Robust Alignment of Incentives*

Andrea Di Giovan Paolo†

November 3, 2026

[Click here for the most recent version]

Abstract

I study robust regulation of a better-informed agent whose incentives are imperfectly aligned with the principal's. The principal faces ambiguity about both the consequences of the agent's actions and his private bias, but knows that small changes in the action cannot have arbitrarily large payoff consequences. She may impose action-contingent penalties or restrict the set of actions available to the agent. I characterize corrective interventions that minimize worst-case regret. Their form depends on the principal's information about the agent's bias. If she can bound its overall magnitude but not its local variation, the optimal intervention is a staircase tax. If she instead knows bounds on the marginal bias, it is a piecewise-linear tax with a single kink.

*I am deeply grateful to Piotr Dworczak, Jeff Ely and Alessandro Pavan for their guidance and advice. For helpful comments and suggestions, I thank Nemanja Antić, Federico Bugni, Yingni Guo, Suraj Malladi, Wojciech Olszewski, Harry Pei, Bruno Strulovici, Yutong Zhang, Alison Zhao, and participants at the Northwestern Theory Lunch. All remaining errors are my own.

†Department of Economics, Northwestern University. Email: andreadigiovanpaolo@u.northwestern.edu.

1. Introduction

Many regulatory and organizational problems share a common structure: a principal cares about the action of a better-informed agent whose incentives are not fully aligned with the principal's objective. A contractor deciding how long to work on a project may know whether additional time will improve quality or simply lead to wasteful delay. At the same time, he may also prefer more time to save on costs or for reasons unrelated to project quality. Similarly, a hiring manager choosing among a pool of candidates may be better informed than the firm about their productivity if hired, while also being biased against a minority group. The manager may therefore be willing to sacrifice productivity to avoid hiring from that group. Whenever the principal can influence the choice through action-contingent taxes she faces the problem of correcting the agent's misaligned incentives while preserving the value of the agent's private information.

If the principal knew exactly how the agent's payoff differs from her own, she could use a corrective tax to exactly offset the misalignment, while preserving the agent's ability to condition his choice on private information about the environment. In practice, however, the principal often has imperfect information about *both* the value of different actions and the extent of misalignment. Any corrective intervention then faces a fundamental trade-off. Aggressive intervention protects against choices driven by bias, but risks distorting choices that reflect valuable information. Conversely, leaving more flexibility to the agent preserves his informational advantage, but leaves greater scope for misaligned incentives to distort behavior.

I formalize and solve this general alignment problem when the principal faces ambiguity about both the mapping from actions to consequences and the agent's misalignment, which together determine a "state of the world". When the environment is complex or the regulatory problem is new, the principal may lack the information needed to credibly assign probabilities to states. At the same time, technological constraints and other known features of the problem often provide a basis for restricting the set of *possible* states.

In particular, I assume the principal can bound the effect of a small change in the action on her payoff, i.e., she knows that nearby actions cannot have arbitrarily different consequences. This implies that the principal is uncertain about the location of the optimum

while still having local information about the cost of small implementation errors. Formally, this is captured by known Lipschitz bounds on the rate at which the principal's payoff can change with the action. In the contractor example, the principal may not know exactly how the value of the project evolves with the completion date, but she may be able to bound the largest gain or loss that an additional hour can generate. In the hiring example, the firm may not know the optimal composition of minority and majority candidates, but may have enough knowledge about the population of candidates to be able to bound the marginal gain or loss in productivity from a small change in that composition.

The second source of ambiguity is the agent's misalignment. The principal does not know the precise magnitude or shape of the bias. I consider two scenarios that differ in how much information she has about it. In both cases, the misalignment is one-directional: the agent prefers weakly higher actions than the principal. In the first scenario, the principal is only able to bound the maximum possible extent of the bias, but does not have enough information to determine how sharply it may change with the action. In the second scenario, the principal has more precise knowledge and knows bounds on the *marginal* bias.

Any form of intervention exposes the principal to a trade-off between over- and under-correction. This provides a natural way to evaluate mechanisms under ambiguity. In any given state and for any mechanism, the principal's regret is the payoff loss relative to the value she would have obtained had she known the state exactly. The principal evaluates each mechanism by the largest regret it can generate over all possible states, and chooses the mechanism that minimizes this worst-case loss relative to the first-best. The optimal corrective intervention therefore balances the losses of over- and under-correction in the worst case.

Results. In the scenario in which the principal can bound the overall magnitude of the agent's bias but not how it changes across actions, the minimax-optimal mechanism takes the form of a staircase tax. The tax is constant over a sequence of intervals and jumps discretely between them. The number and width of the intervals, and the size of the jumps are all determined endogenously by the bounds on the principal's payoff and on the maximum extent of misalignment.

Intuitively, the staircase arises because of the nature of the principal's knowledge about her own payoff and about the agent's bias. Since actions that are close together cannot generate very different payoffs, the trade-off between preserving the agent's information and correcting the bias is resolved locally in favor of the former. The flat parts of the staircase leave complete discretion to the agent, without imposing any marginal cost on slightly higher actions. For actions that are farther apart, however, the potential payoff loss from a biased choice becomes larger. At the same time, because the principal cannot bound how the agent's bias changes with the action, any positive marginal tax risks being ineffective in counteracting the bias. The tax therefore jumps discretely between intervals. These jumps limit the extent to which the bias can push the agent toward higher actions, while the flat regions preserve his ability to use private information within each interval.

The contractor example helps clarify the intuition. In this case, the staircase takes the form of a sequence of soft deadlines. Up to the first deadline, the contractor is free to choose when to deliver without facing any penalty. Crossing the deadline triggers a discrete penalty, after which the schedule is again flat until the next deadline. Suppose first that the state is such that the contractor is not biased, and a small amount of additional time would genuinely improve the project, so that the principal would prefer delivery just after the deadline. The notch may distort the unbiased agent's choice, inducing him to deliver just before the deadline instead. This exposes the principal to regret from over-correction. However, since the two completion dates are close and the value of the improvements that can be made in a short amount of time is bounded, the resulting loss is also bounded. At the same time, the notch protects against under-correction in an alternative state in which delivery before the cutoff is optimal and the contractor prefers a later completion date because of his private bias. Because the principal cannot bound how sharply this private benefit may increase with the completion date, a smooth penalty may provide insufficient correction, while a discrete jump creates a strong incentive not to cross the deadline. The staircase therefore leaves the contractor discretion within each interval, while using jumps to limit the extent to which private incentives can push delivery toward later dates.

The result changes when the principal can bound how the agent's bias varies with the action and the marginal bias is known to lie between a lower and an upper bound. In this case, the minimax-optimal mechanism is a piecewise-linear tax with a single kink. Below

an endogenous cutoff, the tax has slope equal to the lowest possible marginal bias; above the cutoff, its slope jumps to the highest possible marginal bias.

The intuition again comes from the trade-off between over- and under-correction. The lower bound on the marginal bias is the part of the distortion that the principal knows is present, so below the cutoff she corrects only this component, leaving as much scope as possible for the agent to use his private information. The cost is that, if the true marginal bias is closer to its upper bound, correction may be too weak. Above the cutoff, the principal instead protects against the largest possible marginal bias by raising the marginal tax to its upper bound, at the cost of potentially correcting too much when the true bias is small. The cutoff determines how much of the action space is exposed to each of these two errors and is chosen to minimize their worst-case regret.

I consider two applications of the framework that also allow me to illustrate possible microfoundations for the structure of the ambiguity sets. The first one is emissions regulation. An important source of misalignment arises when the agent does not internalize an externality his activity imposes on the environment. I consider a regulator facing a firm whose production generates emissions and which is privately informed about its abatement costs. The regulator does not know the firm's cost function or the exact environmental damage caused by emissions, but may be able to bound marginal abatement costs and marginal damages, for instance through engineering estimates of available abatement technologies and assessments of environmental harm. I show that the analysis for the case of known bounds on the marginal bias applies to this setting, and that a kinked, piecewise-linear emissions tax is minimax optimal.

The second application develops the hiring discrimination problem introduced above. I consider a firm that delegates hiring decisions to a manager who is better informed about applicants' productivity but may have discriminatory preferences. While the firm may be aware of such preferences, assuming that it knows their precise functional form is particularly demanding. For instance, a manager may not only favor hiring workers from his own demographic group but also place additional value on maintaining a majority from that group. Inferring these preferences from observed hiring decisions is difficult because such decisions also reflect the manager's private information about applicants' productivity. To counteract discrimination, the firm may link the manager's compensation to diversity

hiring targets. I show that, even without knowing the productivity distributions within applicant groups, if the firm knows bounds on individual productivity and can place a conservative bound on how much productivity the manager would sacrifice to favor his preferred demographic group, tiered bonuses tied to hiring targets are minimax optimal.

In addition to capturing situations in which the principal operates under considerable uncertainty, the robust approach plays two important roles in my analysis. The first is technical. It allows me to formulate and solve the alignment problem in a general environment. A Bayesian formulation at this level of generality would require a prior over the infinite-dimensional space of possible principal payoff functions and agent biases, making a general characterization of optimal incentive schemes very difficult. Second, this approach yields optimal interventions that take simple and familiar forms, staircase or kinked penalties, providing a new and intuitive rationale for discontinuities (in levels or marginal penalties) in incentive schemes commonly observed in practice.

After introducing the main results, I discuss additional properties of the optimal mechanisms and consider some extensions. First, I show that both mechanisms are undominated: no other feasible mechanism attains weakly lower regret in every state and strictly lower regret in some. Second, I consider how robust they are to additional global information the principal may have about the shape of her payoff functions, relative to the baseline marginal bounds. Knowing that the payoff is single-peaked in the action does not help the principal. By contrast, I show that knowing the payoff is concave can reduce worst-case regret. Under concavity, in the scenario where the principal can bound the absolute, but not the marginal bias, the staircase mechanism remains optimal, but with larger jumps and shorter flat regions. When the principal knows bounds on the marginal bias, instead, the piecewise-linear kinked tax is no longer optimal in general. It is replaced by a logarithmic tax, which may also display a kink depending on parameter values. Finally, I consider a special case of the problem, in which the principal cannot use action-contingent penalties and can only restrict the set of actions available to the agent, as in the literature on optimal delegation. I show that in both ambiguity scenarios, a simple cap on the set of feasible actions is a minimax-optimal delegation set.

Related literature. This paper contributes to a growing literature on robust mechanism design and contracting, which studies the design of incentives when the principal has only partial knowledge of the economic environment (Carroll, 2019). In particular, it contributes to the literature on minimax-regret mechanism design, rooted in statistical decision theory (Wald, 1950; Savage, 1954). This approach has been applied to robust pricing (Bergemann and Schlag, 2011), project choice (Guo and Shmaya, 2023), delegated screening (Malladi, 2020), and monopoly regulation (Guo and Shmaya, 2025). I focus on the problem of correcting the incentives of a partially aligned agent when the principal is uncertain about both her own payoff function and the agent's bias[1].

The economic problem shares important features with delegation. As in the delegation literature, the principal trades off the benefits of allowing an agent to exploit his superior information against the distortions created by preference misalignment (Holmström, 1977; Alonso and Matouschek, 2008; Amador and Bagwell, 2013). Classical delegation models typically assume that the principal knows the agent's preference structure and has a prior over the underlying state. Moreover, with notable exceptions (e.g., Amador and Bagwell, 2013; Ambrus and Egorov, 2017), most of this literature rules out the possibility of using action-contingent penalties to discipline the agent.

A more recent strand of the delegation literature studies robust delegation under uncertainty about the agent's preferences. Frankel (2014) studies ambiguity about the agent's preferences and characterizes mechanisms that maximize the principal's worst-case payoff when the agent makes multiple related decisions. Alonso, Gan, and Hu (2026) combine Bayesian uncertainty about the underlying state with ambiguity about the preferences of the better-informed agent and show that convex delegation sets are maxmin optimal. I consider ambiguity about both the principal's payoff and the agent's bias, and allow the principal to use action-contingent penalties in addition to restricting the agent's discretion. A second difference is the criterion used to evaluate mechanisms. Both Frankel (2014)

[1]The emissions regulation application in Section 4.1 also relates to the classical prices-versus-quantities problem studied by Weitzman (1974). A recent paper by Kang (2026) studies robust regulation of externalities when the regulator is uncertain about the correlation between private demand and marginal externalities, and shows that quantity controls can be optimal when the regulator maximizes worst-case welfare. In my setting, the principal regulates a firm with private information about its abatement costs, while facing uncertainty about both these costs and environmental damages, in order to minimize worst-case regret. I show that a kinked emissions tax is minimax optimal.

and Alonso, Gan, and Hu (2026) maximize the principal's worst-case payoff, whereas I minimize worst-case regret. This distinction is important in my setting: under the type of ambiguity I consider, a maxmin objective is excessively conservative and trivializes the alignment problem. This is an instance of the more general concern noted by Manski (2021) that in some problems the maxmin approach can be too "pessimistic" and lead to decisions that are too insensitive to available information. A recent working paper by He et al. (2026) studies delegation without a prior over the payoff-relevant state and minimizes worst-case regret. In a quadratic-loss environment with a known bias, they characterize an optimal randomized delegation rule consisting of a default action, a gap, and a higher discretionary interval with random boundaries. I focus on deterministic mechanisms and consider ambiguity over a high-dimensional set of functions describing the principal's payoff and the agent's bias. Delegation is a special case of my analysis in which the principal is not allowed to use action-contingent penalties. In my framework a simple cap on the set of actions is a deterministic optimal delegation mechanism.

Finally, the staircase result in particular contributes to the literature on notches and discontinuous incentives. As noted by Kleven (2016), notches are widespread in practice but difficult to rationalize in standard environments with smooth fundamentals, optimizing agents, and sufficiently flexible instruments. Indeed, existing rationales typically rely on exogenous restrictions on the available instruments (Blinder and Rosen, 1985; Gillitzer, Kleven, and Slemrod, 2017). This paper provides a different justification for notches as a robust solution to an alignment problem. Notches arise endogenously within an unrestricted class of deterministic mechanisms as the optimal way to resolve the trade-off between over and under-correction when the principal cannot bound the agent's marginal bias, but can bound the effect of a small change in the action on her own payoff. When the principal can instead bound the local variation of the bias, discontinuities in levels are not needed, but a discontinuity in the marginal tax (a kink) becomes optimal.

2. Model

Environment. A principal contracts with an agent. The set of feasible actions is a compact interval normalized to $X = [0, 1]$. An allocation specifies an action $x \in X$ and a penalty $t \in \mathbb{R}$ imposed on the agent (negative penalties are interpreted as rewards). The principal's

payoff from action x is $V : X \to \mathbb{R}$. The agent's gross payoff from action x is

$$U(x) = V(x) + \delta(x),$$

where $\delta : X \to \mathbb{R}$ captures the misalignment between the principal and the agent's preferences. I assume that δ is weakly increasing, so the agent is biased towards higher actions. When allocation (x, t) is implemented, the agent obtains

$$V(x) + \delta(x) - t.$$

I assume penalties are pure implementation instruments: their role is to align the agent's incentives with the principal's objective but they do not enter the principal's payoff directly. This in particular means that the principal does not face a rent-extraction problem. Her objective is to align the agent's incentives to her own, rather than extracting money from the agent. If we take the penalty to be a monetary tax and the principal as a regulator, one interpretation of the assumption is that payments are purely redistributive and rebated lump-sum to other members of society, so that they have no direct effect on social welfare. For simplicity, I will refer to t as an action-contingent penalty or tax interchangeably.

A state of the world is a pair of functions (V, δ). Together, they determine the principal's and the agent's payoff. The principal does not know the realized state. The agent observes his payoff function[2], U. As will become clear below when considering applications, the term δ does not need to represent an actual behavioral bias. It is meant to capture any source of misalignment between the principal and the agent's payoff. I will often use the term bias for simplicity.

Mechanisms. The principal can elicit information from the agent by committing to a mechanism before the relevant state (V, δ) is realized. A (deterministic) indirect mechanism consists of a message space M and a public mapping $\mu : M \to X \times \mathbb{R}$. After observing his payoff U, the agent chooses a message $m \in M$, and the corresponding allocation $(x(m), t(m))$ is implemented. I impose standard regularity conditions ensuring that the

[2]The results are unchanged if the agent observes V and δ separately. His preferences over allocations depend on the state only through U, so additional information can affect behavior only when he is indifferent among allocations. Because I evaluate mechanisms under tie-breaking against the principal, this cannot improve the principal's guarantee. Thus, either interpretation is possible, depending on the application.

agent's problem admits a solution in every admissible state.[3]

By the Taxation Principle (Rochet, 1985; Guesnerie, 1995), it is without loss of generality to represent such a mechanism by a tax schedule over actions. For each available action x, let $T(x)$ denote the lowest payment at which x can be implemented, and set $T(x) = +\infty$ for actions that are unavailable. The resulting tax schedule is lower semicontinuous and finite at some action. Let

$$\mathcal{T} = \{T : X \to \mathbb{R} \cup \{+\infty\} : T \text{ is lower semicontinuous and finite at some } x \in X\}.$$

The agent's choice can then be written as

$$\Gamma(T; V, \delta) = \arg\max_{x \in X}\{U(x) - T(x)\} = \arg\max_{x \in X}\{V(x) + \delta(x) - T(x)\}. \tag{1}$$

Ambiguity. Rather than assuming that the principal has a Bayesian prior over possible states (V, δ), I assume that she has only coarse information about admissible payoff functions and bias.

Principal's payoff. The first substantive assumption I make is that the principal can bound the marginal effect of a change in the action on her payoff.

Assumption 1. Let $\underline{v}, \overline{v} > 0$ be constants known to the principal. The set of admissible payoff functions for the principal is given by

$$\mathcal{V} = \{V : X \to \mathbb{R} : V(0) = 0,\ -\underline{v}(y - x) \leq V(y) - V(x) \leq \overline{v}(y - x)\ \forall x, y \in X \text{ with } x < y\}.$$

The normalization $V(0) = 0$ is without loss of generality. When V is differentiable, the restriction is equivalent to $-\underline{v} \leq V'(x) \leq \overline{v}$. To simplify notation, I will mostly omit the dependence of $\mathcal{V}$ on $\underline{v}$ and $\overline{v}$. When it needs to be highlighted, I will use the notation $\mathcal{V}(\underline{v}, \overline{v})$.

The assumption captures the idea that, while the principal does not know exactly how the agent's actions affect her payoff, she can bound how much her payoff changes in response

[3] Specifically, I require the induced menu $\mu(M) \subseteq X \times \mathbb{R}$ to be nonempty, closed, and bounded below in transfers.

to small changes in the action.[4] This may be the case, for instance, when x is the time the agent spends working on a project and the principal can bound the value of the progress that can be made over a short interval, as well as the cost of additional delay. Similarly, when x represents the amount of a firm's emissions, a regulator may know bounds on the firm's marginal abatement costs and on marginal environmental damage, even without knowing their exact values. As I will show below, an implication of this assumption is that the principal is in particular able to bound the consequences of small implementation errors.

Agent's misalignment. Next, we have to specify what the principal knows about the agent's misalignment, δ. I consider two scenarios. In the first one, the principal knows the maximal extent of the disagreement with the agent, but not how it varies with the action. In short, we say that the principal can bound the *absolute* bias. Formally, for some $D > 0$, define

$$\mathcal{B}_a = \{\delta : X \to [0, D] : \delta(0) = 0,\ \delta \text{ is nondecreasing and upper semicontinuous}\}.$$

The normalization $\delta(0) = 0$ is again without loss. Upper semicontinuity is a technical assumption made to ensure the agent's problem has a solution in every admissible state.

In the second scenario, the principal has bounds on how the disagreement changes with the action: she can bound the *marginal bias*. Formally, for some known bounds $0 \leq \underline{\delta} < \overline{\delta}$, define

$$\mathcal{B}_m = \left\{\delta : X \to \mathbb{R} : \delta(0) = 0,\ \underline{\delta}(y - x) \leq \delta(y) - \delta(x) \leq \overline{\delta}(y - x),\ \forall\, x < y\right\}.$$

The corresponding ambiguity sets are

$$\mathcal{A}_a = \mathcal{V} \times \mathcal{B}_a \qquad \text{and} \qquad \mathcal{A}_m = \mathcal{V} \times \mathcal{B}_m.$$

I omit the dependence of $\mathcal{A}$ on $\underline{v}$ and $\overline{v}$, since I maintain Assumption 1 throughout the paper. When it is important to highlight this dependence, I will indicate it explicitly.

