on avoiding PCAOB deficiency findings. Third, we separate observed deficiency rates into regulatory policy (monitoring intensity) and auditor behavior (effort choices). Together, these allow us to predict auditor responses to counterfactual policies not yet implemented and to inform the design of effective monitoring regimes.

4 A Dynamic Model of Auditor Effort and Deterrence

We develop a dynamic model of auditor-regulator interactions that captures key economics of deterrence. Our goal is to maintain parsimony while capturing the core economic forces and laying out assumptions under which the model can be estimated. We assume access to data on PCAOB inspections (e.g., detected deficiencies) and the audit process (e.g., audit hours and team sizes). Table 1 summarizes the notations used throughout our study.

Notation	Definition
Primitives	
$a \in \{0,1\}$	Auditor's true effort (unobserved by regulator)
$s \in \{0,1\}$	Public state (0 = clean, 1 = risky)
q_s	Detection/reporting intensity in state s
$\pi(a)$	Deficiency probability given effort a
η	Cost of high effort
θ_R	Disutility per expected reported deficiency
δ	Discount factor
$\varepsilon_t(a)$	Idiosyncratic cost shock for effort a
Derived Objects	
$v(a,s)$	Flow cost: $v(a,s) = \eta a + \theta_R r(a,s)$
$r(a,s)$	Deficiency probability given (a,s): $r(a,s) = q_s \times \pi(a)$
$d_t \in \{0,1\}$	Deficiency outcome (1 = deficiency reported)
$G(s' \mid a,s)$	State transition probability (induced by $s_{t+1} = d_t$)
$w(a,s)$	Choice-specific value function
$V(s)$	Value function
$\mathrm{EV}(a,s)$	Expected continuation value
$p(a \mid s)$	Conditional choice probability (CCP)
Estimation Objects	
$A_t \in \{0,1\}$	Observed effort signal (noisy proxy for a_t)
ρ	Signal precision: $\Pr(A_t = a_t)$
m_s	Signal frequency: $\Pr(A_t = 1 \mid s_t{=}s)$
$R(A,s)$	Observed deficiency probability given (A,s)

Table 1: List of Notations

4.1 Setup

Let $t = 1, 2, \ldots$ index time periods. The unit of observation is the audit firm (not individual engagements). At the start of period t, the auditor chooses true effort $a_t \in \{0,1\}$, where $a_t = 1$ denotes high effort and 0 denotes low effort. This is the auditor's only choice variable. Crucially, true effort is neither contractable nor directly observed by the regulator. This is what makes the inspection problem nontrivial and creates value for PCAOB monitoring. During year t, the regulator reviews the audit firm and may report deficiencies, denoted by $d_t \in \{0,1\}$, where $d_t = 1$ indicates that a deficiency is reported.

4.2 Enforcement Policy

Regulator's Committed Enforcement Policy The regulator uses a dynamic enforcement policy, committing to a monitoring intensity for each possible state $s \in \{0, 1, \ldots, S\}$, where $S \geq 1$ is the total number of regulatory states. For simplicity, we focus on the binary-state case in our main analysis: $S = 1$. This policy is publicly observable and stable over the auditor's decision horizon.[3] Specifically, the regulator commits ex ante to a state-dependent detection/reporting intensity:

$$q_s \in (0, 1], \quad s \in \{0, 1, \ldots, S\}.$$

The parameter q_s is the probability that a deficiency, if present, is detected and reported when the public state is s. We take this policy as given and do not model the regulator's optimization problem; our goal is to characterize auditor responses to the observed enforcement regime.

Decomposing Reported Deficiencies A key feature of our model is the separation of auditor choices from regulator policy. We decompose the probability of a reported deficiency into two components: the underlying quality of the audit (auditor behavior) and the intensity of regulatory monitoring (regulator policy). This decomposition allows us to identify structural parameters from the data, where we observe reported deficiencies but not true effort.

Let $\pi(a_t)$ denote the underlying probability that an inspected audit firm has a deficiency, as a function of effort a_t. We assume that high effort reduces this probability: $0 < \pi(1) \leq \pi(0)$. The reported deficiency probability is

$$r(a_t, s_t) = \underbrace{q_{s_t}}_{\text{regulator policy}} \times \underbrace{\pi(a_t).}_{\text{auditor behavior}} \tag{1}$$

In period t, conditional on inspection, the regulator reports a deficiency ($d_t = 1$) with probability $r(a_t, s_t)$.[4]

[3] We interpret commitment as a rule-based enforcement regime operating under internal guidelines and procedural constraints that make frequent re-optimization unlikely.

[4] In an extension, one could model deficiency rates (total deficiencies divided by audited engagements) rather than binary outcomes, for example by letting $d_t = 1$ if the deficiency rate exceeds a threshold.

State Definition and Transition We do not observe the regulator's internal risk classification. Therefore, we define the public state s_t directly from observable inspection outcomes so that the econometrician observes the state relevant to the auditor's decision problem.[5]

We use a binary state: $s_t = 0$ if the last inspection was clean (no deficiencies) and $s_t = 1$ if the last inspection found at least one deficiency. This specification reflects the PCAOB's risk-based inspection approach, under which auditors with recent deficiencies face heightened scrutiny.[6]

The state updates deterministically to the current inspection outcome:

$$s_{t+1} = d_t \in \{0, 1\}. \tag{2}$$

Together with the deficiency probability $r(a, s)$, this implies the transition law $G(s' \mid a, s) := \Pr(s_{t+1} = s' \mid a_t = a, s_t = s)$: $G(1 \mid a, s) = r(a, s)$ and $G(0 \mid a, s) = 1 - r(a, s)$ on $\{0, 1\}$. Note that effort affects the transition indirectly: effort determines the deficiency probability, and the deficiency outcome determines the next state.

Dynamic Deterrence The auditor's problem can be solved for any enforcement policy q_s. We focus on policies satisfying $q_1 > q_0$, which we term *dynamic deterrence policies.* Under such policies, the regulator monitors more intensively following a deficiency finding. This creates forward-looking incentives: the auditor exerts effort not only to avoid current-period deficiency reports but also to avoid the heightened future scrutiny that a poor regulatory record would bring.

4.3 Auditor payoff and dynamic choice

Incentive Environment In our model, the auditor's effort reflects a trade-off between the cost of exerting effort and the disutility from reported deficiencies. Effort is costly. Auditors must devote substantial resources to achieve high audit quality. Moreover, some clients prefer auditors

[5] Modeling s_t as an observed summary of past inspection outcomes rather than as a latent state is consistent with standard dynamic discrete-choice models (Aguirregabiria and Mira, 2010). Allowing for latent states would add substantial complexity and require additional identifying assumptions, and is beyond our current scope. Our goal is to quantify dynamic deterrence, not to recover the regulator's internal risk classification.

[6] Empirically, Aobdia (2018) and DeFond and Lennox (2017) document that PCAOB monitoring intensity responds to prior inspection outcomes. Our binary specification captures the key feature that recent deficiencies trigger increased scrutiny, while maintaining tractability. Richer state spaces (e.g., incorporating multiple past inspections or deficiency severity) would be a natural extension but would require additional data on the regulator's internal classification.

who issue unqualified opinions without extensive challenges to accounting treatments, in which case more intensive auditing is costly. The static cost of effort captures these considerations. At the same time, auditors face penalties when the PCAOB identifies deficiencies. These penalties include reputational damage, potential sanctions, and heightened future scrutiny. The disutility from deficiencies summarizes these consequences.

