Operations Research

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Audit and Remediation Strategies in the Presence of Evasion Capabilities

Shouqiang Wang, Francis de Véricourt, Peng Sun

To cite this article:
Shouqiang Wang, Francis de Véricourt, Peng Sun (2024) Audit and Remediation Strategies in the Presence of Evasion Capabilities. Operations Research 72(5):1843-1860. https://doi.org/10.1287/opre.2022.0289

informs
https://pubsonline.informs.org/journal/opre

OPERATIONS RESEARCH
Vol. 72, No. 5, September–October 2024, pp. 1843–1860
ISSN 0030-364X (print), ISSN 1526-5463 (online)

Contextual Areas

Audit and Remediation Strategies in the Presence of Evasion Capabilities

Shouqiang Wang,[a,*] Francis de Véricourt,[b] Peng Sun[c]

[a] Naveen Jindal School of Management, The University of Texas at Dallas, Richardson, Texas 75080; [b] ESMT European School of Management and Technology, D-10178 Berlin, Germany; [c] Fuqua School of Business, Duke University, Durham, North Carolina 27708

*Corresponding author

Contact: shouqiang.wang@utdallas.edu, https://orcid.org/0000-0002-0988-7010 (SW); francis.devericourt@esmt.org, https://orcid.org/0000-0003-2832-0738 (FdV); psun@duke.edu, https://orcid.org/0000-0002-3892-5936 (PS)

Received: June 5, 2022
Revised: October 2, 2023; February 23, 2024
Accepted: May 13, 2024
Published Online in Articles in Advance: July 10, 2024

Area of Review: Operations and Supply Chains

https://doi.org/10.1287/opre.2022.0289

Abstract. In this paper, we explore how to uncover an adverse issue that may occur in organizations with the capability to evade detection. To that end, we formalize the problem of designing efficient auditing and remedial strategies as a dynamic mechanism design model. In this setup, a principal seeks to uncover and remedy an issue that occurs to an agent at a random point in time and that harms the principal if not addressed promptly. Only the agent observes the issue's occurrence, but the principal may uncover it by auditing the agent at a cost. The agent, however, can exert effort to reduce the audit's effectiveness in discovering the issue. We first establish that this setup reduces to the optimal stochastic control of a piecewise deterministic Markov process. The analysis of this process reveals that the principal should implement a dynamic cyclic auditing and remedial cost-sharing mechanism, which we characterize in closed form. Importantly, we find that the principal should randomly audit the agent unless the agent's evasion capacity is not very effective, and the agent cannot afford to self-correct the issue. In this latter case, the principal should follow predetermined audit schedules.

Funding: This work was supported by the Deutsche Forschungsgemeinschaft (German Research Foundation) ["Audit Schedules in the Presence of Concealing Effort"; Grant 387250733].

Supplemental Material: The computer code and data that supports the findings of this study and the online appendix are available within this article's supplemental material at https://doi.org/10.1287/opre.2022.0289.

Keywords: dynamic mechanism design • stochastic optimal control • asymmetric information • moral hazard • voluntary disclosure • environmental issue • corporate responsibility

1. Introduction

Organizations sometimes discover an issue that, if not addressed promptly, harms other parties. Rather than fixing the problem, however, these organizations may prefer to conceal the issue and exert effort to evade detection. For instance, when technology executives at the auto manufacturer Volkswagen discovered that their newly designed engine would not comply with the U.S. Clean Air Act, they chose to develop a sophisticated software that evaded emission tests rather than investing in the creation of more expensive but effective emission equipment (Ewing 2016). Similarly, the Quaker Pet Group, one of Walmart's largest suppliers, devised workarounds when the firm realized that one of its own suppliers would not pass Walmart's workplace inspections. Instead of switching to a more expensive but complying supplier, the Quaker Pet Group falsified its order forms, wrongly claiming that it sourced from a Walmart-certified facility (Clifford and Greenhouse 2013). Efforts to evade detection instead of fixing problems are also present in the biotechnology industry (Carreyrou 2018).

Audits serve as a common tool used by firms and regulators to identify and address adverse issues.[1] Although audits can be costly, regulators and businesses often use incentive mechanisms to promote compliance. An example of this is the U.S. Environmental Protection Agency's (EPA) audit policy (EPA 2000), which grants disclosure benefits, typically in the form of reduced penalties, to companies that voluntarily report hazardous incidents. Another notable case involves Levi's, an American blue jeans retailer. Hadler (2013) states that if Levi's uncovers a supplier's failure to meet their requirements and attempts to conceal it, Levi's terminates the contract. However, if a supplier proactively discloses information about a problem, Levi's collaborates with them to find a solution.

Designing an effective audit policy incorporating these incentive mechanisms, however, remains a significant challenge and an open research question. In this paper, we provide managerial insights on how an organization can uncover and remedy an adverse issue that may randomly occur in another entity that can evade detection. To that end, we formalize the problem of designing efficient auditing and remedial strategies as a dynamic principal-agent problem in continuous time, in which the principal seeks to discover an adverse issue. In this setup, the principal uses audits and disclosure benefits to enforce compliance. Importantly, the principal optimizes over a large set of policies, which includes all implementable audit schedules one can reasonably think of.

This formulation reduces to the optimal stochastic control of a piecewise deterministic Markov process (PDP), a class of stochastic processes that generalizes semi-Markov decision processes. We study the control of this PDP analytically, which yields new insights concerning auditing in the presence of evasion capabilities.

Overall, our analysis suggests that the principal should incentivize the agent to always disclose the issue as soon as it occurs, without taking evasive actions or self-correct. For this purpose, the principal audits the agent *randomly* and *periodically* but offers to cover part of the agent's remedial cost when the agent voluntarily reports the issue. In essence, the policy ultimately motivates the agent to come clean. Furthermore, we find that the audits become deterministic when the agent's evasion capacity is not too effective (i.e., it is imperfect and sufficiently costly) *and* the agent cannot afford to self-correct the issue. In this case, the audit follows a *predetermined* schedule. In this sense, the agent's evasion capability affects the very nature of the auditing policy.

To be more specific, we consider a generic situation, in which a principal (e.g., a firm or an employer) seeks to discover and avoid the negative consequences of an adverse event (e.g., a quality or noncompliance issue), the occurrence of which is private information to the agent (e.g., a supplier, or a business unit). A key feature of our setup is that the agent can exert effort to evade the principal's audits (e.g., by falsifying forms, developing circumventing software, establishing hidden accounts, or even self-correct without disclosure). However, evasion may not be perfect, in that an audit may still reveal the issue with positive probability after the agent has exerted evasive efforts. The lower the cost of evasive effort and the resulting detection probability are, the more effective the agent's evasion capability becomes.

When the agent discovers the issue, the agent prefers to conceal the problem from the principal rather than to incur the associated remedial costs for which the agent is liable. This, however, obstructs the timely correction of the problem, which, in turn, harms the principal. To uncover whether such an adverse event has occurred, the principal can decide at any time (possibly randomly) to audit the firm and charge a nondisclosure penalty if the adverse event is detected. Because audits are expensive and the agent can evade them, the principal may also offer to cover part of the remedial costs if the agent voluntarily reports the issue. In our setup, the agent may also decide to self-correct the issue, possibly at a later time, without notifying the principal, in an attempt to avoid potential penalty or to defer the remedial cost.

The principal's objective is to identify an audit schedule and her contribution to the remedial cost that minimizes her total discounted cost. This cost includes the principal's share of the remedial costs, the audit costs and possible damages resulting from failing to address the issue promptly.

This dynamic agency setting involves not only an adverse selection problem due to the agent's private information on the timing of the adverse event, but also a moral hazard problem due to the agent's ability to evade audits or self-correct the issue. Furthermore, audit schedules can take very general forms, as interaudit times can be history dependent and stochastic, following general probability distributions. As such, possible audit times may follow any deterministic schedule, random audits at deterministic times, random time between audits that follow any (well-behaved) probability distributions, or any combinations of these.

We first establish a version of the revelation principle tailored to our dynamic setting, which states that inducing the agent to reveal the adverse event as soon as it occurs is always optimal for the principal. This means that the principal can restrict the search for the optimal strategy to those that remedy the issue without delay. This result allows us to reformulate the problem as an optimal stochastic control problem.

Given this, we first examine situations where the evasion technology is able to render the principal's audits completely ineffective; that is, the evasion capability is perfect. In this case, the problem reduces to the optimal stochastic control of a one-dimensional PDP, which is known to be hard (Davis et al. 1987). Nonetheless, we show that the optimal policy is a cyclic cost sharing and *random* auditing policy. Under this policy, the principal shares part of the remedial costs with the agent, the exact amount of which depends on the timing of the agent's disclosure. The principal adjusts the split of the remedial cost over time, following a cyclic pattern. In the beginning of a cycle, the agent's contribution of the remedial cost increases over time. If the agent's contribution reaches a maximum level (equal to the evasion cost) before any disclosure, the principal runs an audit after an exponentially distributed random time. If the audit does not reveal any issue, the cycle ends, and the agent's contribution is reset to its minimum value to start a new cycle. Overall, the optimal policy alternates between deterministically changing payments and random audits.

We then study situations where the evasion technology is imperfect, so that an audit still reveals the issue with a positive detection probability even after the agent has taken an evasive action. In this case, we show that the optimal policy remains cyclic, but is either *random* or *deterministic*, depending on whether the agent can afford to correct the issue alone. In particular, if the remedial cost is within the agent's limited liability, the optimal policy maintains the same random cyclic structure observed in the case with perfect audits. If the remedial cost is higher than the agent's limited liability, however, the principal sometimes follows a simple deterministic cyclic audit schedule. This happens when the evasion cost is sufficiently high (i.e., higher than a specific threshold, which is decreasing in the postevasion detection probability of audits). In this case, the policy adheres to a similar cyclic pattern, except that the principal runs an audit as soon as the agent's share of the remedial cost reaches its maximum level (the agent's limited liability), rather than after a random period.

Finally, we study the case in which the evasion technology is imperfect but the evasion costs are below the aforementioned threshold. This creates mixed incentives for the agents since the agent's evasive action is imperfect but also inexpensive. In this case, the problem becomes the control of a two-dimensional of PDP governed by two sets of incentive compatible and state contraints. Optimally control this problem is generally intractable not only analytically but also numerically (Chehrazi et al. 2019; also see Remark A.1 in Online Appendix A for more details). In addition, and perhaps more importantly, even if one is able to compute the optimal policy, it may be too complex to implement in practice (see again Remark A.1 in Online Appendix A). We thus restrict the search for efficient auditing policies within a large class of tractable policies, which we refer to as *proportional policies*.

Under a proportional policy, the agent's payment on self-reporting an issue is proportional to the expected penalty of getting caught. In particular, the previous optimal random and cyclic policies belong to this class. Proportional policies, however, do not need to be cyclic or deterministic; they may allow for random audits, a mixture of random and deterministic audits, random audits with different or time-varying rates, and different time interval lengths between audits.

We then show that the optimal proportional policy is a cyclic *random* policy akin to the perfect evasion case, except that the initial cycle is of a different length than the following ones. Furthermore, the parameters of this policy can be obtained by solving a deterministic bivariate constrained optimization problem.

These analytical results allow us to numerically explore the impact of the agent's evasion capacity on the principal's audit policy. Our study reveals a nonmonotone relationship between the evasion detection probability and audit frequency. Specifically, we find that the principal should audit first more and then less frequently as the agent's evasion capability becomes less effective. The total expected audit cost exhibits a similar structure.

Finally, we show that the structure of our policy continues to hold when the principal can inflict different penalties depending on whether the agent does not disclose the issue with or without evasive actions (Section 9.1), a third party, such as non-for-profit organization, can independently uncover the agent's violation (Section 9.2), the principal maximizes social welfare (Section 9.3), and the agent may be uninformed about the event's occurrence (Section 9.4).

2. Literature Review

Stochastic modeling of audits/inspections dates back to the reliability theory literature (Parmigiani 1993, Barlow and Proschan 1996), which mostly focuses on a single decision maker's inspection policy to discover system breakdowns. Extending this framework to a game-theoretical setting, Kim (2015) examines two types of inspection schedules (i.e., a deterministic periodic schedule and an exponential random schedule) to incentivize voluntary disclosure. Wang et al. (2016) adopt a mechanism design framework with costly state verification and show that a deterministic inspection schedule is optimal when used together with subsidies that are decreasing over time between inspections. A key component that distinguishes our paper is that we explicitly account for the agent's opportunistic behavior of evading the principal's audits (i.e., moral hazard). This evasion capability yields fundamentally different results. In particular, we show that the presence of a moral hazard problem can render the previous deterministic schedule suboptimal. Instead, it is sometimes optimal for the principal to alternate between a fixed period with no audit followed by a random period with audit.