[4] A similar assumption appears, in the different context of search under ambiguity, in Malladi (2022). In that setting, items with similar observable characteristics are mapped to similar payoffs to the searcher.

Regret. The presence of ambiguity about both the mapping from actions to consequences, V, and the agent's misalignment, δ, implies that the first best is unattainable. Any form of corrective intervention exposes the principal to two potential mistakes. Because the agent is biased towards higher actions, intervention that is too weak may induce the agent to choose an action above the principal's preferred one. More aggressive taxation can mitigate this problem, but may result in over-correction and induce an action that is too low. I evaluate a mechanism by the regret it induces, i.e., the principal's loss relative to the full-information benchmark. Given a tax schedule T and a state (V, δ), regret is defined as

$$\mathrm{Reg}(T; V, \delta) = \sup_{x \in \Gamma(T;V,\delta)} \left[\max_{x' \in X} V(x') - V(x) \right].$$

Here, the benchmark $\max_{x' \in X} V(x')$ is the payoff the principal could attain if she knew her realized payoff function and could select the action directly. The supremum over the agent's choice correspondence evaluates the mechanism under the tie-breaking rule least favorable to the principal. Equivalently, the regret guarantee must hold for every action that maximizes the agent's payoff. Adverse tie-breaking is consistent with the worst-case approach: the principal aims for a guarantee that remains valid under the least favorable action among those the agent is willing to choose.

For an ambiguity set $\mathcal{A}_i \in \{\mathcal{A}_a, \mathcal{A}_m\}$, the worst-case regret of T is

$$R_i(T) = \sup_{(V,\delta) \in \mathcal{A}_i} \mathrm{Reg}(T; V, \delta),$$

and the principal's problem is

$$\inf_{T \in \mathcal{T}} R_i(T).$$

Interpretation of $\underline{v}$ and $\overline{v}$. To interpret the results below, it is useful to note that, under Assumption 1, the bounds $\underline{v}$ and $\overline{v}$ on the principal's marginal payoff govern the potential cost of under- and over-correcting the agent's behavior, respectively. Let x^* denote a first-best action in any given admissible state (V, δ). If intervention is too aggressive and induces $x < x^*$, the regret from over-correction is bounded by $V(x^*) - V(x) \leq \overline{v}(x^* - x)$. Conversely, if intervention is too weak and the agent chooses $x > x^*$, the regret from under-correction is bounded by $V(x^*) - V(x) \leq \underline{v}(x - x^*)$.

3. Results

I now characterize the minimax-optimal mechanisms in the two ambiguity environments introduced above. Throughout the section, ambiguity about the principal's payoff remains the one described in Assumption 1. I first consider the case in which the principal can bound the absolute (but not the marginal) bias, $\delta \in \mathcal{B}_a$. I then turn to the case in which she instead has bounds on the marginal bias, $\delta \in \mathcal{B}_m$.

3.1. Bounds on absolute bias: the staircase tax

Theorem 1. *Let $\mathcal{A} = \mathcal{A}_a$ and*

$$R_a^\star = \min_{n \in \mathbb{Z}_+} \max \left\{ \frac{nD}{n+1}, \frac{\underline{v}\,\overline{v} + n\underline{v}\,D}{(n+1)(\underline{v}+\overline{v})} \right\}. \tag{2}$$

Let $n^\star$ be any minimizing integer and define

$$\widehat{x}^\star := \min \left\{ \frac{(n^\star+1)R_a^\star}{\underline{v}}, 1 \right\}.$$

(i) $R_a^\star$ is a lower bound on regret, i.e. for any $T \in \mathcal{T}$, $R_a(T) \geq R_a^\star$;

(ii) The mechanism $T_a^\star$, such that $T_a^\star(0) = 0$ and

$$T_a^\star(x) = \begin{cases} j(D - R_a^\star), & \dfrac{jR_a^\star}{\underline{v}} < x \leq \min\left\{ \dfrac{(j+1)R_a^\star}{\underline{v}}, 1 \right\}, \quad j = 0, \ldots, n^\star, \\ n^\star(D - R_a^\star) + D, & \widehat{x}^\star < x \leq 1. \end{cases}$$

achieves $R_a^\star$.

Proof. See Appendix A.1 for the proof of the theorem. Lemma 3 in Appendix A.1 establishes that the minmax problem (2) defining $R_a^\star$ admits solution and characterizes the optimal $n^\star$. □

The first part of the theorem identifies a lower bound on regret, while the second part shows that $T_a^\star$ achieves it. The mechanism takes the shape of a staircase. Starting from zero, it consists of flat regions of length $R_a^\star/\underline{v}$, separated by jumps of size $D - R_a^\star$. There

are $n^\star + 1$ such regions, with the last one truncated if they reach the upper bound 1. If instead

$$\frac{(n^\star + 1)R_a^\star}{\underline{v}} < 1,$$

we have a final jump of size D (so $n^\star + 1$ jumps in total) and a terminal flat region from $\widehat{x}^\star$ to 1. The shape of $T_a^\star$ is illustrated in Figure 1.

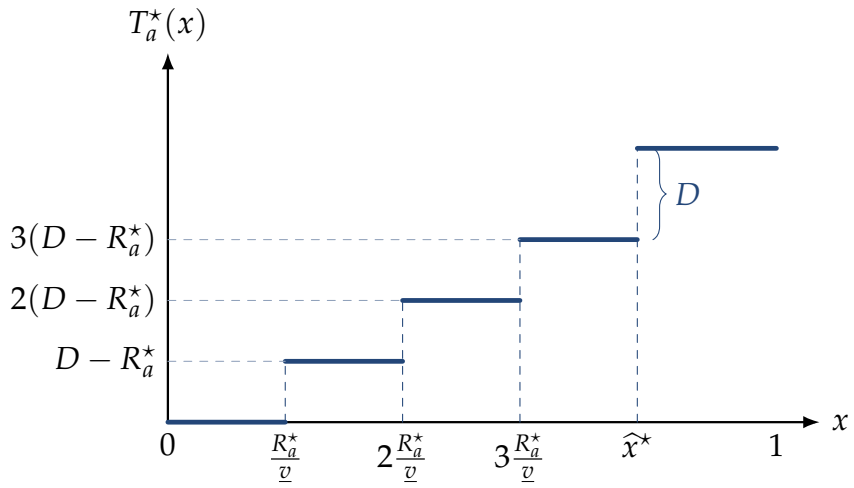


Figure 1: Staircase tax.

The shape arises from the asymmetry between the information the principal has about her payoff and about the agent's bias. The principal can bound how much her payoff changes when the action changes locally, but she cannot do the same for the agent's bias. This implies in particular that regret from small implementation errors is also small, regardless of the location of the first best. Hence, the principal can use a locally flat tax, preserving the value of the agent's information inside each cell without incurring in too much under-correction regret. Importantly, *within* the cell, under-correction is the *only* form of regret the principal is exposed to, since when comparing two actions in the same cell, the agent's choice is undistorted by taxation. Once actions become sufficiently far apart, however, this local guarantee is not enough and some correction is necessary to keep regret low. If the principal had information about the marginal bias, she could fine tune the marginal tax to offset it. In the absence of this information, however, she uses discrete jumps that limit the extent to which even strong bias can induce the agent to move toward higher actions.

Sketch of proof. To make the intuition more precise, suppose the principal wants to guarantee that regret does not exceed some target $r > 0$. Consider a state (V_0, δ_0) with first-best action x_0^*, and suppose that the agent instead chooses some $x_0 > x_0^*$. Then regret is

$$V_0(x_0^*) - V_0(x_0) \leq \underline{v}(x_0 - x_0^*).$$

Hence, whenever

$$x_0 - x_0^* \leq \frac{r}{\underline{v}},$$

regret is bounded by r. This follows from the Lipschitz bound on V, independently of the mechanism chosen by the principal: there is no need to tax differentially actions with a distance $r/\underline{v}$ from each other. Note that this is precisely the size of the steps of the staircase, after replacing the arbitrary r with the optimal regret level $R_a^\star$.

The same guarantee cannot be obtained for actions that are farther apart. To see why, fix $\varepsilon > 0$ and let $r_\varepsilon = r + \varepsilon$. Suppose $x_0 > r_\varepsilon/\underline{v}$, and define

$$x_1^* \equiv x_0 - \frac{r_\varepsilon}{\underline{v}}.$$

Nature can choose the admissible state (V_1, δ_1) given by

$$V_1(z) = \begin{cases} \overline{v}z, & z \leq x_1^*, \\ \overline{v}x_1^* - \underline{v}(z - x_1^*), & z \geq x_1^*, \end{cases} \qquad \delta_1(z) = \begin{cases} 0, & z < x_0, \\ D, & z \geq x_0. \end{cases}$$

Thus V_1 is single-peaked at x_1^*, while the private wedge jumps by its maximal amount D at x_0. By construction,

$$V_1(x_1^*) - V_1(x_0) = \underline{v}(x_0 - x_1^*) = r_\varepsilon > r.$$

Moreover, since V_1 is decreasing to the right of x_1^*, every $z \geq x_0$ generates regret at least $r_\varepsilon > r$. Hence, if T guarantees regret of at most r, no action weakly above x_0 can be chosen in this state. Since x_0 is feasible, the agent must instead choose some $x_1 < x_0$.

Agent optimality then implies

$$V_1(x_1) - T(x_1) > V_1(x_0) + D - T(x_0),$$

where the inequality is strict because x_0 itself cannot be optimal. Rearranging,

$$T(x_0) - T(x_1) > D - [V_1(x_1) - V_1(x_0)].$$

Since x_1^* maximizes V_1,

$$V_1(x_1) - V_1(x_0) \leq V_1(x_1^*) - V_1(x_0) = r_\varepsilon.$$

Therefore,

$$T(x_0) - T(x_1) > D - r_\varepsilon.$$

Thus, for every $\varepsilon > 0$, guaranteeing regret at most r requires the mechanism to generate a tax difference between x_0 and x_1 arbitrarily close to $D - r$. This is the size of the jumps in the staircase, after replacing the arbitrary regret guarantee r with the optimal regret level $R_a^\star$.

The two forces behind the staircase are therefore simple. Over a distance of at most $r/\underline{v}$, the Lipschitz bound on V already limits the principal's loss, so there is no need to tax these actions differently. Once we move beyond this interval, however, Nature can place the first best $r_\varepsilon/\underline{v}$ to the left and let the private wedge jump by D. To preserve the regret guarantee, the mechanism must then induce a lower action, which requires a tax difference of approximately $D - r$. Iterating the same argument along feasible actions generates the staircase: flat cells of length $r/\underline{v}$, separated by jumps of size $D - r$.

The number of intermediate jumps is constrained by the risk of taxing higher actions too much, thereby generating a distortion. However, the principal does not need to extend the sequence of intermediate steps all the way to 1. Once the staircase reaches a sufficiently high threshold $\widehat{x}$, the Lipschitz bound limits the potential additional value of the remaining actions. The principal can therefore impose a final jump of size D, which offsets the maximum possible increase in the agent's private bias when crossing the threshold, without increasing worst-case regret beyond the target level.

3.1.1. Comparative statics

I now describe how changes in the parameters D, $\underline{v}$, and $\overline{v}$ affect the shape of the optimal staircase tax. Throughout this section, I make explicit the dependence of the ambiguity set

and the minimax regret on the primitives, and write $\mathcal{A}_a(D, \underline{v}, \overline{v})$ and $R_a^\star(D, \underline{v}, \overline{v})$, respectively. Let $N^\star(D, \underline{v}, \overline{v})$ denote the set of minimizers of (2), i.e., the set of optimal numbers of jumps before the threshold $\widehat{x}_n^\star$. Let us also define

$$\ell^\star(D, \underline{v}, \overline{v}) := \frac{R_a^\star(D, \underline{v}, \overline{v})}{\underline{v}}, \qquad h^\star(D, \underline{v}, \overline{v}) := D - R_a^\star(D, \underline{v}, \overline{v}).$$

For each $n \in N^\star(D, \underline{v}, \overline{v})$, let $\widehat{x}_n^\star := \min\{(n+1)\ell^\star, 1\}$ denote the corresponding threshold. Then $\ell^\star$ is the length of each full step before $\widehat{x}_n^\star$, while $h^\star$ is the height of each of the n jumps before $\widehat{x}_n^\star$. If $\widehat{x}_n^\star < 1$, the staircase is followed by a final flat step, reached through a final (larger) jump of height D.

Proposition 1. *Let $T_a^\star$ be the optimal mechanism described in Theorem 1. Then,*

- *$R_a^\star(D, \underline{v}, \overline{v})$ is weakly increasing in each of D, $\underline{v}$, and $\overline{v}$;*
- *$\ell^\star(D, \underline{v}, \overline{v})$ is weakly increasing in D and $\overline{v}$, and strictly decreasing in $\underline{v}$;*
- *$h^\star(D, \underline{v}, \overline{v})$ is weakly decreasing in $\underline{v}$ and $\overline{v}$, and strictly increasing in D;*
- *if $D_2 > D_1$, then*

  $$\max N^\star(D_2, \underline{v}, \overline{v}) \leq \min N^\star(D_1, \underline{v}, \overline{v}).$$

  In particular, every selection from $N^\star$ is weakly decreasing in D. Moreover, $N^\star$ is weakly increasing in $\underline{v}$ and $\overline{v}$ in the strong set order.

Proof. Proof in Appendix A.2. □

The first point simply states that the regret guarantee of the principal becomes worse with an increase in any of the parameters. This follows because as $\underline{v}$, $\overline{v}$ or D increase, the set of admissible states (V, δ) also increases, giving more power to Nature to penalize the principal. Formally, for $\overline{v}' \geq \overline{v}$, $\underline{v}' \geq \underline{v}$ and $D' \geq D$, we have $\mathcal{A}_a(D, \underline{v}, \overline{v}) \subseteq \mathcal{A}_a(D', \underline{v}', \overline{v}')$, so, for any $T \in \mathcal{T}$,

$$\sup_{(V,\delta) \in \mathcal{A}_a(D', \underline{v}', \overline{v}')} \mathrm{Reg}(T; V, \delta) \geq \sup_{(V,\delta) \in \mathcal{A}_a(D, \underline{v}, \overline{v})} \mathrm{Reg}(T; V, \delta)$$

As D increases, Nature can select states in which the agent's bias is larger, requiring

stronger correction across flat regions, i.e., larger jumps ($h^\star$ increases). Larger jumps, however, increase the potential loss from over-correcting the agent. The principal therefore finds additional intermediate steps less attractive ($N^\star$ weakly decreases). She may also accept higher regret, allowing wider flat regions ($\ell^\star$ weakly increases). The effect on $\widehat{x}^\star$ is ambiguous, since wider regions extend the staircase while fewer jumps shorten it. For sufficiently large D, the optimal mechanism has no intermediate jumps, and further increases in D leave regret and the length of the flat region unchanged.

As $\underline{v}$ increases, it becomes more costly for the principal to leave discretion to the agent locally. She responds by shortening the flat regions ($\ell^\star$ decreases). With a fixed number of jumps, however, this moves the final jump to the left, increasing the potential loss from over-correcting the agent. She can mitigate this risk by introducing additional intermediate steps, postponing the final jump without lengthening the existing regions ($N^\star$ weakly increases). The principal may also accept higher regret, allowing smaller jumps ($h^\star$ weakly decreases) and partially offsetting the need for shorter flat regions. The net effect on $\ell^\star$ is still negative. The effect on $\widehat{x}^\star$ remains ambiguous, since shorter regions and additional steps work in opposite directions.

As $\overline{v}$ increases, so does the potential value of large actions and, thus, the potential loss from over-correcting the agent. The principal can mitigate this risk by moving the final jump to the right. With a fixed number of jumps, however, this requires longer flat regions, increasing the potential loss from uncorrected deviations within each region. Alternatively, she can introduce additional intermediate steps, postponing the final jump at the cost of greater cumulative tax distortions for a given jump height. As $\overline{v}$ increases, this second option becomes relatively more attractive ($N^\star$ weakly increases). The principal may also accept higher regret, allowing wider flat regions ($\ell^\star$ weakly increases) and smaller jumps ($h^\star$ weakly decreases). The cutoff $\widehat{x}^\star$ moves weakly to the right, since both wider regions and additional intermediate steps extend the staircase.

3.2. Bounds on marginal bias: the kinked tax

Theorem 2. *Let $\mathcal{A} = \mathcal{A}_m$ and define*

$$d \equiv \overline{\delta} - \underline{\delta}, \qquad \alpha \equiv \min\{\underline{v}, d\}, \qquad \beta \equiv \min\{\overline{v}, d\}.$$

Let

$$R_m^\star \equiv \frac{\alpha\beta}{\alpha+\beta}, \qquad c^\star \equiv \frac{\beta}{\alpha+\beta}. \tag{3}$$

Then:

(i) $R_m^\star$ *is a lower bound on regret, i.e., for any* $T \in \mathcal{T}$, $R_m(T) \geq R_m^\star$;

(ii) The mechanism $T_m^\star$ *defined by*

$$T_m^\star(x) = \underline{\delta} x + d(x - c^\star)_+ \tag{4}$$

achieves $R_m^\star$.

Proof. See Appendix A.3. □

The principal faces the same trade-off between over and under-correction as she did when the ambiguity set was $\mathcal{A}_a$, except now she has some knowledge about the agent's marginal bias. In particular, she knows that the marginal bias lies in the positive interval $[\underline{\delta}, \overline{\delta}]$. She uses this local information to target the correction. She first partitions the action set in two intervals separated by a cutoff, $c^\star$. To the left of $c^\star$ she is concerned with minimizing over-correction regret, and therefore sets the marginal tax equal to the minimum marginal bias, $\underline{\delta}$. To the right of $c^\star$, she minimizes under-correction regret, setting the marginal tax equal to the maximum marginal bias, $\overline{\delta}$. This structure of the tax addresses the two sources of regret separately: it focuses on mitigating over-correction below $c^\star$ and under-correction above $c^\star$. The value of the cutoff is chosen to balance these two sources of regret.

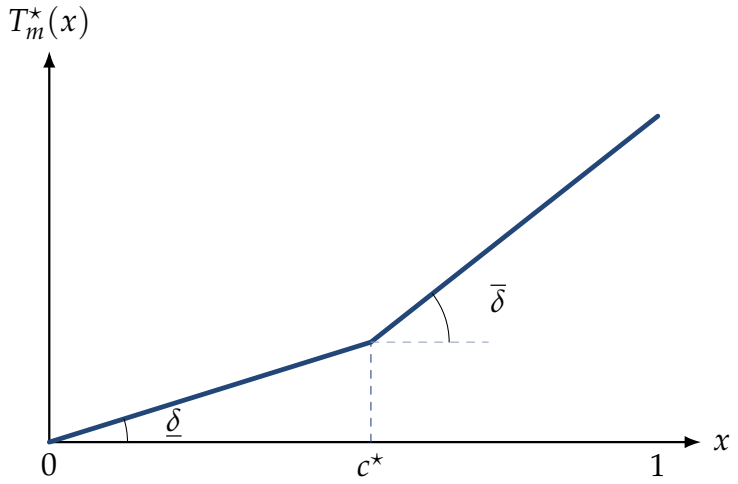


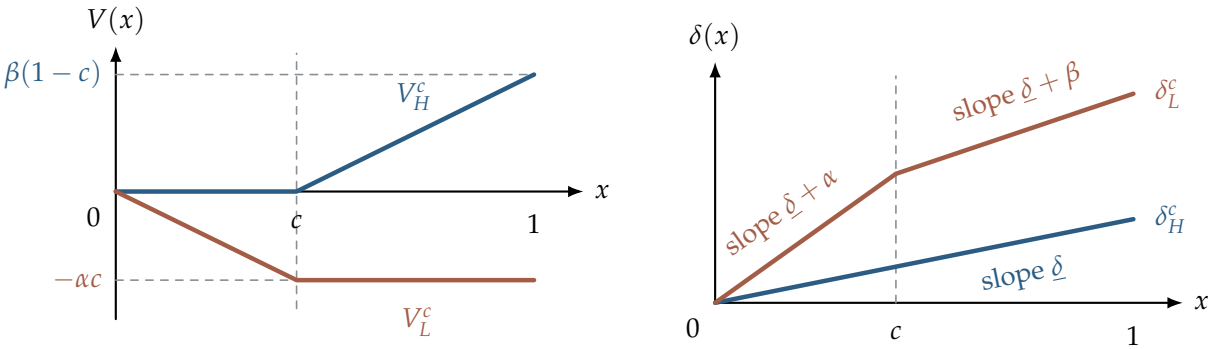
Figure 2: Kinked tax.