Flow Payoff Let $\eta > 0$ denote the per-period cost of high effort, and let $\theta_R > 0$ denote the auditor's disutility from a reported deficiency ($d_t = 1$).[7]

We interpret θ_R as a reduced-form composite capturing both direct costs (e.g., responding to regulatory findings, potential sanctions) and indirect costs (e.g., reputational damage, career concerns within the audit firm). The ratio θ_R/η quantifies the auditor's revealed-preference "price" of avoiding PCAOB findings and is a key object of interest.

The auditor's expected cost is

$$v(a, s) = \eta\, a + \theta_R\, r(a, s), \tag{3}$$

where $r(a, s)$ is defined in (1). The first term is the direct cost of effort; the second is the expected disutility from reported deficiencies.

The auditor also faces idiosyncratic cost shocks. Let $\varepsilon_t(a)$ denote the shock associated with effort level a in period t. We assume that the shocks $\{\varepsilon_t(a) : t \geq 1, a \in \{0, 1\}\}$ independently follow Type-I extreme value distribution, a standard assumption that generates closed-form choice probabilities. The auditor's flow utility from choosing a in state s is[8]

$$-v(a, s) + \varepsilon_t(a). \tag{4}$$

[7] We do not distinguish deficiencies by materiality or frequency; $d_t = 1$ indicates that at least one deficiency is reported, and θ_R reflects the average disutility across deficiency types. Modeling the number of deficiencies would require additional data on how the PCAOB's response varies with deficiency count, which is beyond our current scope.

[8] We assume that auditors are risk-neutral over deficiency outcomes, which is standard in dynamic discrete choice models. This implies that auditors care about the expected deficiency probability $r(a, s)$ rather than higher moments of the deficiency distribution. Randomness in choices arises from the Type-I extreme value shocks $\varepsilon_t(a)$, not from risk preferences.

Dynamic Problem Let $\delta \in (0,1)$ be the discount factor. The choice-specific value function (the expected discounted payoff from choosing a in state s, before observing the cost shock) is

$$w(a,s) = -v(a,s) + \delta \mathrm{EV}(a,s),$$

where $\mathrm{EV}(a,s) := \mathbb{E}[V(s_{t+1}) \mid a_t = a, s_t = s]$ is the expected continuation value. Under the Type-I extreme value assumption, the value function satisfies[9]

$$V(s) = \mathbb{E}\left[\max_a \{-v(a,s) + \varepsilon(a) + \delta \mathrm{EV}(a,s)\}\right] = \log\left(\sum_{a \in \{0,1\}} \exp\{w(a,s)\}\right).$$

Given the transition law $G(s' \mid a, s)$, the expected continuation value is

$$\mathrm{EV}(a,s) = \sum_{s'} G(s' \mid a,s) \log\left(\sum_{a' \in \{0,1\}} \exp\{w(a',s')\}\right). \tag{5}$$

Conditional Choice Probability The auditor's conditional choice probability (CCP) is the probability that the auditor chooses an effort level a conditional on a state s. Under the Type-I extreme value assumption, the auditor's CCP, denoted by $p(a \mid s)$, takes the familiar logit form:

$$p(a \mid s) = \frac{\exp\{w(a,s)\}}{\sum_{a'} \exp\{w(a',s)\}}. \tag{6}$$

We recover the structural parameters by matching model-implied CCPs to their empirical counterparts.

5 A Two-Period Illustration

The above full-fledged dynamic model is suitable for our empirical analysis, but its key economic forces may not be entirely transparent. In this section, we consider a two-period version of the model to clarify those forces in closed form. This simplified setting isolates the dynamic deterrence mechanism and builds intuition for the quantitative results that follow.

We maintain the same setup as the main model, with two modifications: the horizon is two

[9] The expression omits the Euler constant, which does not affect choice probabilities.

periods ($t \in \{1, 2\}$), and the initial state is $s_1 = 0$ (clean record). We focus on dynamic enforcement ($q_1 > q_0$) and compare it to the static benchmark ($q_1 = q_0$).

5.1 Solving for the Auditor's Problem

We solve by backward induction. Since deficiency probabilities depend only on true effort a_t, the noisy signal A_t is irrelevant in characterizing the auditor's optimal behavior.

5.1.1 Period 2: Static Problem

In period 2, the auditor faces a static problem with no continuation value. Let $a := a_2$ and $s := s_2$ denote the period-2 effort and state. Define the period-2 value function in state s by

$$V_2(s) = \mathbb{E}\max\{-\eta - \theta_R q_s \pi(1) + \varepsilon_2(1), -\theta_R q_s \pi(0) + \varepsilon_2(0)\}.$$

The auditor chooses high effort in period 2 in state s if and only if the cost η is offset by the reduction in expected deficiency disutility:

$$\eta \leq (\pi(0) - \pi(1))\theta_R q_s + \underbrace{(\varepsilon_2(1) - \varepsilon_2(0))}_{:=\Delta\varepsilon_2}.$$

In the second period, the CCP is given by

$$p_2(1 \mid s) = \frac{\exp(-\eta + \theta_R q_s(\pi(0) - \pi(1)))}{1 + \exp(-\eta + \theta_R q_s(\pi(0) - \pi(1)))}. \tag{7}$$

The expected utility (omitting the Euler constant) is

$$\begin{aligned} V_2(s) &= \log\Big(\exp\big(-\theta_R q_s \pi(0)\big) + \exp\big(-\eta - \theta_R q_s \pi(1)\big)\Big) \\ &= -\theta_R q_s \pi(0) + \log\Big(1 + \exp\big(-\eta + \theta_R q_s(\pi(0) - \pi(1))\big)\Big). \end{aligned}$$

The value of a clean record. To illustrate the value of a clean record, we differentiate the value function V_2 with respect to monitoring intensity:

$$\begin{aligned}\frac{\partial V_2(s)}{\partial q_s} &= -\theta_R \pi(0) + \theta_R(\pi(0) - \pi(1)) \frac{\exp\big(-\eta + \theta_R q_s(\pi(0) - \pi(1))\big)}{1 + \exp\big(-\eta + \theta_R q_s(\pi(0) - \pi(1))\big)} \\ &\leq -\theta_R \pi(0) + \theta_R(\pi(0) - \pi(1)) \\ &= -\theta_R \pi(1) < 0.\end{aligned}$$

Since $V_2(s)$ is decreasing in q_s and $q_1 \geq q_0$, we have $V_2(0) \geq V_2(1)$, with strict inequality when $q_1 > q_0$. Intuitively, the clean state is more valuable because the auditor faces lower monitoring intensity and hence lower expected disutility.

