Online Appendix: Additional material

ON-A Strategies and outcomes

This part of the appendix contains the formal restrictions on the players' strategy spaces to ensure that any combination of strategies leads to a unique and well-defined outcome.

To this end, it is convenient to have an exogenous underlying stochastic process that governs the arrival of shocks and, given the agent's effort, determines the state of compliance. Let $(\Omega, \mathcal{F}, P)$ be a probability space. Let the marked point process $z = \{z_t\}_{t\geq 0}$ represent the arrival of random shocks, where $z_t = 0$ except at isolated times $t_0 < t_1 < \dots$ which arrive at constant rate $\lambda > 0$. At each random time t_j with $j \in \mathbb{N}$, the value of the shock z_{t_j} is independently and uniformly distributed on $[0, 1]$. Let $\{\mathcal{F}_t\}$ be the natural filtration generated by z. The state of compliance $\{\theta_t\}_{t\geq 0}$ is constant between shocks, and immediately after the arrival of a shock at time t_j, we have $\theta_{t_j} = 1$ if $\alpha\eta_{t_j} \geq z_{t_j}$, and $\theta_{t_j} = 0$ if $\alpha\eta_{t_j} < z_{t_j}$.

A history at time t is a collection of paths

$$h_t = \{\eta_s, \theta_s, \hat{\theta}_s, N^I_s, F_s\}_{s\in[0,t]},$$
$$\text{where } (\eta_s, \theta_s, \hat{\theta}_s, N^I_s, F_s) \in [0, 1] \times \{0, 1\} \times \{0, 1\} \times \mathbb{N}_0 \times \mathbb{R}_+.$$

Throughout, we denote strict histories for which the realization at time t is excluded by h_{t-}. Let H_t be the set of all time-t histories and H_{t-} the set of all strict histories. Let $H = \bigcup_{t\geq 0} H_t$ and $H_- = \bigcup_{t\geq 0} H_{t-}$.

The agent's strategy specifies efforts and reports as functions of histories. A strategy for the agent is defined as a pair $(e, \rho) = (\{e_t, \rho_t\}_{t\geq 0})$ with

$$e_t : H_{t-} \to [0, 1], \quad \rho_t : H_{t-} \times \{0, 1\} \to \{0, 1\},$$

where $e_t(h_{t-})$ is the agent's effort at time t and $\rho_t(h_{t-}, \theta_t)$ is the agent's report at time t after history h_{t-} when the state at time t is θ_t. Note that while the agent submits a report regarding compliance continually, at every $t \geq 0$, this is for notational convenience only. It is equivalent and most natural for most applications to think of the agent as sending messages only sporadically to report transitions in compliance. To capture the principal's uncertainty about the agent's effort choices and the true state of compliance, consider a partition $\mathcal{H}_t^P$ of the history set H_t which comprises all subsets of H_t whose elements are indistinguishable to the principal. Define the partition $\mathcal{H}_{t-}^P$ similarly for strict histories at t. To allow for randomized inspections, we equip the principal with a (private) random signal π, defined on a sufficiently rich probability space with state space Π. A strategy for the principal is defined as a pair $(n, f) = (\{n_t, f_t\}_{t\geq 0})$ of mappings

$$n_t : \Pi \times H_{t-} \times \{0,1\} \to \{0,1\}, \quad f_t : H_{t-} \times \{0,1\}^3 \to \mathbb{R}_+,$$

which are constant on every $H_{t-}^P \in \mathcal{H}_{t-}^P$ for each $t \geq 0$, where f_t is required to be weakly increasing over time. Here, $n_t(\pi, h_{t-}, \hat{\theta}_t)$ is equal to 1 if an inspection is performed at time t and equal to 0 otherwise. By $f_t(h_{t-}, \theta_t, \hat{\theta}_t, \mathrm{d}N_t^I)$ we denote the cumulative fine imposed by the principal at time t. We abuse notation slightly and write $f_t(h_t)$ instead of $f_t(h_{t-}, \theta_t, \hat{\theta}_t, \mathrm{d}N_t^I)$ whenever there is no danger of confusion. The exit decision for each player at any history is a binary variable indicating whether this player decides to exit or not. For the ease of exposition, we do not introduce additional notation for these choices; they translate into lower bounds on the expected payoffs of the players in the equilibrium definition below. The strategies above are to be understood as conditional on no player having exited previously. Actions to be chosen after one player exited are irrelevant.

