Dynamic Contracting and Corporate Tax Strategies*

Juan F. Imbet† Université Paris Dauphine

Marcelo Ortiz ‡ Universitat Pompeu Fabra

Vincent Tena § Université Paris Dauphine

June 18, 2024

Abstract

We use a dynamic moral hazard model in which a firm's owner delegates the corporate tax strategy to an agent in an institutional setting with random inspections by tax officers. The principal cannot observe either the underlying gross profits or the agent's efforts to reduce tax expenses. The inspection date is random and might not occur. Upon inspection, illegal strategies are detected and penalized, but the authorities cannot observe how much as been evaded. The optimal dynamic contract consists of a terminal compensation and a tax reduction strategy. We find that even in contexts without inspections, a risk-aversed owner will not contract tax strategies if the agent's effort costs or the gross profit volatility are too high to compensate for her risk aversion. For contexts with inspections, the optimal compensation includes an additional risk premium for bearing the inspection risk and a lump-sum loss contingent on an inspection. The optimal tax strategy becomes less aggressive over time as the expected penalty increases. The dynamic nature of the model allows for quantitative evaluations. Using calibrated parameters from corporate taxation in the U.S., the model indicates that the agent's compensation for implementing the tax strategy corresponds to 37.8% percent of the expected benefits derived from the tax strategy.

Keywords: Continuous-time contracting, contract theory, tax strategy, tax evasion

JEL Codes: D82, D86, H21, H26, M52.

*Juan Imbet acknowledges support from the 2022 PSL-Junior fellow, the Institute Europlace de Finance, and the QMI Chair at Paris-Dauphine PSL. Marcelo Ortiz acknowledges financial support from the Spanish Ministry of Economy and Competitiveness through the project PID2020-115660GB-I00 and the Severo Ochoa Program for Centers of Excellence in R&D (CEX2019-000915-S). All remaining errors are our own.

†Department of Finance, Université Paris Dauphine and PSL Research University, Pl. du Maréchal de Lattre de Tassigny, 75016 Paris. Email: juan.imbet@dauphine.psl.eu

‡Universitat Pompeu Fabra, UPF - Barcelona School of Management, and Barcelona School of Economics, Spain. Email: marcelo.ortizm@upf.edu

§Department of Finance, Université Paris Dauphine and PSL Research University, Pl. du Maréchal de Lattre de Tassigny, 75016 Paris. Email: vincent.tena@dauphine.psl.eu (Corresponding Author)

1 Introduction

A large body of research on tax planning has documented that corporations pursue various tax-reduction strategies, ranging from benign tax-advantaged investments to aggressive illegal activities. Managers in charge of designing and implementing such strategies will likely trade off the benefit of the additional compensation for this effort with the potential personal economic costs. While some high-profile cases show managers following an active tax minimization strategy, it is puzzling that this phenomenon is not more prevalent given the limited cost of requesting legal tax reductions and the low probability of detecting tax evasion (Weisbach, 2002; Desai and Dharmapala, 2006). We study the extent and intensity of corporate tax strategies from a dynamic contract theory perspective. In particular, we aim to address the following questions: (1) Why would a firm's owner choose not to contract a tax reduction strategy? (2) How does the optimal contract incentivize the manager to avoid taxes across time? (3) How large is the manager's compensation in comparison to the tax reductions?

Our model considers an agent (manager) hired by the principal (firm's owner) to continuously implement a tax strategy for reducing tax expenses over the gross profits generated by a finite-time project. Given that the agent has superior information on the realized gross profits and feasible tax strategies, the principal cannot monitor the agent's effort to reduce tax obligations. The tax strategy splits the gross profit into two components, leaving one untaxed. The untaxed profits are paid back to the principal net of payments to third parties needed to implement the tax strategy (i.e., external legal advisors). The agent incurs private quadratic effort costs for implementing this strategy, reflecting the increasing difficulty of legally reducing gross income or concealing larger tax evasions. The tax authority conducts a random inspection during the project. The inspection only reveals tax reductions at the time of the audit. The inspection does not alter the legal tax strategies. However, in the case of detecting illegal strategies (i.e., tax evasion), it triggers an economic sanction, and taxes cannot be reduced for the rest of the project. The timing of inspections follows an exponential distribution, making the occurrence unpredictable but quantifiable. The agent is compensated at the end of the project after reporting the net profits to the principal.

We start by addressing the optimal dynamic contract in a baseline scenario without tax authority inspections, where the agent's superior information on gross profits and potential tax deductions creates a moral hazard problem. This scenario can be interpreted as companies engaging in lawful tax practices with the confidence that they will not be questioned by tax authorities. It can also be seen as situations with ineffective or corrupt authorities where the lack of inspection threat does not discourage illegal behavior. We analyze the initial conditions under which the owner optimally decides to forgo the benefits of tax reduction strategies, something commonly known as the under-sheltering puzzle (Weisbach, 2002). The marginal benefit of avoiding one dollar must compensate for the marginal tax savings plus an additional factor that increases with the marginal effort costs and the principal's risk premium. Therefore, a risk-aversed principal will not engage in feasible tax avoidance if the managerial effort costs or the volatility of the gross profits is too high. Additionally, while the agent's risk aversion does not affect this initial condition, it influences tax strayegy aggressivity. Specifically, delegating tax strategies to a more risk-averse agent diminishes the optimal tax expense reduction.

Next, we focus on scenarios with inspection risk. In our setting, inspections detect illegal tax strategies with certainty. However, the tax authority is unable to trace back the strategy. When a tax crime is detected, the penalty is calculated as the amount of tax evasion detected on the date of inspection multiplied by the length of the contract. We allow this sanction to vary based on a penalty factor. A lower penalty factor reduces the payment to reflect favorable legal settlements in tax disputes, while a large factor can be interpreted as including additional civil and criminal penalties for tax crimes. Compared to the baseline case, the optimal compensation grows linearly faster to account for the risk premium paid for bearing inspection risk. Additionally, this optimal contract includes a lump-sum loss for the agent contingent on the occurrence of an inspection. The optimal illegal tax strategy is gradually less aggressive over time as the expected penalty increases.

We then analyze two characteristics of the optimal contract: its enforceability and sensibility to different penalty structures. First, Chen and Chu (2005) argue that contracts based on illegal clauses, such as those for contracting tax evasion, are not legally enforceable in court. Therefore, the principal can renege on the contract and refuse to compensate the agent. The agent would refuse the agreement because she knows no court would uphold an illegal contract. This intuition stands in contrast with the widespread practice of tax evasion, even in countries with strict tax enforcement (Slemrod, 2007), and with the evidence of the enduring contractual relationships between large firms and tax consulting companies (ICIJ, 2023; Reuters, 2019). We use our model to establish the theoretical possibility of illegal tax arrangements being self-enforced based on the repetitive nature of tax strategy over successive periods. This repetition fosters a mechanism where the expectation of future interactions provides a basis for compliance beyond legal constraints. Second, we explore the implications of an alternative penalty structure where the agent, instead of the principal, is penalized in case of detecting tax evasion. Crocker and Slemrod (2005) argue that penalties imposed directly on the agent are more effective in curbing tax evasion than those imposed on shareholders. In our setting, the aggressiveness of the tax strategy is unaffected by who carries the penalty to the extent that both the agent and the principal have the same level of risk aversion.

Finally, we rely on the dynamic nature of our model to calibrate some unobservable parameters. In particular, we calibrate the agent's effort cost and the penalty factor to match a set of moments related to corporate tax revenues, penalties, and the fraction of taxes that are expected but not collected from the Internal Revenue Service (IRS). Our analysis suggests that the agent's expected compensation corresponds to 37.8% percent of the expected benefits derived from tax strategies. This fraction largely increases with the tax rates, reflecting the value of the tax strategy for the principal. However, this fraction decreases in tight tax regulatory environments (e.g., higher inspection probability and larger penalties) as these institutional features reduce the principal's benefits from the tax strategy and, thus, are detrimental to optimal compensation.

This paper makes several contributions to the literature on corporate tax strategy and governance. First, we connect with seminal studies using static contract theory to analyze optimal tax avoidance contracts (Chen and Chu, 2005; Crocker and Slemrod, 2005). We build on this literature by developing a dynamic contracting model with random inspection risk. We use this model to explain conditions under which the principal refrains from pursuing tax strategies and to describe the dynamics of optimal compensation and tax aggressiveness. Additionally, we study how illegal clauses can be self-enforced in a dynamic game and that the effectiveness of the penalty structure depends on the relative risk aversion of the agent and the principal. Second, a large amount of empirical evidence suggests that corporate tax strategy benefits shareholders and managers (Desai et al., 2007; Bennedsen and Zeume, 2018). More recently, business commentators have highlighted how companies reward managers for pursuing tax avoidance, often resulting in executives' compensation exceeding corporate tax expenses (The Guardian, 2024; Anderson et al., 2024). In line with this discussion, our calibration analysis suggests that a relevant fraction of the gains derived from the tax strategy is paid to the manager. Lastly, from a methodological perspective, we build on the work of Martin and Villeneuve (2023) by integrating a random exogenous shock into a principal-agent model with moral hazard where a Brownian component drives the firm's output.

2 Model

In this section, we develop a moral hazard model for delegating tax strategy in a setting with a random inspection by the tax authority. Our model embeds the inspections as exogenous shocks (Martin and Villeneuve, 2023) into a classic dynamic moral hazard framework (Holmstrom and Milgrom, 1987).

Time is continuous between $[0, T]$ with $T > 0$, and the firm can be considered as a project with terminal date T. The firm's owner (hereafter, the principal) delegates to an agent (i.e., the manager) the management of the firm's tax strategy. Both the principal and the agent have CARA utility functions on terminal wealth

$$U_P(x) := -\exp(-\gamma_P x) \quad \text{and} \quad U_A(x) := -\exp(-\gamma_A x),$$

with absolute risk aversion parameters $\gamma_P > 0$ and $\gamma_A > 0$.

We depart from the canonical model of Holmstrom and Milgrom (1987) by introducing two novel features: First, we introduce tax evasion incentives. The agent's action strategically splits gross profits into two parts, deducting one from tax obligations. Second, we incorporate a random inspection process imposed by tax authorities that can trigger economic penalties when tax evasion is detected.

2.1 Tax strategy

We define the complete probability space as $(\Omega, \mathcal{G}, \mathbb{P}^0)$ (for a detailed definition, see Appendix A). Firm's gross profits, denoted by $I = (I_t)_{t \leq T}$, evolve exogenously according to the following dynamics:

$$dI_t = \mu dt + \sigma dB_t, \tag{1}$$

$$I_T = I_0 + \int_0^T \mu dt + \int_0^T \sigma dB_t \tag{2}$$

where B represents a standard one-dimensional Brownian motion.[1] This Brownian motion is defined on the filtration $\mathbb{F} = (\mathcal{F}_t)_{t \geq 0}$.