More recently, Varas et al. (2020) study how a principal inspects an agent whose production quality follows a two state Markov chain. Following Board and Meyer-Ter-Vehn (2013), the agent's effort increases the transition rate from low-quality state to high-quality state and reduces the transition rate from high-quality state to low-quality state. The principal and the market form a belief about the quality state over time, which captures the firm's reputation. The firm's payoff is linear in its reputation, and the principal's payoff is convex in the firm's reputation. The principal controls when to conduct costly inspection to fully review the firm's quality state of that moment. Interestingly, their optimal inspection schedule shares a similar structure as ours; each cycle starts with a fixed period of time with no inspection, followed by an exponentially distributed random time before an inspection occurs and finishes the cycle. However, the model and analysis of Varas et al. (2020) is quite different from ours. In particular, with no payment, their model does not rely on the promised utility framework, which is the foundation of our model. Instead, the belief probability is the only state variable in their model.

More generally, there has been an emerging literature in management science that combines incentive management with detections. Bakshi and Gans (2010), for example, study incentive programs that induce firms to improve security, and hence reduce inspection costs, against potential terrorist attacks in cargo shipment. Babich and Tang (2012) compare deferred payment and inspection mechanisms to address the moral hazard issue of "corner cutting" behaviors by suppliers. Hwang et al. (2006) study similar problems but compare the inspection versus certification mechanisms. Cho et al. (2015) study inspection and penalization strategies to combat child labor rather than to ensure product quality. The aforementioned models are generally static in nature and therefore do not address the timing of adverse events as we do. They also do not consider the ability to render inspections ineffective. Levi et al. (2019) study farmers' strategic adulteration behavior in response to quality uncertainty, supply chain dispersion, traceability, and testing sensitivity. Chen et al. (2020) study inspection policies for supply networks with different centrality measures. More recently, Kim and Xu (2023) propose a class of policies that randomize between deterministic and exponential audits to mitigate financial risks and optimize over the policy parameters for given policy structures. An interesting work by Baliga and Ely (2016) examines the use of torture as a means of extracting information from a possibly informed agent who knows the timing of a future attack. With the principal's full commitment, their problem becomes a standard static mechanism design, whereas our principal faces a dynamic adverse selection problem, because our agent knows the timing of the event only when it has happened.

In this stream of research, the only paper that considers deliberate audit evasion is Plambeck and Taylor (2016). One of their key insights is that too high a violation penalty may backfire by creating an incentive for the agent to actively evade the audit, which was first revealed in the economics literature of auditing in the presence of avoidance (Malik 1990). In our setup, the logic of Malik (1990) explains why the principal never requires the agent to incur a cost higher than the effort cost of evading audits. This upper limit on the agent's contribution toward the remedial costs turns out to be the main driver for the optimality of our random audit schedule.

Overall, our work contributes to the longstanding research on the economics of law enforcement initiated by Becker (1968) (see Polinsky and Shavell (2000) for a review). Central to this area of study is the inquiry into the most efficient approach to minimize societal costs while ensuring compliance. What sets our work apart is its focus on a *dynamic* principal-agent framework, which incorporates both *adverse selection* and *moral hazard*.

In a static principal-agent model, Townsend (1979) initiated the paradigm of costly state verification, which, however, is restricted to only deterministic audits. Later, Mookherjee and Png (1989) generalized the analysis to allow random audits and provide conditions for random audits to be optimal. We consider a dynamic setting, which is closer to Ravikumar and Zhang (2012), who examine a tax auditing problem albeit without audit evasion behaviors.

In an information environment with costly state verification *or* multiple periods, Townsend (1988) pointed out that the usual version of the revelation principle (Dasgupta et al. 1979, Myerson 1979) is no longer automatically applicable. He extended the revelation principle separately to these two environments. We contribute to this literature by establishing a version of the revelation principle applicable to a private information environment (adverse selection) with hidden action (moral hazard), and costly state verification.

From a more technical perspective, we leverage existing recursive representation techniques (Spear and Srivastava 1987, Abreu et al. 1990, Ljungqvist and Sargent 2004) to tackle dynamic principal-agent problems (Sannikov 2008, Biais et al. 2010, Li et al. 2013). This approach helps reduce our original principal-agent problem to a stochastic optimal control of a *piecewise deterministic process* (PDP). Optimal control of PDPs, however, are often analytically intractable (Davis et al. 1987). We attack this problem using the verification approach via *quasi-variational inequalities* (Bensoussan and Lions 1984) and obtain a closed-form characterization of the optimal policy.

3. Model

Consider a principal-agent relationship in continuous time. The principal seeks to discover and avoid negative consequences of an adverse issue that occurs at and is privately known to the agent. In the context of environmental regulation, we can conceptualize the EPA as the principal, whereas a firm like Volkswagen assumes the role of the agent. Similarly, when considering supplier compliance matters in the private sector, we can envision influential retailers such as Walmart or Levi's as the principal, with suppliers like the Quaker Pet Group acting as the agent. The adverse issue emerges and comes to the agent's awareness at a random time T, which follows an exponential distribution with rate $\lambda > 0$. If not corrected with appropriate countermeasures, the consequences of this adverse event persist after time T and inflict a cost c per unit of time on the principal. This cost captures the event's detrimental effects imposed on society in the case of environmental violations, or the financial damages and potential harm to the retailer's reputation in the context of supplier noncompliance.[2]

A remedial action can bring an end to these damages at cost r, which covers the expenses involved in restoring the environmental impact, compensating victims, and adopting compliant equipment and repairing the retailer's reputation. This cost is less than the maximum (discounted) negative impact of the event, that is, $r < c/\theta$, where $\theta > 0$ is the discount rate. (Otherwise, no party has any incentive to take the remedial action, and the problem becomes trivial.) Thus, the fact that the agent may need to bear (part of) the remedial cost discourages him from disclosing and fixing the issue, giving rise to a problem of *adverse selection*.

To determine whether the event has occurred, the principal can (possibly randomly) audit the firm at any time and charge a nondisclosure penalty if the issue is uncovered. The agent, however, can exert an effort to evade these audits through deception/falsification without addressing the issue. In other words, the principal also faces a problem of *moral hazard*, in addition to the adverse selection issue. Because of these incentive misalignments and because audits are costly, the principal may alternatively provide the agent with incentives to voluntarily disclose the event.

3.1. Audit and Evasion

Specifically, the principal can conduct an audit with a cost k at any time. Audit schedules can be very general and combine both "impulsive" and/or "intensive" audits. An impulsive audit takes place at time epoch t with probability $q_t^m \in [0,1]$, where we require only finitely many impulsive audit time epochs with $q_t^m > 0$ within any finite time interval. By contrast, an intensive audit occurs in time interval $[t, t+\Delta t)$ with probability $q_t^n \Delta t + o(\Delta t)$, where the audit rate $q_t^n \geq 0$. We denote the principal's *audit schedule* by $Q := (Q_t)_{t\in[0,\infty)}$, where $Q_t := (q_t^m, q_t^n)$. (A rigorous definition is provided in Online Appendix A.) This framework captures any type of reasonable auditing schedules. For example, the principal can decide to follow a deterministic auditing schedule, randomly audit at prespecified times, randomly audit according to an arrival rate, or any combination. This allows the principal to consider all possible scheduling policies one can reasonably imagine. Despite this very rich set of policies, we demonstrate later in the paper that the optimal scheduling policy is easy to understand and implement.

A distinctive feature of our setting is that the agent can exert an *evasive effort*, which is unobservable to the principal, at cost $h > 0$ to render the principal's audits less effective (Lacker and Weinberg 1989). The evasion action and its cost h corresponds to either a one-time occurrence or a continuous level of effort. For instance, the bulk of Volkswagen's evasive efforts consisted in developing a software to avoid detection, which, once developed, was install on each car at negligible cost. When the deceptive mechanism is not automatized and requires continuous effort, evasion cost h represents the agent's cumulative total expected discounted cost of effort.[3]

Evasive actions may not be perfect, however. We denote by $\beta \in [0,1)$ the probability that an audit reveals the issue given that the agent has exerted evasive efforts. Taken together, the pair (h, β) characterizes the agent's evasion technology, such that evasive effort h reduces the detection probability from one to β. Lower values of either h or β provide stronger incentives for the agent to evade the principal's audits.

As an alternative to taking evasive actions, the agent may voluntarily fix the issue at cost r without notifying the principal (i.e., take self-correction actions). Technically, this action is equivalent to a perfect evasion ($\beta = 0$) at cost r but which also terminates the incurrence of cost rate c. The principal may still prefer the agent's disclosing to self-correcting the issue to avoid running unnecessary but costly audits.

When the cost of evasion becomes so significant that it exceeds the remedial cost ($h > r$), the agent lacks any incentive to evade, and the problem simplifies to the one addressed in Wang et al. (2016). Therefore, we focus on the case where evasive actions are meaningful in the sense that

$$h \leq r. \tag{1}$$

3.2. Payment Transfers

If the audit reveals the adverse event at time t, the principal charges the agent a fine $F_t \leq F$, where F is the maximal possible penalty that the principal can inflict on the agent (see Harrington (1988) for a series of justification). In the case of a firm such as Walmart or Levi's, for instance, the penalty may consist in terminating the contract with the supplier, in which case F corresponds to the total opportunity cost associated with this loss of revenue (Hadler 2013). More generally, the agent is protected by limited liability, where F is the maximum penalty that the agent can bear.[4] We focus on the nondegenerate case where

$$h \leq F.$$

Otherwise, the agent cannot afford to evade. Together with (1), we define

$$\overline{h} := r \wedge F, \text{ such that } h \leq \overline{h}, \tag{2}$$

where we use notation $a \wedge b$ to represent $\min\{a, b\}$.

Alternatively, if the agent voluntarily discloses the issue at time t, the principal charges the agent a penalty $P_t \leq F$. We do not assume but will show that $0 \leq P_t \leq \overline{h}$ at optimality. Thus, the policy corresponds to a cost-sharing mechanism, where payments P_t and $r - P_t$ represent a breakdown of remedial cost shared between the agent and the principal, respectively. In particular, if $P_t < r$, then the agent strictly prefers self-disclosure (by paying P_t toward remediation) to self-correction, whereas the agent is indifferent between these two options if $P_t = r$. In practice, such disclosure incentives are typically implemented in the form of penalty reductions (EPA 2000) or remediation assistantship (Hadler 2013).

3.3. Timeline

The sequence of events at any point in time is as follows (Figure 1). The principal first designs and commits to a *policy* $\mathcal{P} := (F_t, P_t, Q_t)_{t\in[0,\infty)}$ that specifies the audit schedule Q_t and the agent's payments (F_t, P_t) on detection and disclosure, respectively. The policy is dynamic in that it is adaptive to the public history $\mathcal{I}_t$, which consists of all previous audit time epochs and audit results up to time t. If the issue occurs, the agent responds to the principal's policy by choosing whether and when to (i) disclose the issue, (ii) evade the audit, or (iii) self-correct. Once a disclosure or an audit detection occurs, the strategic interaction between the principal and the agent ends. Therefore, if time proceeds to time t, no audit must have detected any issue until t, and hence the public history $\mathcal{I}_t$ simply corresponds to all the audits' time epochs that have been run thus far.

3.4. Threat Utility

After taking an evasive action, the agent faces the risk of getting caught if evasion is imperfect with $\beta > 0$. We define a *threat utility* U_t to represent the agent's expected discounted cost from time t onwards after taking an evasive action (conditional on that the issue has occurred, i.e., $T \le t$). Threat utility U_t is fully determined by the audit schedule Q_t, fine F_t, and postevasion detection probability β (see (A.4) of Online Appendix A for a formal definition). In particular, when an evasion is perfect, that is, $\beta = 0$, the agent does not face any threat utility and $U_t = 0$ for $t \ge 0$. In contrast, when $\beta > 0$, the principal can adjust Q_t and F_t to increase U_t and incentivize the agent to comply. Doing so, however, may also increase the principal's auditing costs.