To ensure that any strategy profile results in a unique and well-defined process of actions, we adopt the approach by Kamada and Rao (2023) and require that actions are not changed 'too frequently' on any time interval. To apply this approach, first restrict the strategy spaces for the fine and effort choices. A history $h_t \in H_t$, has an *intervention for the agent* at time t if either $t = 0$, or if $t > 0$ and at least one of the following holds: (i) $\theta_t - \theta_{t-} \neq 0$, (ii) $\hat{\theta}_t - \hat{\theta}_{t-} \neq 0$, (iii) $N_t^I - N_{t-}^I \neq 0$. Similarly, there is an *intervention for the principal* if either $t = 0$, or if $t > 0$ and at least one of the properties (*ii*) and (*iii*) holds. No new information arrives in between interventions. We restrict the principal's fine strategy to reflect this, and require that it be predictable in between inspections. Formally, for any two histories $h_t and h_t'$: $f_t(h_t) \neq f_t(h_t')$ only if there exists $\tau \leq t$ such that τ is an intervention time for the principal and the truncation of the above histories at time τ, h_τ and h_τ', are distinguishable for the principal. In other words, this restriction requires the principal's fines to be specified pathwise; at each intervention, it is fully specified how fines proceed until another intervention arrives. Similarly, we restrict the agent's effort strategy to be predictable in between interventions: For any two histories h_{t-}, h_{t-}': $e_t(h_{t-}) \neq e_t(h_{t-}')$ only if there exists $\tau < t$ such that τ is an intervention time for the agent and $h_\tau \neq h'_\tau$. Based on Kamada and Rao (2023), we require all strategies to fulfil the properties *traceability* and *frictionality* as defined below. Lemma E then shows that any combination of strategies from this class yields a well-defined and unique outcome path. A history h is said to be *consistent* with the agent's strategy (e, ρ) at time t if $\rho_t(h_{t-}, \theta_t) = \hat{\theta}_t$ and $e_t(h_t) = \eta_t$. Similarly, a history h is consistent with the principal's strategy (n, f) at time t if $n_t(\pi, h_{t-}, \hat{\theta}_t) = \mathrm{d}N_t^I$ and $f_t(h_t) = \mathrm{d}F_t$.

Definition 1. *The agent's strategy (ρ, e) is **traceable** if for any time-t history h_t and any principal-action path $\{N_s^I, F_s\}_{s\geq 0}$ that coincides with h_t for all $s < t$, there is a continuation path $\{\hat{\theta}_s, \eta_s\}_{s\geq t}$ that is consistent with (ρ, e). Analogously, The principal's strategy (n, f) is traceable if for any time-t history h_t and any agent-action path $\{\hat{\theta}_s, \eta_s\}_{s\geq 0}$ that coincides with h_t for all $s < t$, there is a continuation path $\{N_s^I, F_s\}_{s\geq t}$ that is consistent with (n, f).*

Definition 2. *The agent's strategy (ρ, e) is **frictional** if for any time-t history h_t, there is conditional probability one that the report path $\{\hat{\theta}_s\}_{s\geq t}$ has only finitely many report changes on any finite interval $[t, u]$ for all paths $\{\eta_s, \hat{\theta}_s\}_{s\geq t}$ such that there is a principal-action path $\{N_s^I, F_s\}_{s\geq t}$ for which the history $\left(h_{t-}, \{N_s^I, F_s\}_{s\geq t}, \{\eta_s, \hat{\theta}_s\}_{s\geq t}\right)$ is consistent with the agent's strategy. Analogously, the principal's strategy (n, f) is frictional if for any time-t history h_t, there is conditional probability one that the inspection path $\{N_s\}_{s\geq t}$ has only finitely many inspections on any finite interval $[t, u]$ for all paths $\{N_s^I, F_s\}_{s\geq t}$ such that there is an action path $\{\eta_s, \hat{\theta}_s\}_{s\geq t}$ for which the history $\left(h_{t-}, \{N_s^I, F_s\}_{s\geq t}, \{\eta_s, \hat{\theta}_s\}_{s\geq t}\right)$ is consistent with the principal's strategy.*

Lemma E (Existence and Uniqueness of consistent Outcome Path)**.** *Given any possible history $h_{u-} = \left\{\pi_0, z_t, \eta_t, \hat{\theta}_t, N_t^I, F_t\right\}_{t\in[0,u)} \cup \{\eta_u\}$, any combination of strategies $((e, \rho), (n, f))$ that are traceable and frictional yields a unique consistent path $\left(\{\eta_t\}_{t\in(u,\infty)}, \{\hat{\theta}_t, N_t^I, F_t, \}_{t\in[u,\infty)}\right)$ almost surely.*

Proof. The proof proceeds in two steps. First we show uniqueness and then existence.