At the initial time $t = 0$, the principal decides whether to contract the agent to perform a *tax strategy*, which involves a series of actions executed continuously over time to reduce tax expenses, represented as $a = (a_t)_{t \leq T}$. The tax strategy can be legal, involving lawful methods to reduce tax expenses, or illegal, involving tax evasion tactics. We define the set of feasible tax strategies as:

$$\mathcal{A} = \{\{a_t\}_{t=0}^{T} : a_t \text{ is } \mathbb{F} - \text{predictable and } a_t \in [0, \bar{a}] \text{ for } \bar{a} > 0\}. \tag{3}$$

The agent incurs a quadratic effort cost for implementing the tax strategy $c(a_t) = \frac{c}{2}a_t^2 \geq 0$ at any date $t \in [0, T]$. This cost reflects the difficulties the agent must overcome to reduce taxable profits. As the tax code options become limited, it becomes progressively harder for managers to increase the amount of tax reductions through legal means. Similarly, as the magnitude of tax evasion that needs to be hidden increases, managers have to put in more effort to conceal it.

2.2 Tax authority and random inspections

We assume the distribution of the tax authority's inspections across time is exogenous and construct the single-jump process N as follows:[2]

$$N_t = \begin{cases} 0 & \text{if no inspection has occurred by time } t \\ 1 & \text{if an inspection has occurred by time } t \end{cases}$$

We assume this single-jump process has an intensity λ. Hence, given that no inspection has occurred before date t, the probability that an inspection will occur between t and $t + dt$ is λdt. Furthermore, the probability that a firm is inspected at time θ before terminal date T is

[1] For tractability, we assume that gross profits are unaffected by the tax strategy. In reality, tax strategies can influence investment decisions and have a heterogeneous impact across firms based on productivity and intangibles (Jacob 2022; Ortiz and Imbet 2023).

[2] Formally, the possibility of multiple audits in the same company could be represented by a Poisson process, reflecting the stochastic nature of audit occurrences. However, since any audit invariably detects illegal activities when they occur, only the first audit is relevant to our model. Therefore, we simplify the representation by considering only the first jump of the Poisson process. This reduction to a single-jump process is justified because subsequent audits become redundant once illegal activity is detected and stopped.

$\mathbb{P}(0 \leq \theta \leq T) = 1 - \exp(-\lambda T)$. Given the process N, we can create the following process M:

$$M_t = N_t - \lambda \int_0^t (1 - N_s) dt \tag{4}$$

which is a $\mathbb{G}$-compensated martingale, where $\mathbb{G} = \{\mathcal{G}_t, t \geq 0\}$ is the smallest complete right-continuous extension of $\mathbb{F}$ that makes θ a stopping time. Hence, $\mathbb{G}$ encompasses all pertinent information: this includes the history of the gross profits, encapsulated within the set $\mathbb{F} = (\mathcal{F}_t)_{t \geq 0}$, and knowledge regarding the occurrence of an inspection.

When a firm pursuing tax evasion is inspected, we assume three consequences. First, the evasion is invariably detected. Second, following the inspection, the company continues operating until T but cannot continue any tax strategy for the remaining duration of the project. Third, the tax authority imposes an economic penalty on the firm.

2.3 Taxable and total profits

The firm reports taxable profits $I^a = (I_t^a)_{t \leq T}$, which depend on a tax strategy $a = (a_t)_{t \leq T}$. These reported profits evolve as:

$$dI_t^a = (\mu - a_t(1 - N_t))dt + \sigma dB_t^a, \tag{5}$$

where B^a is a Brownian motion under the probability measure $\mathbb{P}^a$ equivalent to $\mathbb{P}^0$, formally defined in Appendix A. The deterministic component of the gross profit increases with the expected earnings μ and decreases with tax avoidance a_t. Moreover, σ is a positive constant that denotes the firm's gross profits volatility. A fraction $1 - \delta$ of the savings from the tax strategy is paid to intermediaries. In legal cases, these expenses are driven by the fees of consultants and financial services. In illegal cases, the evaded taxes need to be legitimized before their use, meaning that a fraction of the concealed profits is paid to intermediaries to hide the illegal nature of the money (e.g., money laundering). The total net profits denoted as $X^a = (X_t^a)_{t \leq T}$, represent the after-tax reported earnings combined with the tax savings net of intermediaries' fees. The change in total net profits at time t follows:

$$\begin{aligned} dX_t^a &= (1-\tau)dI_t^a + \delta a_t(1-N_t)dt \\ &= ((1-\tau)\mu + (\delta+\tau-1)a_t(1-N_t))dt + (1-\tau)\sigma dB_t^a. \end{aligned} \tag{6}$$

Where τ is the corporate tax rate. The principal will receive the total net profits from the project, represented by X_T^a, at the final time T. This final *lump-sum* gain includes the benefits from the tax strategy a accumulated up to time T. The principal can only continuously monitor I^a and not X^a, much like the tax authorities. Consequently, the principal only observes the accumulated gains from the tax strategy when the activity terminates or when an inspection occurs.

2.4 Optimal Contracting Problem

The principal compensates the agent with a terminal payment W contingent on the firm's accumulated reported profits I_t^a and the advent of an inspection before the date T. The contract (W, a) comprises a $\mathbb{G}$-measurable wage W that satisfies the usual integrability condition and a tax strategy plan a. For wage W, the agent maximizes:

$$\sup_{a\in\mathcal{A}} E^a\left[U_A\left(W - \frac{c}{2}\int_0^T a_s^2(1-N_s)ds\right)\right]. \tag{7}$$

where $\mathbb{E}^a$ is the expectation under the measure $\mathbb{P}^a$. An incentive-compatible wage W ensures a tax strategy $a^*(W) \in \mathcal{A}$ that satisfies:

$$E^{a^*(W)}\left[U_A\left(W - \frac{c}{2}\int_0^T a_s^*(W)^2(1-N_s)ds\right)\right] \geq E^a\left[U_A\left(W - \frac{c}{2}\int_0^T a_s^2(1-N_s)ds\right)\right] \tag{8}$$

To align the agent's actions with her objectives, the principal offers an incentive-compatible wage scheme, ensuring the agent's best response is $a^*(W)$. We denote the set of best replies to W by $\mathcal{A}^*(W)$ and the set of incentive-compatible wages as $\mathcal{W}_{IC}$. At time $t = 0$, the principal faces a binary decision: to hire the agent to undertake tax strategy throughout the interval $[0, T]$ or to abstain from the tax strategy. The primary goal of the principal is to maximize her expected utility.

This optimization problem can be formally articulated as follows:

$$\sup\left\{E^0\left[U_P\left(X_T^0\right)\right];\sup_{W\in\mathcal{W}_{\mathrm{IC}}}\sup_{a^*\in\mathcal{A}^*(W)}E^{a^*}\left[U_P\left(X_T^{a^*}-W-P(\theta,a)\mathbb{1}_{\theta<T}\right)\right]\right\},\tag{9}$$

s.t.

$$\mathbb{E}^{a^*}\left[U_A\left(W-\frac{c}{2}\int_0^T a^*(W)^2\left(1-N_s\right)ds\right)\right]\geq U_A\left(y_{PC}\right),\tag{10}$$

where Equation (10) is the participation constraint, y_{PC} is the reservation value of the agent, and $P(\theta, a)$ is the penalty incurred by the firm upon being detected during the inspection at date θ.[3] In this formulation, the term $E^0\left[U_P\left(X_T^0\right)\right]$ represents the expected utility derived from choosing not to engage in tax strategy. It is quantified as:

$$E^0\left[U_P\left(X_T^0\right)\right]=-\exp\left[-\gamma_P(1-\tau)T\left(\mu-\frac{\gamma_P\sigma^2(1-\tau)}{2}\right)\right].\tag{11}$$

3 Solution

3.1 No inspection risk

In this section, we focus on cases where the tax authorities do not penalize the firm for its tax strategy, either because the strategy employs lawful means or because of a weak institutional environment. Consequently, corporate tax strategies increase the firm's total profitability, and there is no threat of penalization.

In this case, moral hazard still arises due to the manager's exclusive knowledge of the company's potential tax deductions and the underlying gross profits. The firm's total net profits follow the dynamics:

$$dX_t^a=((1-\tau)\mu+(\delta+\tau-1)a_t)dt+(1-\tau)\sigma dB_t^a,\tag{12}$$

Both the principal and the tax authority can continuously monitor the reported profits I^a given in Equation (5). In contrast, total net profits X^a are not subject to continuous scrutiny by

[3] Further details about the penalty structure are provided in Section 3.2.

any party except the agent, and it is only revealed to the principal at the end of the project or in case of an inspection. The continuous unobservability of the total net profits generates moral hazard friction because it prevents the clear distinction between the tax strategy and the stochastic fluctuations of the firm's profit arising from the Brownian motion B_t^a.

Here, we derive the optimal contract that consists of a triplet $\Pi^* = (y_0^*, \tilde{Z}^*, a^*)$. The contract has three elements: a fixed compensation component (y_0^*), a compensation component sensitive to reported profits $(\tilde{Z}^*)$, and a tax strategy plan (a^*). Let's denote J_t the agent's continuation payoff seen at date t, which is the difference between the compensation W and the cumulative effort costs from date t to termination date T

$$J_t := W - \frac{c}{2}\int_t^T a_s^2 ds. \tag{13}$$

We assume that the agent's continuation payoff is an Itô process and follows the dynamics:

$$dJ_t = \tilde{H}_t dt + \tilde{Z}_t \sigma dB_t. \tag{14}$$

Next, we consider the total utility of the agent seen at date t that can be expressed as:

$$V_A(t) := U_A\left(J_t - \frac{c}{2}\int_0^t a_s^2 ds\right). \tag{15}$$

We apply Itô's formula to $V_A(t)$, then use the fact that $B_t^a = B_t + \int_0^t \frac{a_s}{\sigma} \mathrm{d}s$ and we get

$$\frac{dV_A(t)}{V_A(t)} = \gamma_A\left[\left(-\tilde{H}_t + \frac{\gamma_A \sigma^2}{2}\tilde{Z}_t^2 + c\frac{a_t^2}{2} + a_t \tilde{Z}_t\right)dt - \tilde{Z}_t \sigma dB_t^a\right]. \tag{16}$$

Then, the martingale optimality principle provides a yardstick to determine the agent's ideal strategy. It implies that $\{V_A(t)_{s=0}^T\}$ is a supermartingale under the probability measure $\mathbb{P}^a$ and a martingale under $\mathbb{P}^{a^*}$. Hence,

$$-\tilde{H}_t + \frac{\gamma_A \sigma^2}{2}\tilde{Z}_t^2 + c\frac{(a_t^*)^2}{2} + a_t^* \tilde{Z}_t = 0 \quad \text{a.s.}, \tag{17}$$

only holds when $a_t = a_t^* \quad \forall t \leq T$.