3.5. Agent's Cost

Let stopping time $\sigma(T) \ge T$ denote the instant when the agent takes an action after an issue has emerged at time T. In particular, the agent has not yet taken any action during time interval $[T, \sigma(T))$. If $\sigma(T) = T$, then the agent acts without any delay. At time $\sigma(T)$, the agent's costs for (i) disclosing the issue, (ii) evading audits, or (iii) self-correcting are $P_{\sigma(T)}$, $h + U_{\sigma(T)}$, and r, respectively. Therefore, the agent chooses the lowest among them and incurs a cost

$$\mathfrak{c}_{\sigma(T)} := P_{\sigma(T)} \wedge (h + U_{\sigma(T)}) \wedge r.$$

Next define $\tau(T) > T$ as the time epoch of the first audit after the issue has occurred at time T. If $\sigma(T) \le \tau(T)$, the agent is not audited during $[T, \sigma(T)]$ and incurs a cost $\mathfrak{c}_{\sigma(T)}$ at time $\sigma(T)$. If $\sigma(T) > \tau(T)$, the agent has not yet evaded the audit and will be caught by an audit with certainty at time $\tau(T)$, which results in penalty $F_{\tau(T)}$. Thus, the agent's expected discounted cost of following strategy σ in response to the principal's policy $\mathcal{P}$ is equal to

$$C_a(\mathcal{P}, \sigma) = \mathbb{E}[e^{-\theta\sigma(T)} \mathbb{1}_{\{\sigma(T)\le\tau(T)\}} \mathfrak{c}_{\sigma(T)} + \mathbb{1}_{\{\sigma(T)>\tau(T)\}} F_{\tau(T)} e^{-\theta\tau(T)} \mid \mathcal{P}, \sigma]. \tag{3}$$

3.6. Principal's Problem

Prior to the issue's occurrence at time T, the principal incurs a total discounted auditing cost of $k\int_0^T e^{-\theta t} dN_t$, where N_t represents the counting process for the total number of audits up to time t. After time T, the principal accrues a cost at rate c between T and $\tau(T) \wedge \sigma(T)$, which yields a total discounted cost of $c\int_T^{\sigma(T)\wedge\tau(T)} e^{-\theta t} dt$. If $\sigma(T) > \tau(T)$, the agent is caught by an audit at $\tau(T)$ with certainty, and the principal incurs an audit cost k and the net remedial cost $r - F_{\tau(T)}$. (The fine is the principal's income.) Otherwise (i.e., $\sigma(T) \le \tau(T)$), three situations need to be considered. First, if $\mathfrak{c}_{\sigma(T)} = P_{\sigma(T)}$, the agent discloses the issue and the principal covers the remaining remedial cost, $r - P_{\sigma(T)}$. Second, if $\mathfrak{c}_{\sigma(T)} = r$, the agent self-corrects, and the principal keeps incurring auditing costs (but no damage cost c) indefinitely afterward. We denote the principal's total expected cost from $\sigma(T)$ onward in this case as $W_{\sigma(T)}$, which is determined by control Q_t. Finally if $\mathfrak{c}_{\sigma(T)} = h + U_{\sigma(T)}$, the agent takes an evasive action at

Figure 1. Sequence of Events at Any Moment in Time t ($\Delta t \approx 0$) if the Issue Has Occurred (i.e., for $t \ge T$)

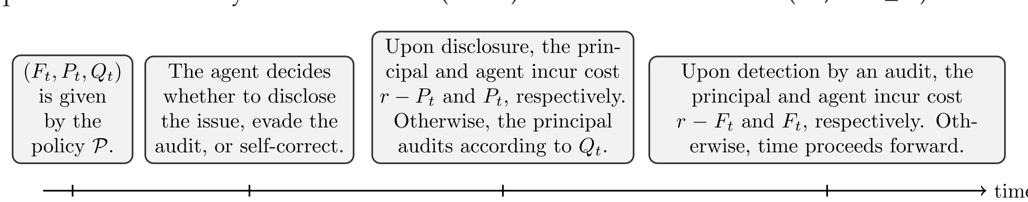

$\sigma(T)$, and we denote the resulting principal's total expected cost from $\sigma(T)$ onward as $V_{\sigma(T)}$, which is determined by control Q_t and F_t. (See (A.6) and (A.5) of Online Appendix A for a formal definition of W_t and V_t, respectively.) Taken together, the principal's total discounted cost is given by

$$\begin{aligned} C(\mathcal{P},\sigma) := \mathbb{E}\Bigg[& k\int_0^T e^{-\theta t} dN_t + c\int_T^{\sigma(T)\wedge\tau(T)} e^{-\theta t} dt \\ & + \mathbb{1}_{\{\sigma(T)>\tau(T)\}} e^{-\theta\tau(T)}(k + r - F_{\tau(T)}) \\ & + \mathbb{1}_{\{\sigma(T)\le\tau(T)\}} e^{-\theta\sigma(T)} \{\mathbb{1}_{\{P_{\sigma(T)}\le\min\{r,h+U_{\sigma(T)}\}\}}(r - P_{\sigma(T)}) \\ & + \mathbb{1}_{\{r<P_{\sigma(T)}, r\le h+U_{\sigma(T)}\}} W_{\sigma(T)} \\ & + \mathbb{1}_{\{h+U_{\sigma(T)}<\min\{r,P_{\sigma(T)}\}\}} V_{\sigma(T)}\} \,|\, \mathcal{P},\sigma \Bigg]. \end{aligned} \tag{4}$$

Overall, the principal's problem consists in designing policy $\mathcal{P}$ that minimizes total expected discounted cost $C(\mathcal{P},\sigma)$ while accounting for the agent's strategic responses. Formally, the principal's *optimal policy* $\mathcal{P}^\star$ is determined as the solution to the following problem,

$$\begin{aligned} C^\star := \min_{\mathcal{P}} C(\mathcal{P},\sigma), \quad \text{subject to} \\ C_a(\mathcal{P},\sigma) \le C_a(\mathcal{P},\sigma') \text{ for all } \sigma', \end{aligned} \tag{5}$$

whereby $C^\star$ denotes the principal's optimal expected total discounted cost. Under the optimal policy $\mathcal{P}^\star$, the agent's optimal total expected discounted cost is given by $C_a^\star := \min_\sigma C_a(\mathcal{P}^\star,\sigma)$, and a best response strategy is a stopping time $\sigma^\star$ (when the agent either discloses or evades) such that $C_a(\mathcal{P}^\star,\sigma^\star) = C_a^\star$.

4. Problem Reformulation

The generality of our framework allows for a large variety of possible auditing policies, which can potentially induce complex disclosure and evasion strategies from the agent. Nonetheless, in this section, we establish that inducing the agent to report the issue without any delay nor evasion is optimal for the principal. In other words, the principal can restrict the search for the optimal policy within the set of incentive-compatible policies that always induce the agent's prompt disclosure. This result extends the classical revelation principle developed for static mechanism design problems (Dasgupta et al. 1979, Myerson 1979) to a dynamic setting with both moral hazard and costly state verification.

Theorem 1 (Optimality of Prompt Disclosure). *For any given policy $\hat{\mathcal{P}} := (\hat{F}_t, \hat{P}_t, \hat{Q}_t)_{t\in[0,\infty)}$ with the agent's best response strategy $\hat{\sigma}^\star$, a policy $\mathcal{P} := (F_t, P_t, Q_t)_{t\in[0,\infty)}$ exists such that*

1. *The fine on detecting the issue through an audit is set to its maximum level, that is, $F_t := F$;*

2. *The agent always prefers disclosing the issue to evading or self-correcting, that is,*

$$P_t \le r \wedge (h + U_t), \quad \textit{for all } t \ge 0, \tag{6}$$

where U_t evolves according to

$$U_t = (1 - q_t^m) U_{t+} + q_t^m \left(\beta F + (1-\beta) U_{t+}^I\right), \;\; \textit{for } q_t^m > 0, \tag{7}$$

$$\frac{dU_t}{dt} = \theta U_t - q_t^n [\beta F + (1-\beta) U_{t+}^I - U_t], \;\; \textit{for } q_t^m = 0, \tag{8}$$

with U_t being reset to U_{t+} (respectively, U_{t+}^I) satisfying (6) right after time t in the absence (respectively, presence) of an audit.

3. *The agent always prefers disclosing without delay, that is, $C_a(\mathcal{P},T) \le C_a(\mathcal{P},\sigma)$ for all σ; or, equivalently, P_t evolves according to*

$$P_t \le (1 - q_t^m) P_{t+} + q_t^m F, \quad \textit{for } q_t^m > 0, \quad \textit{and} \tag{9}$$

$$P_t \le P_{t+}, \;\; \textit{or} \;\; \frac{dP_t}{dt} \ge \theta P_t - q_t^n (F - P_t), \;\; \textit{for } q_t^m = 0, \tag{10}$$

with P_t being reset to P_{t+} (respectively, P_{t+}^I) satisfying (6) right after time t in the absence (respectively, presence) of an audit.

4. *The agent's total discounted expected cost remains the same, whereas the principal is not made worse off, that is, $C_a(\mathcal{P},T) = C_a(\hat{\mathcal{P}},\hat{\sigma}^\star)$ and $C(\hat{\mathcal{P}},\hat{\sigma}^\star) \ge C(\mathcal{P},T) = \mathbb{E}[k\int_0^T e^{-\theta t} dN_t + e^{-\theta T}(r - P_T) | \mathcal{P}]$.*

To establish Theorem 1, we construct a new policy $\mathcal{P}$, under which the agent's payment on disclosure P_t replicates the agent's minimum expected discounted cost under the agent's best response to policy $\hat{\mathcal{P}}$. By doing so, the principal maintains the same expected discounted payment toward remediation and hence equivalent payoff to the agent. This further allows the principal to always remedy the adverse consequences of the issue without delay (i.e., avoid cost c) and uncover the issue through self-reporting (i.e., avoid future unnecessary auditing costs). The principal is thus better off replacing the payment scheme of any arbitrary policy with the one in Theorem 1.

In essence, Theorem 1 states that focusing on policies that always induce prompt disclosure is optimal. Under such a policy, penalty F_t never materializes and only serves as a threat to the agent. As a result, the principal maximizes the penalty to the agent's limited liability F, as stated by the first point of the proposition. We thus refer to a policy $\mathcal{P}$ in the following as a pair $(P,Q) = (P_t, Q_t)_{t\in[0,\infty)}$, and take $F_t = F$ for all t.

This policy must also satisfy the *obedience constraint* (6), per the second point of the theorem. This constraint addresses the moral hazard (i.e., hidden action) incentive that emerges from our setting. The constraint requires that the agent always finds voluntary disclosure (at cost P_t) more economically attractive than either self-correction (at cost r), or evasion and potentially getting caught (at cost $h + U_t$). Therefore, Obedience Constraint (6), together with the limited liability F, imposes an upper bound for the disclosure payment $P_t \le \overline{h}$. Furthermore, Obedience Constraint (6) becomes more stringent for a given Q_t when the evasion becomes relatively easy (i.e., lower value of h). In other words, the principal has to offer the agent higher disclosure incentives (i.e., lower the penalty P_t) to induce no evasive behavior. As a result, cost U_t never materializes on the equilibrium path and again only serves as a threat to the agent. Theorem 1 further explicitly characterizes the dynamic evolution of threat utility U_t in (7) and (8), which are determined by the audit schedule Q_t. Whether the principal should induce a particular action from the agent is in general unclear for moral hazard problems. Theorem 1 shows, however, that it is indeed optimal for the principal to induce the agent and not to evade in our setting.

Finally, *informational IC Constraint* $C_a(\mathcal{P},T) \le C_a(\mathcal{P},\sigma)$ in the last point of the theorem addresses the adverse selection (i.e., hidden information) problem by inducing the agent to disclose the issue as soon as it occurs. Equations (9) and (10) in the proposition express this constraint in a recursive manner for each time instant given that the issue has occurred. They are derived by ensuring the agent's payment P_t (from immediately disclosing the adverse event that has occurred) is no higher than the payment of postponing the disclosure to the next moment. The constraints regulating U_t, (7) and (8), share great similarity with those regulating P_t, (9) and (10). In fact, in the absence of an effective evasion capability (i.e., when $\beta = 1$), the evolution of U_t coincides with the (binding) trajectory of P_t. The presence of an evasion capability, however, requires P_t and U_t to diverge. (See Lemma A.1 and Remark A.1 of Online Appendix A for the feasible range of (P_t, U_t).)

Overall, Theorem 1 allows to reformulate the principal's problem (5) as

$$C^\star = \frac{\lambda}{\theta+\lambda} r + \min_{\mathcal{P}=(P_t,Q_t)_{t\in[0,\infty)}} \mathbb{E}\left[k\int_0^T e^{-\theta t} dN_t - e^{-\theta T} P_T \middle| \mathcal{P}\right], \quad \text{subject to (6)} - \text{(10)}. \tag{11}$$

The principal's objective in (11) explicitly captures the fundamental tradeoff that the principal needs to make. Specifically, the principal can reduce its cost by increasing the agent's payment P_t, whereas larger payments require more frequent audits, resulting in higher audit costs $k\int_0^T e^{-\theta t} dN_t$. Furthermore, (11) does not depend on the impact of the adverse event, c. Indeed, according to Theorem 1, the agent always immediately reports the event, which is then fixed, at optimality. As a result, cost c is irrelevant for the optimal policy. Indeed, because cost c does not appear in the agent's utility, the policy offering sufficient disclosure incentives to the agent does not depend on this cost. Furthermore, enduring cost flow c is never optimal for the principal because paying for the remedial cost r is always preferable in our setup (because $r \le c/\theta$). The principal, however, may incur a cost lower than r by inspecting the agent.