5.1.2 Period 1: Dynamic Problem

Now consider period 1 with initial state $s_1 = 0$. The auditor's payoff from effort a_1 includes the continuation value:

$$-\eta\, a_1 - \theta_R\, q_0 \pi(a_1) + \varepsilon_1(a_1) + \delta\, \underbrace{\mathbb{E}[V_2(s_2) \mid a_1, s_1 = 0]}_{=\mathrm{EV}(a_1,0)}.$$

Because $s_2 = d_1$ and $d_1 \sim \text{Bernoulli}(q_0 \pi(a_1))$, we can compute the continuation value EV explicitly:

$$\mathrm{EV}(a_1, 0) = V_2(0) + q_0 \pi(a_1)\, [V_2(1) - V_2(0)].$$

The continuation value difference is therefore

$$\mathrm{EV}(1, 0) - \mathrm{EV}(0, 0) = q_0(\pi(0) - \pi(1))(V_2(0) - V_2(1)). \tag{8}$$

The auditor chooses high effort in period 1 when

$$\eta \leq (\pi(0) - \pi(1))\theta_R q_0 + \delta(\mathrm{EV}(1, 0) - \mathrm{EV}(0, 0)) + (\varepsilon_1(1) - \varepsilon_1(0)).$$

Therefore, we obtain the following conditional choice probability:[10]

$$p_1(1 \mid s_1 = 0) = \frac{\exp\left(-\eta + \theta_R q_0(\pi(0) - \pi(1)) + \delta(\mathrm{EV}(1,0) - \mathrm{EV}(0,0))\right)}{1 + \exp\left(-\eta + \theta_R q_0(\pi(0) - \pi(1)) + \delta(\mathrm{EV}(1,0) - \mathrm{EV}(0,0))\right)}. \tag{9}$$

5.2 The Value of Dynamic Enforcement

Comparing (9) to (7), we see that period-1 effort incentives include an additional term:

$$\delta\big(\mathrm{EV}(1,0) - \mathrm{EV}(0,0)\big) = \delta q_0(\pi(0) - \pi(1))[V_2(0) - V_2(1)].$$

We call this the *dynamic deterrence term*. The dynamic deterrence effect (i.e., the magnitude of the dynamic deterrence term) is strictly positive when $q_1 > q_0$ (in which case $V_2(0) > V_2(1)$). The term captures the forward-looking benefit of high effort: by reducing the probability of a deficiency today, the auditor is more likely to retain a clean state tomorrow.

The expression immediately implies the following intuitive comparative statics: The dynamic deterrence effect is stronger when

- the discount factor δ is higher (auditors weight the future more),
- the enforcement gap $q_1 - q_0$ is larger (clean records are more valuable), and
- the effort productivity $\pi(0) - \pi(1)$ is higher (effort is more effective at reducing deficiencies).

The two-period model provides closed-form expressions that transparently illustrate the dynamic deterrence mechanism. Our main estimation uses the infinite-horizon specification described in Section 4, which we solve numerically. The qualitative insights from the two-period case carry over: dynamic enforcement induces higher effort by making clean records valuable.

[10] The first-period value function is

$$\begin{aligned} V_1(s) &= \log\Big(\exp\big(-\theta_R q_s \pi(0) + \delta \mathrm{EV}(0,s)\big) + \exp\big(-\eta - \theta_R q_s \pi(1) + \delta \mathrm{EV}(1,s)\big)\Big) \\ &= -\theta_R q_s \pi(0) + \delta \mathrm{EV}(0,s) + \log\Big(1 + \exp\big(-\eta + \theta_R q_s(\pi(0) - \pi(1)) + \delta(\mathrm{EV}(1,s) - \mathrm{EV}(0,s))\big)\Big). \end{aligned}$$

Comparison to static enforcement. Consider a static benchmark where monitoring intensity is constant: $q_1 = q_0$. In this case, $V_2(0) = V_2(1)$, so the dynamic deterrence term vanishes. The period-1 effort threshold collapses to the static condition $\eta \leq \theta_R q_0(\pi(0) - \pi(1))$.

Dynamic enforcement ($q_1 > q_0$) thus induces strictly higher effort in period 1 than static enforcement with the same clean-state intensity q_0. The additional incentive is provided by making a poor record more costly (higher q_1), which makes a clean record more valuable ($V_2(0) - V_2(1) > 0$).

5.2.1 Equilibrium Properties

To evaluate enforcement policies, we define three measures that capture key aspects of equilibrium properties. Let $Q_t = \mathbb{E}[q_{s_t}]$ denote average monitoring intensity, $E_t = \mathbb{E}[a_t]$ denote average effort, and $D_t = \mathbb{E}[d_t]$ denote average deficiency rate in period t, where expectations are taken with initial state $s_1 = 0$.

Although we do not specify the regulator's objective function, it is useful to consider these quantities to understand the trade-off between effort incentives and enforcement intensity. In particular, we suppose that the regulator would like to induce a higher average effort while keeping the average enforcement intensity the same. We now illustrate that the dynamic enforcement policy can achieve this goal.

Period 1. In the first period, monitoring intensity is simply $Q_1 = q_0$, Average effort is $E_1 = p_1(1 \mid 0)$, and the deficiency rate is

$$D_1 = p_1(1 \mid 0)r(1, 0) + p_1(0 \mid 0)r(0, 0).$$

Period 2. In period 2, outcomes depend on the realized state s_2, which equals the period-1 deficiency outcome: $\Pr(s_2 = 1 \mid s_1 = 0) = D_1$. Average monitoring intensity is

$$Q_2 = \Pr(s_2 = 0 \mid s_1 = 0)q_0 + \Pr(s_2 = 1 \mid s_1 = 0)q_1.$$

Average effort is

$$E_2 = \Pr(s_2 = 0 \mid s_1 = 0)p_2(1 \mid 0) + \Pr(s_2 = 1 \mid s_1 = 0)p_2(1 \mid 1).$$

The deficiency rate is

$$\begin{aligned} D_2 &= \Pr(s_2 = 0 \mid s_1 = 0) \Pr(d_2 = 1 \mid s_2 = 0) + \Pr(s_2 = 1 \mid s_1 = 0) \Pr(d_2 = 1 \mid s_2 = 1) \\ &= \Pr(s_2 = 0 \mid s_1 = 0)[p_2(1 \mid 0)r(1, 0) + p_2(0 \mid 0)r(0, 0)] \\ &\quad + \Pr(s_2 = 1 \mid s_1 = 0)[p_2(1 \mid 1)r(1, 1) + p_2(0 \mid 1)r(0, 1)]. \end{aligned}$$

These expressions highlight a key feature of dynamic enforcement. Even though $q_1 > q_0$ raises intensity for auditors in the risky state, the anticipation of this punishment induces higher effort in period 1, reducing D_1 and thus the fraction of auditors who face the higher intensity q_1 in period 2. The net effect on average intensity Q_2 is ambiguous, but effort E_1 unambiguously increases. We quantify these tradeoffs numerically in Section 6.5.

6 Estimation

We start by summarizing the key assumptions maintained throughout estimation.

6.1 Assumptions

1. **Committed enforcement policy.** The regulator commits to a stationary policy (q_0, q_1) that depends only on the public state s. True effort a_t is unobservable. While the regulator may observe proxies for effort (such as audit hours), these signals are noisy and not contracted upon.