Step 1: Uniqueness. Fix a pair of strategies, a history up to u, and any realization of the shock process $\{z_t\}_{t\in[u,\infty)}$. Suppose there are two distinct continuation paths $x = \{\eta_t^x, \hat{\theta}_t^x, N_t^{I^x}, F_t^x\}_{t\in[u,\infty)}$ and $y = \{\eta_t^y, \hat{\theta}_t^y, N_t^{I^y}, F_t^y\}_{t\in[u,\infty)}$ that are consistent with the strategies and the shock path. Let $\underline{t} = \inf\{t \geq u : x_t \neq y_t\}$ be the first time at which the processes differ. Strategy e maps history $h^A_{t^A_k}$ into a deterministic process $\{\eta_s\}_{s\in(t^A_k,\infty)}$ only for times t^A_k at which an intervention for the agent occurs. Likewise, strategy f maps history $h_{t^P_k}$ into a deterministic process $\{F_s\}_{s\in[t^P_k,\infty)}$ for times t^P_k with an intervention for the principal. Therefore, if $\eta^x_s \neq \eta^y_s$ for $s > u$ or $F^x_s \neq F^y_s$ for $s \geq u$, then there must also be a time $t \leq s$ with an intervention at t, i.e. $\exists n \in \mathbb{N}$ s.t. $t = t^A_k$ or $t = t^P_k$. Furthermore, we must have $h^x_t \neq h^y_t$ at this intervention. With probability 1, the realization $\{z_t\}_{t\in[u,\infty)}$ has only finitely many jumps on any closed interval. Hence, by frictionality, there are at most finitely many interventions on any closed interval. Therefore, $\underline{t}$ defined above must be an intervention time and the infimum is attained, i.e., $x_{\underline{t}} \neq y_{\underline{t}}$. We therefore must have $\hat{\theta}^x_{\underline{t}} \neq \hat{\theta}^y_{\underline{t}}$ or $N^{I^x}_{\underline{t}} \neq N^{I^y}_{\underline{t}}$ and, as $\underline{t}$ is the first such time, $h^x_{\underline{t}-} = h^y_{\underline{t}-}$. As $\hat{\theta}^x_{\underline{t}}$ and $\hat{\theta}^y_{\underline{t}}$ both result from the same strategy, this, however, implies that $\hat{\theta}^x_{\underline{t}} = \hat{\theta}^y_{\underline{t}}$, leaving as only possibility that $N^{I^x}_{\underline{t}} \neq N^{I^y}_{\underline{t}}$. This contradicts consistency of both processes with the fixed strategy (as $h^x_{\underline{t}-} = h^y_{\underline{t}-}$). Hence, any pair of traceable and frictional strategies gives at most one consistent outcome.

Step 2: Existence. Existence of a consistent outcome path is shown constructively: Start with arbitrary history $h_{u-} = \left\{\pi_0, z_t, \eta_t, \hat{\theta}_t, N^I_t, F_t\right\}_{t\in[0,u)} \cup \{\eta_u\}$ and fix a realization of the shock process $\{z_t\}_{t\in[u,\infty)}$. We apply the steps below iteratively until they give an outcome path consistent with z and the strategies for $t \geq u$: Define paths $\{\eta^0_t, \hat{\theta}^0_t, N^{I^0}_t, F^0_t\}$ equal to the history up to u and such that for $t > u$: $\eta^0_t = e_t(h_{\max_k t^A_k < u})$, and for $t \geq u$: $\hat{\theta}^0_t = \hat{\theta}_{u-}$, $N^{I^0}_t = N^I_{u-}$ and $\mathrm{d}F^0_t = f_t(h_{\max_k t^A_k < u})$.[36] Let $n = 1$ and $t(1) = u$.