5. Optimal Policy for Perfect Evasion

We first examine the case where the evasion technology is able to render the principal's audits completely ineffective, that is, the detection probability $\beta = 0$. In this case, the agent has the strongest incentive to take the evasive action. This allows us to isolate the effect that the evasion effort h has on the principal's policy. When evasion is perfect, the principal is unable to threaten the agent once an evasive action is taken (i.e., $U_t \equiv 0$, see the proof of Theorem 2 in Online Appendix B). As a result, IC Constraints (7) and (8) become irrelevant, and Obedience Constraint (6) reduces to $P_t \le h$ (because $h \le \overline{h}$ from (2)) and $U_t = 0$. The following result demonstrates that the agent's ability to evade audits induces the principal to use a random (as opposed to deterministic) inspection schedule.

Theorem 2. *If $\beta = 0$, then the principal's optimal policy $(P^\star, Q^\star)$ exhibits a cyclic structure marked by periodic random audits and persists as long as the issue has not been revealed by the agent. Specifically, let $t^\star$ be the unique solution to the following equation in t:*

$$\Gamma(t;h) := (\lambda+\theta)[F-(k+F)e^{-\lambda t}] - \lambda h[1-e^{-(\lambda+\theta)t}] = 0. \tag{12}$$

Then, each cycle $i = 1,2,\ldots$ starts with a deterministic period of length $t^\star$, immediately after the last audit at τ_{i-1} (with $\tau_0 = 0$), during which the principal conducts no audits (i.e., $q_t^{m\star} = q_t^{n\star} \equiv 0$ for $t \in (\tau_{i-1}, \tau_{i-1}+t^\star]$) and charges the agent a payment according to

$$P_t^\star = \Pi(t;t^\star,h) := he^{-\theta(\tau_{i-1}+t^\star-t)}, \text{ for } t \in (\tau_{i-1}, \tau_{i-1}+t^\star]. \tag{13}$$

Starting from $\tau_{i-1}+t^\star$, the principal conducts only intensive audits (i.e., $q_t^{m\star} \equiv 0$) at a finite constant rate while maintaining a constant payment level, respectively, given by

$$q_t^{n\star} \equiv q^{n\star} := \frac{\theta h}{F-h}, \quad \text{and} \quad P_t^\star \equiv h, \quad \text{for } t \in (\tau_{i-1}+t^\star, \tau_i], \tag{14}$$

until the next audit takes place at time τ_i. Namely, conditional on τ_{i-1}, the time until the next audit $\tau_i - \tau_{i-1} - t^\star$ is an independently and identically distributed exponential random variable with rate $q^{n\star}$ given in (14).

Figure 2 illustrates a sample path of payment $P_t^\star$ under the optimal policy characterized by Theorem 2. As depicted by the figure, the optimal policy demonstrates a cyclic structure and alternates between *deterministic and increasing* monetary payments and *random* audits.

Figure 2. Optimal Policy with $F = 10$, $k = 2$, $\lambda = 0.2$, $\theta = 1$, $h = 5$, and $\beta = 0$

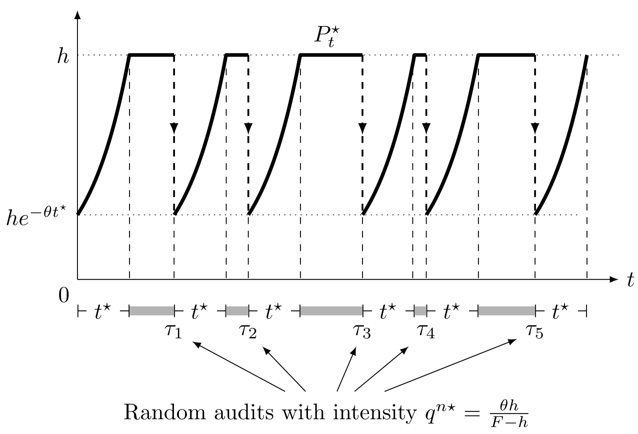

Specifically, the principal starts each cycle by first adjusting payment $P_t^\star$, which increases exponentially from the lower threshold $he^{-\theta t^\star}$ until it reaches the evasion cost h imposed by Obedience Constraint (6). (Recall that $U_t \equiv 0$ if $\beta = 0$.) The increasing curve of $P_t^\star$ ensures that the agent is indifferent between disclosing the issue immediately (by paying $P_t^\star$ toward remediation) and delaying the disclosure to a later time, say $t + \Delta t$, which implies a cost of $(1 - \theta\Delta t)P_{t+\Delta t}^\star$ due to discounting. Namely, the level of $P_t^\star$ is set such that $P_t^\star = (1 - \theta\Delta t)P_{t+\Delta t}^\star$. Taking Δt to zero, we obtain $dP_t^\star/dt = \theta P_t^\star$,[5] which implies the optimal payment trajectory in (13). Furthermore, as long as the increasing trajectory $P_t^\star$ remains below the evasion cost h, the agent has no incentive to take an evasive action. Hence, the monetary instrument provides sufficient incentive for the agent to promptly disclose the issue and the principal does not need to conduct any audit.

Once $P_t^\star$ reaches h (after $t^\star$ units of time since the last audit), the monetary incentives are exhausted. To discourage any evasion, the principal then resorts to audits, while maintaining $P_t^\star$ at the constant level h. The audit is actually *random* with a constant intensity rate $q^{n\star}$. This specific rate ensures that the time-discounting benefit of delaying the disclosure for Δt, which is $(\theta\Delta t)h + o(\Delta t)$, is exactly offset by the net loss of being caught and charged a fine, which is $(q_t^n\Delta t)(F - h) + o(\Delta t)$. That is, audit intensity q_t^n is set such that $(\theta\Delta t)h \approx (q_t^n\Delta t)(F - h)$, which yields the constant audit rate $q^{n\star}$ in (14). It is worth noting that this auditing intensity is time independent and is purely driven by binding IC Constraint (10). Notably, because rate $q^{n\star}$ induces voluntary disclosure *after* the issue has emerged, $q^{n\star}$ does not depend on λ, the rate at which the adverse event may occur. Interestingly, the audit cost k does not impact $q^{n\star}$ either but only affects audits through their frequency, namely the time interval $t^\star$ per (12).

This random inspection structure is in sharp contrast to the *deterministic* audit policy in Wang et al. (2016) for the case where evasive actions are impossible (or equivalently when $h \geq F$). Here, the principal needs to run random audits to account for the agent's moral hazard incentive of evasion in our setting. When adverse selection is the only incentive issue, the agent's payment P_t on disclosure is bounded by the limited liability, F. As a result, deterministic audits are optimal per Wang et al. (2016). In our setting, however, payment P_t needs to stay below $h \leq F$ to prevent the agent from evading audits, which induces audits to be random.

Also, the principal periodically audits the agent at optimality, even though the agent is able to render these audits fully ineffective. This is because the optimal policy is precisely designed to prevent the agent from taking evasive actions (per Obedience Constraint (6), with $U_t = 0$ when $\beta = 0$). As a result, audits and the penalty F serve as credible threats to enforce compliance.

Last, it is important to highlight the remarkable *simplicity* of the optimal policy. In essence, our policy motivates the agent to come clean before a random audit takes place. The principal runs random audits periodically but offers a reduced penalty for self-reporting during a fixed amount of time before each audit. This penalty level gradually increases in a deterministic fashion due to time discounting. When the penalty reaches h, the time until an audit follows an exponential distribution.[6]

Overall, the optimal policy corresponds to a *dynamic cost sharing mechanism* with audits, where the agent is always responsible for a strictly positive portion of the remedial cost $P_t^\star \in (0, r]$, whereas the principal covers the remaining remedial cost $r - P_t^\star \in (0, r]$. Interestingly, the principal can in fact shift the *entire* auditing cost onto the agent, that is, the agent's expected payment is larger than the principal's total expected auditing cost, as stated in the next corollary.

Corollary 1. *Under the optimal policy* $(P^\star, Q^\star)$ *specified in Theorem 2, the total expected discounted costs for the principal and agent are given by* $C^\star = \frac{\lambda}{\lambda+\theta}(r - he^{-\theta t^\star})$ *and* $C_a^\star = \frac{\lambda}{\lambda+\theta}he^{-\theta t^\star} + A^\star$*, respectively, where* $A^\star = \frac{\theta}{\lambda+\theta}\frac{khe^{-\theta t^\star}}{k+F-he^{-\theta t^\star}}$ *is the principal's total auditing expense. Furthermore, we have* $C^\star > 0$ *and* $C_a^\star > A^\star > 0$.

Finally, the total discounted costs incurred by both the principal and the agent do not depend on c, the cost from the adverse event in the absence of remediation measures. This holds true even when the inspection cost k is significantly high. As k increases, the frequency of inspections decreases (i.e., $t^\star$ becomes larger) at optimality, independently of c. In the limit when k approaches infinity, the principal's cost ultimately aligns with the (discounted) remedial cost r.

6. Optimal Policy for Imperfect and Costly Evasions

We see in the last section that when evasion is perfect with $\beta = 0$, the agent has the strongest incentive to evade, and Obedience Constraint (6) always binds at optimality (over some time intervals) per Theorem 2. This is actually true for any evasive cost $h < \overline{h}$, even though a high evasion cost h reduces the agent's incentive to evade and thus should relax the constraint. When the agent's evasive action is imperfect with $\beta > 0$ (and hence $U_t \geq 0$), however, the constraint $P_t \leq h + U_t$ may not always bind anymore. In this section we examine the case when the agent's incentive to evade is sufficiently weak for that constraint to be ignored. We find that this happens when the evasion cost is above a threshold, which we obtain in a closed form. As such, the moral hazard problem is mute, and the optimal policy only binds the upper bound constraint $P_t \leq \overline{h}$.

Theorem 3. *Let t° be the unique positive solution to equation $\Gamma(t;\overline{h}) = 0$ in t, in which function $\Gamma(t,\cdot)$ is defined in (12), and define*

$$\hat{h}(\beta) := \frac{(1-\beta)[F - \overline{h}e^{-\theta t^\circ}]}{F - \overline{h}(1-\beta)e^{-\theta t^\circ}}\overline{h} \in [0, (1-\beta)\overline{h}], \quad \text{for } \beta \in [0,1]. \tag{15}$$

Then, for $\beta > 0$ and $h \geq \hat{h}(\beta)$, the optimal control policy $(P^\star, Q^\star)$ exhibits a cyclic structure similar to Theorem 2. Specifically, each cycle $i = 1, 2, \ldots$ starts with a deterministic period of length t°, immediately after the last audit at τ_{i-1} (with $\tau_0 = 0$), during which $q_t^{m\star} = q_t^{n\star} \equiv 0$ for $t \in (\tau_{i-1}, \tau_{i-1} + t^\circ]$ and $P_t^\star = \Pi(t; t^\circ, \overline{h})$ for $t \in (\tau_{i-1}, \tau_{i-1} + t^\circ]$, in which $\Pi(t;\cdot,\cdot)$ is defined in (13).

Starting from $\tau_{i-1} + t^\circ$, the principal conducts audits in the following fashion:

- *If $r \geq F$, then a deterministic audit occurs at $\tau_i = \tau_{i-1} + t^\circ$ (i.e., $q_{\tau_i}^{m\star} = 1$, $q_t^{m\star} \equiv 0$ for $t \neq \tau_i$, and $q_t^{n\star} \equiv 0$ for all t).*
- *If $r < F$, then the principal conducts an intensive audit (i.e., $q_t^{m\star} \equiv 0$) at a finite constant rate while maintaining a constant payment level, given by*

$$q_t^{n\star} \equiv q^{n^\circ} := \frac{\theta r}{F - r}, \quad \text{and} \quad P_t^\star = r, \quad \text{for } t \in (\tau_{i-1} + t^\circ, \tau_i], \tag{16}$$

respectively, until the next audit takes place at time τ_i.

The corresponding threat utility equals to

$$U_t^\star = \frac{\beta F}{F - (1-\beta)\overline{h}e^{-\theta t^\circ}} P_t^\star, \quad \forall t \geq 0. \tag{17}$$

Based on Theorem 3, we can further obtain the optimal costs of the principal and that agent, which indicate that the agent bears the entire auditing costs.