Because the tax strategy is costly, the principal wants to induce the optimal lowest action, i.e., equality (17) holds. This leads to:

$$\tilde{H} = \inf_a \{\frac{1}{2}ca^2 + a\tilde{Z}\} + \frac{1}{2}\gamma_A \sigma^2 \tilde{Z}^2. \tag{18}$$

Using first-order condition, we get

$$a^* = \begin{cases} 0 & \text{when } \tilde{Z} > 0, \\ \bar{a} & \text{when } \tilde{Z} \leq -\bar{a}c, \\ -\frac{\tilde{Z}}{c} & \text{otherwise.} \end{cases} \tag{19}$$

Therefore, in the interior case where $a^* = -\frac{\tilde{Z}}{c}$,

$$\tilde{H}^{a^*} = \frac{1}{2}\left(\gamma_A \sigma^2 - \frac{1}{c}\right)\tilde{Z}^2, \tag{20}$$

The agent's continuation payoff adheres to the dynamics:

$$dJ_t = \frac{1}{2}\left(\gamma_A \sigma^2 - \frac{1}{c}\right)\tilde{Z}_t^2 dt + \tilde{Z}_t \sigma dB_t, \tag{21}$$

$$\text{with} \quad J_T = W. \tag{22}$$

This equation represents the terminal payment W. The most cost-effective method for ensuring the agent's participation is to set $J_0 = y_{PC}$. As a consequence, the dynamics of the terminal compensation follow:

$$W = y_{PC} + \int_0^T \frac{1}{2}\left(\gamma_A \sigma^2 - \frac{1}{c}\right)\tilde{Z}_t^2 dt + \int_0^T \tilde{Z}_t \sigma dB_t. \tag{23}$$

With the dynamics of W under the optimal contract being defined, our next objective is to characterize the optimal compensation sensitivity of the $\tilde{Z}$. To achieve this, we turn our attention to the principal's problem. The following theorem presents $\tilde{Z}^*$.

Theorem 1. *For the case with no inspection risk, in the interior case where* $a^* = -\frac{\tilde{Z}}{c}$*, the optimal* *compensation sensitivity* $\tilde{Z}^*$ *is given by:*

$$\tilde{Z}^* = \frac{\gamma_P \sigma^2 (1-\tau) - \frac{\delta+\tau-1}{c}}{\sigma^2(\gamma_P+\gamma_A) + \frac{1}{c}}. \tag{24}$$

Proof. See Appendix B.

Hence, the principal will contract the tax strategy if $\tilde{Z}^*$ is negative (i.e., $a^* > 0$ in Equation 19). As the denominator of $\tilde{Z}^*$ is positive, the following conditions should be met for pursuing the tax strategy:

$$\frac{\delta}{(1-\tau)} > 1 + c\sigma^2\gamma_P \tag{25}$$

Altogether, Equations (24) and (25) provide a contractual explanation for the "under sheltering puzzle" initiated by Weisbach (2002) and further explored by Desai and Dharmapala (2006) and McClure (2023). Equation (25) indicates that the principal will engage in tax avoidance if the marginal benefit of avoiding one dollar (δ) equals the marginal tax savings $(1-\tau)$ plus an additional factor that increases with the marginal effort costs (c) times the principal's risk premium. Interestingly, Equation (25) implies that even in contexts without inspection risk, a risk-aversed principal will not engage in feasible tax avoidance if the managerial effort costs or the gross profit volatility are too high. Additionally, while the agent's risk aversion, γ_A, does not affect the threshold for engaging in the tax strategy, it influences the intensive margin by decreasing the compensation sensitivity, as shown in Equation (24). Specifically, delegating tax strategies to a more risk-averse agent diminishes the optimal level of tax avoidance.

Under $\mathbb{P}^{a^*}$, the dynamic of the optimal terminal compensation is[4]

$$W^{\pi^*} = y_{PC} + \frac{1}{2}\left(\gamma_A\sigma^2 + \frac{1}{c}\right)(Z^*)^2\, T + \sigma Z^* B_T^{a^*}. \tag{26}$$

The terminal compensation includes a fixed component that ensures the agent's participation. The second and third elements in Equation 26 represent the compensation sensitivity to reported profits, which depend on the agent's risk aversion, effort cost, and the underlying gross profit

[4] We show in Appendix A that $B_t^{a^*} = B_t + \int_0^t \frac{a_s^*}{\sigma}(1-N_s)\,\mathrm{d}s, t \in [0,T]$ is a $\mathbb{G}$ -Brownian motion under $\mathbb{P}^{a^*}$, hence the change of sign in front of the term $\frac{1}{c}$ that appears between Eq. (23) and Eq. (26).

volatility.

3.2 Inspection risk

As discussed before, tax crimes are identified during an inspection, triggering an economic sanction. This penalty is defined as:

$$P(\theta, a) = \rho\tau\theta a_\theta, \tag{27}$$

where θ denotes the inspection time, and a_θ is the current level of tax evasion detected in the inspection. The constant ρ is a penalty factor. Low positive values of ρ reflect cases where the authorities offer a favorable settlement to stop a tax dispute, while high values indicate legal environments that penalize harder tax crimes. We assume the penalty is paid immediately after the inspection, and afterward, the firm cannot reduce taxes anymore. Hence, the principal's utility associated with the remaining profits at any date t after the inspection satisfies

$$U_P(L(T-t)), \tag{28}$$

$$\text{where } L = (1-\tau)\left(\mu - \frac{\gamma_P(1-\tau)\sigma^2}{2}\right). \tag{29}$$

Our model revolves around time and two state variables: the firm's total net profits, denoted by $X = \{X_t^a\}_{s=0}^T$ and the agent's continuation value, denoted by $W = \{W_t\}_{s=0}^T$. Hence, the principal's value function $v(t, x, w)$, is the maximum expected utility the principal can attain at time t, given a firm's total net profits x and an agent's continuation value w.

Before the inspection, while the agent implements the tax strategy, the firm's total net profits satisfy

$$dX_t^a = ((1-\tau)\mu + (\delta + \tau - 1)a_t(1 - N_t))dt + (1-\tau)\sigma dB_t^a, \tag{30}$$

And the agent's continuation value follows

$$W_t = y_{PC} + \int_0^t \tilde{Z}_s (1 - N_s) \sigma dB_s^a + \int_0^t \tilde{K}_s (1 - N_s) dN_s \\ + \int_0^t \left(\frac{\gamma_A}{2} \sigma^2 \tilde{Z}_s^2 + \frac{c}{2} a_s^2 (1 - N_s) + \frac{\lambda}{\gamma_A} \left[\exp\left(-\gamma_A \tilde{K}_s \right) - 1 \right] \right) (1 - N_s) ds. \tag{31}$$

This continuation value depends on the participation constraint and makes the agent's compensation sensitive to both the firm's reported profits through $\tilde{Z}$ and the contingent inspection via $\tilde{K}$. We will focus on the interior solution $a^* = -\tilde{Z}/c$. This illegal tax strategy is aligned with the principal's expectations (incentive compatibility) and maximizes the principal's expected value under the measure $\mathbb{P}^{a^*}$.

We introduce the Hamilton–Jacobi–Bellman (HJB) equation to delineate the principal's value function $v(t, x, w)$, which depends on the date t, the current level of reported profits x, and the current level of the agent's continuation value w. We leverage the separability property of the exponential utility function, enabling us to express $v(t, x, w)$ in the simplified form $U_P(x - w)\phi_0(t)$, where ϕ_0 is a function of time. Using this simplified form in the HJB equation, we find that ϕ_0 satisfies the following second-order partial differential equation (PDE) :

$$\phi_0'(t) = \inf_{(Z_t)_{t \leq T \wedge \theta}, (K_t)_{t \leq T \wedge \theta}} \Bigg\{ \gamma_P \phi_0(t) \Bigg((1 - \tau)\mu - \frac{\gamma_P}{2}(1 - \tau)^2 \sigma^2 + \frac{\lambda}{\gamma_A} \\ + \left(\frac{-(\delta + \tau - 1)}{c} + \gamma_P (1 - \tau)\sigma^2 \right) Z_t^* - \left(\sigma^2 (\gamma_A + \gamma_P) + \frac{1}{c} \right) \frac{(Z_t^*)^2}{2} \\ - \frac{\lambda}{\gamma_A} \exp\left\{ -\gamma_A K_t^* \right\} \Bigg) - \lambda \exp\left\{ \gamma_P \left(K_t^* - \frac{\rho \tau Z_t^* t}{c} - L(T - t) \right) \right\} \Bigg\}, \tag{32}$$

where $\phi_0(.)$ must meet the boundary conditions[5]:

$$\phi_0(T) = 1, \phi_0(0) = \exp(-\gamma_P * L * T). \tag{33}$$

This implies that under the condition $Z \leq -\bar{a} \times c$ (implying $a^* = -\frac{Z}{c}$), the optimal contract

[5] K_0 is zero since no jump has occurred at $t = 0$.

yields:

$$K_t^* = \frac{\log \phi_0(t) + \gamma_P(\frac{\rho\tau Z_t^* t}{c} + L(T-t))}{\gamma_A + \gamma_P}, \tag{34}$$

and $Z^* = (Z_t^*)_{t \leq T\wedge\theta}$ is such that:

$$\begin{aligned} 0 =&\phi_0(t)\left(-\frac{(\delta+\tau-1)}{c} + \gamma_P(1-\tau)\sigma^2 - \left(\sigma^2(\gamma_A+\gamma_P) + \frac{1}{c}\right) Z_t^*\right) \\ &+ \frac{\lambda\rho\tau t}{c}\phi_0(t)^{\frac{\gamma_P}{\gamma_A+\gamma_P}} \exp\left\{-\frac{\gamma_A\gamma_P}{\gamma_A+\gamma_P}\left(\frac{\rho\tau Z_t^* t}{c} + L(T-t)\right)\right\}. \end{aligned} \tag{35}$$

The following theorem summarizes our main result.

Theorem 2. *Assume that $-\bar{a}c < Z < 0$, so $a^* = -\frac{Z}{c}$, and let the controls (K^*, Z^*) satisfy the equations (34) and (35). Then $\pi^* = (y_{PC}, a^*, Z^*, K^*)$ parameterizes the optimal contract $\left(W_T^{\pi^*}, a^*\right)$.*

Proof. See Appendix B.