Sketch of the proof. The proof of optimality of the kinked tax first establishes a lower bound on the regret attainable by any mechanism, and then shows that the kinked tax achieves it. The full proof is in Appendix A.3. Here I illustrate the construction of the lower bound, as it provides intuition for the shape of the optimal tax. In particular, it identifies two states that are particularly problematic for the principal, since they are "maximally different" from her perspective (she prefers opposite actions), but the agent has the same preferences in both. Thus, any mechanism cannot correct the incentives of the agent in one state without affecting the incentives in the other state, generating an unavoidable amount of regret.

Fix an arbitrary cutoff $c \in (0,1)$ and consider two admissible states (V_H^c, δ_H^c) and (V_L^c, δ_L^c) defined as follows

$$\begin{aligned} V_H^c(x) &= \beta(x-c)_+, & \delta_H^c(x) &= \underline{\delta} x, \\ V_L^c(x) &= -\alpha \min\{x, c\}, & \delta_L^c(x) &= \underline{\delta} x + \alpha \min\{x, c\} + \beta(x-c)_+. \end{aligned}$$

The principal's payoffs and the agent's biases in the two states are illustrated in Figure 3.



Principal payoffs Agent biases

Figure 3: Lower-bound construction for the kinked tax.

By construction, the agent's payoff before transfers is identical in the two states:

$$V_H^c(x) + \delta_H^c(x) = V_L^c(x) + \delta_L^c(x) = \underline{\delta}x + \beta(x-c)_+.$$

Thus, any mechanism induces the same set of optimal actions in both states, even though the principal prefers opposite actions. If the agent chooses $x \leq c$, regret in the high state is at least $\beta(1-c)$. If instead the agent chooses $x \geq c$, regret in the low state is at least αc. Therefore, any mechanism T must generate regret

$$R_m(T) \geq \min\{\beta(1-c), \alpha c\}.$$

Since this holds for any mechanism and every $c \in (0,1)$, we can maximize this lower bound over c. Since $\beta(1-c)$ is decreasing in c, while αc is increasing, their minimum is maximized when $\beta(1-c) = \alpha c$. This gives $c^* = \beta/(\alpha+\beta)$ and regret $R_m^\star = \alpha\beta/(\alpha+\beta)$.

This explains why the kinked tax is a natural candidate for the minimax-optimal mechanism. In state H, the agent has the smallest possible marginal bias, so that aggressive taxation risks inducing an action that is too low. In state L, instead, the agent has a large bias, so that insufficient taxation risks inducing an action that is too high. Since the principal cannot distinguish between these states, she cannot eliminate both sources of regret simultaneously. The kinked tax resolves this trade-off by addressing the two risks separately. Below the cutoff, the principal sets the marginal tax equal to the smallest possible marginal bias, $\underline{\delta}$. Above the cutoff, she instead sets the marginal tax equal to the

largest possible marginal bias, $\overline{\delta}$.

3.2.1. Comparative statics

Fixing $\underline{\delta}$, the shape of the kinked penalty is completely pinned down by the value of the cutoff $c^\star(d, \underline{v}, \overline{v})$ and the slope d. The proposition below illustrates how the cutoff changes with the parameters.

Proposition 2. *Let $T_m^\star$ be the optimal mechanism described in Theorem 2. Then,*

- *$R_m^\star(d, \underline{v}, \overline{v})$ is weakly increasing in each of d, $\underline{v}$ and $\overline{v}$;*
- *$c^\star(d, \underline{v}, \overline{v})$ is weakly decreasing in $\underline{v}$ (strictly if $\underline{v} < d$), weakly increasing in $\overline{v}$ (strictly if $\overline{v} < d$). Moreover,*
  - *It is constant in d if $\underline{v}, \overline{v} > d$ or $\underline{v}, \overline{v} < d$;*
  - *It is strictly increasing in d if $\underline{v} < d < \overline{v}$;*
  - *It is strictly decreasing in d if $\overline{v} < d < \underline{v}$.*

Proof. The result follows immediately by differentiating $R_m^\star$ and $c^\star$ with respect to d, $\underline{v}$, and $\overline{v}$ on each region where they are differentiable, and using continuity at the boundaries $d = \underline{v}$ and $d = \overline{v}$. □

As for the staircase tax, worst-case regret $R_m^\star(d, \underline{v}, \overline{v})$ is (weakly) increasing in all the parameters (see Section 3.1.1). The cutoff is weakly decreasing in $\underline{v}$ (strictly if $\underline{v} < d$) and weakly increasing in $\overline{v}$ (strictly if $\overline{v} < d$). This is because a larger $\underline{v}$ makes under-correction more costly, inducing the principal to expand the high-tax region by shifting the cutoff to the left. The opposite is true for $\overline{v}$. The effect of an increase in $d = \overline{\delta} - \underline{\delta}$ depends on its relative size with respect to $\underline{v}$ and $\overline{v}$. When d is either smaller or larger than both bounds, the cutoff remains unchanged. If instead $\underline{v} < d < \overline{v}$, an increase in d strictly increases $c^\star$. To see why, note that in this case the potential regret from under-correction is already limited by the Lipschitz bound $\underline{v}$, while over-correction regret is limited by the maximum rate at which bias increases, d, since this determines the tax rate in the high-tax region. Thus, increasing d makes over-correction more costly without increasing the potential loss

from under-correction. The principal responds by shifting the cutoff to the right, reducing the high-tax region. The opposite occurs when $\overline{v} < d < \underline{v}$, in which case increasing d shifts the cutoff to the left.

3.3. Discussion

In this section, I examine additional properties of the optimal mechanisms. First, I show that both the staircase tax of Theorem 1 and the kinked tax of Theorem 2 are undominated in their respective ambiguity sets. Undominatedness provides a natural refinement of robust optimality (Börgers, 2017; Dworczak and Pavan, 2022; Li, Börgers, and Wang, 2024). When a minimax optimal mechanism is dominated, there exists another mechanism that is also minimax optimal but performs weakly better in every admissible state and strictly better in at least one state. Thus, it is reasonable to assume that a rational decision maker would restrict attention to undominated mechanisms. I start by formally defining dominance in this setting.

Definition 1. Let $T_1, T_2 \in \mathcal{T}$ be feasible mechanisms. We say T_1 *dominates* T_2 in $\mathcal{A}$ if

$$\mathrm{Reg}(T_1; V, \delta) \leq \mathrm{Reg}(T_2; V, \delta)$$

for all $(V, \delta) \in \mathcal{A}$, with a strict inequality for at least one state (V, δ). We say that a mechanism T is *undominated* if there does not exist another feasible mechanism T' that dominates it.

Then, we have the following result.

Proposition 3. *Let $\mathcal{A} = \mathcal{A}_a$ and let $T_a^\star$ be the staircase penalty characterized in Theorem 1. Then $T_a^\star$ is undominated among all feasible mechanisms.*

And the corresponding result for the kinked penalty.

Proposition 4. *Let $\mathcal{A} = \mathcal{A}_m$ and let $T_m^\star$ be the kinked mechanism characterized in Theorem 2. Then $T_m^\star$ is undominated among all feasible mechanisms.*

The proofs of Propositions 3 and 4 are in Appendix A.4 and A.5, respectively.

Intuitively, the undominatedness results rely on the richness of the state space. When

comparing the optimal mechanism (staircase or kinked tax) with an alternative that changes the relative incentives it provides across any pair of actions, there exists an admissible state in which the optimal mechanism induces the principal's first-best action, while the alternative induces a sub-optimal one. Because differences in principal's payoffs can be arbitrarily small, even a small modification of the mechanism can have this effect[5].

A second observation is that we have been assuming the principal has local information about her payoff V: she knows that actions that are close together result in similar payoffs. It is natural to ask how the mechanisms we identified perform when the principal has additional, global information about the shape of V. The following corollary shows that the staircase and the kinked tax continue to be minimax optimal when V is also restricted to be (weakly) single-peaked.

Corollary 1. *Let $\mathcal{V}_{\text{sp}} \subseteq \mathcal{V}$ denote the set of single-peaked payoff functions. For $i \in \{a, m\}$, let $\mathcal{A}_{i,\text{sp}} = \mathcal{V}_{\text{sp}} \times \mathcal{B}_i$ and denote by $R_{i,\text{sp}}(T)$ the worst-case regret of mechanism T over this restricted ambiguity set. Then, $\inf_{T \in \mathcal{T}} R_{a,\text{sp}}(T) = R_{a,\text{sp}}(T_a^\star) = R_a^\star$ and $\inf_{T \in \mathcal{T}} R_{m,\text{sp}}(T) = R_{m,\text{sp}}(T_m^\star) = R_m^\star$..*

The proof follow immediately from the arguments used to establish Theorems 1 and 2 in Appendix A.1 and A.3, informally described in Sections 3.1 and 3.2. There, I first identify a lower bound on regret that any mechanism must induce, and then show that the staircase and the kinked tax achieve that bound, respectively. Constructing the lower bounds requires identifying critical states that are particularly problematic for any mechanism. Both in the staircase and in the kinked tax case, these states only involve weakly single-peaked payoff functions for the principal. Hence, even when we restrict V to be single-peaked, the lower bounds on regret continue to hold. Moreover, restricting the state space cannot increase the worst-case regret of a given mechanism, so the staircase and the kinked tax continue to achieve these bounds and remain minimax optimal.

In Section 3.4 I consider what happens when V is also known to be concave. I show that, under symmetric slope bounds, a staircase remains optimal, although its parameters

[5]This same argument suggests that, in this framework, undominatedness does not uniquely select a minimax-optimal mechanism. Nevertheless, it rules out certain minimax-optimal alternatives. For instance, when the optimal staircase has a nonempty terminal region, replacing its final jump of size D with a hard cap at $\widehat{x}^\star$ preserves minimax optimality but results in a dominated mechanism.

may change, whereas the kinked tax generally fails to achieve the minimax regret. I then characterize the optimal mechanism in this case.

3.4. Concave V

In the baseline model, the principal only knows bounds on her marginal payoff. In some applications, however, she may have additional information about the shape of her payoff function. In this section, I consider the case where the principal also knows that V is concave.[6] This additional information restricts the ambiguity set and can therefore only weakly reduce the principal's minimax regret. As discussed in Section 3.3, knowing that V is single-peaked is not enough to reduce regret. I now show that concavity, by contrast, provides valuable information that can strictly reduce minimax regret. Moreover, although jumps and kinks remain features of optimal mechanisms under these smaller ambiguity sets, the additional assumption has different implications for the shape of optimal mechanisms depending on the scenario considered. When the principal only knows a bound on the absolute bias, a staircase penalty remains optimal under symmetric Lipschitz bounds. By contrast, when the principal knows bounds on the marginal bias, the piecewise-linear kinked tax introduced in Section 3.2 is no longer generally optimal. Instead, a minimax-optimal mechanism takes the form of a piecewise-logarithmic tax, which has a kink when $d > \underline{v} + \overline{v}$. Let us begin by strengthening Assumption 1 as follows.

Assumption 2. Let $\mathcal{V}$ be the ambiguity set specified in Assumption 1. Then, V belongs to the ambiguity set

$$\mathcal{V}_{\text{cv}} := \{V \in \mathcal{V} : V \text{ is concave}\}$$

Then, define $\mathcal{A}_{i,\text{cv}} = \mathcal{V}_{\text{cv}} \times \mathcal{B}_i$, and the corresponding worst-case regret $R_{i,\text{cv}}(T)$, for $i \in \{a, m\}$.

I first consider the ambiguity about the agent's bias studied in Section 3.1, where $\delta \in \mathcal{B}_a$: the principal knows bounds on the absolute bias but not on the marginal bias. I further assume that the Lipschitz bounds are symmetric, i.e., $\underline{v} = \overline{v} = v$.

[6] For instance, in the hiring application considered in Section 4.2, the marginal productivity of workers selected from a given group is weakly decreasing in the mass hired from that group. As a result, the firm's payoff is concave in the composition of its workforce.

Theorem 3. *Let $\underline{v} = \overline{v} = v$ and $\mathcal{A} = \mathcal{A}_{a,\text{cv}}$. Then the minimax regret is*

$$R^\star_{a,\text{cv}} := \inf_{T \in \mathcal{T}} R_{a,\text{cv}}(T) = \min\left\{\frac{v}{2}, \frac{vD}{v+D}\right\}. \tag{5}$$

Let

$$n^\star_{\text{cv}} := \left\lfloor \frac{v}{D} \right\rfloor, \qquad \ell^\star_{\text{cv}} := \frac{R^\star_{a,\text{cv}}}{v}, \qquad h^\star_{\text{cv}} := D - R^\star_{a,\text{cv}}, \qquad \widehat{x}^\star_{\text{cv}} := (n^\star_{\text{cv}} + 1)\ell^\star_{\text{cv}}. \tag{6}$$

Then $\widehat{x}^\star_{\text{cv}} \leq 1$, and the mechanism $T^\star_{a,\text{cv}}$, such that $T^\star_{a,\text{cv}}(0) = 0$ and

$$T^\star_{a,\text{cv}}(x) = \begin{cases} jh^\star_{\text{cv}}, & j\ell^\star_{\text{cv}} < x \leq (j+1)\ell^\star_{\text{cv}}, \quad j = 0, \ldots, n^\star_{\text{cv}}, \\ n^\star_{\text{cv}} h^\star_{\text{cv}} + D, & \widehat{x}^\star_{\text{cv}} < x \leq 1, \end{cases} \tag{7}$$

achieves $R^\star_{a,\text{cv}}$.

Proof. See Appendix A.7. □

The theorem shows that even with this additional knowledge about the shape of V, the optimal corrective intervention still takes the form of a staircase, as in Theorem 1. The $n^\star_{\text{cv}} + 1$ flat regions before $\widehat{x}^\star_{\text{cv}}$ each have length $\ell^\star_{\text{cv}}$, and are separated by jumps of size $h^\star_{\text{cv}}$. If $\widehat{x}^\star_{\text{cv}} < 1$, the staircase is followed by a final jump of size D and a terminal flat region. It is then interesting to analyze the effect of concavity on the individual elements of the staircase: the length of the flat parts and the size of the vertical jumps. This is the subject of the following corollary.

Corollary 2. *Suppose $\underline{v} = \overline{v} = v > 0$. Let $R^\star_a$, $\ell^\star$, and $h^\star$ be defined as in Theorem 1, and $R^\star_{a,cv}$, $\ell^\star_{cv}$, and $h^\star_{cv}$ as in Theorem 3. Then,*

$$R^\star_{a,cv} \leq R^\star_a, \qquad \ell^\star_{cv} \leq \ell^\star, \qquad h^\star_{cv} \geq h^\star.$$

All three inequalities are strict if and only if $v > D$ and $v/D \notin \mathbb{N}$.

Proof. See Appendix A.8. □

Concavity can reduce over-correction regret by allowing the principal to infer the value of higher actions from the agent's refusal to cross an earlier tax threshold. To see this,

suppose the agent chooses the right endpoint x of a flat region rather than the right endpoint $z = x + \ell$ of the next one, where ℓ is the step length and $h = T(z) - T(x)$ is the intermediate jump height. Incentive compatibility and the monotonicity of δ imply $V(z) - V(x) \leq h$.

Consider now a higher action $y > z$. Without concavity, the principal's payoff may increase at rate v beyond z. Combining Lipschitzness with incentive compatibility gives

$$V(y) - V(x) \leq \min\{T(y) - T(x),\, h + v(y - z)\}.$$

Under concavity, instead, the principal's payoff cannot increase more rapidly beyond z than it did between x and z. Since the agent's choice implies that the payoff gain over this interval of length ℓ is at most h, concavity gives

$$V(y) - V(x) \leq \min\left\{T(y) - T(x),\, h + \frac{h}{\ell}(y - z)\right\}.$$

The concavity bound is tighter whenever $h/\ell < v$ and $T(y) - T(x) > h + (h/\ell)(y - z)$. Thus, concavity can limit the regret from over-correction even when higher actions are taxed more heavily. This allows the principal to impose larger intermediate jumps and shorten the flat regions of the staircase, thereby reducing under-correction regret.

Next, I consider the case where the principal knows bounds on the marginal bias, i.e. $\delta \in \mathcal{B}_m$.

Theorem 4. *Let $\mathcal{A}_{m,\mathrm{cv}}$ be the ambiguity set $\mathcal{A}_m$ restricted to concave V, and suppose $\underline{v}, \overline{v} > 0$. Write $d := \overline{\delta} - \underline{\delta}$. Define*

$$R^{\star}_{m,\mathrm{cv}} := \max_{c \in [0,1]} \min\{\underline{v}c, \overline{v}(1-c), dc(1-c)\}, \tag{8}$$

and let $c^{\star}_{\mathrm{cv}}$ be a maximizer. Then

$$\inf_{T \in \mathcal{T}} R_{m,\mathrm{cv}}(T) = R^{\star}_{m,\mathrm{cv}},$$

and the following tax achieves $R^\star_{m,\mathrm{cv}}$:

$$T^\star_{m,\mathrm{cv}}(x) = \underline{\delta} x + \begin{cases} R^\star_{m,\mathrm{cv}} \log \dfrac{1}{1-x}, & 0 \le x \le c^\star_{\mathrm{cv}}, \\ R^\star_{m,\mathrm{cv}} \log \dfrac{1}{1-c^\star_{\mathrm{cv}}} + d(x - c^\star_{\mathrm{cv}}) - R^\star_{m,\mathrm{cv}} \log \dfrac{x}{c^\star_{\mathrm{cv}}}, & c^\star_{\mathrm{cv}} < x \le 1. \end{cases} \tag{9}$$

Proof. See Appendix A.9. □

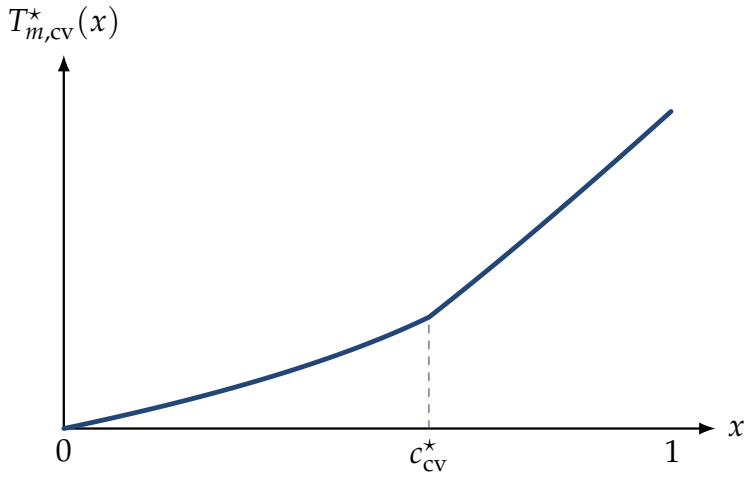


Figure 4: Kinked piecewise-logarithmic tax.

The intuition behind the optimality of the logarithmic tax is, once again, that concavity allows the principal to use the agent's local incentives to bound her payoff from actions farther away from the one chosen. To see this, normalize $\underline{\delta} = 0$, so that $d = \overline{\delta}$, and suppose the agent chooses x, while x^* maximizes V. Consider an arbitrary tax T and assume, for simplicity, that it admits one-sided derivatives at x. Incentive compatibility gives bounds on the principal's marginal payoff at x. Since V is concave, its marginal value is nonincreasing, so these local bounds also restrict the payoff gain from moving to any other

action. In particular, we obtain[7]

$$V(x^*) - V(x) \leq \begin{cases} (1-x)\min\{\overline{v}, T'_+(x)\}, & x^* > x, \\ x\min\{\underline{v}, d - T'_-(x)\}, & x^* < x. \end{cases}$$

The principal can exploit this information to adjust the marginal correction to the action chosen by the agent. Suppose she wants to guarantee regret no greater than some target $r > 0$. To address over-correction, a sufficient condition is

$$(1-x)T'_+(x) \leq r \quad \Longrightarrow \quad T'_+(x) \leq \frac{r}{1-x}.$$

In particular, the principal can choose the largest marginal tax compatible with the bound, $T'(x) = r/(1-x)$. As x increases, the remaining distance to higher actions shrinks, allowing the principal to impose a higher marginal tax without increasing the bound on over-correction regret. Conversely, a sufficient condition to control under-correction regret is

$$x[d - T'_-(x)] \leq r \quad \Longrightarrow \quad T'_-(x) \geq d - \frac{r}{x}.$$

The principal can then choose the smallest marginal tax compatible with this constraint, $T'(x) = d - r/x$. As x increases, the distance to lower actions grows, requiring stronger correction to keep under-correction regret bounded by r. As shown in Appendix A.9, the principal can combine these two rules by dividing the action space into two regions separated by a cutoff. Below the cutoff, the marginal tax is determined by the bound on over-correction regret; above the cutoff, it is determined by the bound on under-correction regret. The cutoff $c^\star_{\text{cv}}$ is chosen so that both sources of regret remain bounded by $r = R^\star_{m,\text{cv}}$ on either side. Integrating the resulting marginal tax rates gives the piecewise-logarithmic schedule. Note that the kink at c^*_{cv} arises only in a specific parameter region, where $d > \underline{v} + \overline{v}$, otherwise the tax is smooth.