2. **Deficiency technology.** High effort (weakly) reduces the probability of deficiency: $r(a, s) = q_s \times \pi(a)$ and $\pi(1) \leq \pi(0)$ with $\pi(0) > 0$.

3. **Inspection timing.** The auditor is inspected in each period t. For large audit firms inspected annually, t corresponds to a calendar year. For small audit firms inspected triennially, t represents one inspection cycle (three years).

The first assumption fixes the enforcement environment. We model the regulator as committing to a stationary, state-dependent policy q_s that is publicly understood. In our institutional setting, this corresponds to the generally accepted auditing standards (GAAS) and the publicly available inspection policies set by the PCAOB. This assumption pins down the auditor's dynamic problem and keeps it time-invariant. If major changes in PCAOB policy occurred during the sample period, we would split the sample accordingly to maintain the stationarity assumption.

The second assumption defines the mapping from effort to reported deficiencies.[11] Even though the regulator does not observe true effort a_t, effort affects the underlying deficiency probability $\pi(a)$, which in turn affects the distribution of outcomes the regulator observes. The decomposition $r(a, s) = q_s \times \pi(a)$ separates monitoring intensity (regulator policy) from audit quality (auditor behavior). This is the key restriction that allows us to map reported deficiencies in the data to objects that enter the auditor's incentives.

The third assumption aligns the model's time index with the data. Since our unit of observation is the audit firm, we index time by inspection opportunities.

Information Sets Let $A_t \in \{0, 1\}$ denote a noisy signal of effort constructed from observable audit inputs (e.g., audit hours H_t), where $\Pr(A_t = a_t) = \rho$ for some accuracy parameter $\rho > 0.5$. The information sets are as follows:

- **Auditor:** $\mathcal{I}_t^A = \mathcal{I}_{t-1}^A \cup \{s_t, \varepsilon_t(a)\}$, where $\varepsilon_t(a)$ is the idiosyncratic cost shock for effort a in period t. The auditor also knows the monitoring policy (q_s), the transition rule, and the distribution of the cost shocks as common knowledge primitives.

- **Regulator:** $\mathcal{I}_t^R = \mathcal{I}_{t-1}^R \cup \{s_t, A_t\}$. The regulator observes the public state and the effort signal, but neither the true effort a_t nor the idiosyncratic shocks $\varepsilon_t(a)$. Importantly, the regulator's enforcement policy conditions only on s_t.[12]

[11] Auditors may learn from inspection results and adjust their procedures. We do not explicitly model this learning channel. To the extent that learning correlates with PCAOB inspection outcomes, our estimates capture part of this effect. Similarly, we do not model within-firm spillovers across engagement teams, though such effects would be partially captured by our audit-firm-level estimation.

[12] Although the regulator can observe audit inputs, PCAOB monitoring intensity is primarily determined by past inspection outcomes and risk-based factors. Our model captures this by having the enforcement policy condition on the public state s_t, which reflects prior inspection results, rather than on current-period input signals. Moreover, as discussed above, observable inputs do not reveal whether effort is appropriate given engagement characteristics, which limits their usefulness for targeting inspections.

- **Econometrician:** $\mathcal{I}_t^E = \mathcal{I}_{t-1}^E \cup \{s_t, d_t, A_t, H_t\}$, where H_t denotes the audit inputs (e.g., audit hours) from which the noisy signal A_t is constructed. The econometrician observes the public state, the deficiency outcomes, the effort signal, and the underlying audit inputs.

6.2 Effort and Observable Signals

6.2.1 Recovering Structural Objects from Observable Data

True effort a_t is unobservable, but the econometrician observes a noisy signal A_t constructed from audit inputs. This subsection shows how to recover the structural objects that depend on a_t, the conditional choice probabilities $p(a \mid s)$, and deficiency probabilities $r(a, s)$, from observable quantities that depend on A_t.

Effort Observability Observable audit inputs, such as hours, imperfectly reveal true audit quality. Moreover, even if effort were perfectly observed, determining whether it was *appropriate* requires knowing engagement-specific characteristics. The same hours may suffice for a simple engagement but not a complex one. PCAOB inspections address this by evaluating whether audit procedures were sufficient given engagement requirements, not merely whether inputs were expended. Our model abstracts from engagement-level heterogeneity but captures the key feature that inspection outcomes d_t reveal quality information that inputs alone cannot provide.

Signal Structure We model A_t as a binary signal of true effort with precision $\rho \in (1/2, 1)$:

$$\Pr(A_t = a_t \mid a_t, s_t) = \rho, \qquad \Pr(A_t \neq a_t \mid a_t, s_t) = 1 - \rho.$$

Note that the signal A_t plays no role in equilibrium, because payoffs and strategies depend only on a_t and not on A_t. However, A_t is important for estimation, because it is what we observe. We now discuss how to recover structural objects from this noisy signal.

Recovering Choice Probabilities The conditional choice probability $p(a \mid s)$ cannot be directly estimated because we do not observe true effort. However, we can recover it from the observable signal frequency. Let $m_s := \Pr(A_t = 1 \mid s_t = s)$ denote the frequency of $A_t = 1$ in state s. By the law of iterated expectations,

$$m_s = \sum_{a_t} \Pr(A_t = 1 \mid a_t, s) \Pr(a_t \mid s) = \rho\, p(1 \mid s) + (1-\rho)\, p(0 \mid s).$$

Solving for $p(1 \mid s)$ gives

$$p(1 \mid s) = \frac{m_s - (1-\rho)}{2\rho - 1}. \tag{10}$$

This formula links the unobservable probability $p(1 \mid s)$ to the observed signal frequency m_s.

Recovering Deficiency Probabilities To link the unobservable deficiency probabilities $r(a, s)$ to observable quantities, we first derive the posterior probability of true effort given the signal:

$$\Pr(a_t = 1 \mid A_t = 1, s) = \frac{\rho\, p(1 \mid s)}{m_s}, \quad \Pr(a_t = 1 \mid A_t = 0, s) = \frac{(1-\rho)\, p(1 \mid s)}{1 - m_s}. \tag{11}$$

The deficiency probability conditional on the *observed* signal (A, s) is

$$\begin{aligned} R(A, s) &\coloneqq \Pr(d_t = 1 \mid A_t, s) \\ &= \sum_a \underbrace{\Pr(d_t = 1 \mid a, s)}_{=q_s \pi(a)} \Pr(a \mid A_t = A, s) \qquad (12) \\ &= q_s \left[\pi(0) + \big(\pi(1) - \pi(0)\big) \Pr(a_t = 1 \mid A_t = A, s)\right]. \qquad (13) \end{aligned}$$

Observe that

$$\begin{aligned} R(1, s) - R(0, s) &= q_s\big(\pi(1) - \pi(0)\big)(\Pr(a_t = 1 \mid A_t = 1, s) - \Pr(a_t = 1 \mid A_t = 0, s)) \\ &= \big(r(1, s) - r(0, s)\big)(\Pr(a_t = 1 \mid A_t = 1, s) - \Pr(a_t = 1 \mid A_t = 0, s)). \end{aligned}$$

After rearranging, we have

$$r(1, s) - r(0, s) = \frac{R(1, s) - R(0, s)}{p(1 \mid s)\left[\rho/m_s - (1-\rho)/(1 - m_s)\right]}. \tag{14}$$

This expression relates the unobservable object $r(1, s) - r(0, s)$ to the observable object $R(1, s) - R(0, s)$. By substituting (12) into (14), we have

$$r(0, s) = \frac{\rho(1 - m_s)R(0, s) - (1 - \rho)m_s R(1, s)}{\rho(1 - m_s) - (1 - \rho)m_s}. \tag{15}$$

These expressions recover the unobservable deficiency probabilities $r(0, s)$ and $r(1, s)$ from objects that depend only on observable data: $R(A, s)$ and m_s. As a consistency check, note that when $\rho = 1$ (perfect signal), we have $r(a, s) = R(a, s)$.