Corollary 2. *Under the optimal policy (t°, F) specified in Theorem 3, the total expected discounted costs for the principal and agent are given by $C^\star = \frac{\lambda}{\lambda+\theta}(r - \overline{h}e^{-\theta t^\circ})$ and $C_a^\star = \frac{\lambda}{\lambda+\theta}\overline{h}e^{-\theta t^\circ} + A^\star$, respectively, where $A^\star = \frac{\theta}{\lambda+\theta}\frac{k\overline{h}e^{-\theta t^\circ}}{k+F-\overline{h}e^{-\theta t^\circ}}$ is the principal's total auditing expense. Furthermore, we have $C^\star > 0$ and $C_a^\star > A^\star > 0$.*

In essence, Theorem 3 shows that a cyclic policy akin to the one identified in Theorem 2 remains optimal when evasions are imperfect ($\beta > 0$) and sufficiently costly ($h \geq \hat{h}(\beta)$). More specifically, each cycle still features a deterministic no-audit period (with length t°) followed by an audit that resets the cycle. In particular, the audit is *random* with a constant rate if remedial cost r is lower than limited liability F (Figure 3(a)). The random nature of the audit stems again from binding Obedience Constraint (6). Specifically, Equation (16) indicates that the agent's payment P_t never goes beyond the remidial cost r to mitigate the agent's self-correction incentive. However, when remedial costs more than the agent's limited liability (i.e., $r \geq F$), the entire policy becomes deterministic and periodic (Figure 3(b)). In this case, Obedience Constraint (6) is not binding under the optimal policy and the setting reduces to the one studied in Wang et al. (2016). Formally, we define the class of deterministic cyclic policies as follows.

Definition 1. A *deterministic cyclic policy* $(\overline{t}, \overline{p})$, with periodicity $\overline{t} > 0$ and maximum payment $\overline{p} \in (0, F]$, is a policy (P, Q) such that a deterministic audit occurs at every time epoch $\tau_i = \overline{t} \times i$ (i.e., $q_{\tau_i}^m = 1$, $q_t^m = 0$ for $t \neq \tau_i$, and $q_t^{n\star} := 0$ for all t) for $i = 1, 2, \ldots$, as long as the agent does not reveal the issue. The payment between two consecutive audits follows the same trajectory $P_t = \overline{p}e^{-\theta(\tau_i - t)}$ for $t \in (\tau_{i-1}, \tau_i]$ and $i = 1, 2, \ldots$.

In particular, the optimal policy for $r \geq F$ in Theorem 3 is a deterministic cyclic policy (t°, F).

When the evasion cost h is sufficiently high ($h \geq \hat{h}(\beta)$), the principal can ignore the agent's incentive to evade, and hence threat utility U_t. Problem (11) then reduces to a single-dimensional stochatic control on state P_t.

Figure 4 illustrates this point and depicts threshold $\hat{h}(\beta)$ in the space of evasion capabilities (β, h) following (15). Threshold $\hat{h}(\beta)$ is below the line $h = (1-\beta)\overline{h}$. Thus, when evasion is imperfect, there always exist some evasion capabilities (β, h) with $\hat{h}(\beta) \leq h < \overline{h}$ that the principal can safely ignore. More generally, the optimal policy in Figure 4's shaded area is the same as descried in Theorem 3, regardless of evasion capabilities (β, h).

In contrast, threshold $\hat{h}(\beta)$ converges to the evasion cost's upper bound $\overline{h}$ as evasion becomes more effective, that is, β approaches zero. At the limit $\beta = 0$, the principal can never ignore the agent's evasion

Figure 3. Optimal Policy for $h \geq \hat{h}(\beta)$, with $F = 10$, $k = 2$, $\lambda = 0.2$, and $\theta = 1$

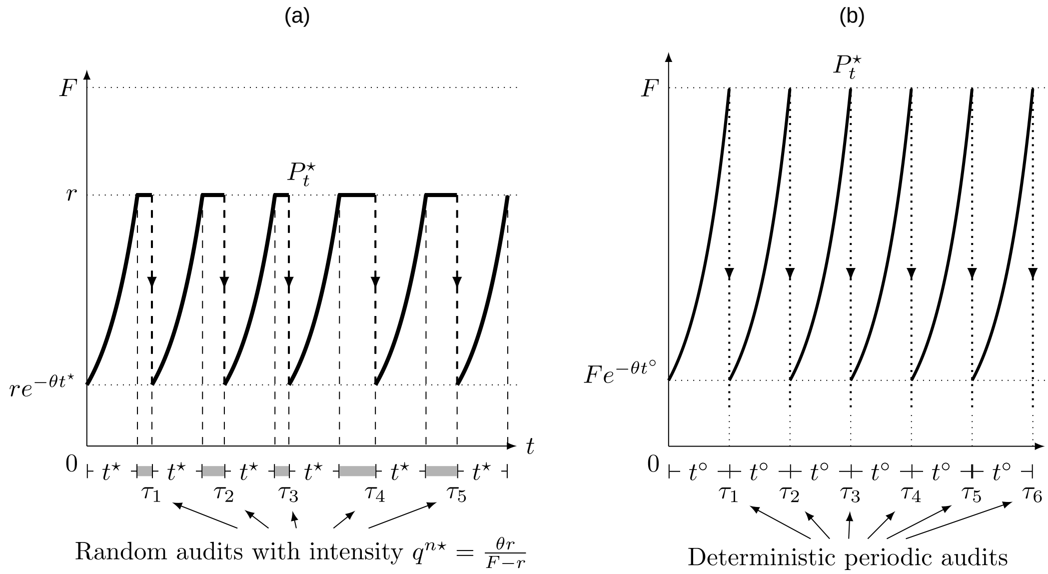

Notes. (a) $r = 7 < F$. (b) $r = 12 > F$.

capability and associated Obedience Constraint (6), for *all* values of the evasion cost $h \in [0, \bar{h}]$. The optimal audit policy becomes random in this case per Theorem 2. In general, the principal needs to explicitly account for Obedience Constraint (6) in Figure 4's unshaded area, for which $h < \hat{h}(\beta)$ and $\beta > 0$. We explore these cases in the next section, where we leverage the fact that payment $P_t^\star$ for self-disclosure is proportional to threat utility $U_t^\star$ at optimality per Equation (17).

7. Policies for Imperfect and Inexpensive Evasion

We now explore situations where the agent has mixed incentives to evade. That is, the agent's evasive action is imperfect with $\beta > 0$, but also inexpensive with $h < \hat{h}(\beta)$ (as defined in (15)), corresponding to the unshaded area of Figure 4. Recall that when evasion is either perfect ($\beta = 0$) or relatively costly ($h \geq \hat{h}(\beta)$), the principal's problem (11) reduces to the stochastic control of a one-dimensional piecewise deterministic process (in P_t). This is because threat utility U_t reduces to zero when $\beta = 0$ or does not affect the agent's evasive action when $\beta > 0$ and $h \geq \hat{h}(\beta)$. When $\beta > 0$ and $h < \hat{h}(\beta)$, however, Problem (11) becomes a genuine *two*-dimensional control of piecewise deterministic process in (P_t, U_t) governed by IC Constraints (7)–(10) and State Constraint (6). Problems of this sort are known to be generally intractable not only analytically but also numerically (Chehrazi et al. 2019; also see Remark A.1 in Online Appendix A for more details). In addition, and perhaps more importantly, the optimal policy solving this problem may be too complex to be implementable in practice (see again Remark A.1 in Online Appendix A).

Figure 4. (Color online) Shaded Area Represents Parametric Range of (β, h) with $h \geq \hat{h}(\beta)$, for Which Theorem 3 Holds ($r = 7$, $F = 10$, $k = 2$, $\lambda = 0.2$, and $\theta = 1$)

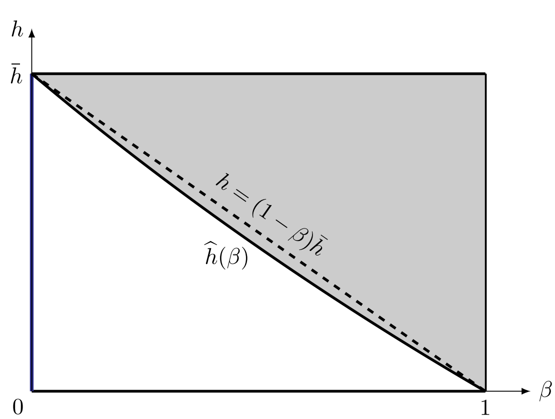

In the following, we restrict the search for efficient auditing policies within a large class of tractable policies. Under all the optimal policies we have seen thus far, the threat utility U_t is always proportional to the disclosure incentive P_t. Motivated by this property, we define the class of *proportional* policies $\mathcal{P}_\gamma = (P_t, Q_t)_{t \in [0,\infty)}$, for some $\gamma \geq 0$, such that policy $\mathcal{P}_\gamma$ induces a threat utility U_t proportional to P_t, that is, $U_t = \gamma P_t$ for all $t \geq 0$. For instance, a proportional policy is optimal when evasion is perfect or sufficiently costly ($\beta = 0$ or $h \geq \hat{h}(\beta)$) with $\gamma = 0$ and $\gamma = \beta F / [F - (1-\beta)\bar{h} e^{-\theta t^\circ}]$, per Theorem 2 and Equation (17) of Theorem 3, respectively.

In other words, under a proportional policy, the principal rewards the agent for self-reporting an issue by having the agent pay a fraction of the expected penalty of getting caught evading audits. Proportional policies do not need to be cyclic or deterministic; they may allow for random audits, a mixture of random and deterministic audits, random audits with different or time-varying rates, and different time interval lengths between audits. Although the class of proportional policies remains fairly large, we show in the following the optimal proportional policy takes a cyclical structure as in the previous cases.

Finding the best proportional policy corresponds to solving the following optimization problem:

$$\hat{C}^\star := \min_{\gamma\ge 0,\,(P_t,\,Q_t)_{t\in[0,\infty)}} \mathbb{E}\left[k\int_0^T e^{-\theta t}dN_t - e^{-\theta T}P_T\middle|\mathcal{P}\right],$$
$$\text{subject to } (6)-(10) \text{ and } U_t=\gamma P_t,\quad \forall t\ge 0. \tag{18}$$

In contrast to the principal's problem in (11), Problem (18) involves an additional decision variable, γ, and a new *proportionality* constraint, $U_t=\gamma P_t$. This additional restriction renders the principal's problem more tractable both mathematically and numerically for $h<\hat{h}(\beta)$. Indeed, the main result of this section shows that the infinite-dimensional optimization problem (18) can be solved through the following two-dimensional *static* optimization:

$$K^\star := \min_{\substack{\beta\le\gamma\le 1-(h/\bar{h})\\ 0\le x\le 1}} \left(\frac{1-\gamma}{h}\right)^\rho \frac{k+\beta[F/\gamma - hx/(1-\gamma)]}{\rho[(\beta F/h)(1-\gamma)/\gamma+(1-\beta)x]-(\rho-1)-x^\rho}$$
$$\text{subject to } \frac{1-(\beta F/h)(1-\gamma)/\gamma}{1-\beta} \le x \le \frac{F(\gamma-\beta)(1-\gamma)}{h(1-\beta)\gamma}\wedge\frac{r(1-\gamma)}{h}, \tag{19}$$

with $\rho:=(\theta+\lambda)/\theta$, where the constraint is mandated by the proportional policy's feasibility. (See Remark C.1 of Online Appendix C for additional explanations about this reframing of Stochastic Control Problem (18) into Static Optimization Problem (19).)

The following theorem shows that the optimal solution to (19) fully characterizes the optimal proportion and the optimal policy for Problem (18).