i) By traceability, there are paths $\{\eta^n_t, \hat{\theta}^n_t\}_{t\geq 0}$ such that, for $t < t(n)$: $\{\hat{\theta}^n_t, \eta^n_t\} = \{\eta^{n-1}_t, \hat{\theta}^{n-1}_t\}$ and that $\{\eta^n_t, \hat{\theta}^n_t, N^{I^{n-1}}_t, F^{n-1}_t\}_{t\geq 0}$ is consistent with the agent's strategy and process z for $t \geq t(n)$. Set $\{\eta^n_t, \hat{\theta}^n_t\}$ equal to these processes. Similarly, traceability implies that there exist paths $\{N^{I^n}_t, F^n_t\}$ with $(N^{I^n}_t, F^n_t) = (N^{I^{n-1}}_t, F^{n-1}_t)$ for $t < t(n)$ and such that $\{\eta^n_t, \hat{\theta}^n_t, N^{I^n}_t, F^n_t\}_{t\geq 0}$ is consistent with the principal's strategy on $t \geq u$. Set $\{N^{I^n}_t, F^n_t\}$ equal to these processes and continue to step (ii).

ii) If $\{\eta^n_t, \hat{\theta}^n_t, N^{I^n}_t, F^n_t\}$ is consistent with the strategies for all $t \in [u, \infty)$, stop the pro-

[36] That is, report and inspections are held constant from u onward and fines and effort are chosen according to the strategies (depending only on the last intervention before u) for the case that no further interventions occur.

cedure. The proof is complete. Otherwise, redefine $n = n + 1$ and set $t(n + 1)$ equal to the largest time v such that there is an intervention at v and $\{\eta_t^n, \hat{\theta}_t^n, N_t^{I^n}, F_t^n\}$ is consistent with the strategies for all $t \in [u, v)$, go to step (i).

If the above procedure stops after finite n, that's because of having given a consistent process and the proof is complete. In the case in which it does not stop after finitely many iterations,

$$\lim_{n\to\infty} \{\eta_t^n, \hat{\theta}_t^n, N_t^{I^n}, F_t^n\}_{t\geq 0}$$

is consistent with the strategies on $[u, \infty)$ with probability one. To see this, note that for every n, $t(n + 2) > t(n)$. Given that, with probability one, any finite interval has only finitely many interventions, $\lim_{n\to\infty} t(n) = \infty$ which implies consistency of the resulting process for all $t \in [u, \infty)$. □

ON-B Martingale representation of promised utility

Proof of Lemma A. Denote by $\mathcal{F}$ the filtration generated by the random processes θ, $\hat{\theta}$ and ν^I. Define

$$W_t := \int_0^t e^{-rs}(-\,\mathrm{d}F_s - c\eta_s\,\mathrm{d}s) + e^{-rt}U_t.$$

The corresponding representation in differential form is

$$dW_t = e^{-rt}(-\,\mathrm{d}F_t - c\eta_t\,\mathrm{d}t) - re^{-rt}U_t + e^{-rt}\,\mathrm{d}U_t. \tag{23}$$

The process $\{W_t\}$ is an $\mathcal{F}$-martingale by construction. By the martingale representation theorem for marked point processes (Last and Brandt, 1995, Theorem 1.13.2), there exist $\mathcal{F}$-predictable functions $\tilde{\Delta}_t^\theta$, $\tilde{\Delta}_t^{\hat{\theta}}$ and $\tilde{\Delta}_t^I$ such that

$$dW_t = \sum_{a\in\{\theta,\hat{\theta},I\}} \tilde{\Delta}_t^a(\mathrm{d}N_t^a - \mathrm{d}\nu_t^a) \tag{24}$$

Replacing $\tilde{\Delta}_t^a = e^{-rt}\Delta_t^a$ and then equating (23) and (24) yields

$$\mathrm{d}U_t = rU_t\,\mathrm{d}t + \mathrm{d}F_t + c\eta_t\,\mathrm{d}t + \sum_{a\in\{\theta,\hat{\theta},I\}} \Delta_t^a(\mathrm{d}N_t^a - \mathrm{d}\nu_t^a).$$

This is the representation of the evolution of promised utilities shown in the lemma. □