[7]The agent's optimality condition implies $V(y) - V(x) \leq T(y) - T(x) + \delta(x) - \delta(y)$ for every y. Since δ is nondecreasing, taking $y \downarrow x$ gives $V'_+(x) \leq T'_+(x)$. By concavity, $V(x^*) - V(x) \leq V'_+(x)(x^* - x)$ whenever $x^* > x$. Combining these inequalities with the upper slope bound on V and $x^* \leq 1$ gives the first bound. Similarly, using $\delta(x) - \delta(y) \leq d(x-y)$ for $y < x$ and taking $y \uparrow x$ gives $-V'_-(x) \leq d - T'_-(x)$. Concavity and the lower slope bound on V then imply the second bound.

3.5. Optimal delegation mechanisms

So far, we have allowed the principal to use action-contingent penalties to correct the agent's incentives. In some settings, however, such instruments may not be available, and the principal can only intervene by restricting the set of actions available to the agent. This problem has been the main focus of the literature on optimal delegation, starting with Holmström (1977). Formally, suppose that the principal can commit to a delegation set before the state (V, δ) is realized, and the agent chooses his preferred action among those that are available. A delegation mechanism is a nonempty compact set $S \subseteq X$. Let

$$\mathcal{S} = \{S \subseteq X : S \text{ is nonempty and compact}\}$$

denote the set of feasible delegation mechanisms. Given a delegation set S and a state (V, δ), the agent chooses from

$$\Gamma(S; V, \delta) = \arg\max_{x \in S} \{V(x) + \delta(x)\}.$$

Since S is compact and the agent's payoff is upper semicontinuous, the choice correspondence is nonempty in every admissible state.

For an ambiguity set $\mathcal{A}_i$, $i \in \{a, m\}$, worst-case regret is

$$R_i(S) = \sup_{(V,\delta) \in \mathcal{A}_i} \operatorname{Reg}(S; V, \delta),$$

and the principal solves

$$R_i^{\star,\mathrm{del}} = \inf_{S \in \mathcal{S}} R_i(S).$$

Note that delegation mechanisms are a special case of the mechanisms considered in the baseline model. Indeed, any delegation set $S \in \mathcal{S}$ can be represented by the tax schedule

$$T_S(x) = \begin{cases} 0, & x \in S, \\ +\infty, & x \notin S. \end{cases}$$

Hence, restricting the principal's instruments to delegation mechanisms weakly increases minimax regret: $\inf_{T \in \mathcal{T}} R_i(T) \leq R_i^{\star,\mathrm{del}}$.

As the following result shows, under both ambiguity sets, the principal can achieve minimax regret by imposing a simple cap on the set of feasible actions.

Proposition 5. *Suppose the principal is restricted to delegation mechanisms $S \in \mathcal{S}$. Then:*

(i) ***Bound on absolute bias.*** *Let $\mathcal{A} = \mathcal{A}_a$. Then the minimax regret under delegation is $R_a^{\star,\text{del}} = \min\left\{D, \frac{\underline{v}\overline{v}}{\underline{v}+\overline{v}}\right\}$. The delegation set $S_a^\star = [0, c_a^{\star,\text{del}}]$, where*

$$c_a^{\star,\text{del}} = \begin{cases} 1, & D \leq \dfrac{\underline{v}\overline{v}}{\underline{v}+\overline{v}}, \\ \dfrac{\overline{v}}{\underline{v}+\overline{v}}, & D > \dfrac{\underline{v}\overline{v}}{\underline{v}+\overline{v}}, \end{cases}$$

achieves $R_a^{\star,\text{del}}$.

(ii) ***Bound on marginal bias.*** *Let $\mathcal{A} = \mathcal{A}_m$. Then the minimax regret under delegation is $R_m^{\star,\text{del}} = \frac{\min\{\underline{v},\overline{\delta}\}\overline{v}}{\min\{\underline{v},\overline{\delta}\}+\overline{v}}$. The delegation set $S_m^\star = [0, c_m^{\star,\text{del}}]$, where*

$$c_m^{\star,\text{del}} = \frac{\overline{v}}{\min\{\underline{v},\overline{\delta}\}+\overline{v}},$$

achieves $R_m^{\star,\text{del}}$.

Proof. See Appendix A.6. □

Recall that in both ambiguity scenarios, $\mathcal{A}_a$ and $\mathcal{A}_m$, the agent is (weakly) biased upwards with respect to the principal. This means that, in the absence of action-contingent penalties, whenever a first-best action is feasible, the agent will never choose a *lower* action that is strictly worse for the principal. On the other hand, if the principal caps the set of available actions below her first best, then she may distort the choice of the agent even if he is unbiased. Even in this case, any action selected by the agent must give the principal at least as high a payoff as the highest available action. This suggests that, for any delegation set S, its largest element plays a central role in determining a lower bound on worst-case regret.

To formalize the intuition, fix some arbitrary $S \in \mathcal{S}$ and consider $\bar{s} = \max S$. Whenever $\bar{s} < 1$, the principal risks capping the action of an agent who is unbiased, when the state is

actually such that higher actions are optimal for the principal. In particular, consider the admissible state (V_H, δ_H) in which $\delta_H(x) = 0$ and $V_H(x) = \overline{v}x$. Here, the agent would pick the highest available action, $\bar{s}$, and regret is $\overline{v}(1 - \bar{s})$. Setting a high $\bar{s}$ mitigates this risk, but at the same time it exposes the principal to the opposite problem: the agent might be biased and the first best lie far below the highest action. Consider the state $V_L(x) = -kx$ and $\delta_L(x) = k \min\{x, \bar{s}\}$, where $k := \min\{\underline{v}, D/\bar{s}\}$. This choice of k ensures that the pair (V_L, δ_L) is admissible. Here the first best is $x = 0$, but the agent's payoff is such that he is indifferent among all actions in S: $V_L(x) + \delta_L(x) = 0$ for all $x \in S$. Under adversarial tie-breaking, Nature can then select $\bar{s}$, generating a regret of $k\bar{s}$. Putting the two cases together, this means that any delegation set S with highest action $\bar{s}$ generates a regret of at least $\max\{\overline{v}(1 - \bar{s}), k\bar{s}\}$ Since the principal controls $\bar{s}$, she can set it equal to the value that minimizes this lower bound on regret, which is

$$c_a^{\star,\text{del}} = \begin{cases} \frac{\overline{v}}{\underline{v}+\overline{v}} \text{ if } D > \frac{\underline{v}\overline{v}}{\underline{v}+\overline{v}} \\ 1 \text{ if } D \leq \frac{\underline{v}\overline{v}}{\underline{v}+\overline{v}} \end{cases}$$

But then, the capped delegation set $[0, c_a^{\star,\text{del}}]$ is a natural candidate for optimality. Appendix A.6 completes the argument showing that such set it actually achieves the lower bound.

4. Applications

4.1. Corrective taxation: the case of externalities

A potential source of misalignment between the principal's and the agent's payoffs is that the agent does not internalize the environmental harm generated by his actions. In regulating emissions, the principal may have limited information about the agent's private production or abatement costs (Dasgupta, Hammond, and Maskin, 1980), while also being unable to precisely quantify the social damages associated with emissions (Weitzman, 1974; Rudik, 2020). The complexity of this type of regulatory problems makes them particularly suited to a robust approach (e.g., Kang, 2026). In this section, I consider an application of my framework to an emissions regulation problem.

Consider an environmental regulator who faces a firm whose production process generates

a baseline emission level of $\bar{x}$. Emissions have a negative effect on the environment, generating a social cost $H(x)$ when the firm emits x. The firm can reduce emissions by undertaking costly abatement. Let $C(a)$ denote the cost of abating a units of emissions, so that choosing an emission level $x \in [0, \bar{x}]$ entails an abatement cost $C(\bar{x} - x)$. The private benefit of higher emissions is therefore the saved abatement cost, so the firm's payoff is $U(x) = C(\bar{x}) - C(\bar{x} - x)$. In what follows, I normalize $\bar{x} = 1$.

The regulator can set a tax T on the firm as a function of the (observed) level of emissions, x. Her objective is to achieve the efficient level of pollution, that maximizes welfare

$$V(x) = [C(\bar{x}) - C(\bar{x} - x)] - H(x)$$

In this example, the misalignment between payoffs coincides with the external harm, $\delta(x) = H(x)$.

The regulator does not observe the firm's abatement-cost function and cannot precisely determine the damages associated with each level of emissions. This implies that she faces the same trade-off between over and under-correction described in Section 2. Too weak interventions may cause the firm to pollute too much, while excessive taxation may generate excessive social costs in terms of costly abatement.

Assumption 3. The regulator knows constants $0 < \underline{k} < \overline{k}$ such that, for all $0 \leq a < a' \leq 1$,

$$\underline{k}(a' - a) \leq C(a') - C(a) \leq \overline{k}(a' - a)$$

Moreover, she also knows constants $0 < \underline{h} < \overline{h}$ such that, for all $0 \leq x < x' \leq 1$,

$$\underline{h}(x' - x) \leq H(x') - H(x) \leq \overline{h}(x' - x)$$

The first inequality bounds the marginal cost of abatement on the regulated range, while the second bounds the marginal environmental harm. Such bounds could in principle be constructed starting from conservative engineering estimates of the costs of feasible abatement technologies and estimates of the relationship between emissions and monetary value of the resulting environmental and health damages (Carlson et al., 2000; N. Z. Muller and Mendelsohn, 2009).

Let $\mathcal{C}$ and $\mathcal{H}$ denote the sets of all functions satisfying the respective bounds in Assumption 3, with the normalizations $C(0) = H(0) = 0$. Let $\mathcal{A}_{\text{em}}$ be the set of induced pairs (V, δ) as (C, H) ranges over $\mathcal{C} \times \mathcal{H}$. The primitive restrictions imply, for all $x < y$,

$$\begin{aligned} -(\overline{h} - \underline{k})(y - x) \leq V(y) - V(x) \leq (\overline{k} - \underline{h})(y - x), \\ \underline{h}(y - x) \leq \delta(y) - \delta(x) \leq \overline{h}(y - x). \end{aligned} \tag{10}$$

I also make the following assumption

$$\overline{k} > \underline{h}, \qquad \underline{k} < \overline{h} \tag{11}$$

This is to ensure the regulator's problem does not have a trivial solution. Indeed, if the first condition is violated, all the admissible value functions V are monotonically decreasing in x: the highest admissible marginal abatement cost is lower than the smallest marginal damage from emissions, and the principal can achieve zero regret by mandating the highest possible level of abatement. Similarly, if the second condition is violated, then all admissible V are increasing in x, and the optimum is not to regulate the firm. So, we can focus on the parameter region where (11) holds.

Note that, by Assumption 3, the functions $V(x) = C(1) - C(1 - x) - H(x)$ and $\delta(x) = H(x)$ satisfy the bounds of Theorem 2, with

$$\underline{v} = \overline{h} - \underline{k}, \qquad \overline{v} = \overline{k} - \underline{h}, \qquad \underline{\delta} = \underline{h}, \qquad \overline{\delta} = \overline{h}. \tag{12}$$

Thus, the set of feasible states $\mathcal{A}_{\text{em}}$ is a subset of the ambiguity set $\mathcal{V} \times \mathcal{B}_m$ in Theorem 2. However, $\mathcal{A}_{\text{em}}$ is non-rectangular: in addition to (12), admissible pairs (V, δ) must also satisfy

$$\underline{k}(y - x) \leq [V(y) + \delta(y)] - [V(x) + \delta(x)] \leq \overline{k}(y - x).$$

This means that, while Theorem 2 still provides an upper bound on the regret of the kinked tax, in principle, its lower bound need not automatically remain valid for this smaller ambiguity set. The following lemma verifies that it does.

Lemma 1. *Suppose Assumption 3 and condition (11) hold. Define* $d = \overline{h} - \underline{h}$, $\alpha = \min\{\overline{h} - \underline{k}, d\}$ *and* $\beta = \min\{\overline{k} - \underline{h}, d\}$. *Then the minimax regret over* $\mathcal{A}_{\text{em}}$ *is* $R^{\star}_{\text{em}} = \frac{\alpha\beta}{\alpha+\beta}$, *and a minimax-*

optimal emissions tax is

$$T^{\star}_{\text{em}}(x) = \underline{h}x + d(x - c^{\star}_{\text{em}})_{+}, \qquad c^{\star}_{\text{em}} = \frac{\beta}{\alpha + \beta}. \tag{13}$$

The lower-bound construction follows the same intuition as in Theorem 2. Since the firm does not internalize environmental damages, we can consider two states with the same abatement-cost function but different environmental damages. In one state, damages are sufficiently low that the socially optimal choice is not to abate, while in the other they are sufficiently high that full abatement is optimal. However, since the firm's private payoff is identical in the two states, any tax induces the same emissions choices in both. The regulator therefore faces the same trade-off between over and under-correction as in the baseline model.

The abatement-cost function used in the proof has a clear technological interpretation (see, for instance, Burtraw et al., 2001; Fowlie and N. Muller, 2019). The firm can choose between two abatement technologies. The first (the status quo) has no fixed cost and marginal abatement cost $k_0 = \min\{\overline{k}, \overline{h}\}$, while the second (the "green" technology) requires a fixed adoption cost K_{ξ} but has a lower marginal abatement cost $k_{\text{G}} = \max\{\underline{k}, \underline{h}\}$. The firm chooses the cheaper technology for each level of abatement a, resulting in

$$C^{(\xi)}(a) = \min\{k_0 a, K_{\xi} + k_{\text{G}} a\}, \qquad K_{\xi} = (k_0 - k_{\text{G}})(1 - \xi).$$

The green technology is adopted when abatement exceeds $1 - \xi$, since the savings in variable abatement costs then outweigh its fixed adoption cost.

4.2. Hiring discrimination

Consider a firm that must fill a fixed mass $q \in (0, 1)$ of positions from a unit mass of applicants. The firm delegates the choice of workers to a hiring manager. Each applicant belongs to one of two groups, $g \in \{A, B\}$, and has productivity $w \in [\underline{w}, \overline{w}]$, with $\underline{w} < \overline{w}$. Productivity can be interpreted as the expected contribution of the applicant to firm output (net of wages), conditional on the information available to the hiring manager. The masses of applicants belonging to the two groups are μ_A and μ_B, respectively, with $\mu_A + \mu_B = 1$. Assume that $q \leq \min\{\mu_A, \mu_B\}$, so that the firm can potentially fill all its positions using

applicants from either group. The firm wants to maximize the productivity of its workforce, irrespective of the group composition.

Let F_g denote the distribution of productivity conditional on group g. The manager observes the productivity of individual applicants, while the firm does not know F_g. The firm only knows possible bounds on a worker's productivity: $\underline{w} \leq w \leq \overline{w}$. Formally, the firm's ambiguity set is

$$\mathcal{F} = \left\{(F_A, F_B) : F_A \text{ and } F_B \text{ are probability distributions supported on } [\underline{w}, \overline{w}]\right\}.$$

The action $x \in X = [0, 1]$ in this case is the fraction of positions filled with applicants from group A. At composition x, the manager therefore hires a mass qx of group-A applicants and $q(1-x)$ of group-B applicants.

I assume that the manager shares the firm's objective of maximizing productivity, but also has a preference for hiring applicants from group A. Conditional on the group composition, however, the manager selects the most productive applicants within each group. This is consistent with a discriminatory preference that depends on the group composition, rather than on the identity of individual applicants within each group. Formally, the manager obtains an additional utility $\delta(x)$ from filling a fraction x of positions with group-A applicants.

While the firm may know the direction of the manager's discriminatory preferences, identifying their precise magnitude and functional form is considerably more demanding. Observed hiring choices reflect both the manager's preferences and his private information about applicants' productivity. Moreover, the manager's preferences need not vary linearly, or even smoothly with the composition of hires. For instance, a manager may not only prefer hiring applicants from his own demographic group but also place additional value on having that group constitute a majority in the team. To capture this, I assume the firm has minimal information about the shape of the bias. Formally, it is only willing to place a conservative upper bound D on how much productivity the manager would be willing to sacrifice to favor group A, i.e. $\delta \in \mathcal{B}_a$.

For each $g \in \{A, B\}$, let

$$Q_g(u) := \inf\{w \in [\underline{w}, \overline{w}] : F_g(w) \geq u\}, \qquad p_g(z) := Q_g\left(1 - \frac{z}{\mu_g}\right),$$

where Q_g is the generalized quantile function and $p_g(z)$ is the productivity of the marginal group-g applicant when a mass z is hired from that group. Since the most productive applicants are selected first, p_g is (weakly) decreasing. The total productivity of a mass z of selected group-g applicants is

$$P_g(z) := \int_0^z p_g(s)\, ds.$$

For a given pair $F = (F_A, F_B)$, total productivity at composition x is $\widetilde{V}_F(x) = P_A(qx) + P_B(q(1-x))$. Normalizing the all-$B$ allocation to zero, the firm's payoff is

$$V_F(x) = P_A(qx) + P_B(q(1-x)) - P_B(q). \tag{14}$$

The manager's payoff is

$$V_F(x) + \delta(x) - T(x)$$

Set $L := q(\overline{w} - \underline{w})$ and let

$$\mathcal{V}_{\text{cv}}(L, L) := \big\{V \in \mathcal{V}(L, L) : V \text{ is concave}\big\},$$

Note that $\mathcal{V}_{\text{cv}}(L, L)$ is the ambiguity set introduced in Assumption 2, with the symmetric Lipschitz bound used in Theorem 3. The next lemma shows that ambiguity about the conditional distributions of applicants' productivity is equivalent to ambiguity over normalized concave payoff functions with Lipschitz constant L.

Lemma 2. *Under the assumptions above,*

$$\big\{V_F : (F_A, F_B) \in \mathcal{F}\big\} = \mathcal{V}_{\text{cv}}(L, L).$$

Proof. See Appendix A.11. □

The lemma provides a microfoundation for the Lipschitz bounds in Assumptions 1 and 2. An increase in x replaces a mass $q\, dx$ of group-B workers with the same mass of group-

A workers. The productivity difference between any added and removed worker is at most $\overline{w} - \underline{w}$ in absolute value. Hence the firm can bound the payoff consequences of a small change in the composition without knowing exactly which composition maximizes aggregate productivity. Concavity follows from the assumption that, within each group, the manager selects applicants in decreasing order of productivity.

Together with the assumption that $\delta \in \mathcal{B}_a$, Lemma 2 implies that the firm's problem corresponds exactly to the ambiguity set of Theorem 3, with $v = L = q(\overline{w} - \underline{w})$. Therefore, a staircase incentive scheme is minimax optimal. In particular, the firm can offer tiered bonuses for reaching progressively higher hiring targets for group B. These incentives preserve the manager's discretion within each interval, while limiting the distortions caused by potentially sharp changes in his discriminatory preferences.

5. Conclusion

This paper studies how a principal can align the incentives of a better-informed agent when she faces ambiguity about both the consequences of the agent's actions and the extent of his private bias. Rather than assuming that the principal can form a prior over the economic environment, I consider a setting in which she only knows bounds on how her payoff changes with the agent's action and has partial information about the nature of the agent's misalignment. I characterize corrective interventions that minimize worst-case regret relative to the full-information benchmark, balancing the risk of insufficiently correcting the agent's bias against that of distorting choices that reflect valuable private information.

The optimal intervention depends on what the principal knows about the agent's bias. When she can bound its overall magnitude but not how it changes locally across actions, the minimax-optimal mechanism is a staircase tax, with flat regions that preserve the agent's discretion and discrete jumps that limit the distortions induced by his bias. When the principal instead knows bounds on the marginal bias, the optimal tax is piecewise linear, with a kink separating a region of low marginal taxation from one of high marginal taxation. I also show how additional information about the shape of the principal's payoff can affect the optimal mechanisms, and that simple caps are optimal when the principal

is restricted to delegation mechanisms. Applications to emissions regulation and hiring discrimination illustrate how the informational restrictions characterizing the ambiguity sets can arise in different economic environments.