Nonparametric Estimation We estimate m_s and $R(A, s)$ non-parametrically. Let $\hat{m}_s$ be the sample average of A in state s and $\hat{R}(A, s)$ be the sample average of d given (A, s). Specifically,

$$\hat{m}_s = \frac{\sum_{i,t} \mathbf{1}\{(A_{it}, s_{it}) = (1, s)\}}{\sum_{i,t} \mathbf{1}\{s_{i,t} = s\}}, \quad \hat{R}(A, s) = \frac{\sum_{i,t} \mathbf{1}\{(d_{i,t}, A_{it}, s_{it}) = (1, A, s)\}}{\sum_{i,t} \mathbf{1}\{(A_{it}, s_{it}) = (A, s)\}}, \tag{16}$$

where the summation is over audit firms (i) and periods (t). Hence, we can construct the estimates of the choice probabilities as $\hat{p}(1 \mid s) = \frac{\hat{m}_s - (1-\rho)}{2\rho - 1}$ from (10):

$$\hat{p}(1 \mid s) = \frac{\hat{m}_s - (1 - \rho)}{2\rho - 1}.$$

Similarly, we can construct the estimates of the true deficiency probabilities as $\hat{r}(a, s)$ by substituting $\hat{m}_s$ and $\hat{R}(A, s)$ into (14) and (15).

6.2.2 Constructing the Effort Signal from Data

The model uses a binary signal $A_t \in \{0, 1\}$ that is informative about latent effort a_t. We construct A_t from observed audit inputs such as total audit hours. The goal is to capture the auditor's effort choice rather than predictable variation in workload driven by observable characteristics such as the number and size of client engagements, industry mix, or time trends. For example, an auditor with more clients mechanically logs more hours, and we want to isolate effort beyond what these observables would predict.

Simple Approach Suppose we observe total audit hours H_t for an audit firm in period t. We set

$$A_t = \mathbf{1}\{H_t \geq c\},$$

where c is a cutoff such as the within-period median across audit firms (or a fixed percentile). This approach is transparent but potentially confounds effort with workload because total hours mechanically rise with the number and size of engagements.

Residualized Approach To isolate effort from predictable workload variation, we benchmark observed hours against predicted hours given observable determinants. Let i index audit firms and H_{it} denote total hours in period t. We estimate

$$\log H_{it} = \mu_i + \lambda_t + X_{it}'\beta + u_{it}, \tag{17}$$

where μ_i are audit-firm fixed effects, λ_t are period fixed effects, and X_{it} captures portfolio characteristics (e.g., number of issuer engagements, average client size, and industry mix). We then compute the fitted value $\widehat{\log H}_{it} = \hat{\mu}_i + \hat{\lambda}_t + X_{it}'\hat{\beta}$ and the residual $\hat{u}_{it} = \log H_{it} - \widehat{\log H}_{it}$.

We construct the effort signal from this residual:

$$A_{it} = \mathbf{1}\{\hat{u}_{it} \geq c_u\},$$

where c_u is a percentile cutoff (e.g., the median, $c_u = 0$). This construction classifies a firm as high-effort ($A_{it} = 1$) when it supplies more inputs than predicted given observables. When engagement-level data are available, we apply the same approach at the engagement level and aggregate to the audit-firm level. We also report robustness to alternative cutoffs and to alternative input measures (e.g., fees, team size, or reporting lag).

Signal Precision The signal precision $\rho = \Pr(A_t = a_t)$ governs how informative the observed effort proxy is about true effort. This parameter determines the mapping between the unobservable structural objects and the observables through (14) and (15).

We cannot estimate ρ jointly with (η, θ_R) without additional restrictions, because ρ affects the mapping from observables to structural objects rather than entering the auditor's decision problem directly. Instead, we calibrate ρ and examine sensitivity to this choice. We report results for different values of ρ.

6.3 Parameter Estimation

We estimate the preference parameters (η, θ_R) by solving for the auditor's dynamic problem for each set of candidate parameters (Rust, 1987). This approach proceeds as follows: for any candidate (η, θ_R), we solve the auditor's dynamic choice problem (i.e., deriving the $\mathrm{EV}(a, s)$ given the parameters). This delivers the model-implied conditional choice probabilities (CCPs), $p(a \mid s; \eta, \theta_R)$. We then choose (η, θ_R) to minimize the distance between model-implied and empirical CCPs.

The discount factor and the signal precisions are calibrated. As a starting point, we set $\delta = 0.95$ and $\rho = 0.7$. We then examine how the model's fit varies with δ and ρ and select appropriate values.

Matching the Model-Implied CCPs Taking as given the calibrated parameters (δ, ρ) and the recovered object $\hat{r}(a, s)$ from the previous subsection, we solve for $\mathrm{EV}(a, s)$ as the fixed point of

$$\mathrm{EV}(a, s) = \sum_{s'} \hat{G}(s' \mid a, s) \log \left(\sum_{a' \in \{0,1\}} \exp\left\{ -\eta a' - \theta_R \hat{r}(a', s') + \delta \mathrm{EV}(a', s') \right\} \right), \tag{18}$$

where $\hat{G}(s' \mid a, s)$ is constructed from (2): $\hat{G}(1 \mid a, s) = \hat{r}(a, s)$ and $\hat{G}(0 \mid a, s) = 1 - \hat{r}(a, s)$.

With the expected value function EV, the model-implied CCPs are as follows (Rust, 1987):

$$p(a \mid s; \eta, \theta_R) = \frac{\exp\left\{ -\eta a - \theta_R \hat{r}(a, s) + \delta \mathrm{EV}(a, s) \right\}}{\sum_{a'} \exp\left\{ -\eta a' - \theta_R \hat{r}(a', s) + \delta \mathrm{EV}(a', s) \right\}}. \tag{19}$$

We estimate (η, θ_R) by minimizing the weighted squared distance between the empirical probabilities ($\hat{p}(1 \mid s)$) and the model-implied probabilities:

$$\min_{(\eta,\theta_R)} \sum_{s\in\{0,1\}} W_s \left(\hat{p}(1 \mid s) - p(1 \mid s; \eta, \theta_R)\right)^2, \tag{20}$$

where W_s is the sample share of observations in state s. This weighting ensures that states with more observations receive proportionally greater influence in estimation.[13]

We compute standard errors using the nonparametric bootstrap. Specifically, we resample audit firm-period observations with replacement, re-estimate $(\hat{\eta}, \hat{\theta}_R)$ for each bootstrap sample, and report the standard deviation of the bootstrap distribution as the standard error.