Finally, the optimal wage satisfies[6]

$$\begin{aligned} W^{\pi^*} =&y_{PC} + \int_0^{T\wedge\theta}\left[\frac{1}{2}\left(\sigma^2\gamma_A + \frac{1}{c}\right)(Z_s^*)^2\, ds + \sigma Z_s^* dB_s^*\right] \\ &+ \int_0^{T\wedge\theta}\left[\frac{\lambda}{\gamma_A}\left[\exp\left(-\gamma_A K_s^*\right) - 1\right]\right] ds + K_{\theta^-}^* \mathbb{1}_{\theta\leq T}, \end{aligned} \tag{36}$$

Comparing this optimal compensation to the case without inspection risk in Equation 26, we see that both contracts consider a minimum wage y_{PC} that ensures the agent's participation. Additionally, both include a component sensitive to reported profits, but now, the size of this component depends on the occurrence of an inspection before the terminal date. If that is the case, then this component will stop increasing after the inspection date. Furthermore, in Equation 36, the optimal compensation includes two additional elements to account for the inspection risk, $\int_0^{T\wedge\theta} \lambda\left[\exp\left(-\gamma_A K_s^*\right) - 1\right] ds + K_{\theta^-}^* \mathbb{1}_{\theta\leq T}$[7]. The first element is the compensation for bearing inspection risk. The second represents the contingent loss in the compensation triggered by an inspection. Both additional elements are jointly determined in the optimal contract sensitivity to

[6]The notation t^- is the limit approaching t from the left. It is such that $t^- = \lim_{s\uparrow t} s$, where $s < t$, and denotes the moment immediately preceding t, where no time has elapsed between t^- and t, but any event occurring at time t has not yet happened at t^-.

[7]The notation $T\wedge\theta$ represents the minimum of T and θ, indicating that the evasion activity ceases upon the audit's occurrence. By construction, we have $\int_0^{T\wedge\theta} a^*(s)ds = \int_0^T a^*(s)(1-N_s)ds$

inspection (K_t^*).

3.3 Analysis of the Optimal Policies

3.3.1 Tax Strategy

Figure 1 illustrates the dynamics of tax strategies for different tax rates. High tax rates encourage contracts that promote intense tax strategies. When there are no inspection risks (legal strategies or corrupt authorities), the aggressiveness of the strategy remains stable over time. However, for cases with inspection risk, the optimal tax aggressiveness decreases as time passes because the expected penalties increase over time. To further understand the dynamics of the illegal tax strategy, Figure 2 illustrates the evolution of the two forces behind the expected penalty cost: the probability of being inspected and the respective penalty. On the one hand, a longer evasion is penalized more severely. On the other hand, as shown in Figure 1, optimal evasion decreases over time. While the first factor dominates initially, the latter becomes more prominent at the end of the contract. The combination of both factors generates a concave dynamic for the penalty.

3.3.2 Wage Dynamics

Figure 3 illustrates the dynamics of the optimal sensitivity of the compensation to the inspection risk, denoted by $K^* = (K_t^*)_{0 \leq t \leq T}$. Its analysis reveals several important aspects. Initially, at time $t = 0$, the optimal sensitivity to inspection, $K^*(t)$, must be zero to maintain the integrity of the model for two reasons. First, the fixed compensation y_{PC} at the onset of the contract implies that $K^*(0)$ is equal to zero to meet the participation constraint. Second, as the penalty is proportional to the contract length, at $t = 0$, any evasion detected at this point received a null penalty. Then, the evolution of K^* reflects the risk premium paid to the manager according to the evolution of the penalty $P(a_t^*, t)$. The concave nature of the penalty over time, as illustrated in Figure 2, is mirrored in Figure 3 by the convex shape of the compensation sensitivity to audit risk K^*.

Figure 4 illustrates the process of the optimal terminal compensation for a firm with and without inspection risk. Up to time $\theta \wedge T$, the compensation for the agent facing inspection risk is larger to offset the burden of bearing inspection risk. If an inspection occurs at $\theta < T$, the contingent loss for the inspection is triggered, and its compensation remains constant until the end of the project because no additional tax strategy can be pursued.

3.4 Additional analyses of the optimal dynamic contract

3.4.1 Self-Enforceability of the Optimal Contract

This section shows that the optimal dynamic contract can be self-enforceable in our setting. Consider an initial contract denoted by Π_0^* and ending at $t = 0$. The principal's decision to compensate this agent hinges on the following condition

$$E_0[U_p(X_T^{\Pi_1^*} - W_T^{\Pi_1^*} - P(a^{\Pi_1^*}, \theta)\mathbb{1}_{0\leq\theta\leq T} - W_{0^-}^{\Pi_0^*})] \geq E_0[U_p(X_T^0)]. \tag{37}$$

Here, $X_T^{\Pi_1^*}$ and $W_T^{\Pi_1^*}$ represent the terminal net profit and wage under a new contract Π_1^*, respectively. $P(a^{\Pi_1^*}, \theta)$ denotes the penalty function, and $W_{0^-}^{\Pi_0^*}$ symbolizes the wage promised for the just-concluded contract Π_0^*. This condition implies that the principal will fulfill the contract's financial obligations if it is expected to be more beneficial than the alternative of defaulting on payment and being unable to implement tax strategies in the upcoming period. Assuming stationarity on this condition, i.e., all agents are identical for the purpose of the contract, then Π_1^* and Π_0^* are also identical (hereafter denoted by Π^*). Using CARA utility function properties, we simplify this condition by focusing on the principal's certainty equivalent. This leads us to the following statement:

Condition 1. *The contract is self-enforceable if:*

$$W_{0^-} \leq E_0(\chi_1^*) - \frac{\gamma_P}{2} Var(\chi_1^*) - \left(E_0(\chi_0^*) - \frac{\gamma_P}{2} Var(\chi_0^*)\right). \tag{38}$$

where

$$\chi_1^* := X_T^{\Pi^*} - W_T^{\Pi^*} - P(a^{\Pi^*}, \theta)\mathbb{1}_{0\leq\theta\leq T}, \tag{39}$$

$$\chi_0^* := X_T^0. \tag{40}$$

Chen and Chu (2005) suggests that the folk theorem's effectiveness for contracting tax evasion hinges on an endless sequence of interactions between the principal and the agent. Our approach differs as it considers the principal engaging with a succession of agents, where each is functionally indistinguishable from the others in terms of their role and impact on the contract. This arrangement effectively creates a scenario that mirrors the conditions of an indefinite series of interactions, as postulated by the folk theorem, thereby preserving the theorem's relevance and applicability in our model.

3.4.2 Alternative Penalisation

Next, we explore the implications of an alternative penalty structure where the agent, instead of the principal, is penalized for tax evasion. We modify the agent's continuation value to reflect this change and adjust the Hamilton–Jacobi–Bellman (HJB) equation accordingly.

In the context where the agent bears the penalty, the continuation value W_t is modified as follows:

$$\begin{aligned} W_t = y_{PC} &+ \int_0^t Z_s\,(1-N_s)\,\sigma dB_s^a + \int_0^t (K_s - \rho\tau a_s s)\,(1-N_s)\,dN_s \\ &+ \int_0^t \left(\frac{\gamma_A}{2}\sigma^2 Z_s^2 + \frac{c}{2}a_s^2\,(1-N_s) + \frac{\lambda}{\gamma_A}\left[\exp\left(-\gamma_A(K_s - \rho\tau a_s s)\right) - 1\right]\right)(1-N_s)ds. \end{aligned} \tag{41}$$

Employing similar techniques as used in the main model, the HJB equation for this scenario is:

$$\begin{aligned} \phi_0'(t) = \inf_{(Z_t)_{t\le T\wedge\theta},(K_t)_{t\le T\wedge\theta}} &\Bigg\{\gamma_P\phi_0(t)\Bigg((1-\tau)\mu - \frac{\gamma_P}{2}(1-\tau)^2\sigma^2 + \frac{\lambda}{\gamma_A} \\ &+ \left(\frac{-(\delta+\tau-1)}{c} + \gamma_P(1-\tau)\sigma^2\right)Z_t^* - \left(\sigma^2(\gamma_A+\gamma_P) + \frac{1}{c}\right)\frac{(Z_t^*)^2}{2} \\ &- \frac{\lambda}{\gamma_A}\exp\left\{-\gamma_A(K_t^* + \frac{\rho\tau Z_t^* t}{c})\right\}\Bigg) - \lambda\exp\left\{\gamma_P\left(K_t^* - L(T-t)\right)\right\}\Bigg\}, \end{aligned} \tag{42}$$

The optimal controls in this scenario are determined as follows:

$$K_t^* = \frac{\log\phi_0(t) + (\gamma_A\frac{\rho\tau Z_t^* t}{c} + \gamma_P L(T-t))}{\gamma_A+\gamma_P}, \tag{43}$$

$$\begin{aligned} 0 = &\left(-\frac{(\delta+\tau-1)}{c} + \gamma_P(1-\tau)\sigma^2 - \left(\sigma^2(\gamma_A+\gamma_P) + \frac{1}{c}\right)Z_t^*\right) \\ &+ \frac{\lambda\rho\tau t}{c}\phi_0(t)^{\frac{-\gamma_A}{\gamma_A+\gamma_P}}\exp\left\{-\frac{\gamma_A\gamma_P}{\gamma_A+\gamma_P}\left(\frac{\rho\tau Z_t^* t}{c} + L(T-t)\right)\right\}. \end{aligned} \tag{44}$$

Under the condition that $\gamma_A = \gamma_P$, the optimal contract in this scenario is the same as in our main case. This equivalence suggests that penalizing the agent or the firm's owner yields similar outcomes in terms of the optimal contract structure.

It is noteworthy that this result contrasts with the findings in Crocker and Slemrod (2005). They argue that penalties imposed directly on the agent are more effective in curbing tax evasion than those imposed on shareholders. This is because the agent's decision to engage in legal or illegal activities creates an information asymmetry outside the principal's control, leading to the "non-equivalency" of penalties.

In contrast, our model presents a different scenario. Here, the principal delegates tax evasion tasks to the agent, with the information asymmetry concerning the extent of unreported profits rather than the illegal nature of the activity itself. This variation leads to a divergent outcome: if both the agent and the principal share equivalent levels of risk aversion, penalizing the agent becomes the same as penalizing the principal.

4 Calibration

In this section, we calibrate the model to address our last research question regarding the magnitude of the optimal compensation relative to the size of tax reductions. For this analysis, we borrow some parameter estimates from recent studies. From Bertomeu et al. (2022) we obtain the risk aversion coefficients (γ_A and γ_P) as their mid-range estimate (0.61/2). From Nikolov et al. (2021) we obtain the profit-generating process parameters as a fraction of total assets for Compustat firms ($\mu = 0.066$ and $\sigma = 0.034$). From Hoopes et al. (2012) we obtain the average percentage of Compustat firms inspected by the Internal Revenue Service (IRS) in a year (29%), which we use to set the inspection frequency λ.[8] Based on McClure (2023), we set the intermediation expenses to be $1 - \delta = 0.064$. The corporate income tax rate is $\tau = 0.2581$, representing the average combined federal and state statutory rate in 2022 (Enache 2022). Finally, we assume perfect competition in the labor market ($y_{pc} = 0$), and set an upper bar on a, which is never binding. Finally, to match the duration of the parameters, we solve the model for $T = 1$.