More broadly, this paper illustrates the value of a robust approach to incentive alignment. On the one hand, such approach is realistic in environments characterized by substantial uncertainty and high complexity, in which the principal may lack the information needed to form a credible prior over the consequences of the agent's actions and the extent of his bias. On the other hand, taking this limited information as a starting point makes it possible to formulate and solve the alignment problem at a level of generality that would be difficult to achieve while maintaining tractability under a traditional Bayesian approach. Finally, the results offer a new rationale for discontinuities observed in real-world incentive schemes that does not rely on exogenous restrictions on the instruments available to the principal, but rather on realistic limitations in her knowledge of the environment she is regulating.

References

Alonso, Ricardo, Tan Gan, and Ju Hu (2026). "Robust delegation". CEPR Discussion Paper No. 21382.

Alonso, Ricardo and Niko Matouschek (2008). "Optimal delegation". *Review of Economic Studies* 75(1): 259–293.

Amador, Manuel and Kyle Bagwell (2013). "The theory of optimal delegation with an application to tariff caps". *Econometrica* 81(4): 1541–1599.

Ambrus, Attila and Georgy Egorov (2017). "Delegation and nonmonetary incentives". *Journal of Economic Theory* 171: 101–135.

Bergemann, Dirk and Karl Schlag (2011). "Robust monopoly pricing". *Journal of Economic Theory* 146(6): 2527–2543.

Blinder, Alan S. and Harvey S. Rosen (1985). "Notches". *American Economic Review* 75(4): 736–747.

Börgers, Tilman (2017). "(No) Foundations of dominant-strategy mechanisms: a comment on Chung and Ely (2007)". *Review of economic design* 21(2): 73–82.

Burtraw, Dallas, Karen Palmer, Ranjit Bharvirkar, and Anthony Paul (2001). "Cost-effective reduction of NOx emissions from electricity generation". *Journal of the Air & Waste Management Association* 51(10): 1476–1489.

Carlson, Curtis, Dallas Burtraw, Maureen Cropper, and Karen L. Palmer (2000). "Sulfur dioxide control by electric utilities: What are the gains from trade?" *Journal of Political Economy* 108(6): 1292–1326.

Carroll, Gabriel (2019). "Robustness in mechanism design and contracting". *Annual Review of Economics* 11: 139–166.

Dasgupta, Partha, Peter Hammond, and Eric Maskin (1980). "On imperfect information and optimal pollution control". *Review of Economic Studies* 47(5): 857–860.

Dworczak, Piotr and Alessandro Pavan (2022). "Preparing for the Worst but Hoping for the Best: Robust (Bayesian) Persuasion". *Econometrica* 90(5): 2017–2051.

Fowlie, Meredith and Nicholas Muller (2019). "Market-based emissions regulation when damages vary across sources: What are the gains from differentiation?" *Journal of the Association of Environmental and Resource Economists* 6(3): 593–632.

Frankel, Alexander (2014). "Aligned delegation". *American Economic Review* 104(1): 66–83.

Gillitzer, Christian, Henrik Jacobsen Kleven, and Joel Slemrod (2017). "A characteristics approach to optimal taxation: Line drawing and tax-driven product innovation". *Scandinavian Journal of Economics* 119(2): 240–267.

Guesnerie, Roger (1995). *A contribution to the pure theory of taxation*. Econometric Society Monographs 25. Cambridge: Cambridge University Press.

Guo, Yingni and Eran Shmaya (2023). "Regret-minimizing project choice". *Econometrica* 91(5): 1567–1593.

Guo, Yingni and Eran Shmaya (2025). "Robust monopoly regulation". *American Economic Review* 115(2): 599–634.

He, Wei, Fei Li, Jiangtao Li, and Hanqing Ye (2026). "Delegation without priors". arXiv preprint 2609.35156.

Holmström, Bengt (1977). "On incentives and control in organizations". PhD thesis. Stanford University.

Kang, Zi Yang (2026). "Prices vs. quantities: Robust regulation". arXiv preprint 2603.15832.

Kleven, Henrik Jacobsen (2016). "Bunching". *Annual Review of Economics* 8: 435–464.

Li, Jiangtao, Tilman Börgers, and Kexin Wang (2024). "Undominated mechanisms". SSRN Working Paper No. 5007659.

Malladi, Suraj (2020). "Delegated screening and robustness". SSRN Working Paper No. 3740206.

Malladi, Suraj (2022). "Searching in the dark and learning where to look". SSRN Working Paper No. 4084113.

Manski, Charles F. (2021). "Econometrics for decision making: Building foundations sketched by Haavelmo and Wald". *Econometrica* 89(6): 2827–2853.

Muller, Nicholas Z. and Robert Mendelsohn (2009). "Efficient pollution regulation: Getting the prices right". *American Economic Review* 99(5): 1714–1739.

Rochet, Jean-Charles (1985). "The taxation principle and multi-time Hamilton–Jacobi equations". *Journal of Mathematical Economics* 14(2): 113–128.

Rudik, Ivan (2020). "Optimal climate policy when damages are unknown". *American Economic Journal: Economic Policy* 12(2): 340–373.

Savage, Leonard J. (1954). *The foundations of statistics*. New York: Wiley.

Wald, Abraham (1950). *Statistical decision functions*. New York: Wiley.

Weitzman, Martin L. (1974). "Prices vs. quantities". *Review of Economic Studies* 41(4): 477–491.

A. Proofs

A.1. Proof of Theorem 1

We begin with a lemma showing that the staircase is indeed well-defined and characterize the optimal number of jumps, n^*.

Let us define

$$\rho_n := \max\left\{\frac{nD}{n+1}, \frac{\underline{v}\overline{v} + n\underline{v}D}{(n+1)(\underline{v}+\overline{v})}\right\}, \qquad n \in \mathbb{Z}_+. \tag{15}$$

Then,

$$R_a^\star = \min_{n\in\mathbb{Z}_+} \rho_n$$

Lemma 3. $\min_{n\in\mathbb{Z}_+} \rho_n$ *is attained.*

If $\overline{v} < D$*, the unique minimizer is* $n^\star = 0$*. If* $\overline{v} = D$*,*

$$\arg\min_{n\in\mathbb{Z}_+} \rho_n = \left\{0, \ldots, \left\lfloor \frac{\underline{v}}{D} \right\rfloor\right\}.$$

If $\overline{v} > D$*, let*

$$k := \left\lfloor \frac{\underline{v}}{D} \right\rfloor.$$

Then

$$\arg\min_{n\in\mathbb{Z}_+} \rho_n = \begin{cases} \{k\}, & \frac{\underline{v}(\overline{v}+kD)}{(k+1)(\underline{v}+\overline{v})} < \frac{(k+1)D}{k+2}, \\ \{k, k+1\}, & \frac{\underline{v}(\overline{v}+kD)}{(k+1)(\underline{v}+\overline{v})} = \frac{(k+1)D}{k+2}, \\ \{k+1\}, & \frac{\underline{v}(\overline{v}+kD)}{(k+1)(\underline{v}+\overline{v})} > \frac{(k+1)D}{k+2}. \end{cases}$$

Proof. Write

$$\rho_n = \max\left\{\frac{nD}{n+1}, \frac{\underline{v}}{\underline{v}+\overline{v}}\left(D + \frac{\overline{v}-D}{n+1}\right)\right\}.$$

The first term is strictly increasing in n. The second is strictly increasing, constant, or strictly decreasing depending on whether $\overline{v} < D$, $\overline{v} = D$, or $\overline{v} > D$.

If $\overline{v} < D$, ρ_n is strictly increasing, so $n^\star = 0$ is the unique minimizer.

If $\overline{v} = D$, the second term equals $\underline{v}D/(\underline{v}+D)$ for every n, and

$$\frac{nD}{n+1} \leq \frac{\underline{v}D}{\underline{v}+D} \quad \iff \quad nD \leq \underline{v}.$$

Hence the stated argmin set follows.

Finally, suppose $\overline{v} > D$. In this case, for $n \leq k$, ρ_n equals the decreasing term, while for $n \geq k+1$, it equals the increasing term. Hence, to determine the optimal value of $n^\star$ we have to compare $(k+1)D/(k+2)$ and $\underline{v}(\overline{v}+kD)/(k+1)(\underline{v}+\overline{v})$. If the former is strictly larger (resp. smaller), $n^\star = k$, (resp. $n^\star = k+1$). Finally, if they're equal any $n^\star \in \{k, k+1\}$ is optimal. □

Now we can proceed with the proof of Theorem 1. The proof is organized as follows. First we show that Equation 2 is a lower bound on regret for any mechanism T. Next, we show that the staircase achieves this bound.

Lower bound. Fix any mechanism T and call $r \equiv R_a(T)$. We want to show that

$$r \geq R_a^\star = \min_{n \in \mathbb{Z}_+} \rho_n \tag{16}$$

If $r \geq \min\{D, \overline{v}\}$, then we are done. Indeed, for $n = 0$,

$$\rho_0 = \frac{\overline{v}\underline{v}}{\underline{v}+\overline{v}} < \overline{v}$$

Moreover,

$$\frac{nD}{n+1} < D,$$

for every finite $n \geq 0$, and

$$\frac{\underline{v}\overline{v} + n\underline{v}D}{(n+1)(\underline{v}+\overline{v})} \to \frac{\underline{v}}{\underline{v}+\overline{v}}D < D$$

as $n \to \infty$. Hence,

$$R_a^\star = \min_{n \in \mathbb{Z}_+} \rho_n < D$$

Both imply

$$R_a^\star < \min\{D, \overline{v}\} \leq r$$

and the lower bound holds. So let us assume instead $r < \min\{D, \overline{v}\}$.

Normalize $\min_{x \in X} T(x) = 0$, so that $T(x) \geq 0$ for all feasible x. Fix $\varepsilon > 0$ such that

$$r_\varepsilon \equiv r + \varepsilon < \min\{D, \overline{v}\}$$

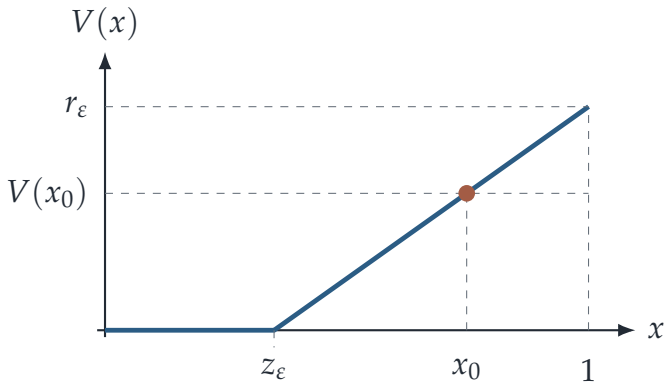
and set

$$z_\varepsilon \equiv 1 - \frac{r_\varepsilon}{\overline{v}}.$$

Consider the state

$$V_0(x) = \overline{v}(x - z_\varepsilon)_+, \qquad \delta_0(x) \equiv 0$$

The agent is unbiased and the principal's payoff is flat until a cutoff z_ε, at which point it increases linearly at rate $\overline{v}$. Note that (V_0, δ_0) is admissible, i.e., $V_0 \in \mathcal{V}$ and $\delta_0 \in \mathcal{B}_a$. The function V_0 is represented below.



Take any $x_0 \in \Gamma(T; V_0, \delta_0)$. We must have $x_0 > z_\varepsilon$, otherwise

$$r = R_a(T) \geq V_0(1) - V_0(x_0) = r_\varepsilon > r,$$

a contradiction. Let z_0 be an action such that $T(z_0) = 0$. Since z_0 is feasible, and x_0 is chosen, agent's IC requires

$$V_0(x_0) - T(x_0) \geq V_0(z_0) \geq 0$$

Thus,

$$\begin{aligned} T(x_0) \leq V_0(x_0) &= V_0(1) - \overline{v}(1 - x_0) \\ &= r_\varepsilon - \overline{v}(1 - x_0) \end{aligned}$$

Rearranging

$$\overline{v}(1 - x_0) + T(x_0) \leq r_\varepsilon \tag{17}$$

Now, consider the following recursive construction. Start from any action x_i that is chosen under T in some admissible state. If

$$x_i \leq \frac{r_\varepsilon}{\underline{v}},$$

stop. Otherwise, if $x_i > r_\varepsilon / \underline{v}$, set

$$p_{i+1} \equiv x_i - \frac{r_\varepsilon}{\underline{v}} > 0$$

and consider the state

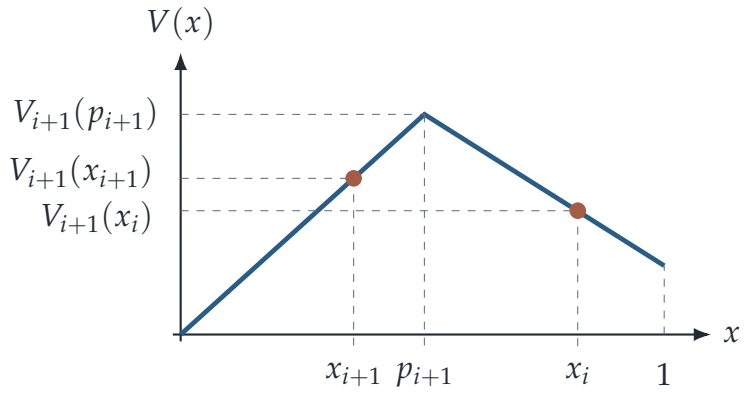
$$V_{i+1}(x) \equiv \begin{cases} \overline{v}x, & 0 \leq x \leq p_{i+1}, \\ \overline{v}p_{i+1} - \underline{v}(x - p_{i+1}), & p_{i+1} \leq x \leq 1, \end{cases}$$

and

$$\delta_{i+1}(x) \equiv \begin{cases} 0, & x < x_i, \\ D, & x \geq x_i. \end{cases}$$

The state (V_{i+1}, δ_{i+1}) is admissible. Moreover, by construction,

$$V_{i+1}(p_{i+1}) - V_{i+1}(x_i) = \underline{v}(x_i - p_{i+1}) = r_\varepsilon.$$



Note that every action chosen in the state (V_{i+1}, δ_{i+1}) lies strictly below x_i. Suppose not, and there exists some $x' \in \Gamma(T; V_{i+1}, \delta_{i+1})$ such that $x' \geq x_i$. Since $x_i > p_{i+1}$ and V_{i+1} is decreasing on $[p_{i+1}, 1]$,

$$V_{i+1}(x') \leq V_{i+1}(x_i).$$

Therefore,

$$\begin{aligned} r = R_a(T) &\geq V_{i+1}(p_{i+1}) - V_{i+1}(x') \\ &\geq V_{i+1}(p_{i+1}) - V_{i+1}(x_i) = r_\varepsilon > r, \end{aligned}$$

a contradiction.

Then, there must exist some action $x_{i+1} \in \Gamma(T; V_{i+1}, \delta_{i+1})$ such that $x_{i+1} < x_i$. Since x_i is feasible and is not chosen, IC requires

$$V_{i+1}(x_{i+1}) + \delta_{i+1}(x_{i+1}) - T(x_{i+1}) > V_{i+1}(x_i) + \delta_{i+1}(x_i) - T(x_i)$$

Using the definition of δ_{i+1} and rearranging we have

$$\begin{aligned} T(x_i) - T(x_{i+1}) &> V_{i+1}(x_i) - V_{i+1}(x_{i+1}) + D \\ &= D - r_\varepsilon + V_{i+1}(p_{i+1}) - V_{i+1}(x_{i+1}) \\ &\geq D - r_\varepsilon \end{aligned} \tag{18}$$

where the equality in the second line follows by adding and subtracting $V_{i+1}(p_{i+1})$ and noting that $V_{i+1}(p_{i+1}) - V_{i+1}(x_i) = r_\varepsilon$ by construction. The last inequality follows from the fact that V_{i+1} is maximized at p_{i+1}.

Rearranging (18) we get

$$V_{i+1}(p_{i+1}) - V_{i+1}(x_{i+1}) < T(x_i) - T(x_{i+1}) - (D - r_\varepsilon) \tag{19}$$

If $x_{i+1} \leq p_{i+1}$, then, using the definition of V_{i+1},

$$V_{i+1}(p_{i+1}) - V_{i+1}(x_{i+1}) = \overline{v}(p_{i+1} - x_{i+1})$$

So (19) gives

$$p_{i+1} - x_{i+1} \leq \frac{T(x_i) - T(x_{i+1}) - (D - r_\varepsilon)}{\overline{v}}$$

Also,

$$x_i - x_{i+1} = \underbrace{x_i - p_{i+1}}_{=r_\varepsilon/\underline{v}} + p_{i+1} - x_{i+1}$$

Then,

$$x_i - x_{i+1} \leq \frac{r_\varepsilon}{\underline{v}} + \frac{T(x_i) - T(x_{i+1}) - (D - r_\varepsilon)}{\overline{v}} \tag{20}$$

If instead $x_{i+1} > p_{i+1}$, then

$$x_i - x_{i+1} < x_i - p_{i+1} = \frac{r_\varepsilon}{\underline{v}}.$$

Moreover, by (18),

$$T(x_i) - T(x_{i+1}) - (D - r_\varepsilon) > 0.$$

So (20) holds also in this case.

Starting from x_0, repeat this construction until reaching the first action x_n such that

$$x_n \leq \frac{r_\varepsilon}{\underline{v}}.$$

By (18), at every step T must decrease by more than

$$D - r_\varepsilon > 0.$$

Since $T(x_0) \leq r_\varepsilon$ and $T(x) \geq 0$ for every feasible action x, the procedure must stop at some finite n.

Summing (18) across all steps we have

$$n(D - r_\varepsilon) \leq \sum_{i=0}^{n-1} [T(x_i) - T(x_{i+1})] = T(x_0) - T(x_n) \leq T(x_0) \leq r_\varepsilon.$$

Rearranging,

$$r_\varepsilon \geq \frac{nD}{n+1}. \tag{21}$$

Next, write

$$1 = x_n + \sum_{i=0}^{n-1} (x_i - x_{i+1}) + (1 - x_0).$$

By construction,

$$x_n \leq \frac{r_\varepsilon}{\underline{v}}.$$

Moreover, summing (20) gives

$$\sum_{i=0}^{n-1} (x_i - x_{i+1}) \leq \frac{nr_\varepsilon}{\underline{v}} + \frac{T(x_0) - T(x_n) - n(D - r_\varepsilon)}{\overline{v}}.$$

Finally, using (17)

$$1 - x_0 \leq \frac{r_\varepsilon - T(x_0)}{\overline{v}}.$$

Putting everything together,

$$\begin{aligned} 1 &\leq \frac{(n+1)r_\varepsilon}{\underline{v}} + \frac{T(x_0) - T(x_n) - n(D - r_\varepsilon)}{\overline{v}} + \frac{r_\varepsilon - T(x_0)}{\overline{v}} \\ &= \frac{(n+1)r_\varepsilon}{\underline{v}} + \frac{(n+1)r_\varepsilon - nD - T(x_n)}{\overline{v}} \\ &\leq \frac{(n+1)r_\varepsilon}{\underline{v}} + \frac{(n+1)r_\varepsilon - nD}{\overline{v}}, \end{aligned}$$

where the last inequality follows from $T(x_n) \geq 0$. Multiplying by $\underline{v}\overline{v}$ and rearranging gives

$$r_\varepsilon \geq \frac{\underline{v}\overline{v} + n\underline{v}D}{(n+1)(\underline{v} + \overline{v})}. \tag{22}$$

Combining (21) and (22),

$$r_\varepsilon \geq \max\left\{\frac{nD}{n+1}, \frac{\underline{v}\overline{v} + n\underline{v}D}{(n+1)(\underline{v}+\overline{v})}\right\} = \rho_n.$$

Therefore,

$$r_\varepsilon \geq \rho_n \geq \min_{m\in\mathbb{Z}_+} \rho_m = R_a^\star.$$

This holds for every sufficiently small $\varepsilon > 0$. Taking $\varepsilon \to 0$ gives

$$r \geq R_a^\star.$$

Since T was arbitrary, this proves the lower bound.

Upper bound. Let $n^\star$ minimize ρ_n and set

$$\widehat{x}^\star \equiv \min\left\{1, \frac{(n^\star+1)R_a^\star}{\underline{v}}\right\}.$$

Consider the staircase $T_a^\star$ defined by $T_a^\star(0) = 0$ and

$$T_a^\star(x) = j(D - R_a^\star)$$

whenever

$$\frac{jR_a^\star}{\underline{v}} < x \leq \min\left\{\frac{(j+1)R_a^\star}{\underline{v}}, 1\right\}, \qquad j = 0, \ldots, n^\star.$$

If $\widehat{x}^\star < 1$, set

$$T_a^\star(x) = n^\star(D - R_a^\star) + D \qquad \forall x > \widehat{x}^\star.$$

Thus, if the staircase ends before 1, it is followed by a final jump of size D and a flat tail.