Summary of Estimation Procedure The complete estimation procedure is as follows:

1. Fix calibrated parameters: the discount factor δ and the precision of the effort signal ρ.
2. Construct sample averages $\hat{R}(A, s)$ and $\hat{m}_s$ from the data.
3. Recover empirical CCPs: $\hat{p}(1 \mid s) = \frac{\hat{m}_s-(1-\rho)}{2\rho-1}$.
4. Recover deficiency probabilities $\hat{r}(a, s)$ using (14)–(15).
5. Construct $\hat{G}(s' \mid a, s)$ using (2): $\hat{G}(1 \mid a, s) = \hat{r}(a, s)$ and $\hat{G}(0 \mid a, s) = 1 - \hat{r}(a, s)$.
6. Estimate (η, θ_R) by by solving (20).

Heterogeneity A parsimonious approach to accommodate heterogeneity in preferences is to stratify the sample by audit firm characteristics and estimate parameters separately for each stratum. For example, we partition firms by size (large vs. small) and estimate $(\eta_g, \theta_{R,g})$ for each group g. This allows effort costs and sensitivity to deficiencies to vary across firm types while maintaining tractability.

An alternative approach would be to estimate a single model with group-specific parameters or random coefficients. We adopt stratification because it is more transparent and because we are

[13] Alternative weighting schemes, such as equal weights or optimal minimum distance weights, should yield similar results.

primarily interested in comparing incentive responses across firm types rather than estimating a distribution of heterogeneity.

6.4 Estimation Results and Interpretation

This subsection describes how we interpret the estimated parameters and describes our counterfactual exercises.

Dynamic Deterrence To interpret the estimates, consider the relative benefit of high effort versus low effort in state s:

$$\Delta w(s) = -\eta - \theta_R [r(1, s) - r(0, s)] + \delta [\mathrm{EV}(1, s) - \mathrm{EV}(0, s)]. \tag{21}$$

The first term is the direct cost of effort. The second term is the current-period benefit: high effort reduces deficiency probability, which is valued at θ_R per expected deficiency. The third term is the continuation value benefit: high effort increases the probability of reaching the clean state ($s = 0$) tomorrow.

Under state-dependent enforcement ($q_1 > q_0$), the clean state is more valuable than the risky state: $V(0) > V(1)$. This follows because the auditor faces lower monitoring intensity and hence lower expected disutility in the clean state. Moreover, because high effort increases the probability of transitioning to the clean state, the continuation value of high effort exceeds that of low effort: $\mathrm{EV}(1, s) > \mathrm{EV}(0, s)$. This positive continuation term captures the dynamic deterrence channel: effort today improves tomorrow's regulatory standing.

The value of a clean record, $V(0) - V(1)$, summarizes the strength of dynamic incentives. A larger value indicates that auditors place greater weight on avoiding the risky state, amplifying the deterrence effect of state-dependent enforcement.

Counterfactual Analyses A key advantage of the structural approach is the ability to evaluate policies not observed in the data. We design counterfactual exercises to address two policy-relevant questions. First, how much does dynamic enforcement contribute to deterrence beyond what a constant-intensity policy would achieve? This quantifies the value of the PCAOB's state-dependent monitoring approach. Second, could the regulator improve audit quality by adjusting the enforcement intensity gap between clean and risky states? This sheds light on whether the current policy is optimally calibrated.

We use the estimated model to evaluate alternative enforcement policies. Throughout, we fix the preference parameters $(\hat{\eta}, \hat{\theta_R})$ at their calibrated values and vary the enforcement policy (q_0, q_1).

We consider two counterfactual exercises:

1. **Static enforcement benchmark.** We highlight the value of dynamic enforcement by comparing it to a static counterfactual. In particular, we set $q_0 = q_1 = \bar{q}$, where $\bar{q}$ is chosen so that the long-run average enforcement intensity matches the baseline policy.[14] This benchmark shuts down the dynamic deterrence channel and isolates the role of the continuation-value wedge $V(0) - V(1)$. By analyzing how the average effort and deficiency probability change under the static counterfactual, we can quantify the value of dynamic enforcement.

2. **Alternative dynamic policies.** We vary the enforcement gap $q_1 - q_0$ on a grid and, for each gap, scale (q_0, q_1) so that the long-run average enforcement intensity is held fixed at the baseline level. This analysis helps us understand whether and how much the regulator can improve the average effort and deficiency probability by using an alternative dynamic enforcement policy.

Our counterfactuals vary PCAOB monitoring intensity while holding other features of the environment fixed. In practice, alternative detection channels exist, including SEC enforcement, private litigation, and client-initiated auditor switches. Our estimates capture the marginal effect of PCAOB monitoring conditional on these other mechanisms.

6.5 Numerical Simulations

We illustrate the model's predictions through numerical simulations of the two-period example. These exercises demonstrate how dynamic enforcement can improve outcomes relative to static policies.

[14] Equivalently, since payoffs depend on enforcement only through $r(a, s) = q_s \pi(a)$, we can implement this benchmark by equalizing $r(a, 0)$ and $r(a, 1)$ for each a and scaling to match the same long-run average intensity.

6.5.1 Parameterization

Table 2 reports the baseline parameter values. We set $\eta = 0.5$ and $\theta_R = 2.0$, implying that the disutility from a deficiency is four times the cost of high effort. We use $\delta = 0.95$ as the discount factor, which corresponds to an annual discount rate of approximately 5%, a standard choice in the literature. The returns on efforts are $\pi(0) = 0.6$ and $\pi(1) = 0.2$, so high effort reduces the deficiency probability by two-thirds. These parameters are chosen for illustration.

Parameter	Value	Definition
η	0.5	Cost of high effort
θ_R	2.0	Disutility per expected deficiency
$\pi(0)$	0.6	Deficiency prob. under low effort
$\pi(1)$	0.2	Deficiency prob. under high effort
δ	0.95	Discount factor

Table 2: Baseline Simulation Parameters

6.5.2 Dynamic Enforcement Can Dominate Static Enforcement

A key prediction of the model is that dynamic enforcement can outperform static enforcement: the regulator can induce more effort while using fewer monitoring resources. We demonstrate this by comparing two policies designed to achieve similar long-run monitoring intensity.

- **Static policy:** $q_0 = q_1 = 0.27$ (constant intensity regardless of record)
- **Dynamic policy:** $q_0 = 0.23$, $q_1 = 0.90$ (state-dependent intensity)

Table 3 reports the discounted sums of monitoring intensity ($Q = Q_1 + \delta Q_2$), effort ($E = E_1 + \delta E_2$), and deficiency rates ($D = D_1 + \delta D_2$).

Policy	Q	E	D
Dynamic ($q_0 = 0.23$, $q_1 = 0.90$)	0.511	0.846	0.215
Static ($q_0 = q_1 = 0.27$)	0.526	0.837	0.225
Difference (Dynamic − Static)	−0.015	+0.009	−0.010

Table 3: Comparing Dynamic and Static Enforcement

The dynamic policy dominates on all three dimensions: it requires less monitoring (Q: 0.511 vs. 0.526), induces more effort (E: 0.846 vs. 0.837), and produces fewer deficiencies (D: 0.215 vs. 0.225).