Theorem 4. *For $\beta>0$ and $h<\hat{h}(\beta)$, a unique solution $(\gamma^\star,x^\star)$ to (19) exists, such that the optimal proportion solving (18) is equal to $\gamma^\star$. Furthermore, for an initial period of length $t_0^\star=-(1/\theta)\ln[(1-\gamma^\star)p_0^\star/h]$ with $p_0^\star=(\rho K^\star)^{-\theta/\lambda}<h/(1-\gamma^\star)$, the optimal policy $(P^\star,Q^\star)$ for (18) conducts no audits (i.e., $q_t^{m\star}=q_t^{n\star}\equiv 0$ for $t\in[0,t_0^\star]$) and charges the agent a payment according to*

$$P_t^\star = p_0^\star e^{\theta t},\ \textit{for } t\in[0,t_0^\star]. \tag{20}$$

Starting from $t_0^\star$, the optimal policy $(P^\star,Q^\star)$ for (18) exhibits a cyclic structure marked by periodic random audits and persists as long as the issue has not been revealed by the agent. Specifically, each cycle $i=1,2,\ldots$ starts with only intensive audits (i.e., $q_t^{m\star}\equiv 0$) at a finite constant rate while maintaining the constant payment level, respectively given by

$$q_t^{n\star}\equiv q^{n\star} := \frac{\theta}{\beta\frac{F(1-\gamma^\star)}{h\gamma^\star}+(1-\beta)x^\star-1},\quad \textit{and}\quad P_t^\star\equiv\frac{h}{1-\gamma^\star},$$
$$\textit{for } t\in(t_0^\star,\tau_1]\bigcup_{i=1}^{\infty}(\tau_i+t^\star,\tau_{i+1}], \tag{21}$$

where $t^\star=(-1/\theta)\ln x^\star$ and τ_i is the ith audit. Immediately after the last audit at τ_i, the principal applies no audits (i.e., $q_t^{m\star}\equiv q_t^{n\star}:=0$) for a deterministic period of length $t^\star$ and charges the agent a payment according to

$$P_t^\star=\frac{h}{1-\gamma^\star}e^{-\theta(\tau_i+t^\star-t)},\ \textit{for } t\in\bigcup_{i=1}^{\infty}(\tau_i,\tau_i+t^\star]. \tag{22}$$

Figure 5 illustrates the policy structure identified by Theorem 4 for a sample path of audits. This structure is similar to the prefect evasion case with $\beta=0$ (Theorem 2), except for the initial period of length $t_0^\star$. Specifically, Obedience Constraint (6) restricts the payment P_t to be upper bounded by $h/(1-\gamma^\star)\le\bar{h}$. Between two consecutive intensive auditing episode, IC Constraint (10) is binding, and therefore the payment P_t follows a deterministic exponential trajectory as in (20) and (22). Once the payment P_t reaches its upper bound $h/(1-\gamma^\star)$, the principal exhausts the monetary incentive, and the

Figure 5. Optimal Proportional Policy with $F=10$, $k=2$, $\lambda=0.2$, $\theta=1$, $h=1$, and $\beta=0.6$

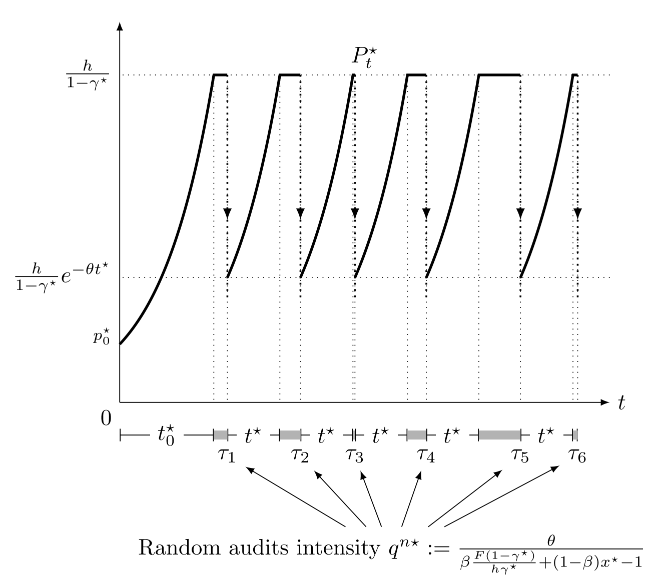

payment P_t remains constant at this level per (21). The principal then switches to random audits to incentivize the agent, with a constant audit rate also defined in (21).[7] In contrast to the perfect evasion case $\beta = 0$, however, IC Constraint (10) may no longer bind in this case.[8]

Importantly, the optimal proportional policy in Theorem 4 always relies on random audits. The audits would be deterministic if the constant audit rate in (21) was infinite, which would happen if $x^\star$ was equal to its lower bound in (19). However, Theorem 4 states that this never occurs (see also Lemma C.3 in Online Appendix C).[9] This further implies that a deterministic cyclic policy cannot be optimal (over the whole set of feasible policies) when $h < \hat{h}(\beta)$, as shown by the following proposition.

Proposition 1. *A deterministic cyclic policy is optimal if and only if $r \geq F$ and $h \geq \hat{h}(\beta)$.*

Finally, the cyclic structure of the optimal policy in Theorem 4 yields a cost decomposition that is similar to the one characterized in Corollaries 1 and 2. (With a slight abuse of notation and for simplicity, we still use notations $A^\star$, $C^\star$, and $C_a^\star$ to denote the corresponding costs under the optimal proportional policy.)

Corollary 3. *Under the principal's optimal proportional policy $(P^\star, Q^\star)$ prescribed in Theorem 4, the total expected discounted costs for the principal and agent are given by $C^\star = \frac{\lambda}{\lambda+\theta}(r - p_0^\star)$ and $C_a^\star = \frac{\lambda}{\lambda+\theta} p_0^\star + A^\star$, respectively, where $A^\star = \frac{\theta}{\lambda+\theta} \frac{k p_0^\star}{k+\beta[F/\gamma^\star - hx^\star/(1-\gamma^\star)]}$ is the principal's total audit expense. Furthermore, we have $C^\star > 0$ and $C_a^\star > A^\star > 0$.*

8. Effect of Evasion Capability on Costs

The previous analysis allows exploring the impact of the agent's evasion capacity on the principal's audit policy. To that end, we vary detection probability β and numerically evaluate the resulting audit frequency and the associated expected costs.

In our setup, the mean sojourn time between two consecutive audits is equal to $t^\star + 1/q^{n\star}$. Indeed, the auditing policies in Theorems 2–4 alternate between a payment phase of fixed length $t^\star$ (or t°) and a random audit phase of average length $1/q^{n\star}$ (which is zero if the audit is deterministic). Thus, a *lower* value of $t^\star + 1/q^{n\star}$ indicates *more frequent* audits.

Our results allow us to evaluate this sojourn time and the associated expected costs. In particular, we obtain the sojourn time from Equations (12) and (14) of Theorem 2 when $\beta = 0$, and from t° and (16) of Theorem 3 when $h \geq \hat{h}(\beta)$. When $\beta > 0$ and $h < \hat{h}(\beta)$, we have $t^\star = (-1/\theta)\ln x^\star$, in which $x^\star$ is the optimal solution of (19), and $q^{n\star}$ as defined in (21). The expected costs are evaluated from Corollaries 2 and 3, depending on whether $h \geq \hat{h}(\beta)$ or not.

Figure 6 depicts the principal's audit frequency and the corresponding expected costs as a function of detection probability β. Specifically, Figure 6(a) plots the mean sojourn time, whereas Figure 6(b) depicts the expected auditing cost $A^\star$, the principal's overall cost $C^\star$, and the agent's cost $C_a^\star$. Recall also that the agent bears the audit cost, such that $C_a^\star - A^\star > 0$ corresponds to the agent's expected contribution to the remedial cost.

Figure 6(a) demonstrates a nonmonotone (and possibly discontinuous) relationship between detection probability β and audit frequency. Specifically, the principal audits first more and then less frequently as detection probability β increases (i.e., evasion becomes less effective) in the shaded region ($h < \hat{h}(\beta)$). In the unshaded region (i.e., $h \geq \hat{h}(\beta)$), the audit frequency remains constant per Theorem 3. As shown in Figure 6(b), auditing cost $A^\star$ exhibits a similar unimodal structure, albeit less pronounced. Furthermore, as evasion becomes less effective (i.e., β increases), the principal can transfer a higher proportion of the remedial costs onto the agent. The principal's overall cost $C^\star$ decreases in probability β as a result.

To understand this unimodal structure, probability β has two countervailing effects. On one hand, as evasion becomes less and thus audits more effective, the principal can rely less on financial incentives and more on auditing to enforce compliance. In this case, the mean sojourn time decreases, and audit cost $A^\star$ increases. In addition, the principal decreases the financial incentives by transferring a higher proportion of the remedial cost to the agent, increasing the agent's cost but decreasing the principal's. Conversely, because audits are more effective, the principal can audit less frequently to enforce compliance and thus reduce audit cost $A^\star$. The first effect dominates the second one when the evasive action is more effective, but the second effect dominates the first one when the action is more effective, yielding the overall unimodal impact of probability β we observe in Figure 6(b).

9. Extensions

In this section, we extend our base model and analysis in four directions to reflect additional considerations from practice. We will demonstrate that the results we obtained from previous sections still remain valid.

9.1. Postevasion Penalty

In our base model (Section 3), we assume that the principal charges the agent the same penalty F_t on detection by an audit, irrespective of whether the agent has taken an evasive action or not. In the following, we allow the principal to use different penalties for the two types of violation. Specifically, we assume that when an audit uncovers an adverse event (which happens with probability β), it also reveals whether the agent has exerted evasive efforts. Notations $\overline{F}_t$ and F_t denote then the penalties on detecting an advert event with or with

Figure 6. (Color online) Effect of Postevasion Detection Probability β on the Audit Frequency and Different Cost Components with $h=2$, $r=F=10$, $k=2$, $\lambda=0.2$, and $\theta=1$

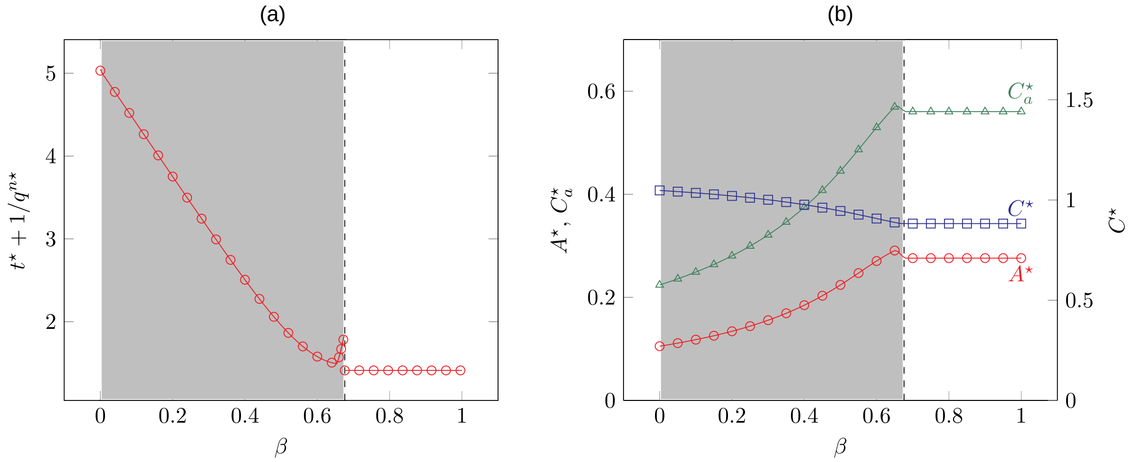

Notes. The shaded (respectively, unshaded) parameter ranges correspond to $h < \hat{h}(\beta)$ (respectively, $h \geq \hat{h}(\beta)$). (a) Audit frequency. (b) Costs.

evasion, respectively. When the agent voluntarily discloses the issue, the principal continues to charge the agent with payement P_t. All three penalties are bounded by maximum fine F.

With this extension, the principal's *policy*, $\mathcal{P} := (F_t, \overline{F}_t, P_t, Q_t)_{t\in[0,\infty)}$, now specifies three payments charged to the agent (depending on the discovery channel and the presence of evasion) and the audit schedule Q_t. In response to such a policy, the agent's strategy σ is again to choose whether and when to self-report the issue or to conceal it through the evasive action. Notably, *threat utility* U_t (which is the agent's expected discounted cost from time t onward after taking an evasive action), is now determined by $(\overline{F}_t, Q_t)$ (see (D.1) of Online Appendix D for a formal definition). The sequence of events remains the same as in our base model (Figure 1).

The following theorem establishes a key result for this extension. Similar to Theorem 1, we can again restrict our search for the optimal policy among all policies that induce the agent to self-report the issue without any delay or evasion. Consequently, only the payment P_t will be induced, and all the other two payments $(F_t, \overline{F}_t)$ act as *off-equilibrium* threats.

Theorem 5 (Optimality of Prompt Disclosure: Extended). *For any given policy* $\hat{\mathcal{P}} := (\hat{F}_t, \hat{\overline{F}}_t, \hat{P}_t, \hat{Q}_t)_{t\in[0,\infty)}$, *there always exists a policy* $\mathcal{P} := (F_t, \overline{F}_t, P_t, Q_t)_{t\in[0,\infty)}$, *under which*

1. *The fine on detecting the issue or the evasion through an audit is set to its maximum level, that is,* $F_t = \overline{F}_t := F$;

2. *The agent always prefers disclosure of the issue to evasion or self-correction, that is,* (6) *holds with* U_t *still evolving according to* (7) *and* (8);

3. *The agent always prefers to disclose without delay, so that* (9) *and* (10) *hold;*

4. *The agent's total discounted expected cost remains the same, whereas the principal is not made worse off, compared with under policy* $\hat{\mathcal{P}}$.