By definition of $R_a^\star$,

$$R_a^\star \geq \frac{n^\star D}{n^\star + 1},$$

or equivalently

$$n^\star(D - R_a^\star) \leq R_a^\star. \tag{23}$$

Moreover, if $\widehat{x}^\star < 1$,

$$R_a^\star \geq \frac{\underline{v}\overline{v} + n^\star\underline{v}D}{(n^\star+1)(\underline{v}+\overline{v})}.$$

Using $\widehat{x}^\star = (n^\star + 1)R_a^\star/\underline{v}$ and rearranging,

$$n^\star(D - R_a^\star) + \overline{v}(1 - \widehat{x}^\star) \leq R_a^\star. \tag{24}$$

Fix any admissible state (V, δ) and any $x \in \Gamma(T_a^\star; V, \delta)$. Since all $y \in X$ are feasible, by agent's IC we must have

$$V(y) - V(x) \leq T_a^\star(y) - T_a^\star(x) + \delta(x) - \delta(y), \ \forall y \in X. \tag{25}$$

We want to show that

$$V(y) - V(x) \leq R_a^\star \qquad \forall y \in X.$$

First suppose $x \leq \widehat{x}^\star$. Consider first some $y \leq \widehat{x}^\star$. If $y \geq x$, (25), monotonicity of δ and (23) give

$$V(y) - V(x) \leq T_a^\star(y) - T_a^\star(x) \leq n^\star(D - R_a^\star) \leq R_a^\star.$$

If instead $y < x$ and x, y belong to the same flat step, then

$$x - y \leq \frac{R_a^\star}{\underline{v}},$$

so, by Lipschitzness of V,

$$V(y) - V(x) \leq \underline{v}(x - y) \leq R_a^\star.$$

Finally, if $y < x$ and x, y belong to different steps, then

$$T_a^\star(x) - T_a^\star(y) \geq D - R_a^\star.$$

Since $\delta(x) - \delta(y) \leq D$, (25) gives

$$V(y) - V(x) \leq -(D - R_a^\star) + D = R_a^\star.$$

It remains to consider $y > \widehat{x}^\star$. This case is possible only if $\widehat{x}^\star < 1$. Since $\widehat{x}^\star$ is feasible and $\widehat{x}^\star \geq x$, agent's IC and monotonicity of δ give

$$V(\widehat{x}^\star) - V(x) \leq T_a^\star(\widehat{x}^\star) - T_a^\star(x) \leq n^\star(D - R_a^\star).$$

Moreover, by Lipschitzness of V,

$$V(y) - V(\widehat{x}^\star) \leq \overline{v}(y - \widehat{x}^\star) \leq \overline{v}(1 - \widehat{x}^\star).$$

Hence, using (24),

$$V(y) - V(x) \leq n^\star(D - R_a^\star) + \overline{v}(1 - \widehat{x}^\star) \leq R_a^\star.$$

Now suppose instead $x > \widehat{x}^\star$. Then necessarily $\widehat{x}^\star < 1$, and

$$T_a^\star(x) = n^\star(D - R_a^\star) + D.$$

For any $y \leq \widehat{x}^\star$,

$$T_a^\star(y) \leq n^\star(D - R_a^\star)$$

and $\delta(x) - \delta(y) \leq D$. Hence, by (25),

$$V(y) - V(x) \leq n^\star(D - R_a^\star) - [n^\star(D - R_a^\star) + D] + D = 0.$$

Thus,

$$V(y) \leq V(x) \qquad \forall y \leq \widehat{x}^\star. \tag{26}$$

Consider now some $y > \widehat{x}^\star$. If $y \geq x$, then $T_a^\star(y) = T_a^\star(x)$ and monotonicity of δ gives

$$V(y) - V(x) \leq 0.$$

If instead $\widehat{x}^\star < y < x$, then by (26),

$$V(\widehat{x}^\star) \leq V(x).$$

Therefore,

$$V(y) - V(x) \leq V(y) - V(\widehat{x}^\star) \leq \overline{v}(y - \widehat{x}^\star) \leq \overline{v}(1 - \widehat{x}^\star) \leq R_a^\star,$$

where the last inequality follows from (24).

Thus, for every admissible state, every action x chosen by the agent, and every $y \in X$,

$$V(y) - V(x) \leq R_a^\star.$$

Therefore,

$$R_a(T_a^\star) \leq R_a^\star.$$

Together with the lower bound,

$$R_a(T_a^\star) = R_a^\star = \inf_T R_a(T).$$

Hence $T_a^\star$ is minimax optimal. $\square$

A.2. Proof of Proposition 1

(i) The first point follows because as the parameters $\underline{v}$, $\overline{v}$ and D increase, the set of admissible states (V, δ) also increases, giving more power to Nature to penalize the principal. Formally, for $\overline{v}' \geq \overline{v}$, $\underline{v}' \geq \underline{v}$ and $D' \geq D$, we have $\mathcal{A}_a(D, \underline{v}, \overline{v}) \subseteq \mathcal{A}_a(D', \underline{v}', \overline{v}')$, so, for any $T \in \mathcal{T}$,

$$\max_{(V,\delta)\in\mathcal{A}_a(D',\underline{v}',\overline{v}')} \mathrm{Reg}(T; V, \delta) \geq \max_{(V,\delta)\in\mathcal{A}_a(D,\underline{v},\overline{v})} \mathrm{Reg}(T; V, \delta)$$

(ii) The fact that $\ell^\star(D, \underline{v}, \overline{v}) \equiv R_a^\star(D, \underline{v}, \overline{v})/\underline{v}$ is weakly increasing in D and $\overline{v}$ follows directly from part (i). Next, write

$$\ell^\star(D, \underline{v}, \overline{v}) = \frac{R_a^\star(D, \underline{v}, \overline{v})}{\underline{v}} = \min_{n\in\mathbb{Z}_+} \max \left\{ \frac{nD}{\underline{v}(n+1)},\ \frac{\overline{v} + nD}{(n+1)(\underline{v} + \overline{v})} \right\}.$$

Note first that this minmax problem has a solution, i.e., for each value of $(D, \underline{v}, \overline{v})$ there exists a finite $n^\star(D, \underline{v}, \overline{v})$ that solves the problem. This follows from the same argument used in Lemma 3. Moreover, for each fixed $n \geq 0$, the maximum is strictly decreasing in $\underline{v}$. These two observations imply that $\ell^\star(D, \underline{v}, \overline{v})$ is also strictly decreasing in $\underline{v}$.

(iii) The fact that $h^\star(D, \underline{v}, \overline{v}) \equiv D - R_a^\star(D, \underline{v}, \overline{v})$ is weakly decreasing in $\underline{v}$ and $\overline{v}$ follows directly from part (i). Replacing the formula (2) for $R_a^\star$ into h we get

$$h^\star = D - R_a^\star = \max_{n\in\mathbb{Z}_+} \min \left\{ \frac{D}{n+1},\ \frac{D[\underline{v} + (n+1)\overline{v}] - \underline{v}\overline{v}}{(n+1)(\underline{v} + \overline{v})} \right\}.$$

For every fixed n, both terms inside the minimum are strictly increasing in D. Hence the minimum is strictly increasing in D. Moreover, the maximum over n is attained for every $(D, \underline{v}, \overline{v})$ (see Lemma 3). Therefore $h^\star(D, \underline{v}, \overline{v})$ is strictly increasing in D.

(*iv*) Define

$$\rho_n(D, \underline{v}, \overline{v}) := \max\left\{\frac{nD}{n+1},\ \frac{\underline{v}\overline{v} + n\underline{v}D}{(n+1)(\underline{v}+\overline{v})}\right\}.$$

so that

$$R_a^\star(D, \underline{v}, \overline{v}) = \min_{n\in\mathbb{Z}_+} \rho_n(D, \underline{v}, \overline{v}).$$

Moreover, define the difference

$$\Delta_n(D, \underline{v}, \overline{v}) := \rho_{n+1}(D, \underline{v}, \overline{v}) - \rho_n(D, \underline{v}, \overline{v}).$$

Since

$$\rho_n(D, \underline{v}, \overline{v}) = \begin{cases} \dfrac{\underline{v}(\overline{v}+nD)}{(n+1)(\underline{v}+\overline{v})}, & nD \le \underline{v}, \\ \dfrac{nD}{n+1}, & nD \ge \underline{v}, \end{cases}$$

we have

$$\Delta_n = \begin{cases} \dfrac{\underline{v}(D-\overline{v})}{(\underline{v}+\overline{v})(n+1)(n+2)}, & (n+1)D \le \underline{v}, \\ \dfrac{(n+1)D}{n+2} - \dfrac{\underline{v}(\overline{v}+nD)}{(n+1)(\underline{v}+\overline{v})}, & nD \le \underline{v} \le (n+1)D, \\ \dfrac{D}{(n+1)(n+2)}, & \underline{v} \le nD. \end{cases}$$

The three expressions agree at the relevant boundaries.

Consider first D. In the three regions above,

$$\frac{\partial \Delta_n}{\partial D} = \begin{cases} \dfrac{\underline{v}}{(\underline{v}+\overline{v})(n+1)(n+2)}, \\ \dfrac{\underline{v} + (n+1)^2\overline{v}}{(n+1)(n+2)(\underline{v}+\overline{v})}, \\ \dfrac{1}{(n+1)(n+2)}, \end{cases}$$

which is strictly positive throughout. Hence Δ_n is strictly increasing in D for every n. For any $m > n$,

$$\rho_m - \rho_n = \sum_{j=n}^{m-1} \Delta_j,$$

so ρ has strict increasing differences in (n, D).

Let $D_2 > D_1$, $m \in N^\star(D_2, \underline{v}, \overline{v})$, and $n \in N^\star(D_1, \underline{v}, \overline{v})$. Suppose by contradiction that $m > n$, then optimality implies

$$\rho_m(D_1) - \rho_n(D_1) \geq 0 \qquad \text{and} \qquad \rho_m(D_2) - \rho_n(D_2) \leq 0,$$

whereas strict increasing differences imply

$$\rho_m(D_2) - \rho_n(D_2) > \rho_m(D_1) - \rho_n(D_1),$$

a contradiction. Thus $m \leq n$ for every such pair, and therefore

$$\max N^\star(D_2, \underline{v}, \overline{v}) \leq \min N^\star(D_1, \underline{v}, \overline{v}).$$

Next consider $\overline{v}$. In the three regions,

$$\frac{\partial \Delta_n}{\partial \overline{v}} = \begin{cases} -\dfrac{\underline{v}(\underline{v} + D)}{(\underline{v} + \overline{v})^2 (n+1)(n+2)}, \\ -\dfrac{\underline{v}(\underline{v} - nD)}{(n+1)(\underline{v} + \overline{v})^2}, \\ 0, \end{cases}$$

which is weakly negative. Thus ρ has decreasing differences in $(n, \overline{v})$.

To see that this implies the desired strong-set-order result, let $\overline{v}_2 > \overline{v}_1$, take

$$n_1 \in N^\star(D, \underline{v}, \overline{v}_1), \qquad n_2 \in N^\star(D, \underline{v}, \overline{v}_2),$$

and suppose $n_1 > n_2$. Decreasing differences and optimality give

$$0 \leq \rho_{n_1}(\overline{v}_2) - \rho_{n_2}(\overline{v}_2) \leq \rho_{n_1}(\overline{v}_1) - \rho_{n_2}(\overline{v}_1) \leq 0.$$

Hence all inequalities are equalities, so

$$n_2 \in N^\star(D, \underline{v}, \overline{v}_1) \qquad \text{and} \qquad n_1 \in N^\star(D, \underline{v}, \overline{v}_2).$$

This is precisely

$$N^\star(D, \underline{v}, \overline{v}_2) \geq_s N^\star(D, \underline{v}, \overline{v}_1).$$

Finally, consider $\underline{v}$. If $\overline{v} < D$, Lemma 3 gives

$$N^\star(D, \underline{v}, \overline{v}) = \{0\}$$

for every $\underline{v}$, so the claim is immediate. Suppose instead that $\overline{v} \geq D$. In the three regions,

$$\frac{\partial \Delta_n}{\partial \underline{v}} = \begin{cases} -\dfrac{\overline{v}(\overline{v} - D)}{(\underline{v} + \overline{v})^2 (n+1)(n+2)}, \\ -\dfrac{\overline{v}(\overline{v} + nD)}{(n+1)(\underline{v} + \overline{v})^2}, \\ 0, \end{cases}$$

which is weakly negative. Hence ρ has decreasing differences in $(n, \underline{v})$. The same argument as above therefore yields

$$N^\star(D, \underline{v}_2, \overline{v}) \geq_s N^\star(D, \underline{v}_1, \overline{v}) \qquad \text{whenever } \underline{v}_2 > \underline{v}_1.$$

□

A.3. Proof of Theorem 2

First, we show that $R_m^\star$ is a lower bound on regret for any mechanism T. Next, we show that the kinked tax achieves this bound. Recall the definitions

$$d \equiv \overline{\delta} - \underline{\delta}, \qquad \alpha \equiv \min\{\underline{v}, d\}, \qquad \beta \equiv \min\{\overline{v}, d\}.$$

Lower bound. Fix any mechanism T and call $r \equiv R_m(T)$. We want to show that $r \geq R_m^\star$.

Fix an arbitrary cutoff $c \in [0, 1]$. Consider two admissible states, which we denote by (V_H^c, δ_H^c) and (V_L^c, δ_L^c).

In the first state, let

$$V_H^c(x) = \beta(x - c)_+, \qquad \delta_H^c(x) = \underline{\delta} x.$$

The principal's payoff is constant up to c and then increases at rate β. The agent's marginal bias is at its lower bound.

In the second state, let

$$V_L^c(x) = -\alpha \min\{x, c\}$$

and

$$\delta_L^c(x) = \underline{\delta} x + \alpha \min\{x, c\} + \beta(x - c)_+$$

Here, the principal's payoff decreases at rate α up to c and is constant for $x > c$.

Both states are admissible. Indeed, since $\alpha \leq \underline{v}$ and $\beta \leq \overline{v}$, both V_L^c and V_H^c satisfy Assumption 1 for all $c \in [0, 1]$. The biases are also admissible. The marginal bias in the first state is $\underline{\delta}$ everywhere. In the second state, it is $\underline{\delta} + \alpha$ below c and $\underline{\delta} + \beta$ above c. Since $\alpha, \beta \leq d$, both slopes belong to $[\underline{\delta}, \overline{\delta}]$.

Note that by construction the agent has exactly the same payoff in the two states,

$$V_H^c(x) + \delta_H^c(x) = V_L^c(x) + \delta_L^c(x) = \underline{\delta} x + \beta(x - c)_+ \qquad \forall x \in X. \tag{27}$$

Thus, under any mechanism T, the agent has the same set of optimal actions in the two states. Denote this set by $A_c(T)$.

In state H, the first-best payoff is achieved at $x = 1$. The regret generated by any action x is

$$r_H^c(x) = V_H^c(1) - V_H^c(x) = \begin{cases} \beta(1 - c), & x \leq c, \\ \beta(1 - x), & x > c. \end{cases}$$

In the second state, the first-best payoff is attained at $x = 0$, so regret is

$$r_L^c(x) = V_L^c(0) - V_L^c(x) = \begin{cases} \alpha x, & x \leq c, \\ \alpha c, & x > c. \end{cases}$$

Since Nature can choose either state, and the agent's choice set is the same in both, we have

$$r \geq \sup_{x \in A_c(T)} \max\{r_H^c(x), r_L^c(x)\}.$$

Now, consider any $x \in A_c(T)$. If $x \leq c$, then in the first state regret is

$$r_H^c(x) = \beta(1-c).$$

If instead $x \geq c$, then in the second state regret is

$$r_L^c(x) = \alpha c.$$

Therefore, regardless of the action chosen by the agent, at least one of the two states generates regret of at least

$$\min\{\beta(1-c), \alpha c\}.$$

It follows that

$$r \geq \min\{\beta(1-c), \alpha c\}. \tag{28}$$

Importantly, this holds for every $c \in [0,1]$. Hence,

$$r \geq \max_{c\in[0,1]} \min\{\beta(1-c), \alpha c\}.$$

The minimum is maximized at the cutoff $c^\star$ at which the two terms are equal:

$$\beta(1-c^\star) = \alpha c^\star.$$

Solving,

$$c^\star = \frac{\beta}{\alpha+\beta}.$$

Substituting into (28),

$$r \geq \alpha c^\star = \beta(1-c^\star) = \frac{\alpha\beta}{\alpha+\beta} = R_m^\star.$$

Since T was arbitrary, this proves the lower bound.

Upper bound. Consider the kinked tax

$$T_m^\star(x) = \underline{\delta} x + d(x-c^\star)_+,$$

where $c^\star$ is defined in (3).

Fix any admissible state (V, δ) and any $x \in \Gamma(T_m^\star; V, \delta)$. Define

$$\widetilde{\delta}(x) \equiv \delta(x) - \underline{\delta} x.$$

By the bounds on the marginal bias, for every $y > x$,

$$0 \leq \widetilde{\delta}(y) - \widetilde{\delta}(x) \leq d(y - x). \tag{29}$$

Thus, $\widetilde{\delta}$ is nondecreasing and d-Lipschitz.

Since the agent chooses x, incentive compatibility requires

$$V(x) + \delta(x) - T_m^\star(x) \geq V(y) + \delta(y) - T_m^\star(y) \qquad \forall y \in X.$$

Substituting the definitions of $\widetilde{\delta}$ and $T_m^\star$ and rearranging, we obtain

$$V(y) - V(x) \leq d\left[(y - c^\star)_+ - (x - c^\star)_+\right] + \widetilde{\delta}(x) - \widetilde{\delta}(y) \qquad \forall y \in X. \tag{30}$$

We want to show that

$$V(y) - V(x) \leq R_m^\star \qquad \forall y \in X.$$

Suppose first that $y \geq x$. We distinguish three cases.

If $y \leq c^\star$, then both actions lie below the kink, where the tax corrects only the known component of the bias. By (30),

$$V(y) - V(x) \leq \widetilde{\delta}(x) - \widetilde{\delta}(y) \leq 0.$$

Thus, the agent cannot choose an action below the cutoff when a higher action below the cutoff has strictly greater social value.

If instead $x \geq c^\star$, both actions lie above the kink. Using (30) and monotonicity of $\widetilde{\delta}$,

$$V(y) - V(x) \leq d(y - x) + \widetilde{\delta}(x) - \widetilde{\delta}(y) \leq d(y - x).$$

Moreover, by Assumption 1,

$$V(y) - V(x) \leq \overline{v}(y - x).$$

Combining the two inequalities,

$$V(y) - V(x) \leq \min\{\overline{v}, d\}(y - x) = \beta(y - x) \leq \beta(1 - c^\star).$$

Finally, suppose $x < c^\star < y$. Since $c^\star$ is feasible and the agent chooses x, using $y = c^\star$ in (30) gives

$$V(c^\star) - V(x) \leq \widetilde{\delta}(x) - \widetilde{\delta}(c^\star) \leq 0.$$

Hence,

$$V(y) - V(x) \leq V(y) - V(c^\star) \leq \overline{v}(y - c^\star),$$

where the last inequality follows from Assumption 1.

On the other hand, using (30) directly and monotonicity of $\widetilde{\delta}$,

$$\begin{aligned} V(y) - V(x) &\leq d(y - c^\star) + \widetilde{\delta}(x) - \widetilde{\delta}(y) \\ &\leq d(y - c^\star). \end{aligned}$$

Combining the two bounds,

$$V(y) - V(x) \leq \min\{\overline{v}, d\}(y - c^\star) = \beta(y - c^\star) \leq \beta(1 - c^\star).$$

Thus, whenever $y \geq x$,

$$V(y) - V(x) \leq \beta(1 - c^\star) = R_m^\star. \tag{31}$$

Now suppose instead that $y < x$. Again, we distinguish three cases.

If $x \leq c^\star$, both actions lie below the cutoff. Using (30) and the bound on $\widetilde{\delta}$,

$$V(y) - V(x) \leq \widetilde{\delta}(x) - \widetilde{\delta}(y) \leq d(x - y).$$

Moreover, by Assumption 1,

$$V(y) - V(x) \leq \underline{v}(x - y).$$

Therefore,

$$V(y) - V(x) \leq \min\{\underline{v}, d\}(x - y) = \alpha(x - y) \leq \alpha c^\star.$$

If instead $y \geq c^\star$, both actions lie above the cutoff. By (30),

$$\begin{aligned} V(y) - V(x) &\leq \widetilde{\delta}(x) - \widetilde{\delta}(y) - d(x-y) \\ &\leq 0, \end{aligned}$$

where the last inequality follows from (29).