The intuition is as follows. Under dynamic enforcement, the regulator inspects *less* intensively in period 1 ($q_0 = 0.23$ versus 0.27), yet the auditor exerts *more* effort ($E_1 = 0.433$ versus 0.429). This occurs because auditors anticipate harsher future enforcement if the regulator finds a deficiency in today's inspection. Put simply, the threat of future punishment substitutes for current monitoring.

6.5.3 Policy Tradeoffs: Contour Plots

To understand the policy design tradeoffs, Figure 1 presents contour plots of Q, E, and D over the (q_0, q_1) plane. Each contour traces policy combinations of (q_0, q_1) that yield the same outcome.

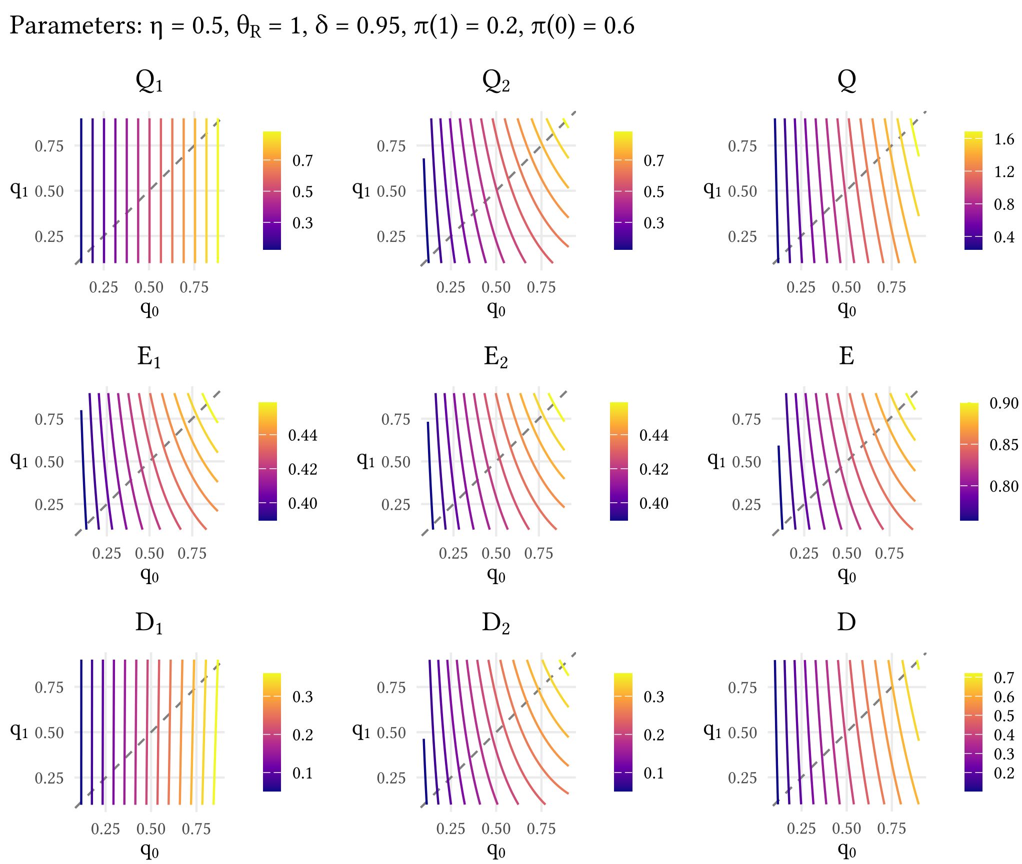

Figure 1: Contour Plots of Equilibrium Outcomes

The dashed diagonal marks $q_1 = q_0$ (static enforcement). Points above the diagonal correspond to dynamic enforcement ($q_1 > q_0$).

The diagonal line $q_1 = q_0$ represents a static enforcement benchmark where the regulator applies the same intensity regardless of past performance. Points above the diagonal are dynamic policies ($q_1 > q_0$), where the regulator monitors risky firms more intensively.

At any point on the diagonal, the marginal rate of substitution between q_0 and q_1 is greater than 1 (i.e., $|\partial q_1/\partial q_0| > 1$). This means that holding Q constant, a small decrease in q_0 requires a large increase in q_1. The same pattern holds for E and D.

To see why, note that the clean-record state is more likely on the equilibrium path. Starting from $s_1 = 0$, the auditor is always in the clean state in period 1 and remains there with probability $1 - D_1$ in period 2. Consequently, q_0 governs the bulk of monitoring activity. A change in q_0 affects all of period 1 and most of period 2, whereas a change in q_1 affects only the fraction D_1 of auditors who transition to the risky state. When D_1 is small (as it typically is under reasonable effort levels), q_1 receives little direct weight in the aggregate outcomes.

Regulator's Objectives and Welfare Our proposed analysis uses effort (E) and deficiency rates (D) as outcome measures. We assume that higher effort and lower deficiencies are desirable. This is consistent with the view that the regulator's monitoring activities are costly (Polinsky and Shavell, 2000). However, we deliberately refrain from constructing a formal welfare function for two reasons. First, doing so would require specifying the social value of audit quality relative to its cost, a difficult empirical question beyond our scope. Second, as Dye (1993) emphasizes, increased auditor effort is not unambiguously welfare improving if the marginal cost to auditors (and potentially their clients, through higher fees) exceeds the marginal social benefit in terms of improved financial reporting quality.

Our counterfactual results should therefore be interpreted as characterizing the regulator's *ability* to influence auditor behavior through policy design, not as welfare rankings of alternative policies. A complete welfare analysis would require modeling the demand side of the audit market, the pass-through of compliance costs to clients, and the downstream effects of audit quality on investor decisions. We leave these important extensions for future work.

6.6 Limitations

We note several simplifications. First, we treat audit fees and client behavior as exogenous; endogenizing these would require modeling auditor-client bargaining and audit market structure. Second, we assume the signal precision ρ is constant across firms and states; in practice, the informativeness of audit hours may vary. Third, our two-state specification captures the key dynamic incentive but abstracts from richer state spaces that might better reflect the PCAOB's internal risk classification. Fourth, we do not model auditor entry, exit, or market structure effects. These extensions represent promising directions for future research. Fifth, our model assumes the state affects payoffs only through monitoring intensity. If a clean regulatory record also directly improves business outcomes (e.g., client retention), our estimates of dynamic deterrence are conservative.

7 Conclusion

This paper proposes a structural approach to evaluate the PCAOB's dynamic enforcement policy. We develop a dynamic model of auditor behavior in which the regulator commits to state-dependent monitoring intensity, and auditors make forward-looking effort choices that account for both current-period inspection risk and the future consequences of a poor regulatory record.

Our proposed analysis makes three contributions. First, we quantify the value of dynamic enforcement relative to static monitoring by estimating how much additional deterrence state-dependent policies provide. Second, we recover structural preference parameters (the cost of effort and the disutility from deficiencies) that have a direct economic interpretation. Third, we decompose observed deficiency rates into regulatory policy and auditor behavior, enabling counterfactual evaluation of alternative enforcement regimes.