To calibrate the remaining parameters, c and ρ , we match two simulated moments with their empirical counterparts. The calibration is conducted by minimizing a weighted sum of squares

[8]To match the 29% probability of inspection in a year, we define $\lambda = -\log(1 - 0.29)$.

between the simulated moments of the model and the empirical moments. In Appendix C, we describe how to solve the model numerically. The simulation is conducted as follows. We solve the model for a given combination of c and ρ and use respective policy functions to simulate a panel of firms between 0 and T.

We match two moments. The first one is the expected value of the tax expenses under a tax strategy relative to the case without any tax deduction, $\mathbb{E}[(\tau I_T^a)/(\tau I_T)]$. We use the IRS' corporate tax gap's estimates and tax collection data for the empirical match. In particular, we consider τI_T^a as the actual IRS's tax collection and τI_T to be equal to this collection plus the tax gap the IRS expected to receive but did not collect. The second moment is the expected value of the penalty relative to the tax collection from a firm pursuing a tax strategy, $\mathbb{E}[(\rho-1)\tau\theta a_\theta/(\tau I_T^a)]$. Notice that by using $(\rho - 1)$, we capture the fraction of the penalty payment that is additional to the owed taxes in our model. Consistently, we use the IRS's penalties collected from corporations for the empirical match. To compute the empirical moments, we collected these variables from the IRS Data book for the available years between 2000 and 2022. We obtain that, on average, 87.9% of the expected tax collection was actually paid and that the penalties represented 0.643% of the actual tax paid. With these two empirical moments, the calibrated parameters are $c = 7.604$ and $\rho = 1.504$.[9]

4.1 Agent compensation relative to the evaded taxes

We now use the calibrated parameters to quantify the magnitude of the agent compensation relative to the intensity of the illegal tax strategy. The expected benefit for the principal driven by the tax evasion is calculated as follows:

$$\begin{aligned} E^{a^*}\left[X_T^{a^*} - W^{\pi^*} - P(\theta, a^*)\mathbb{1}_{\theta<T} + W^{\pi^*} - X_T^0\right] = &\int_0^T \left(\int_0^{T\wedge t} (\delta+\tau-1)a_s^* ds\right) \cdot \lambda e^{-\lambda t}\, dt \\ &- \int_0^T \rho\tau a_s^* \cdot s \cdot \lambda e^{-\lambda s}\, ds \\ &+ e^{-\lambda T}\int_0^T (\delta+\tau-1)a_s^* dt, \end{aligned} \tag{45}$$

[9] A structural estimation is not feasible in our setup given that our model does not consider other observable corporate policies (cash, investment, financing), limiting our capacity to simulate observable moments. Moreover, data on tax evasion and penalties at the firm level are not available. These facts preclude us from obtaining more moment conditions to perform a structural estimation.

The first three terms on the left-hand side represent the total net profits received by the principal minus the agent's compensation and penalties. To capture the magnitude of the benefits driven by the illegal tax strategy, we subtract from these residual profits the net profits that the principal would receive with a legal strategy. The expected value of the agent's compensation, as seen from the initial date, is represented by:

$$\begin{aligned}\mathbb{E}^{a^*}[W^{\pi^*}] &= y_{PC} + \int_0^T \left(\int_0^{t\wedge T} \frac{1}{2}\left(\sigma^2\gamma_A + \frac{1}{c}\right)(Z_s^*)^2\, ds \right) \lambda e^{-\lambda t}\, dt \\ &\quad + \int_0^T K_t^* \lambda e^{-\lambda t}\, dt + \int_0^T \left(\int_0^{t\wedge T} \frac{\lambda}{\gamma_A}\left(\exp(-\gamma_A K_s^*) - 1\right) ds \right) \lambda e^{-\lambda t}\, dt \\ &\quad + e^{-\lambda T}\Big\{ \int_0^T \frac{1}{2}\left(\sigma^2\gamma_A + \frac{1}{c}\right)(Z_s^*)^2\, dt + \int_0^T \left(\frac{\lambda}{\gamma_A}\left(\exp(-\gamma_A K_s^*) - 1\right) \right) dt \Big\},\end{aligned} \tag{46}$$

We analyze the agent's expected compensation relative of the expected tax evasion benefits using the following ratio:

$$\frac{\mathbb{E}^{a^*}[W^{\pi^*}]}{\mathbb{E}^{a^*}[X_T^{a^*} - P(\theta, a^*)\mathbb{1}_{\theta<T} - X_T^0]}. \tag{47}$$

We solve numerically the above expressions using our baseline calibration and obtain that 37.8% of the expected marginal benefit derived from tax evasion corresponds to the agent's compensation. This calibration highlights the magnitude of the economic incentives for managers in pursuing tax strategies. Indeed, recent public debates on fair taxation have noticed that managers' compensation exceeds tax payments by top U.S. corporations because managers are rewarded by implementing tax avoidance strategies The Guardian (2024); Anderson et al. (2024). To further understand the managerial economic incentives, we show in Figure 5 the comparative statics of the agent's share of the tax strategy benefits. For simplicity, we refer to the ratio in Equation 47 as $W/\Delta X$. Intermediation expenses $(1-\delta)$ decrease the expected compensation as a fraction of tax savings. The intuition behind this is that as the benefit of the tax strategy is captured by intermediaries, the agent's effort in reducing tax expenses is less valuable for the principal. Then we look at tax institutional features. We find that the agent's compensation increases with the tax rate τ. This is because it becomes more profitable for the principal to hire the agent to implement a tax strategy and also because the manager faces higher expected penalties in case of inspection, increasing the risk premium in her compensation. The parameters defining the penalty, λ and ρ, play a similar role in discouraging tax avoidance, resulting in a lower agent's compensation. However, the compensation is more sensitive to changes in the probability of inspection than to changes in the penalty factor. For example, if the inspection rate increases from 0.3 to 0.4 (a 10 percentage point increase), it implies a reduction in the compensation ratio of almost 4 percentage points. Similarly, a 10 percentage point increase in the penalty factor from 1.5 to 1.6 is linked to a reduction in the compensation ratio of less than 1 percentage point.

5 Conclusion

This paper examines corporate tax strategies from a dynamic contracting approach. Our model provides a theoretical explanation for why not all companies engage in tax strategies, even in cases where legal avoidance is feasible. We derive the optimal compensation for the agent implementing the tax strategy, contrasting scenarios with and without inspections from tax authorities. Inspection risk increases the agent's compensation. In the event of an audit, the agent's compensation suffers a lump-sum loss. Tax aggressiveness decreases over time as the expected penalty increases. Our calibration analysis suggests that a relevant fraction of the gains derived from tax strategies are paid to the agent.

References

Aksamit, A. and Jeanblanc, M. (2017). *Enlargement of Filtration with Finance in View*. SpringerBriefs in Quantitative Finance. Springer.

Anderson, S., Tashman, Z., and Rice, W. (2024). More for them, less for us: Corporations that pay their executives more than uncle sam.

Bennedsen, M. and Zeume, S. (2018). Corporate tax havens and transparency. 31(4):1221–1264.

Bertomeu, J., Jung, J., and Marinovic, I. (2022). Moral hazard and the value of information: A structural approach. *SSRN Electronic Journal*.

Chen, K.-P. and Chu, C. C. (2005). Internal control versus external manipulation: A model of corporate income tax evasion. *RAND Journal of Economics*, 36(1):151–164.

Crocker, K. J. and Slemrod, J. (2005). Corporate tax evasion with agency costs. *Journal of Public Economics*, 89(9):1593–1610.

Desai, M. A. and Dharmapala, D. (2006). Corporate tax avoidance and high-powered incentives. *Journal of financial Economics*, 79(1):145–179.

Desai, M. A., Dyck, A., and Zingales, L. (2007). Theft and taxes. *Journal of Financial Economics*, 84(3):591–623.

Enache, C. (2022). Corporate tax rates around the world, 2022.

Holmstrom, B. and Milgrom, P. (1987). Aggregation and linearity in the provision of intertemporal incentives. *Econometrica*, 55(2):303.

Hoopes, J. L., Mescall, D., and Pittman, J. A. (2012). Do irs audits deter corporate tax avoidance? *The Accounting Review*, 87(5):1603–1639.

ICIJ (2023). Behind the scenes of the PwC tax leak scandal with Neil Chenoweth. https://www.icij.org/inside-icij/2023/10/behind-the-scenes-of-the-pwc-tax-leak-scandal-with-neil-chenoweth/[Accessed: May 2024].

Jacob, M. (2022). Real effects of corporate taxation: A review. *European Accounting Review*, 31(1):269–296.

Jeanblanc, M. and Rutkowski, M. (2000). Modelling of shutdown risk. *Mathematical Tools*.

Martin, J. and Villeneuve, S. (2023). Risk-sharing and optimal contracts with large exogenous risks. *Decisions in Economics and Finance*, pages 1–43.

McClure, C. G. (2023). How costly is tax avoidance? evidence from structural estimation. *The Accounting Review*, 98(6):353–380.

Nikolov, B., Schmid, L., and Steri, R. (2021). The sources of financing constraints. *Journal of Financial Economics*, 139(2):478–501.

Ortiz, M. and Imbet, J. F. (2023). Private firms and offshore finance. Research Paper 4557679, Université Paris-Dauphine.

Reuters (2019). Google shifted $23 billion to tax haven Bermuda in 2017. https://www.reuters.com/article/us-google-taxes-netherlands-idUSKCN1OX1G9 [Accessed: May 2024].

Sannikov, Y. (2008). A continuous-time version of the principal-agent problem. *Review of Economic Studies*, 75(3):957–984.

Slemrod, J. (2007). Cheating ourselves: The economics of tax evasion. *Journal of Economic Perspectives*, 21(1):25–48.

The Guardian (2024). Companies paid top executives more than they paid in US taxes – report. https://www.theguardian.com/business/2024/mar/13/top-us-executives-salaries-versus-tax-payments [Accessed: May 2024].

Weisbach, D. A. (2002). Ten truths about tax shelters. *Tax Law Review*, 55:215.

Figures

Figure 1: Optimal Tax Strategy

Note: This figure illustrates the dynamics of the optimal tax strategy $a^*(t)$ for the model with and without inspection risk.