It remains to consider $y < c^\star < x$. Using (30), we obtain

$$\begin{aligned} V(y) - V(x) &\leq \widetilde{\delta}(x) - \widetilde{\delta}(y) - d(x-c^\star) \\ &\leq d(x-y) - d(x-c^\star) \\ &= d(c^\star - y). \end{aligned}$$

We can obtain a second bound by comparing the agent's chosen action x with the cutoff $c^\star$. Incentive compatibility gives

$$\begin{aligned} V(c^\star) - V(x) &\leq \widetilde{\delta}(x) - \widetilde{\delta}(c^\star) - d(x-c^\star) \\ &\leq 0. \end{aligned}$$

Hence $V(c^\star) \leq V(x)$. Using Assumption 1,

$$\begin{aligned} V(y) - V(x) &\leq V(y) - V(c^\star) \\ &\leq \underline{v}(c^\star - y). \end{aligned}$$

Combining the two bounds,

$$V(y) - V(x) \leq \min\{\underline{v}, d\}(c^\star - y) = \alpha(c^\star - y) \leq \alpha c^\star.$$

Therefore, whenever $y < x$,

$$V(y) - V(x) \leq \alpha c^\star = R_m^\star. \tag{32}$$

Combining (31) and (32), we have shown that, for every admissible state (V, δ) and every $x \in \Gamma(T_m^\star; V, \delta)$,

$$V(y) - V(x) \leq \max\{\alpha c^\star, \beta(1-c^\star)\} = R_m^\star \qquad \forall y \in X.$$

Taking the maximum over y, the supremum over the agent's optimal actions, and the supremum over admissible states gives

$$R_m(T_m^\star) \leq R_m^\star.$$

Together with the lower bound established above,

$$R_m(T_m^\star) = R_m^\star = \inf_{T \in \mathcal{T}} R_m(T).$$

This proves the optimality of the kinked penalty.

A.4. Proof of Proposition 3

Suppose by contradiction that T dominates $T_a^\star$. In every state where $T_a^\star$ achieves zero regret, T must also achieve zero regret.

Let $\ell^\star = R_a^\star / \underline{v}$ and $h^\star = D - R_a^\star$. Fix $x \in (0, \widehat{x}^\star]$ and consider

$$\delta_x(z) = T_a^\star(x)\mathbf{1}_{\{z \geq x\}}, \qquad V_0(z) = -\varepsilon z, \qquad V_x(z) = \varepsilon(x - |z - x|),$$

where $0 < \varepsilon < \min\{\underline{v}, \overline{v}\}$. Since $T_a^\star$ is nondecreasing and $T_a^\star(x) \leq n^\star h^\star \leq R_a^\star < D$, the bias is admissible and $\delta_x \leq T_a^\star$, with equality at 0 and x. Thus $T_a^\star$ uniquely implements the first-best actions 0 and x in the respective states (V_0, δ_x) and (V_x, δ_x).

To achieve zero regret in these states T must (uniquely) implement the same actions. In particular, the agent must prefer 0 to x in state (V_0, δ_x):

$$V_0(0) + \delta_x(0) - T(0) > V_0(x) + \delta_x(x) - T(x) \quad \Rightarrow \quad T(x) - T(0) > T_a^\star(x) - \varepsilon x$$

The opposite must be true in state (V_x, δ_x):

$$V_x(x) + \delta_x(x) - T(x) > V_x(0) + \delta_x(0) - T(0) \quad \Rightarrow \quad T(x) - T(0) < T_a^\star(x) + \varepsilon x$$

Putting the two conditions together we get

$$T_a^\star(x) - \varepsilon x \leq T(x) - T(0) \leq T_a^\star(x) + \varepsilon x.$$

As $\varepsilon \downarrow 0$, the states remain well-defined and admissible. Taking the limit yields

$$T(x) = T_a^\star(x) + T(0) \qquad \forall x \le \widehat{x}^\star. \tag{33}$$

If $\widehat{x}^\star = 1$, the proof is complete.

So suppose instead that $\widehat{x}^\star < 1$. We first show that

$$\overline{v}z - T_a^\star(z) \le \overline{v}\widehat{x}^\star - T_a^\star(\widehat{x}^\star) \qquad \forall z \le \widehat{x}^\star. \tag{34}$$

This is immediate if $n^\star = 0$, since in that case $T_a^\star$ is constant in this region. In instead $n^\star \ge 1$, then optimality of $n^\star$ implies $\overline{v} \ge D$, since both terms of ρ_n in Equation (15) are strictly increasing in n when $\overline{v} < D$. Moreover, $\widehat{x}^\star = (n^\star + 1)\ell^\star < 1$ and, since $R_a^\star = \rho_{n^\star}$,

$$n^\star D \le (n^\star + 1)R_a^\star < \underline{v},$$

so the second term of $\rho_{n^\star}$ binds. Hence,

$$R_a^\star = \frac{\underline{v}(\overline{v} + n^\star D)}{(n^\star + 1)(\underline{v} + \overline{v})}.$$

So, we have

$$\overline{v}\ell^\star - h^\star = \frac{\overline{v} - D}{n^\star + 1} \ge 0.$$

For z in step $j \in \{0, \dots, n^\star\}$, we have $z \le (j+1)\ell^\star$ and $T_a^\star(z) = jh^\star$. Since $\widehat{x}^\star = (n^\star + 1)\ell^\star$ and $T_a^\star(\widehat{x}^\star) = n^\star h^\star$, it follows that $\widehat{x}^\star - z \ge (n^\star - j)\ell^\star$. Therefore,

$$\begin{aligned}[\overline{v}\widehat{x}^\star - T_a^\star(\widehat{x}^\star)] - [\overline{v}z - T_a^\star(z)] &= \overline{v}(\widehat{x}^\star - z) - (n^\star - j)h^\star \\ &\ge (n^\star - j)(\overline{v}\ell^\star - h^\star) \ge 0,\end{aligned}$$

where the last inequality follows from $\overline{v}\ell^\star \ge h^\star$. This proves (34).

Now fix $x > \widehat{x}^\star$ and consider two states:

$$\begin{aligned}V_{\widehat{x}^\star}(z) &= \overline{v}\min\{z, \widehat{x}^\star\} - \varepsilon(z - \widehat{x}^\star)_+, \\ \delta_{\widehat{x}^\star, x}(z) &= b\mathbf{1}_{\{z \ge \widehat{x}^\star\}} + (D - b)\mathbf{1}_{\{z \ge x\}},\end{aligned}$$

and

$$\begin{aligned}V_x^+(z) &= \overline{v}\min\{z, \widehat{x}^\star\} + \varepsilon(z - \widehat{x}^\star)_+ - 2\varepsilon(z - x)_+, \\ \delta_x^+(z) &= D\mathbf{1}_{\{z \ge x\}},\end{aligned}$$

where $0 < b < D$ and $0 < \varepsilon < \min\{\underline{v}, \overline{v}\}$. Both states are admissible, with unique first-best actions $\widehat{x}^\star$ and x, respectively.

By (34), in both states the agent's utility under $T_a^\star$ at any $z < \widehat{x}^\star$ is at most $\overline{v}\widehat{x}^\star - n^\star h^\star$. At $\widehat{x}^\star$ in the first state, utility equals $\overline{v}\widehat{x}^\star + b - n^\star h^\star$; at x in the second, it equals $\overline{v}\widehat{x}^\star + \varepsilon(x - \widehat{x}^\star) - n^\star h^\star$. Since $T_a^\star(z) = n^\star h^\star + D$ for $z > \widehat{x}^\star$, it's easy to check that every other action yields strictly lower utility. Thus $T_a^\star$ uniquely implements the first best in both states.

Dominance therefore requires T to implement $\widehat{x}^\star$ in the first state and x in the second. The IC conditions for these two actions give

$$T(x) - T(\widehat{x}^\star) > D - b - \varepsilon(x - \widehat{x}^\star),$$
$$T(x) - T(\widehat{x}^\star) < D + \varepsilon(x - \widehat{x}^\star).$$

Letting $b, \varepsilon \downarrow 0$ yields

$$T(x) - T(\widehat{x}^\star) = D = T_a^\star(x) - T_a^\star(\widehat{x}^\star).$$

Since $x > \widehat{x}^\star$ was arbitrary, together with (33) this implies

$$T(x) = T_a^\star(x) + T(0) \qquad \forall x \in X.$$

Hence T and $T_a^\star$ induce identical choices and regret in every state, contradicting dominance. □

A.5. Proof of Proposition 4

Suppose by contradiction that some mechanism T dominates $T_m^\star$. In every state where $T_m^\star$ achieves zero regret, T must also achieve zero regret.

Recall that

$$T_m^\star(z) = \underline{\delta} z + d(z - c^\star)_+, \qquad d = \overline{\delta} - \underline{\delta}.$$

Fix $0 < \varepsilon < \min\{\underline{v}, \overline{v}\}$. For any $y \in X$, consider

$$V_y(z) = \varepsilon(y - |z - y|),$$

which is admissible and uniquely maximized at y.

First, suppose $y \leq c^\star$ and consider the admissible bias $\delta(z) = \underline{\delta} z$. The agent's payoff under

$T_m^\star$ is

$$V_y(z) + \delta(z) - T_m^\star(z) = V_y(z) - d(z - c^\star)_+,$$

which is uniquely maximized at y. Thus $T_m^\star$ achieves zero regret, and dominance requires T to uniquely implement y.

Now suppose $y \geq c^\star$ and consider instead $\delta(z) = \overline{\delta} z$. The agent's payoff under $T_m^\star$ is

$$V_y(z) + \delta(z) - T_m^\star(z) = V_y(z) + d \min\{z, c^\star\}.$$

This is strictly increasing up to y and strictly decreasing thereafter, so $T_m^\star$ again uniquely implements y. Since T dominates $T_m^\star$, also T must uniquely implement y.

Fix any $x < c^\star$ and consider the two states with principal's payoffs V_x and $V_{c^\star}$ and the common bias $\delta(z) = \underline{\delta} z$. The incentive constraints under T imply

$$T(x) - T(c^\star) < \underline{\delta}(x - c^\star) + \varepsilon(c^\star - x),$$
$$T(x) - T(c^\star) > \underline{\delta}(x - c^\star) - \varepsilon(c^\star - x).$$

Since these inequalities hold for arbitrarily small $\varepsilon > 0$,

$$T(x) - T(c^\star) = \underline{\delta}(x - c^\star) \qquad \forall x \leq c^\star. \tag{35}$$

Similarly, fix any $x > c^\star$. Consider the two states with social values V_x and $V_{c^\star}$ and common bias $\delta(z) = \overline{\delta} z$. Dominance requires T to uniquely implement x and $c^\star$, respectively. Their incentive constraints give

$$T(x) - T(c^\star) < \overline{\delta}(x - c^\star) + \varepsilon(x - c^\star),$$
$$T(x) - T(c^\star) > \overline{\delta}(x - c^\star) - \varepsilon(x - c^\star).$$

Letting $\varepsilon \downarrow 0$ yields

$$T(x) - T(c^\star) = \overline{\delta}(x - c^\star) \qquad \forall x \geq c^\star. \tag{36}$$

Combining (35) and (36), we obtain

$$T(x) - T(c^\star) = T_m^\star(x) - T_m^\star(c^\star) \qquad \forall x \in X.$$

Therefore,

$$T(x) = T_m^\star(x) + T(c^\star) - T_m^\star(c^\star) \qquad \forall x \in X.$$

Hence T and $T_m^\star$ differ only by an additive constant and induce identical choice sets and regret in every admissible state. This contradicts strict dominance. Therefore $T_m^\star$ is undominated. □

A.6. Proof of Proposition 5

We first establish a lower bound on regret for any delegation set $S \in \mathcal{S}$. Next, we show that the capped delegation sets in the proposition achieve these bounds.

Lower bound. Fix any $S \in \mathcal{S}$ and let $\bar{s} = \max S$.

Consider first $\mathcal{A}_a$. In the state $V_H(x) = \overline{v}x$, $\delta_H(x) = 0$, the agent chooses $\bar{s}$, while the principal's first best is 1. Regret is therefore $\overline{v}(1 - \bar{s})$.

Next, let $k = \min\{\underline{v}, D/\bar{s}\}$ if $\bar{s} > 0$, and $k = 0$ otherwise. Consider the state $V_L(x) = -kx$, $\delta_L(x) = k \min\{x, \bar{s}\}$. This state is admissible since $k \leq \underline{v}$ and δ_L is nondecreasing, continuous, and bounded above by $k\bar{s} \leq D$. Here the principal's first best is 0, while the agent is indifferent among all actions in S, since $V_L(x) + \delta_L(x) = 0$ for every $x \in S$. Under adverse tie-breaking, the agent can choose $\bar{s}$, generating regret $k\bar{s} = \min\{\underline{v}\bar{s}, D\}$. Hence, for any $S \in \mathcal{S}$,

$$R_a(S) \geq \max\{\min\{\underline{v}\bar{s}, D\}, \overline{v}(1 - \bar{s})\}. \tag{37}$$

Note that if $\underline{v}\overline{v}/(\underline{v} + \overline{v}) < D$, then the lower bound is minimized at $\bar{s} = \overline{v}/(\underline{v} + \overline{v})$. If $D \leq \underline{v}\overline{v}/(\underline{v} + \overline{v})$, instead, the lower bound is at least D for every $\bar{s}$. At $\bar{s} = 1$ the lower bound is exactly $\min\{\underline{v}, D\} = D$. Hence its minimum is D, attained at $\bar{s} = 1$.

Now consider $\mathcal{A}_m$. In the state $V_H(x) = \overline{v}x$, $\delta_H(x) = \underline{\delta}x$, the agent chooses $\bar{s}$ and regret is $\overline{v}(1 - \bar{s})$. In the state $V_L(x) = -\min\{\underline{v}, \overline{\delta}\}x$, $\delta_L(x) = \overline{\delta}x$, the agent's payoff is $(\overline{\delta} - \min\{\underline{v}, \overline{\delta}\})x$, which is nondecreasing. Thus $\bar{s}$ is agent-optimal, while the principal's first best is 0, generating regret $\min\{\underline{v}, \overline{\delta}\}\bar{s}$. Both states are admissible, so

$$R_m(S) \geq \max\{\min\{\underline{v}, \overline{\delta}\}\bar{s}, \overline{v}(1 - \bar{s})\}. \tag{38}$$

The right-hand side is minimized when $\min\{\underline{v}, \overline{\delta}\}\bar{s} = \overline{v}(1-\bar{s})$, giving $\bar{s} = \overline{v}/(\min\{\underline{v}, \overline{\delta}\} + \overline{v})$ and regret $\min\{\underline{v}, \overline{\delta}\}\overline{v}/(\min\{\underline{v}, \overline{\delta}\} + \overline{v})$.

Since S was arbitrary, this establishes the lower bounds on minimax regret in both ambiguity scenarios.

Upper bound. Consider a capped delegation set $S_c = [0, c]$ for some $c \in [0, 1]$. Fix any admissible state (V, δ) and any $x \in \Gamma(S_c; V, \delta)$. We want to bound $V(y) - V(x)$ for every $y \in X$.

Suppose first that $y \leq c$. Since y is feasible, agent optimality implies $V(y) - V(x) \leq \delta(x) - \delta(y)$. If $y \geq x$, monotonicity of δ implies $V(y) - V(x) \leq 0$. If instead $y < x$, Assumption 1 gives

$$V(y) - V(x) \leq \min\{\underline{v}(x - y), \delta(x) - \delta(y)\}.$$

Under $\mathcal{A}_a$, the right-hand side is bounded by $\min\{\underline{v}c, D\}$, while under $\mathcal{A}_m$ it is bounded by $\min\{\underline{v}, \overline{\delta}\}(x - y) \leq \min\{\underline{v}, \overline{\delta}\}c$.

Now suppose $y > c$. Since c is feasible and $x \leq c$, agent optimality and monotonicity of δ imply $V(c) \leq V(x)$. Therefore, by Assumption 1,

$$V(y) - V(x) \leq V(y) - V(c) \leq \overline{v}(1 - c).$$

Since both the state and the agent-optimal action x were arbitrary, we have shown that

$$R_a(S_c) \leq \max\{\min\{\underline{v}c, D\}, \overline{v}(1 - c)\},$$

and

$$R_m(S_c) \leq \max\{\min\{\underline{v}, \overline{\delta}\}c, \overline{v}(1 - c)\}.$$

Choosing $c = c_a^{\star,\text{del}}$ and $c = c_m^{\star,\text{del}}$, respectively, the right-hand sides coincide with the lower bounds established above. Hence, both capped delegation sets are minimax optimal. □

A.7. Proof of Theorem 3

We first show that (5) is a lower bound on regret for any mechanism $T \in \mathcal{T}$. Next, we show that the staircase in (7) achieves this bound.

Lower bound. Fix any $T \in \mathcal{T}$ and define

$$q := \min\{D, v\}, \qquad c := \frac{q}{v+q}.$$

Notice that

$$R^\star_{a,\mathrm{cv}} = \min\left\{\frac{v}{2}, \frac{vD}{v+D}\right\} = \frac{vq}{v+q} = vc = q(1-c).$$

Consider the admissible state

$$V_0(z) = qz, \qquad \delta_0(z) = 0.$$

Let $x \in \Gamma(T; V_0, 0)$ be any agent-optimal action. If $x \le c$, then, since the first-best action under V_0 is 1,

$$R_{a,\mathrm{cv}}(T) \ge q(1-x) \ge q(1-c) = R^\star_{a,\mathrm{cv}}.$$

Suppose instead that $x > c$, and let $y := x - c > 0$. Consider the state

$$V_1(z) = \begin{cases} qz, & z \le y, \\ qy - v(z-y), & z > y, \end{cases} \qquad \delta_1(z) = \begin{cases} 0, & z < x, \\ D, & z \ge x. \end{cases} \tag{39}$$

The function V_1 is concave, v-Lipschitz and satisfies $V_1(0) = 0$, since $0 < q \le v$. The bias δ_1 is nondecreasing and u.s.c., belongs to $[0, D]$, and $\delta_1(0) = 0$. Hence $(V_1, \delta_1) \in \mathcal{A}_{a,\mathrm{cv}}$. Moreover, V_1 is maximized at y and

$$V_1(y) - V_1(x) = v(x-y) = vc = R^\star_{a,\mathrm{cv}}.$$

Note that, by construction x is optimal for the agent also in state (V_1, δ_1), i.e., $x \in$

$\Gamma(T; V_1, \delta_1)$. To see this, consider

$$\begin{aligned} V_1(x) + D - T(x) &= qy - v(x-y) + D - T(x) \\ &= qx + D - q - T(x) \\ &\geq qx - T(x) \end{aligned}$$

where we used the definition of $c = q/(v+q)$ and $y = x - c$ and $q = \min\{D, v\}$. Next, consider $z < x$,

$$V_1(z) + \delta_1(z) - T(z) \leq qz - T(z)$$

since $\delta_1(z) = 0$ and $V_1(z) = V_0(z) = qz$ by construction for all $z \leq y$ and $V_1(z) < V_0(z)$ for $z > y$. Using $x \in \Gamma(T; V_0, \delta_0)$ we have

$$qx - T(x) \geq qz - T(z), \ \forall z \in [0, 1]$$

Putting the inequalities together we get

$$V_1(x) + \delta_1(x) - T(x) \geq V_1(z) + \delta_1(z) - T(z), \ \forall z < x$$

If instead $z \geq x$, note that

$$T(z) - T(x) \geq q(z - x) \geq 0$$

so $T(z) \geq T(x)$. But since V_1 is decreasing above y and $z \geq x > y$, we have

$$V_1(z) - V_1(x) = -v(z - x) \leq 0$$

so $V_1(z) \leq V_1(x)$. Finally, $\delta_1(z) = \delta_1(x) = D$. Then,

$$V_1(x) + \delta_1(x) - T(x) \geq V_1(z) + \delta_1(z) - T(z) \ \forall z \geq x$$

This completes the proof that $x \in \Gamma(T; V_1, \delta_1)$. But then, regret from T is at least

$$V_1(y) - V_1(x) = vc = R^{\star}_{a,\text{cv}}$$

This proves the lower bound.