Theorem 5 shows that to induce the agent's timely disclosure (i.e., the equilibrium outcome), the principal should simply maximize the threat (i.e., the off-equilibrium outcome) by setting the fine on detection to the agent's limited liability $F_t = \overline{F}_t := F$ just as in our base model, regardless of whether an evasion is revealed. As a result, the threat to the agent U_t still follows the same evolution as in (7) and (8). As such, just like our base model, Theorem 5 essentially allows us to again reduce any policy $(F_t, \overline{F}_t, P_t, Q_t)$ simply to (P_t, Q_t), which satisfies (6), (9), and (10). Therefore, the principal's problem remains the same as (11) and all the results in Sections 5–8 still apply. In other words, the principal does not benefit from penalizing the nondisclosure of an issue differently when the agent also took evasive actions. In this sense, our base model is without loss of generality.

9.2. Third-Party Discovery

In our base model (Section 3), the principal has two information channels to uncover the adverse event, namely the agent's voluntary disclosure and the audit's detection. Other channels may exist, for instance, when a nongovernmental organization or an independent news media uncovers the adverse event. In this section, we extend our model to incorporate such a setting. More specifically, we assume that the issue can be revealed by a third party with a constant probability rate μ after its occurrence regardless of whether an evasion has been taken. (This means that a not-yet-uncovered issue will come to the principal's awareness after an exponentially distributed period of time.) We focus on the case such that $\mu < \overline{\mu} := \min\{\lambda, \theta r/(F-r)^+\}$, that is, the third-party discovery rate is not too high.[10] When μ takes sufficiently high values such that $\mu \geq \overline{\mu}$, the incentive problem discussed in this paper is arguably less relevant. In this case, the principal relies more on third parties to uncover the problem, rather than its own audit policy.[11]

The following proposition shows that the optimal policy in this extended setting can still be identified within our current framework.

Proposition 2. *Let* $(P_t^\star, Q_t^\star)_{t\in[0,\infty)}$ *be the solution to* (11) *where the discount rate* θ *is replaced by* $\theta+\mu$*, the limited liability* F *by* $\frac{\theta}{\theta+\mu}F$*, the remedial cost* r *by* $r-\frac{\mu}{\theta+\mu}F>0$*, the hazard rate* λ *by* $\lambda-\mu>0$*, and the auditing cost* k *by* $k(\lambda-\mu)/\lambda$*. Then, the optimal policy* $\tilde{\mathcal{P}}^\star := (\tilde{F}_t^\star, \tilde{P}_t^\star, \tilde{Q}_t^\star)_{t\in[0,\infty)}$ *with an exogenous discovery rate* μ *is given by*

$$\tilde{F}_t^\star = F, \quad \tilde{P}_t^\star = P_t^\star + \frac{\mu}{\theta+\mu}F, \quad and \quad \tilde{Q}_t^\star = Q_t^\star. \tag{23}$$

In essence, the exogenous discovery channel plays two roles. First, it acts as a costless random audit (with constant rate μ and perfect detection probability), which helps the principal to reduce the agent's disclosure benefit and inflate the penalty $P_t^\star$ by $\frac{\mu}{\theta+\mu}F$ (i.e., the expected discounted penalty due to the exogenous discovery). Second, it acts to speed up the discounting as an exogenous discovery would immediately terminates the strategic interaction between the principal and the agent. Thus, the discount rate θ is inflated to $\theta+\mu$.

Proposition 2 immediately implies that the optimal policies we obtained for perfect evasions (Theorem 2) and for imperfect but sufficiently costly evasions (Theorem 3) can be reparameterized as the optimal policies in the presence of exogenous discovery channel. For imperfect and inexpensive evasions (Section 7), we consider the class of *proportional* policies in the form $U_t = \gamma P_t$ for some constant γ. Under the reparameterization identified in Proposition 2, we can generalize it to and optimize within the class of *affine* policies in the form $\tilde{U}_t - \frac{\mu}{\theta+\mu}F = \gamma\left(\tilde{P}_t - \frac{\mu}{\theta+\mu}F\right)$ in the presence of exogenous discovery channel.

9.3. Social Welfare

In our current setting, the principal's objective is to minimize its total discounted cost in (4). This objective is reasonable for many settings when the principal is a self-interested party such as a private enterprise (e.g., Walmart). Yet, when the principal is a regulatory agency such as EPA, it may also care about the cost incurred by the agent in (3) and aims to minimize the total *social cost* (i.e., the principal is a social welfare maximizer). In this case, we follow the mainstream literature on regulation economics (Baron and Myerson 1982, Laffont and Tirole 1993), public economics (Dahlby 2008), and environmental regulations (Boyer and Laffont 1999, Lyon and Maxwell 2003, Wang et al. 2016) to assume that any cost incurred by the principal is $\alpha>0$ times more expensive than that of the agent, where the fact α corresponds to the deadweight loss of applying public funds, and captures economic frictions created by regulation (e.g., by raising distortionary taxes).[12] As a result, the principal's problem (5) can be revised as

$$\begin{aligned}&\min_{\mathcal{P}} (1+\alpha)C(\mathcal{P},\sigma) + C_a(\mathcal{P},\sigma),\\ &\text{subject to} \quad C_a(\mathcal{P},\sigma) \le C_a(\mathcal{P},\sigma') \text{ for all } \sigma'.\end{aligned} \tag{24}$$

Proposition 3. *The solution to* (24) *is the same as that to* (11) *with auditing cost* k *replaced by* $(1+1/\alpha)k$.

Proposition 3 shows that the socially optimal policy can essentially be identified by reparameterizing the principal's problem as the one in our base model. This is because the agent's problem and hence the IC constraint in (24) remains unchanged. As a result, we can still focus on the class of policies inducing the agent's prompt disclosure according to Theorem 1. In the principal's objective function, the principal's auditing cost is amplified by the factor α, and the monetary transfer P_t is not completely canceled due to the deadweight loss.

9.4. Imperfectly Informed Agent

In our base model, the agent is assumed to be perfectly *informed* about the adverse issue once it occurs. However, it is plausible that the agent is genuinely unaware of the occurrence of the issue. Assume that the agent can observe the adverse event's occurrence only with a probability $\delta\in(0,1)$ and cannot find it with probability $1-\delta$. In the latter case, the agent cannot disclose the issue even if it has occurred. However, we assume that an audit can still uncover the issue and whether it was observable to the agent. Let $\underline{F}$ be the penalty that the principal levies on audit detection on the agent who did not observe the issue.[13] As in our base model, it is optimal for the principal to lever both disclosure penalty P_t and audits Q_t to induce the agent's prompt disclosure without evasion or self-correction.

In this alternative setting, the principal faces a tradeoff concerning audits, which need to detect unobservable issues to the agent while properly incentivizing the agent to disclose them when they are observable. Therefore, the principal's problem becomes

$$\begin{aligned}\min_{\mathcal{P}=(P_t,Q_t)_{t\in[0,\infty)}} \quad & \delta\left\{\frac{\lambda}{\theta+\lambda}r + \mathbb{E}\left[k\int_0^T e^{-\theta t}dN_t - e^{-\theta T}P_T\middle|\mathcal{P}\right]\right\}\\ &+(1-\delta)\mathbb{E}\left[k\int_0^T e^{-\theta t}dN_t\right.\\ &\left.+c\int_T^{\tau(T)} e^{-\theta t}dt + e^{-\theta\tau(T)}(k+r-\underline{F})\middle|\mathcal{P}\right]\\ \text{subject to} \quad & (6)\text{–}(10).\end{aligned} \tag{25}$$

Proposition 4. *If* $\underline{F} = k+r-c/\theta$*, then the solution to* (25) *is the same as that to* (11) *with auditing cost* k *replaced by* k/δ.

The sufficient condition $\underline{F} = k + r - c/\theta$ in Proposition 4 essentially charges the uninformed agent for the cost of detecting and repairing the issue but deducts the cost of harm to the principal caused by the issue due to its delayed detection. This condition acts to render the effects that the uninformed agent inflicts on the principal independent of the delay of detection $\tau(T) - T$. Otherwise, the principal's problem would be of a fundamentally different nature and needs a separate analytical treatment that we leave for future research.[14]

10. Conclusion

This paper studies the impact of audit evasion capabilities on the efficient audit and remedial strategies. We represent the evasion capability as a costly effort that reduces the audits' detection probability. To evaluate the impact of this capacity on auditing schedules, we allow very general classes of control policies, instead of restricting to a few specific structures. In particular, we do not assume a priori whether the policy is deterministic or random.

The presence of this evasion capability gives rise to a moral hazard problem in which the agent may self-repair or opt to exert effort aimed at avoiding detection. Furthermore, the adverse issue's occurrence is the agent's private information. Even if the agent does not actively evade audits, the agent may nonetheless decide not to disclose or remedy the problem immediately, which would cause damage to the principal. This gives rise to an adverse selection problem. And because the time at which the issue occurs is random, the problem is dynamic. As such, audits act as a threat and deter the agent from both taking evasive actions and delaying the disclosure of a violation.

Taken together, this situation corresponds to a dynamic principal-agent problem with costly state verification, adverse selection and moral hazard. We reformulate this problem as the stochastic optimal control of a piecewise deterministic process. The analysis of this dynamic stochastic control problem yields two important new managerial insights.

First, the presence of an evasion capability may require the principal to run audits randomly. This contrasts with the deterministic audit schedules that are optimal when the agent cannot hide the issue from audits (but may still not disclose the issue voluntarily) (Wang et al. 2016). More specifically, the principal should randomly audit the agent, unless the agent's evasion capacity is not very effective, and the agent cannot afford to self-correct the issue. In this later case, the principal should follow predetermined audit schedules. In this sense, our findings provide a novel rationale for why audits are sometimes random in practice. Technically, the key driver for random audits in our setup is the upper limit that the moral hazard problem imposes on the agent's contribution toward the remedial costs.

Second, as we increase the audit's probability of detection, the principal should audit the agent first more and then less frequently. This means, in particular, that an improvement in the agent's evasion capability can actually decrease the principal's audit costs (but always increases the principal's total cost).

Overall, our analysis yields a policy that is easy to understand and implement: The policy runs a series of random audits but always motivates the agent to come clean. After each audit, the principal first offers a penalty reduction, which is discounted over time according to basic accounting principles. After a fixed amount of time, the penalty reduction stops changing and stays constant until the next inspection, which takes a simple exponentially distributed random time to occur.

Importantly, this policy outperforms any implementable audit schedules (including nonexponential inspection times, combinations of prescheduled audits with random inspections, etc.). In addition, this structure continues to hold when (1) different levels of penalties can be inflicted depending on whether the violation is accompanied with evasive actions, (2) a third party can independently uncover the violation, (3) the principal maximizes social welfare, (4) the agent may not be able to observe the issue's occurrence, (5) the penalty associated with a violation can take any finite value, (6) the agent's evasive actions can aggravate the environmental impact, and (7) the cost of effort is either a lump sum or a flow overtime.

Our model can potentially be extended in a few other directions. In particular, our setup could be generalized to account for different effort levels, such that higher evasive effort reduces audits' efficacy. Our present model corresponds to a case with two effort levels (the evasive effort and the self-repair effort, respectively). We suspect that conditions exist for more than two effort levels such that the agent either evades at the highest intensity or not at all, in which case our results should hold. If not, the threat utility becomes multidimensional, which requires different and novel analytical approaches.

Another potential direction is to consider additional sources of adverse selection in our model. For instance, the hiding cost could be the agent's private information. Alternatively, the agent may privately know upfront whether the agent will be able to observe the event's occurrence (similar to Baliga and Ely 2016). Accounting for these extensions require introducing agents of different types, which, following the revelation principle, requires the principal to offer a *menu* of dynamic contracts. The design of these contracts, in turn, requires representing the agent's dynamic optimal responses to the contract of each type. Overall, this yields a highly nontrivial problem,[15] which we leave for future research.

From a technical perspective, although the optimal control of general PDP is notoriously difficult, we solve this problem in closed form for a given class of PDP (Theorems 2 and 3). When the problem becomes intractable (as in Section 7), we optimize over a subset of policies and again solve the problem in closed form (Theorem 4). We accomplish this result by reducing the stochastic dynamic optimization problem into a deterministic one. Also, the subset of policies we consider is quite general and focuses on policies that are implementable in practice.