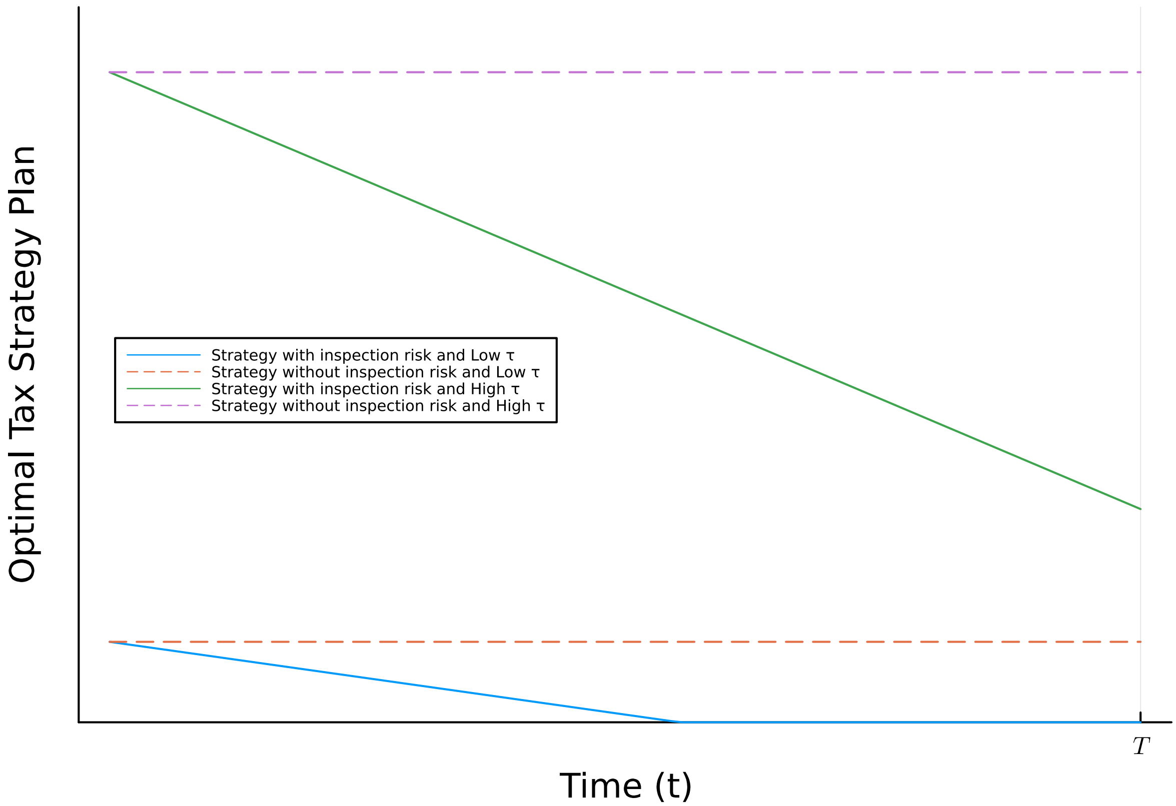

Figure 2: Penalty Dynamics

Note: This figure illustrates the dynamics of the penalty $P(a^*(t), t)$ imposed on detected tax-evading firms

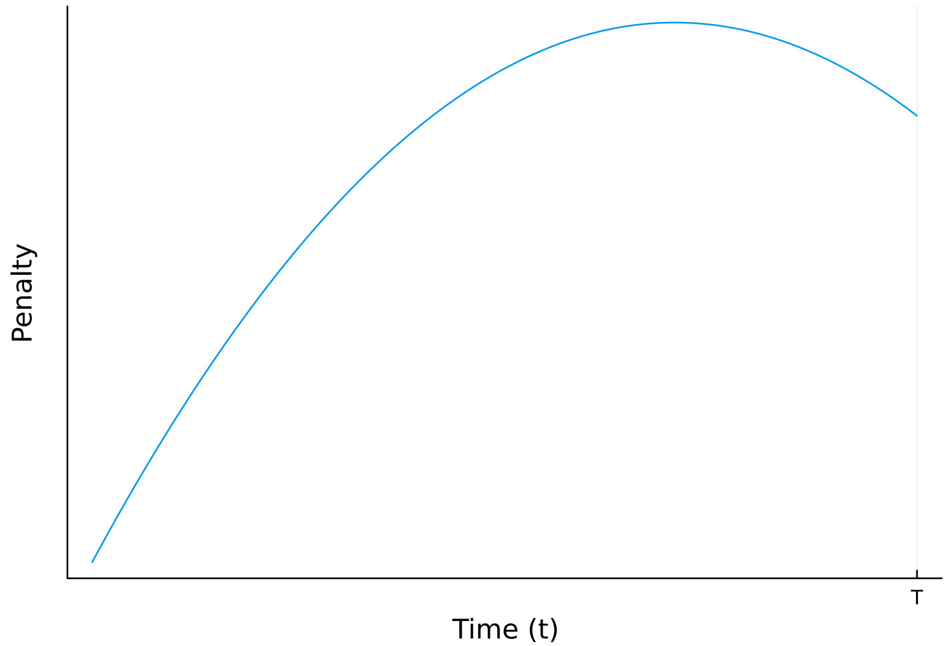

Figure 3: Compensation Sensitivity to Inspection Risk

Note: This figure illustrates the dynamics of the optimal sensitivity of the compensation to the inspection risk K_t^*.

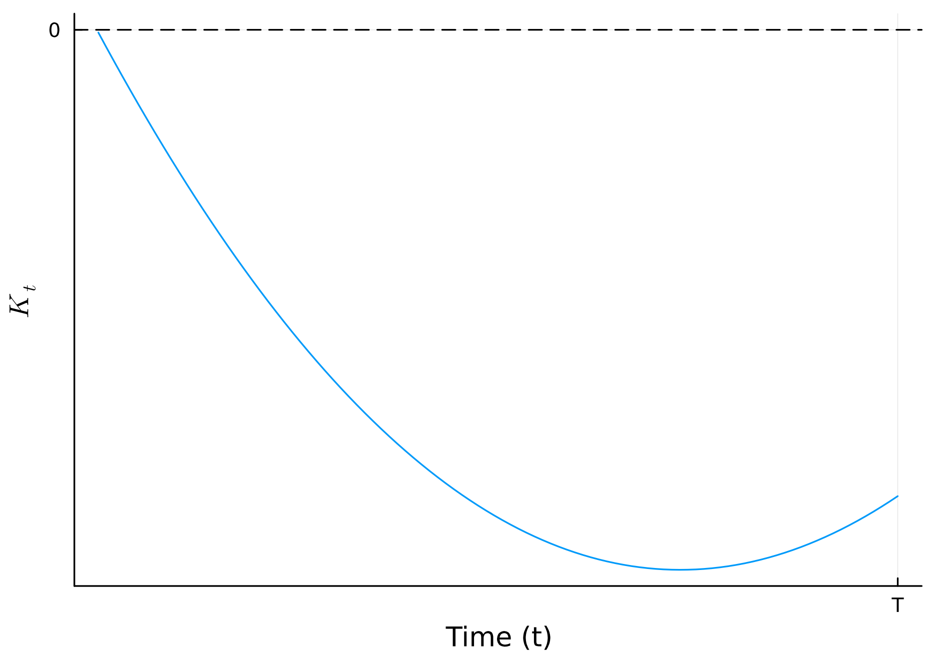

Figure 4: The Dynamics of the Optimal Compensation

Note: This figure illustrates the dynamics of the terminal compensation process for the model with and without inspection risk.

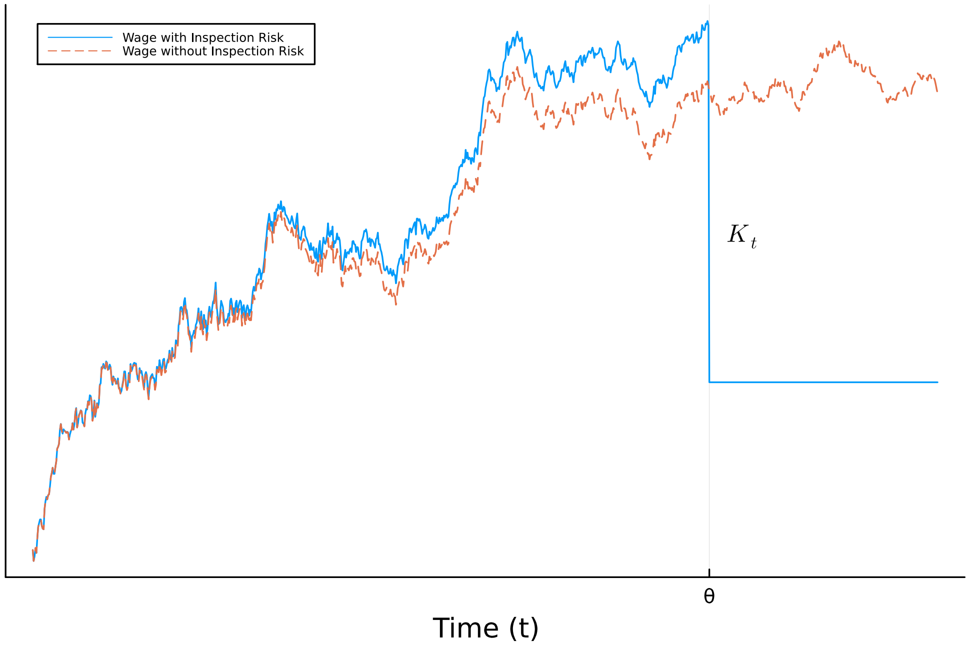

Figure 5: Agent's compensation relative to the gains derived from the tax strategy $\frac{\mathbb{E}^{a^*}[W^{\pi^*}]}{\mathbb{E}^{a^*}[X_T^{a^*} - P(\theta, a^*)\mathbb{1}_{\theta<T} - X_T^0]}$

Note: This figure provides the comparative statics for the optimal agent's compensation, based on the parameter values in Section 4. In the plot, we refer to this ratio by $\frac{W}{\Delta X}$.