ON-C Proof of Lemma C with left-open support

In this section, we show how the arguments in the proof of Lemma C extend to the case in which $t^0 \notin \mathcal{T}$. First, note that if its infimum t^0 is not contained in the set $\mathcal{T}$, then for any $\delta > 0$, we can find an $\epsilon \in (0, \delta)$ such that $t^0 + \epsilon \in \mathcal{T}$. Further, by choosing δ small enough, we can ensure that the expected inspection probability $\int_{t^0}^{t^0+\delta} \mathrm{d}\nu_s^I$ becomes arbitrarily small. In the first case with $U_{t^0}^0 > -B$, there exists an $\epsilon > 0$ small enough such that $t^0 + \epsilon \in \mathcal{T}$ and also $U_{t^0+\epsilon}^0 > -B$ by right-continuity of U_t^0. In this case we can apply the argument above to schedule a predictable inspection at time $t^0 + \epsilon$. To satisfy the agent's incentive constraints, this modification is paired either with an additional fine after a high report at $t^0 + \epsilon$ or with an additional transition fine for any transition at times $s \in [t^0, t^0 + \epsilon)$, depending on the sign of $\Delta_{t^0+\epsilon}^I$. In the second case with $U_s^0 = -B$, on $[t^0, t^0 + \delta)$ for some $\delta > 0$, then by $\Delta_s^I > -B - U_s^0$, we have that $\Delta_{t^0+\epsilon}^I > 0$. In this case, we can proceed in a similar way as above and introduce an additional fine to compensate for the increase in the agent's expected payoff caused by performing the inspection with probability 1 and keep the path of persistent payoffs U_s^1 unchanged for $s \leq t^0$. However, to ensure that the obedience and honesty constraints are also satisfied on $(t^0, t^0 + \epsilon]$, the fine is increased gradually on the interval $(t^0, t^0 + \epsilon)$. Specifically, construct the fine such that the honesty constraint (H) binds (with $U_s^0 = -B$):

$$0 = -rB\,\mathrm{d}t - \lambda\alpha(U_s^1 + B)\,\mathrm{d}t + \mathrm{d}F_t + c\,\mathrm{d}t. \tag{25}$$

In the promise-keeping constraint (Pk), substituting for $\mathrm{d}F_s$ with the binding honesty constraint (25) and inserting $U_s^0 = -B$ determines the evolution of U_s^1 on $(t^0, t^0 + \epsilon)$ via the differential equation

$$\hat{u}_s' = (r + \lambda)(\hat{u}_s + B).$$

We keep the persistent utility at t^0 unchanged, so the initial condition for the ODE is $\hat{u}_{t^0} = U_{t^0}^1$, which leads to the solution

$$\hat{u}_s = U_{t^0}^1 e^{(r+\lambda)(s-t^0)} + B\left(e^{(r+\lambda)(s-t^0)} - 1\right),$$

for $s \in [t^0, t^0 + \epsilon)$. To ensure, that this trajectory of persistent utility is feasible, we verify that the fine $\mathrm{d}Fs$ is positive and that the solution $\hat{u}_{t^0+\epsilon}$ does not exceed $U_{t^0+\epsilon}^1 + \Delta_{t^0+\epsilon}^I$ from the original equilibrium. The latter is necessary to reach $U_{t^0+\epsilon}^1 + \Delta_{t^0+\epsilon}^I$ as the continuation payoff after inspection at $t^0 + \epsilon$. For the fine, (25) with $U_s^1 = \hat{u}_s$ gives

$$\frac{\mathrm{d}F_s}{\mathrm{d}t} = -c + rB + \lambda\alpha(\hat{u}_s + B) = -c + rB + \lambda\alpha(U_{t^0}^1 + B)e^{(r+\lambda)(s-t^0)}.$$

This term is decreasing in s and therefore smallest at $s = t^0$, where it is positive if

$$(r + \lambda\alpha)B + \lambda\alpha U_{t^0}^1 \geq c.$$

For the original equilibrium to satisfy the obedience constraint we must have $U_{t^0}^1 \geq -B + \frac{c}{\lambda\alpha}$, so that the above inequality must be satisfied and the fines are positive. To check that $\hat{u}_{t^0+\epsilon}$ constructed above does not lie above $U_{t^0+\epsilon}^1 + \Delta_{t^0+\epsilon}^I$ from the original equilibrium, note that the inspections in the original equilibrium had no effect on the honesty constraint (H) as, by assumption, we are in the case $U_s^0 = -B$. Therefore, as the original equilibrium satisfied the honesty constraints, the evolution of $\hat{u}_s$, which was constructed by making the honesty constraint binding, must lie weakly below the original U_s^1 and therefore $\hat{u}_{t^0+\epsilon} \leq U_{t^0+\epsilon}^1 + \Delta_{t^0+\epsilon}^I$ since $\Delta_{t^0+\epsilon}^I$ is positive by $\Delta_{t^0+\epsilon}^I > -B - U_{t^0}^0 = 0$. Hence, the newly constructed equilibrium includes a fine at inspection time $t^0 + \epsilon$ of $\hat{U}_{t^0+\epsilon}^1 - (U_{t^0+\epsilon}^1 + \Delta_{t^0+\epsilon}^I)$ so that the persistent utility increases to the one from the original continuation equilibrium after inspection at time $t^0 + \epsilon$. □