Besides the problem of evading detection that we address in this paper, the optimal control of piecewise deterministic Markov processes provides a fruitful framework to address other types of issues related to auditing. For example, in different situations, the agent does not exert effort to evade audits but rather directly influences the likelihood of an adverse issue occurring. This can be modeled as the agent's effort level determining the arrival rate. Variations of our analytical framework could help study this and other settings related to audit and remedial strategies. The rise of sustainability and corporate social responsibility concerns is conferring increasing importance on these questions.

Acknowledgments

The authors thank the area editor, associate editor, and three anonymous reviewers for many sharp comments and constructive suggestions that significantly improved the paper, both in content and exposition.

Endnotes

[1] For example, the U.S. EPA has developed systematic audit protocols under various legislations (https://www.epa.gov/compliance/audit-protocols).

[2] The cost inflicted on the principal by the persistence of the issue can be indiscernible, making the principal unable to infer the emergence of the issue from such a cost. This situation is pervasive. For example, such costs may be confounded with other factors like demand or price fluctuations. These costs may also represent any risk that may materialize to the principal *in the future*, such as the reputation damage associated with a third party publicly revealing the issue.

[3] Modeling the agent's evasion as a binary decision lends analytical tractability. This modeling choice has also been widely adopted in dynamic moral hazard literature (Biais et al. 2010, Myerson 2012, Sun and Tian 2018).

[4] Environmental economists have argued against the principal enjoying a surplus beyond the remedial costs, that is, $F > r$, which can be politically and legally prohibited (Harrington 1988). In this paper, we make no assumptions regarding whether F is larger or smaller than r. In this case, inflicting F may force bankruptcy on the agent (Clifford and Greenhouse 2013).

[5] This corresponds to a binding informational IC constraint (10) at any point in time.

[6] In practice, a random audit with a constant rate could be implemented in the following way. Say, the audit occurs with probability x% each day during the random audit phase. Then the policy is to audit the agent when the first two digits after the decimal place of a commonly observable stock index's opening price is no larger than x.

[7] Threat utility U_t is then also constant by ensuring $dU_t/dt = 0$ according to (8).

[8] IC Constraint (10) is binding if and only if x^* binds its upper bound $\frac{F(\gamma^*-\beta)(1-\gamma^*)}{h(1-\beta)\gamma^*}$ in (19).

[9] As commented in Section 3, only in the limiting case with $h = 0$, the auditing rate becomes infinite and hence cyclic deterministic audits (except for the initial period) is optimal.

[10] In particular, we have $\overline{\mu} = \lambda$ if $F \le r$.

[11] Technically, it can be optimal for the principal to induce self-correction if $\mu \ge \frac{\theta r}{(F-r)^+}$. For $\lambda \le \mu < \frac{\theta r}{(F-r)^+}$, the principal's problem can be reformulated as (11) in our base model, albeit with different discount rates for the principal and agent. Dynamic contract design problems with different discount rates between the principal and the agent involves more complex control and is beyond the scope of this paper (Cao et al. 2023).

[12] A reasonable estimate for α is significantly positive in the magnitude of 0.3 for the U.S. economy (Ballard et al. 1985, Jones et al. 1990) for empirical estimations of α.

[13] If the principal is not able to distinguish between whether the issue is observable or not to the agent, the agent can always claim to be uninformed even upon audit detection, effectively lowering the informed agent's limited liability to $\min\{F, \underline{F}\}$.

[14] If $\underline{F} \ne k + r - c/\theta$, the objective function in (25) involves evaluating $\mathbb{E}[e^{-\theta(\tau(T)-T)} | T, \mathcal{P}]$ as shown in the proof of Proposition 4, where $\mathbb{P}[\tau(T) - T > t | T, \mathcal{P}] = \mathbb{P}[N_{t+T} - N_T = 0 | T, \mathcal{P}] = e^{-\int_T^{t+T} q_s^m ds} \prod_{s=T}^{t+T} (1 - q_s^n)$. Hence, the objective function in (25) can no longer be expressed as an expectation of the integral with respect to dN_t.

[15] We can consider an example with pure dynamic moral hazard, arguably simpler than our setting (with both static moral hazard and dynamic adverse selection), studied in Sun and Tian (2018). The working paper (Tian et al. 2023) attempts to generalize Sun and Tian (2018) to include an upfront static adverse selection such that the principal does not know the agent's effort cost. The results in Tian et al. (2023) are much more complex than those in Sun and Tian (2018). In fact, the optimal menu itself does not seem to possess tractable, let alone implementable, structures. Therefore, Tian et al. (2023) resorts to restricting contract spaces and approximating the optimal solution by upper and lower bounding its performance.

References

Abreu D, Pearce D, Stacchetti E (1990) Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica* 58(5):1041–1063.

Babich V, Tang C (2012) Managing opportunistic supplier adulteration: Deferred payments, inspection, and combined mechanisms. *Manufacturing Service Oper. Management* 14(2):301–314.

Bakshi N, Gans N (2010) Securing the containerized supply chain: Analysis of government incentives for private investment. *Management Sci.* 56(2):219–233.

Baliga S, Ely JC (2016) Torture and the commitment problem. *Rev. Econom. Stud.* 83(4):1406–1439.

Ballard C, Shoven J, Whalley J (1985) General equilibrium computations of the marginal welfare costs of taxes in the United State. *Amer. Econom. Rev.* 75(1):128–138.

Barlow RE, Proschan F (1996) *Mathematical Theory of Reliability* (SIAM, Philadelphia).

Baron DP, Myerson RB (1982) Regulating a monopolist with unknown costs. *Econometrica* 50(4):911–930.

Becker GS (1968) Crime and punishment: An economic approach. *J. Political Econom.* 72(2):169–217.

Bensoussan A, Lions JL (1984) *Impulse Control and Quasivariational Inequalities.* Translation of *Contrôle Impulsionnel et Inéquations Quasi-variationnelles* (Gauthier-Villars, Paris).

Biais B, Mariotti T, Rochet J-C, Villeneuve S (2010) Large risks, limited liability, and dynamic moral hazard. *Econometrica* 78(1):73–118.

Board S, Meyer-Ter-Vehn M (2013) Reputation for quality. *Econometrica* 81:2381–2462.

Boyer M, Laffont J-J (1999) Toward a political theory of the emergence of environmental incentive regulation. *RAND J. Econom.* 30(1):137–157.

Cao P, Sun P, Tian F (2023) Punish underperformance with suspension: Optimal dynamic contracts in the presence of switching cost. *Management Sci.* 70(5):3020–3037.

Carreyrou J (2018) *Bad Blood: Secrets and Lies in a Silicon Valley Startup* (Alfred A. Knopf, New York).

Chehrazi N, Glynn PW, Weber TA (2019) Dynamic credit-collections optimization. *Management Sci.* 65(6):2737–2769.

Chen J, Ai A, Dawande D (2020) Supplier centrality and auditing priority in socially responsible supply chains. *Manufacturing Service Oper. Management* 22(6):1199–1214.

Cho S-H, Fang X, Tayur S (2015) Combating strategic counterfeiters in licit and illicit supply chains. *Manufacturing Service Oper. Management* 17(3):273–289.

Clifford S, Greenhouse S (2013) Fast and flawed inspections of factories abroad. Accessed October 24, 2016, http://www.nytimes.com/2013/09/02/business/global/superficial-visits-and-trickery-undermine-foreign-factory-inspections.html.

Dahlby D (2008) *The Marginal Cost of Public Funds* (MIT Press, Cambridge, MA).

Dasgupta P, Hammond P, Maskin E (1979) The implementation of social choice rules. *Rev. Econom. Stud.* 46(2):185–216.

Davis MHA, Dempster MAH, Sethi SP, Vermes D (1987) Optimal capacity expansion under uncertainty. *Adv. Appl. Probability* 19(1):156–176.

EPA (2000) *Incentives for Self-Policing: Discovery, Disclosure, Correction and Prevention of Violations* (U.S. Environmental Protection Agency, Washington, DC).

Ewing J (2016) VW presentation in '06 showed how to foil emissions tests. Accessed October 24, 2016, http://www.nytimes.com/2016/04/27/business/international/vw-presentation-in-06-showed-how-to-foil-emissions-tests.html?_r=0.

Hadler T (2013) Reward supplier honesty, not performance. Accessed June 8, 2023, https://nbs.net/reward-supplier-honesty-not-performance/.

Harrington W (1988) Enforcement leverage when penalties are restricted. *J. Public Econom.* 37(1):29–53.

Hwang I, Radhakrishnan S, Su L (2006) Vendor certification and appraisal: Implications for supplier quality. *Management Sci.* 52(10):1472–1482.

Jones LP, Tandon P, Vogelsang I (1990) *Selling Public Enterprise* (MIT Press, Cambridge, MA).

Kim S-H (2015) Time to come clean? Disclosure and inspection policies for green production. *Oper. Res.* 63(1):1–20.

Kim Y, Xu Y (2023) Operational risk management: Optimal inspection policy. *Management Sci.* 70(6):4087–4104.

Lacker JM, Weinberg JA (1989) Optimal contracts under costly state falsification. *J. Political Econom.* 97(6):1345–1363.

Laffont J-J, Tirole J (1993) *A Theory of Incentives in Procurement and Regulation* (MIT Press, Cambridge, MA).

Levi R, Singhvi S, Zheng Y (2019) Economically motivated adulteration in farming supply chains. *Management Sci.* 66(1):209–226.

Li H, Zhang H, Fine CH (2013) Dynamic business share allocation in a supply chain with competing suppliers. *Oper. Res.* 61(2):280–297.

Ljungqvist L, Sargent TJ (2004) *Recursive Macroeconomic Theory* (MIT Press, Cambridge, MA).

Lyon TP, Maxwell JW (2003) Self-regulation, taxation and public voluntary environmental agreements. *J. Public Econom.* 87(7–8):1453–1486.

Malik AS (1990) Avoidance, screening and optimum enforcement. *RAND J. Econom.* 21(3):341–353.

Mookherjee D, Png I (1989) Optimal auditing, insurance, and redistribution. *Quart. J. Econom.* 104(2):399–415.

Myerson R (1979) Incentive compatibility and the bargaining problem. *Econometrica* 47(1):61–73.

Myerson RB (2012) A model of moral-hazard credit cycles. *J. Political Econom.* 120(5):847–878.

Parmigiani G (1993) Scheduling inspections in reliability. Basu AP, ed. *Advances in Reliability* (North Holland, Amsterdam), 303–319.

Plambeck E, Taylor T (2016) Supplier evasion of a buyer's audit: Implications for motivating supplier social and environmental responsibility. *Manufacturing Service Oper. Management* 18(2):184–197.

Polinsky AM, Shavell S (2000) The economic theory of public enforcement of law. *J. Econom. Literature* 38(1):45–76.

Ravikumar B, Zhang Y (2012) Optimal auditing insurance in a dynamic model of tax compliance. *Theoretical Econom.* 7(2):241–282.

Sannikov Y (2008) A continuous-time version of the principal-agent problem. *Rev. Econom. Stud.* 75(3):957–984.

Spear S, Srivastava S (1987) On repeated moral hazard with discounting. *Rev. Econom. Stud.* 54(4):599–617.

Sun P, Tian F (2018) Optimal contract to induce continued effort. *Management Sci.* 64(9):4193–4217.

Tian F, Zhang F, Sun P, Duenyas I (2023) Dynamic moral hazard with adverse selection – Probation, sign-on bonus and delayed payment. Working paper, Duke University, Durham, NC.

Townsend RM (1979) Optimal contracts and competitive markets with costly state verification. *J. Econom. Theory* 21(2):265–293.

Townsend RM (1988) Information constrained insurance: The revelation principle extended. *J. Monetary Econom.* 21(2–3):411–450.

Varas F, Marinovics I, Skrzypacz A (2020) Random inspections and periodic reviews: Optimal dynamic monitoring. *Rev. Econom. Stud.* 87:2893–2937.

Wang S, Sun P, de Véricourt F (2016) Inducing environmental disclosures: A dynamic mechanism design approach. *Oper. Res.* 64(2):371–389.

Shouqiang Wang is an associate professor of operations management at Naveen Jindal School of Management, the University of Texas at Dallas. He has broad research interests in strategic decision-making problems in various business sectors, as well as areas of public interests. He draws methodological tools from and contributes to the intersection between operations research/management science and microeconomics.

Francis de Véricourt is a professor of management science and the Joachim Faber chair in business and technology at ESMT, Berlin, Germany. His research interest is in the area of analytics, operations and decision science, with applications in healthcare, sustainability and human–artificial intelligence interaction.

Peng Sun is a J.B. Fuqua professor of business administration at the Fuqua School of Business, Duke University. His research interests in recent years are focused on mathematical analysis of dynamic incentive management in policies and businesses.