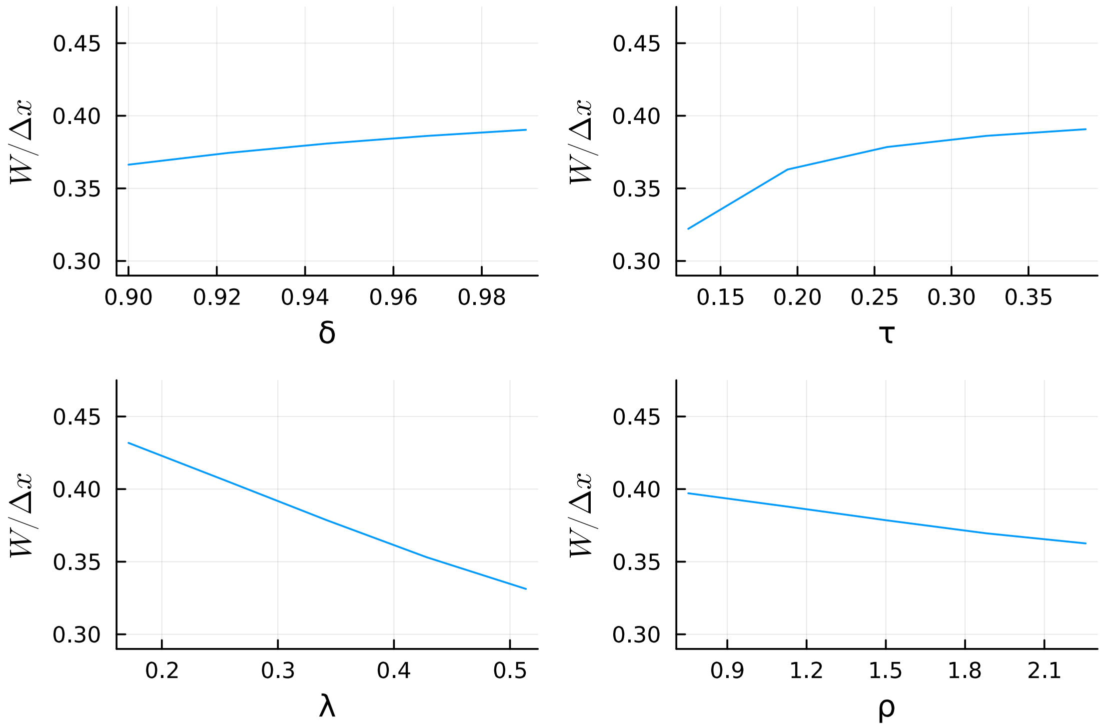

A Appendix: Probability Background

Let $(\Omega, \mathcal{G}, \mathbb{P}^0)$ be a complete probability space, and consider the following stochastic processes:

- B a standard $\mathbb{F} = (\mathcal{F}_t)_{t\geq 0}$-Brownian motion,
- N a right-continuous single-jump process defined as $N_t = \mathbb{1}_{\theta \leq t}, \quad t \in [0,T]$ where θ is some positive random variable independent of B that models the random inspection time.

We expand the information set $\mathbb{F}$ to incorporate details regarding inspections, specifically the timing of any past inspections. More specifically, we progressively enlarge the filtration associated with the Brownian motion and we consider $\mathbb{G} = (\mathcal{G}_t)_{t\geq 0}$ the smallest complete right-continuous enlargement of $\mathbb{F}$ such that θ is a $\mathbb{G}$-stopping time (see Aksamit and Jeanblanc (2017) and Jeanblanc and Rutkowski (2000)). The single-jump process N is assumed to have fixed intensity λ, and

$$M_t = N_t - \int_0^t \lambda\,(1 - N_s)\,ds, \quad t \in [0,T]$$

is a $\mathbb{G}$-compensated martingale.

Furthermore, the probability that an inspection occurs during a period $[t,T]$, with $t < T$, is:

$$\mathbb{P}(t \leq \theta \leq T) = 1 - \exp\{-\lambda(T-t)\}.$$

We denote by $\mathbb{P}^0$ (respectively $\mathbb{P}^a$) the probability distribution of the firm's income in the absence of action from the agent (respectively with the action strategy $a = (a_t)_{0\leq t\leq T}$ is realized by the agent). The probability measures $\mathbb{P}^0$ and $\mathbb{P}^a$ are equivalent and the Radon-Nikodym derivative of $\mathbb{P}^a$ with respect of $\mathbb{P}^0$ is such that

$$\left.\frac{d\mathbb{P}^a}{d\mathbb{P}^0}\right|_{\mathcal{G}_T} = \exp\left(\int_0^T -\frac{a_s}{\sigma}\,(1-N_s)\,\mathrm{d}\boldsymbol{B}_s - \frac{1}{2}\int_0^T \left|\frac{a_s}{\sigma}\right|^2 (1-N_s)\,\mathrm{d}s\right)$$

and $B_t^a = B_t + \int_0^t \frac{a_s}{\sigma}\,(1-N_s)\,\mathrm{d}s, t \in [0,T]$ is a $\mathbb{G}$ -Brownian motion under $\mathbb{P}^a$.

B Appendix: Omitted Proofs

Proof of Theorem 1. We provide here the technical details and derivations for the optimal process $\tilde{Z}^*$ that were omitted in the main text. Using the representation of W, the principal's expected utility can be written

$$\mathbb{E}^{a^*}\left[U_P\left(X_T^{a^*} - W\right)\right]$$

we know that

$$\begin{aligned}
X_T^{a^*} &= \int_0^T \Big((1-\tau)\mu - (\delta+\tau-1)\frac{\tilde{Z}_t}{c}\Big)dt + (1-\tau)\sigma\int_0^T dB_t^{a^*} \\
W &= y_0 + \int_0^T \frac{1}{2}\Big(\gamma_A\sigma^2 + \frac{1}{c}\Big)\tilde{Z}_t^2 dt + \sigma\int_0^T \tilde{Z}_t dB_t^{a^*} \\
X_T^{a^*} - W &= \int_0^T \Big((1-\tau)\mu - \frac{\delta+\tau-1}{c}\tilde{Z}_t - \frac{1}{2}(\gamma_A\sigma^2 + \frac{1}{c})\tilde{Z}_t^2\Big)dt \\
&\quad + \sigma\int_0^T \Big((1-\tau) - \tilde{Z}_t\Big)dB_t^{a^*} - y_0
\end{aligned} \tag{48}$$

or in differential form

$$\begin{aligned}
dX_t^{a^*} - dW_t &= h(\tilde{Z}_t)dt + \sigma((1-\tau) - \tilde{Z}_t)dB_t^{a^*} \\
h(\tilde{Z}_t) &= (1-\tau)\mu - \frac{\delta+\tau-1}{c}\tilde{Z}_t - \frac{1}{2}(\gamma_A\sigma^2 + \frac{1}{c})\tilde{Z}_t^2 \\
X_0^{a^*} - W_0 &= -y_0
\end{aligned} \tag{49}$$

note that the dynamics of $dX_t^{a^*} - dWt$ do not depend on t, which suggests that the optimal action is constant (e.g. $\{\tilde{Z}_t\}_{s=0}^T = \tilde{Z}$). Using the properties of the exponential utility function, we have that

$$\mathbb{E}^{a^*}\left[U_P\left(X_T^{a^*} - W\right)\right] = -\exp\Big(-\gamma_P\big(-y_0 + h(\tilde{Z}_t)T\big) + \frac{1}{2}\gamma_P^2\sigma^2((1-\tau) - \tilde{Z})^2 T\Big) \tag{50}$$

The first-order condition implies that

$$\frac{\delta+\tau-1}{c} + (\gamma_A\sigma^2 + \frac{1}{c})\tilde{Z} - \gamma_P\sigma^2((1-\tau) - \tilde{Z}) = 0 \tag{51}$$

solving for $\tilde{Z}$ yields

$$\tilde{Z} = \frac{\gamma_P \sigma^2 (1-\tau) - \frac{\delta+\tau-1}{c}}{\sigma^2(\gamma_A+\gamma_P)+\frac{1}{c}} \tag{52}$$

□

Proof of Theorem 2. We start by proving the following proposition:

Proposition 1 (Sufficient Conditions for Participation)**.** *Consider a contract defined by* $\Pi = (y_0, a, \tilde{Z}_t, \tilde{K}_t)$*, where* y_0 *is the initial value, a is a parameter,* $\tilde{Z}_t$ *and* $\tilde{K}_t$ *are time-dependent variables. This contract offers a wage, denoted as* W_T^Π*, which is the value at time* T *of a process governed by the following dynamics:*

$$dW_t^\Pi = \left(\frac{1}{2}(\gamma_A\sigma^2 - \frac{1}{c})\tilde{Z}_t^2 + \frac{\lambda}{\gamma_A}(\exp(-\gamma_A\tilde{K}_t)-1)+\lambda\tilde{K}_t\right)dt + \tilde{K}_t dM_t + \tilde{Z}_t\sigma dB_t, \tag{53}$$

$$\textit{with } W_0^\Pi = y_0, \tag{54}$$

The contract satisfies the participation constraint if $y_0 \geq y_{PC}$*.*

We define a wage process that contains sensitivities $\tilde{Z}_t$ and $\tilde{K}_t$ with respect to the Brownian component and the inspection process, as well as a deterministic payment $g(Z_t, K_t)$. We find the processes $\tilde{Z}_t, \tilde{K}_t$, and the function $f(t, W_t) := U_A(W_t - \frac{c}{2}\int_0^t a_s^2 ds)$ in equilibrium. For a contract $\pi = (y_0, a, \beta, H)$ consider the following wage process

$$dW_t^\Pi = g(\tilde{Z}_t, \tilde{K}_t)dt + \tilde{Z}_t\sigma dB_t + \tilde{K}_t dM_t \tag{55}$$

with initial condition $W_0^\pi = y_0$. Applying Itô's lemma for jump-diffusion processes, and assuming the time of the inspection θ has not yet occurred, we have:

$$\frac{\mathrm{d}U_A(W_t^{\Pi} - \frac{c}{2}\int_0^t a_s^2 ds)}{U_A(W_t^{\Pi} - \frac{c}{2}\int_0^t a_s^2 ds)} = \gamma_A \frac{c}{2} a_t^2 dt - \gamma_A (g(\tilde{Z}_t, \tilde{K}_t) - \lambda \tilde{K}_t) dt - \gamma_A \tilde{Z}_t \sigma (-\frac{a_t}{\sigma} dt + dB_t^a)$$
$$+ \frac{1}{2}\gamma_A^2 \sigma^2 \tilde{Z}_t^2 dt + \big(\exp(-\gamma_A \tilde{K}_t) - 1) dN_t$$
$$= \gamma_A \Big(\frac{c}{2} a_t^2 - g(\tilde{Z}_t, \tilde{K}_t) + \frac{1}{2}\gamma_A \sigma^2 \tilde{Z}_t^2 + \frac{\lambda}{\gamma_A}(\exp(-\gamma_A \tilde{K}_t) - 1)) + \lambda \tilde{K}_t + a_t \tilde{Z}_t \Big) dt$$
$$- \gamma_A \tilde{Z}_t \sigma dB_t^a + (\exp(-\gamma_A K_t) - 1) dM_t \tag{56}$$

The utility process has to be a $\mathbb{P}^a$-martingale (Proposition 1 in Sannikov 2008). We choose the minimum compensation g such that the drift of the process is equal to zero.