The structural approach complements existing reduced-form evidence. While prior studies document that PCAOB oversight affects auditor behavior, reduced-form methods cannot distinguish whether auditors respond to current-period inspection risk or to future consequences, nor can they predict responses to policies not yet implemented. Our framework addresses both limitations, enabling quantitative evaluation of alternative monitoring regimes. We acknowledge, however, that our approach also has limitations. In particular, we treat audit fees and client behavior as exogenous and abstract from auditor entry and exit. These simplifications isolate the dynamic deterrence mechanism that is our primary focus. Despite these limitations, our analysis can inform regulatory design by quantifying the incentive effects of state-dependent enforcement.

References

Aguirregabiria, V. and P. Mira (2010). Dynamic discrete choice structural models: A survey. Journal of Econometrics 156(1): 38–67.

Aobdia, D. (2018). The impact of the pcaob individual engagement inspection process—preliminary evidence. The Accounting Review 93(4): 53–80.

Aobdia, D. (2020). The economic consequences of audit firms' quality control system deficiencies. Management Science 66(7): 2883–2905.

Aobdia, D. and N. Shroff (2017). Regulatory oversight and auditor market share. Journal of Accounting and Economics 63(2-3): 262–287.

Becker, G. S. (1968, March). Crime and Punishment: An Economic Approach. Journal of Political Economy 76(2): 169–217.

Bertomeu, J., Y. Liang, I. Marinovic, et al. (2023). A primer on structural estimation in accounting research. Foundations and Trends® in Accounting 18(1–2): 1–137.

Blundell, W., G. Gowrisankaran, and A. Langer (2020). Escalation of scrutiny: The gains from dynamic enforcement of environmental regulations. American Economic Review 110(8): 2558–2585. Publisher: American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

Boone, J. P., I. K. Khurana, and K. Raman (2015). Did the 2007 pcaob disciplinary order against deloitte impose actual costs on the firm or improve its audit quality? The Accounting Review 90(2): 405–441.

Breuer, M., E. Labro, H. Sapra, and A. A. Zakolyukina (2024). Bridging theory and empirical research in accounting. Journal of Accounting Research 62(3): 1121–1139.

CAQ (2021). Guide to pcaob inspections. https://www.thecaq.org/wp-content/uploads/2021/03/caq_guide-to-PCAOB-inspections_2021-03.pdf. 2021.

Cheynel, E., L. Y. Liu, L. Tong, and F. Zhou (2025). Estimating the value of auditing services for private firms. Available at SSRN 5366462.

Cheynel, E. and F. S. Zhou (2024). Auditor tenure and misreporting: Evidence from a dynamic oligopoly game. Management Science 70(8): 5557–5585.

Christensen, B. E., N. G. Lundstrom, and N. J. Newton (2021). Does the disclosure of pcaob inspection findings increase audit firms' litigation exposure? The Accounting Review 96(3): 191–219.

Christensen, H. B., L. Y. Liu, and M. Maffett (2020). Proactive financial reporting enforcement and shareholder wealth. Journal of Accounting and Economics 69(2-3): 101267.

DeFond, M. L. and C. S. Lennox (2011). The effect of sox on small auditor exits and audit quality. Journal of Accounting and Economics 52(1): 21–40.

DeFond, M. L. and C. S. Lennox (2017). Do pcaob inspections improve the quality of internal control audits? Journal of Accounting Research 55(3): 591–627.

Dye, R. A. (1993). Auditing standards, legal liability, and auditor wealth. Journal of political Economy 101(5): 887–914.

Fung, S. Y. K., K. Raman, and X. K. Zhu (2017). Does the pcaob international inspection program improve audit quality for non-us-listed foreign clients? Journal of Accounting and Economics 64(1): 15–36.

Hanlon, M. and N. Shroff (2022). Insights into auditor public oversight boards: Whether, how, and why they "work". Journal of Accounting and Economics 74(1): 101497.

Kang, K. and B. S. Silveira (2021). Understanding disparities in punishment: Regulator preferences and expertise. Journal of Political Economy 129(10): 2947–2992.

Lamoreaux, P. T. (2016). Does pcaob inspection access improve audit quality? an examination of foreign firms listed in the united states. Journal of Accounting and Economics 61(2-3): 313–337.

Lamoreaux, P. T., M. J. Mowchan, and W. Zhang (2022). Does pcaob regulatory enforcement deter low quality audits? Available at SSRN 4166137.

Lennox, C. and J. Pittman (2010). Auditing the auditors: Evidence on the recent reforms to the external monitoring of audit firms. Journal of accounting and economics 49(1-2): 84–103.

Polinsky, A. M. and S. Shavell (1998). On offense history and the theory of deterrence. International Review of Law and Economics 18(3): 305–324.

Polinsky, A. M. and S. Shavell (2000). The economic theory of public enforcement of law. Journal of economic literature 38(1): 45–76. Publisher: American Economic Association.

Rust, J. (1987). Optimal replacement of gmc bus engines: An empirical model of harold zurcher. Econometrica: Journal of the Econometric Society: 999–1033.

Shroff, N. (2020). Real effects of pcaob international inspections. The Accounting Review 95(5): 399–433.

A Additional Institutional Details

This appendix provides additional details on the PCAOB's quality control evaluation and enforcement program, supplementing the institutional background in Section 2.

A.1 PCAOB Quality Control Evaluation

After identifying potential audit-specific deficiencies, the PCAOB evaluates the firm as a whole to assess its quality control system. The PCAOB maintains separate quality control standards that mandate policies and procedures covering several areas.[15] These include

- The firm's management structure and processes, including the tone at the top;
- Practices for partner management;
- Policies and procedures for client acceptance and retention;
- Internal inspection programs and the firm's response to deficiencies in audit quality; and
- Independence policies and procedures.

Quality control deficiencies are handled differently from engagement-specific deficiencies. When an inspection report includes quality control criticisms, these findings are not initially included in the public copy of the report. Instead, firms are given 12 months after the issuance of the final inspection report to demonstrate that they have resolved any issues. If the auditor successfully addresses quality control criticisms within this remediation period, the findings remain confidential. Otherwise, the PCAOB publicly releases these criticisms as "Part II findings."

A.2 PCAOB Enforcement Program

In addition to its inspection function, the PCAOB operates an enforcement program. In the event of serious violations of standards, laws, or rules beyond routine deficiencies, the PCAOB will investigate firms and individuals. The PCAOB's stated monitoring policy is to prioritize enforcement that "address[es] those issues that pose the greatest risk to investors and are most likely to deter improper conduct."[16]

[15] The quality control standards can be accessed on the PCAOB's website: https://pcaobus.org/oversight/standards/qc-standards.

[16] See https://pcaobus.org/oversight/enforcement.

If violations are found, the PCAOB may impose a range of sanctions, including

- Censures,
- Monetary penalties, and
- Limitations on a firm's or an individual's ability to audit public companies or broker-dealers.

Our model focuses on the inspection process and Part I findings rather than enforcement actions. While enforcement represents an important component of the PCAOB's oversight, the inspection regime affects a broader set of auditors and engagements. Thus, it is the primary channel through which most auditors experience regulatory monitoring.