$$g(\tilde{Z}_t, \tilde{K}_t) = \frac{1}{2}\gamma_A \sigma^2 \tilde{Z}_t^2 + \frac{\lambda}{\gamma_A}(\exp(-\gamma_A \tilde{K}_t) - 1)) + \lambda \tilde{K}_t + \inf_a \{ \frac{c}{2} a_t^2 + a_t \tilde{Z}_t \} \tag{57}$$

which yields an optimal action $a_t^* = -\frac{\tilde{Z}_t}{c}$ and a compensation function g equal to

$$g(\tilde{Z}_t, \tilde{K}_t) = \frac{1}{2}\left(\gamma_A \sigma^2 - \frac{1}{c}\right) \tilde{Z}_t^2 + \frac{\lambda}{\gamma_A}(\exp(-\gamma_A \tilde{K}_t) - 1) + \lambda \tilde{K}_t \tag{58}$$

so under $\mathbb{P}^a$

$$dW_t = \left(\frac{1}{2}\left(\gamma_A \sigma^2 - \frac{1}{c}\right) Z_t^2 + \frac{\lambda}{\gamma_A}(\exp(-\gamma_A \tilde{K}_t) - 1) - a_t \tilde{Z}_t \right) dt + \tilde{K}_t dN_t + \tilde{Z}_t \sigma dB_t^a, \tag{59}$$

and under $\mathbb{P}^{a^*}$:

$$dW_t = \left(\frac{1}{2}(\gamma_A \sigma^2 + \frac{1}{c}) \tilde{Z}_t^2 + \frac{\lambda}{\gamma_A}(\exp(-\gamma_A \tilde{K}_t) - 1) \right) dt + \tilde{K}_t dN_t + \tilde{Z}_t \sigma dB_t^{a^*}, \tag{60}$$

Next, given the dynamics of dW_t we compute $dU_A(W_t)/U_A(W_t)$ under $\mathbb{P}^0$

$$dU_A(W_t) = -\frac{1}{2}\gamma_A \frac{\tilde{Z}_t^2}{c} U_A(W_t) dt - \gamma_A \tilde{Z}_t \sigma U_A(W_t) dB_t + (\exp(-\gamma_A \tilde{K}_t) - 1) U_A(W_t) dM_t \tag{61}$$

Now we use Lemma 3.1. in Martin and Villeneuve (2023) to show that if $(U_A(W_t), -\gamma_A \tilde{Z}_t \sigma U_A(W_t), (\exp(-\gamma_A \tilde{K}_t) - 1) U_A(W_t))$ is a solution satisfying the participation constrain.

Setting $-\gamma_A U_A(W_t)\tilde{Z}_t = Z_t^{(W,a)}$ and $(\exp(-\tilde{K}_t\gamma_A)-1)U_A(W_t) = -K_t^{(W,a)}$ let us identify that $U_A(W_t)$ is a solution to the BDSE. This solution is unique, and the participation constraint is satisfied only if

$$U_A(W_0) = U_A(y_0) = U_0^{(W,a)} = \mathbb{E}\Big[U_A(W - \frac{c}{2}\int_0^T a_s^2 ds)\Big] \geq U_A(y_{PC}) \tag{62}$$

Hence, the wage process given by the wage dynamics in Proposition 1 will satisfy the participation constraint if $y_0 \geq y_{PC}$. The principal binds the contract at the lowest initial value for the wage process that satisfies the participation constraint, so $y_0 = y_{PC}$.

Now, we move to prove Theorem 2. We do so by following the proof of Lemma 4.1 in Martin and Villeneuve (2023). We consider a family of stochastic processes $R^a(W) := (R_t^a)_{t\in[0,T]}$ indexed by a that satisfies:

1. $R_T^a = U_A\left(W - \frac{c}{2}\int_0^T a_s^2\left(1-N_s\right)\mathrm{d}s\right)$,
2. R^a is a $\mathbb{P}^a$-supermartingale for any effort a,
3. R_0^a is independent of a.
4. There exists a^* such that R^{a^*} is a $\mathbb{P}^{a^*}$-martingale.

Then,

$$R_0^{a^*} = \mathbb{E}^{a^*}\left[U_A\left(W - \frac{c}{2}\int_0^T (a_s^*)^2(1-N_s)ds\right)\right] \geq \mathbb{E}^a\left[U_A\Big(W - \frac{c}{2}\int_0^T a_s^2(1-N_s)ds\Big)\right],$$

meaning that a^* is the optimal agent's action in response to the contract W.

We define the family $R^a(W) := (R_t^a)_{t\in[0,T]}$ as

$$R_t^a := U_A\left(W_t^{\Pi} - \frac{c}{2}\int_0^t a_s^2\left(1-N_s\right)\mathrm{d}s\right)$$

Applying Itô's formula to R_t^a under the probability measure $\mathbb{P}^0$ leads to:

$$\begin{aligned}\frac{\mathrm{d}R_t^a}{R_t^a} &= (1-N_t)\Big[-\gamma_A\tilde{Z}_t\sigma\mathrm{d}B_t + \left(e^{-\gamma_A\tilde{K}_t}-1\right)\mathrm{d}N_t + \\ &+ \gamma_A\left\{\frac{1}{2}\gamma_A\sigma^2\tilde{Z}_t^2 - g\left(\tilde{Z}_t,\tilde{K}_t\right) + \frac{c}{2}a_t^2 + \lambda\tilde{K}_t + \frac{\lambda}{\gamma_A}\left(e^{-\gamma_A\tilde{K}_t}-1\right)\right\}\mathrm{d}t\Big].\end{aligned} \tag{63}$$

where g is defined as Proposition 1.

Assuming that the tax strategy maximizes equation (9), the principal's problem becomes

$$V_P := \sup_{y \geq y_{pc}} v(0, x, w)$$

where

$$v(0, x, w) = \sup_{\pi} \mathbb{E}^* \left[U_P \left(X_T^a - W_T^\pi \right) \right]$$

To construct an optimal contract, we will develop a smooth solution for the Hamilton-Jacobi-Bellman (HJB) equation of the principal's problem:

$$\begin{aligned} 0 =& \partial_t v(t, x, w) + \sup_{Z,K} \Bigg\{ \partial_x v(t, x, w) \left((1-\tau)\mu - (\delta + \tau - 1) \frac{Z_t}{c} \right) + \partial_y v(t, x, w) \left[\frac{Z_t^2}{2} \left(\frac{1}{c} + \sigma^2 \gamma_A \right) \right. \\ & \left. + \frac{\lambda}{\gamma_A} \left[\exp\left(-\gamma_A K_t \right) - 1 \right] \right] + \lambda \left[U_P \left(x - y - K_t - \frac{\rho \tau Z_t t}{c} + L \cdot (T - t) \right) - v(t, x, w) \right] \\ & + \partial_{xx} v(t, x, w) \frac{(1-\tau)^2 \sigma^2}{2} + \partial_{yy} v(t, x, w) \frac{Z_t^2 \sigma^2}{2} + \partial_{xy} v(t, x, w) Z_t (1 - \tau) \sigma^2 \Bigg\}. \end{aligned} \tag{64}$$

The function $U_P(x - w)\phi_0(t)$ serves as a classical solution to our HJB equation. We find that :

$$\begin{aligned} \phi_0'(t) = \inf_{(Z_t)_{t \leq T \wedge \theta}, (K_t)_{t \leq T \wedge \theta}} \Bigg\{ & \gamma_P \phi_0(t) \Bigg((1-\tau)\mu - \frac{\gamma_P}{2} (1-\tau)^2 \sigma^2 + \frac{\lambda}{\gamma_A} \\ & + \left(\frac{-(\delta + \tau - 1)}{c} + \gamma_P (1-\tau) \sigma^2 \right) Z_t \\ & - \left(\sigma^2 (\gamma_A + \gamma_P) + \frac{1}{c} \right) \frac{(Z_t)^2}{2} \\ & - \frac{\lambda}{\gamma_A} \exp\left\{ -\gamma_A K_t \right\} \Bigg) - \lambda \exp \left\{ \gamma_p \left(K_t - \frac{\rho \tau Z_t t}{c} - L(T - t) \right) \right\} \Bigg\} \end{aligned} \tag{65}$$

Taking the first order condition with respect to the controls K and Z give

$$K_t^* = \frac{\log \phi_0(t) + \gamma_P \left(\frac{\rho \tau Z_t^* t}{c} + L(T - t) \right)}{\gamma_A + \gamma_P} \tag{66}$$

and $Z^* = (Z_t^*)_{t \leq T \wedge \theta}$ such that:

$$\begin{aligned} 0 =& \phi_0(t) \left(-\frac{(\delta + \tau - 1)}{c} + \gamma_P (1-\tau)\sigma^2 - \left(\sigma^2(\gamma_A + \gamma_P) + \frac{1}{c} \right) Z_t^* \right) \\ &+ \frac{\lambda \rho \tau t}{c} \phi_0(t)^{\frac{\gamma_P}{\gamma_A + \gamma_P}} \exp\left\{ -\frac{\gamma_A \gamma_P}{\gamma_A + \gamma_P} \left(\frac{\rho \tau Z_t^* t}{c} + L(T-t) \right) \right\}. \end{aligned} \tag{67}$$

These optimal controls are bounded and independent of y. Finally, it follows that:

$$\begin{aligned} \phi_0'(t) = \gamma_P \phi_0(t) \Bigg(& (1-\tau)\mu - \frac{\gamma_P}{2}(1-\tau)^2\sigma^2 + \frac{\lambda}{\gamma_A} + \left(\frac{-(\delta + \tau - 1)}{c} + \gamma_P(1-\tau)\sigma^2 \right) Z_t^* \\ & - \left(\sigma^2(\gamma_A + \gamma_P) + \frac{1}{c} \right) \frac{(Z_t^*)^2}{2} - \frac{\lambda}{\gamma_A} \exp\{-\gamma_A K_t^*\} \Bigg) \\ & - \lambda \exp\left\{ \gamma_p \left(K_t^* - \frac{p Z_t^* t}{c} - L(T-t) \right) \right\} \end{aligned}$$

□

C Numerical Solution

To compute $\tilde{K}_t$ and $\tilde{Z}_t$ as well as ϕ_0 in equation (42) for the case of illegal tax-evasion, we perform the following iterative procedure. Processes $\tilde{K}_t^i$ and $\tilde{Z}_t^i$ and the function ϕ_0^i are indexed by the iteration i.

1. Initialize $\tilde{Z}_t^0$ with Equation (24).
2. Compute K_t^0 based on Equation (34).
3. Solve for ϕ_0^0 in Equation (35) using finite-differences based on the boundary condition in (33).
4. In iteration n compute $\tilde{Z}_t^n$ using Equation (35) given ϕ_0^{n-1}.
5. Go to step 2, and stop when $\tilde{Z}_t$, $\tilde{K}$ and ϕ_0 converge.