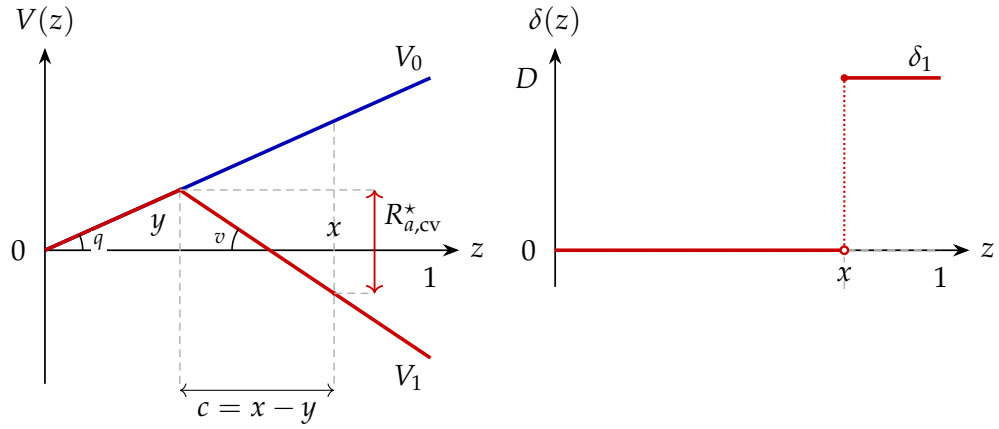


Figure 5: Lower-bound construction for the concave case.

Upper bound. Consider the staircase $T^\star_{a,\mathrm{cv}}$ defined in (7). Recall that

$$\ell^\star_{\mathrm{cv}} = \frac{R^\star_{a,\mathrm{cv}}}{v}, \qquad h^\star_{\mathrm{cv}} = D - R^\star_{a,\mathrm{cv}}, \qquad n^\star_{\mathrm{cv}} = \left\lfloor \frac{v}{D} \right\rfloor, \qquad \widehat{x}^\star_{\mathrm{cv}} = (n^\star_{\mathrm{cv}} + 1)\ell^\star_{\mathrm{cv}}.$$

The following inequalities will be useful:

$$0 < \widehat{x}^\star_{\mathrm{cv}} \le 1, \qquad n^\star_{\mathrm{cv}} h^\star_{\mathrm{cv}} \le R^\star_{a,\mathrm{cv}}, \qquad v(1-\widehat{x}^\star_{\mathrm{cv}}) \le R^\star_{a,\mathrm{cv}}, \qquad \frac{h^\star_{\mathrm{cv}}}{\ell^\star_{\mathrm{cv}}}(1-\ell^\star_{\mathrm{cv}}) = R^\star_{a,\mathrm{cv}} \quad \text{if } n^\star_{\mathrm{cv}} \ge 1. \tag{40}$$

To see this, suppose first that $D \le v$. In this case,

$$R^\star_{a,\mathrm{cv}} = \frac{vD}{v+D}, \qquad \ell^\star_{\mathrm{cv}} = \frac{D}{v+D}, \qquad h^\star_{\mathrm{cv}} = \frac{D^2}{v+D}, \qquad n^\star_{\mathrm{cv}} D \le v.$$

Consequently, $n^\star_{\mathrm{cv}} h^\star_{\mathrm{cv}} \le R^\star_{a,\mathrm{cv}}$, $\widehat{x}^\star_{\mathrm{cv}} \le 1$, and

$$v(1-\widehat{x}^\star_{\mathrm{cv}}) = R^\star_{a,\mathrm{cv}}\left(\frac{v}{D} - n^\star_{\mathrm{cv}}\right) < R^\star_{a,\mathrm{cv}}, \qquad \frac{h^\star_{\mathrm{cv}}}{\ell^\star_{\mathrm{cv}}}(1-\ell^\star_{\mathrm{cv}}) = R^\star_{a,\mathrm{cv}}.$$

If $D > v$, then $n^\star_{\mathrm{cv}} = 0$, $R^\star_{a,\mathrm{cv}} = v/2$, and $\ell^\star_{\mathrm{cv}} = \widehat{x}^\star_{\mathrm{cv}} = 1/2$, so that (40) also holds.

Fix any admissible state $(V,\delta) \in \mathcal{A}_{a,\mathrm{cv}}$, any $x \in \Gamma(T^\star_{a,\mathrm{cv}}; V, \delta)$, and any $y \in \arg\max_{z\in X} V(z)$. For every feasible action z, agent optimality implies

$$V(z) - V(x) \le T^\star_{a,\mathrm{cv}}(z) - T^\star_{a,\mathrm{cv}}(x) + \delta(x) - \delta(z). \tag{41}$$

We want to show that $V(y) - V(x) \le R^\star_{a,\mathrm{cv}}$.

Suppose first that $x \le \widehat{x}^\star_{\mathrm{cv}}$.

Consider first $y \le \widehat{x}^\star_{\mathrm{cv}}$, so that y is feasible. If $y \ge x$, (41), monotonicity of δ, and (40) give

$$V(y) - V(x) \le T^\star_{a,\mathrm{cv}}(y) - T^\star_{a,\mathrm{cv}}(x) \le n^\star_{\mathrm{cv}} h^\star_{\mathrm{cv}} \le R^\star_{a,\mathrm{cv}}.$$

If instead $y < x$ and both actions belong to the same flat region, then $x - y \le \ell^\star_{\mathrm{cv}}$, so Lipschitzness gives

$$V(y) - V(x) \le v(x - y) \le v\ell^\star_{\mathrm{cv}} = R^\star_{a,\mathrm{cv}}.$$

Finally, if $y < x$ and the actions belong to different regions, then $T^\star_{a,\mathrm{cv}}(x) - T^\star_{a,\mathrm{cv}}(y) \ge h^\star_{\mathrm{cv}}$. Since $\delta(x) - \delta(y) \le D$, (41) gives

$$V(y) - V(x) \le D - h^\star_{\mathrm{cv}} = R^\star_{a,\mathrm{cv}}.$$

Hence the required bound holds whenever $x, y \le \widehat{x}^\star_{\mathrm{cv}}$.

It remains to consider $x \le \widehat{x}^\star_{\mathrm{cv}} < y$. Since V is concave and y maximizes V, the function V is nondecreasing on $[0, y]$.

Let $i \in \{0, \dots, n^\star_{\mathrm{cv}}\}$ index the flat region containing x (taking $i = 0$ if $x \in [0, \ell^\star_{\mathrm{cv}}]$), and let $z_i := (i+1)\ell^\star_{\mathrm{cv}}$ be its right endpoint. Thus $x \le z_i \le \widehat{x}^\star_{\mathrm{cv}} < y$ and $T^\star_{a,\mathrm{cv}}(z_i) = T^\star_{a,\mathrm{cv}}(x)$. Since V and δ are nondecreasing on $[x, z_i]$, we have $V(z) \ge V(x)$ and $\delta(z_i) \ge \delta(x)$ while x is optimal for the agent, $V(z_i) + \delta(z_i) \le V(x) + \delta(x)$. This implies $V(z_i) = V(x)$ and $\delta(z_i) = \delta(x)$ and z_i is also optimal for the agent.

If $i = n^\star_{\mathrm{cv}}$, then $z_i = \widehat{x}^\star_{\mathrm{cv}}$, and Lipschitzness and (40) imply

$$V(y) - V(x) = V(y) - V(\widehat{x}^\star_{\mathrm{cv}}) \le v(1 - \widehat{x}^\star_{\mathrm{cv}}) \le R^\star_{a,\mathrm{cv}}.$$

Suppose instead that $i < n^\star_{\mathrm{cv}}$. The next right endpoint $z_i + \ell^\star_{\mathrm{cv}}$ is feasible and

$$T^\star_{a,\mathrm{cv}}(z_i + \ell^\star_{\mathrm{cv}}) - T^\star_{a,\mathrm{cv}}(z_i) = h^\star_{\mathrm{cv}}.$$

Agent optimality of z_i therefore implies

$$V(z_i + \ell^\star_{\mathrm{cv}}) - V(z_i) \le h^\star_{\mathrm{cv}} - [\delta(z_i + \ell^\star_{\mathrm{cv}}) - \delta(z_i)] \le h^\star_{\mathrm{cv}}.$$

Since $z_i < z_i + \ell^\star_{\text{cv}} < y$, concavity implies that secant slopes are weakly decreasing, so

$$\frac{V(y) - V(z_i)}{y - z_i} \leq \frac{V(z_i + \ell^\star_{\text{cv}}) - V(z_i)}{\ell^\star_{\text{cv}}} \leq \frac{h^\star_{\text{cv}}}{\ell^\star_{\text{cv}}}.$$

Using $V(x) = V(z_i)$, $z_i \geq \ell^\star_{\text{cv}}$, and (40),

$$V(y) - V(x) = V(y) - V(z_i) \leq \frac{h^\star_{\text{cv}}}{\ell^\star_{\text{cv}}}(y - z_i) \leq \frac{h^\star_{\text{cv}}}{\ell^\star_{\text{cv}}}(1 - \ell^\star_{\text{cv}}) = R^\star_{a,\text{cv}}.$$

Thus the bound holds whenever $x \leq \widehat{x}^\star_{\text{cv}}$.

Now suppose instead that $x > \widehat{x}^\star_{\text{cv}}$. If $y \leq \widehat{x}^\star_{\text{cv}}$, agent optimality implies

$$\begin{aligned} V(x) - V(y) &\geq T^\star_{a,\text{cv}}(x) - T^\star_{a,\text{cv}}(y) - [\delta(x) - \delta(y)] \\ &= n^\star_{\text{cv}} h^\star_{\text{cv}} - T^\star_{a,\text{cv}}(y) + D - [\delta(x) - \delta(y)] \\ &\geq 0, \end{aligned}$$

where the last inequality follows from $T^\star_{a,\text{cv}}(y) \leq n^\star_{\text{cv}} h^\star_{\text{cv}}$ and $\delta(x) - \delta(y) \leq D$. Hence $V(y) - V(x) = 0$, since y maximizes V.

If instead $y > \widehat{x}^\star_{\text{cv}}$, Lipschitzness and (40) imply

$$V(y) - V(x) \leq v|y - x| \leq v(1 - \widehat{x}^\star_{\text{cv}}) \leq R^\star_{a,\text{cv}}.$$

Hence the bound also holds whenever $x > \widehat{x}^\star_{\text{cv}}$.

Thus, for every admissible state, every action x chosen by the agent, and every first-best action y,

$$V(y) - V(x) \leq R^\star_{a,\text{cv}}.$$

Therefore,

$$R_{a,\text{cv}}(T^\star_{a,\text{cv}}) \leq R^\star_{a,\text{cv}}.$$

Together with the lower bound, this proves the theorem. □

A.8. Proof of Corollary 2

Since the concave ambiguity set is a subset of the baseline ambiguity set, $R^\star_{a,cv} \leq R^\star_a$.

Suppose first that $v \leq D$. By Theorem 3, $R^\star_{a,cv} = v/2$. Setting $n = 0$ in (2) gives $R^\star_a \leq v/2$, so $R^\star_a = R^\star_{a,cv}$.

Now suppose $v > D$. By Theorem 3, $R^\star_{a,cv} = vD/(v+D)$, while Theorem 1 gives

$$R^\star_a = \min_{n \in \mathbb{Z}_+} \max \left\{ \frac{nD}{n+1}, \frac{v+nD}{2(n+1)} \right\}.$$

For any $n \in \mathbb{Z}_+$ with $nD \leq v$,

$$\max \left\{ \frac{nD}{n+1}, \frac{v+nD}{2(n+1)} \right\} - R^\star_{a,cv} = \frac{(v-D)(v-nD)}{2(n+1)(v+D)}.$$

For $nD \geq v$, instead,

$$\max \left\{ \frac{nD}{n+1}, \frac{v+nD}{2(n+1)} \right\} - R^\star_{a,cv} = \frac{D(nD-v)}{(n+1)(v+D)}.$$

Both expressions are nonnegative and equal to zero if and only if $nD = v$. Since the minimum in (2) is attained, $R^\star_a = R^\star_{a,cv}$ if and only if $v/D \in \mathbb{N}$. Otherwise, $R^\star_{a,cv} < R^\star_a$.

Finally, by the definitions in Theorems 1 and 3,

$$\ell^\star = \frac{R^\star_a}{v}, \qquad \ell^\star_{cv} = \frac{R^\star_{a,cv}}{v},$$

and

$$h^\star = D - R^\star_a, \qquad h^\star_{cv} = D - R^\star_{a,cv}.$$

The inequalities follow immediately. □

A.9. Proof of Theorem 4

Lower bound. Fix any $c \in (0,1)$ and define

$$r(c) := \min\{\underline{v}c, \overline{v}(1-c), dc(1-c)\}.$$

Consider the following two states:

$$V_H^c(x) = \frac{r(c)}{1-c}x, \qquad \delta_H^c(x) = \underline{\delta}x,$$
$$V_L^c(x) = -\frac{r(c)}{c}x, \qquad \delta_L^c(x) = \left(\underline{\delta} + \frac{r(c)}{c(1-c)}\right)x.$$

Both payoff functions are affine, hence concave, and satisfy the slope bounds by the definition of $r(c)$. The biases are also admissible, since $r(c)/c(1-c) \leq d$. Moreover, the two states induce the same payoff for the agent:

$$V_H^c(x) + \delta_H^c(x) = V_L^c(x) + \delta_L^c(x) = \left(\underline{\delta} + \frac{r(c)}{1-c}\right)x \qquad \forall x \in X.$$

In state H, the first best is 1, and the regret from choosing action x is

$$R_H^c(x) = \frac{r(c)}{1-c}(1-x).$$

Similarly, in state L, the first best is 0, and regret is

$$R_L^c(x) = \frac{r(c)}{c}x.$$

Therefore, for any $x \leq c$,

$$R_H^c(x) \geq R_H^c(c) = r(c),$$

while for any $x \geq c$,

$$R_L^c(x) \geq R_L^c(c) = r(c).$$

Hence, $\max\{R_H^c(x), R_L^c(x)\} \geq r(c)$ for all x with equality at $x = c$.

Since the two states are indistinguishable to the agent, every mechanism T incurs worst-case regret of at least $r(c)$. Indeed,

$$\begin{aligned} R_{m,\text{cv}}(T) &\geq \max\left\{\sup_{x\in A_T} R_H^c(x), \sup_{x\in A_T} R_L^c(x)\right\} \\ &= \sup_{x\in A_T} \max\{R_H^c(x), R_L^c(x)\} \\ &\geq r(c). \end{aligned}$$

Moreover, since, this construction is feasible for every $c \in (0,1)$, Nature can choose the pair that generates the largest lower bound:

$$R_{m,\mathrm{cv}}(T) \geq R^\star_{m,\mathrm{cv}} := \max_{c\in[0,1]} \min\{\underline{v}c, \overline{v}(1-c), dc(1-c)\} \qquad \forall T \in \mathcal{T}. \tag{42}$$

which is well-defined since $r(c)$ is continuous and $r(0) = r(1) = 0$.

This establishes the lower bound.

Let $c^\star_{\mathrm{cv}} \in (0,1)$ be a maximizer of $r(c)$. By the definition of $R^\star_{m,\mathrm{cv}}$ and $c^\star_{\mathrm{cv}}$,

$$R^\star_{m,\mathrm{cv}} \leq \underline{v}c^\star_{\mathrm{cv}}, \qquad R^\star_{m,\mathrm{cv}} \leq \overline{v}(1-c^\star_{\mathrm{cv}}), \qquad R^\star_{m,\mathrm{cv}} \leq dc^\star_{\mathrm{cv}}(1-c^\star_{\mathrm{cv}}). \tag{43}$$

Moreover, since $R^\star_{m,\mathrm{cv}}$ is the maximum of $r(c)$,

$$\min\{\underline{v}x, \overline{v}(1-x), dx(1-x)\} \leq R^\star_{m,\mathrm{cv}} \qquad \forall x \in X. \tag{44}$$

Upper bound. Consider the tax $T^\star_{m,\mathrm{cv}}$ in (9). For every admissible δ, define $\widetilde{\delta}(x) := \delta(x) - \underline{\delta}x$. Then

$$0 \leq \widetilde{\delta}(y) - \widetilde{\delta}(x) \leq d(y-x) \qquad \forall x < y.$$

Similarly, define the residual tax $\widetilde{T}^\star(x) := T^\star_{m,\mathrm{cv}}(x) - \underline{\delta}x$. The agent's payoff can then be written as $V + \widetilde{\delta} - \widetilde{T}^\star$.

The residual tax $\widetilde{T}^\star$ is finite and continuous, with marginal rates

$$(\widetilde{T}^\star)'(x) = \begin{cases} \dfrac{R^\star_{m,\mathrm{cv}}}{1-x}, & x < c^\star_{\mathrm{cv}}, \\ d - \dfrac{R^\star_{m,\mathrm{cv}}}{x}, & x > c^\star_{\mathrm{cv}}. \end{cases} \tag{45}$$

Using the third condition in (43),

$$0 \leq \frac{R^\star_{m,\mathrm{cv}}}{1-x} \leq \frac{R^\star_{m,\mathrm{cv}}}{1-c^\star_{\mathrm{cv}}} \leq d \qquad \text{for } x < c^\star_{\mathrm{cv}},$$

and

$$0 \leq d - \frac{R^\star_{m,\mathrm{cv}}}{c^\star_{\mathrm{cv}}} \leq d - \frac{R^\star_{m,\mathrm{cv}}}{x} \leq d \qquad \text{for } x > c^\star_{\mathrm{cv}}.$$

Thus all marginal rates of $\widetilde{T}^\star$ belong to $[0,d]$.

Moreover, they are increasing on either side of $c^\star_{\mathrm{cv}}$, and

$$(\widetilde{T}^\star)'_+(c^\star_{\mathrm{cv}}) - (\widetilde{T}^\star)'_-(c^\star_{\mathrm{cv}}) = d - \frac{R^\star_{m,\mathrm{cv}}}{c^\star_{\mathrm{cv}}(1-c^\star_{\mathrm{cv}})} \geq 0.$$

by (43). Thus $\widetilde{T}^\star$ is convex.

Fix an arbitrary admissible state (V,δ) and any action $x \in \Gamma(T^\star_{m,\mathrm{cv}}; V, \delta)$. Let $p \in \arg\max_{y\in X} V(y)$. Agent optimality implies

$$V(z) - V(x) \leq \widetilde{T}^\star(z) - \widetilde{T}^\star(x) + \widetilde{\delta}(x) - \widetilde{\delta}(z) \qquad \forall z \in X. \tag{46}$$

Suppose first that $p > x$. For any $x < z < p$, concavity of V, (46), and monotonicity of $\widetilde{\delta}$ imply

$$\frac{V(p)-V(x)}{p-x} \leq \frac{V(z)-V(x)}{z-x} \leq \frac{\widetilde{T}^\star(z)-\widetilde{T}^\star(x)}{z-x}.$$

Letting $z \downarrow x$ and using the upper slope bound on V, we obtain

$$V(p) - V(x) \leq \min\{\overline{v}, (\widetilde{T}^\star)'_+(x)\}(1-x). \tag{47}$$

Similarly, if $p < x$, concavity and (46), together with the upper slope bound on $\widetilde{\delta}$, imply, for $p < z < x$,

$$\frac{V(p)-V(x)}{x-p} \leq \frac{V(z)-V(x)}{x-z} \leq d + \frac{\widetilde{T}^\star(z)-\widetilde{T}^\star(x)}{x-z}.$$

Letting $z \uparrow x$ and combining the resulting inequality with the lower slope bound on V, we obtain

$$V(p) - V(x) \leq \min\{\underline{v}, d - (\widetilde{T}^\star)'_-(x)\}(x-p).$$

Since $(\widetilde{T}^\star)'_-(x) \leq d$, the coefficient on the right-hand side is nonnegative. Thus, using $x - p \leq x$,

$$V(p) - V(x) \leq \min\{\underline{v}, d - (\widetilde{T}^\star)'_-(x)\}x. \tag{48}$$

We now show that both bounds are at most $R^\star_{m,\mathrm{cv}}$.

Suppose $x < c^\star_{\mathrm{cv}}$. If $p > x$, by (45),

$$(\widetilde{T}^\star)'_+(x)(1-x) = R^\star_{m,\mathrm{cv}},$$

so (47) implies that regret is at most $R^\star_{m,\mathrm{cv}}$.

If instead $p < x$, (48) gives

$$V(p) - V(x) \le \min\left\{\underline{v}, d - \frac{R^\star_{m,\mathrm{cv}}}{1-x}\right\} x.$$

Notice that

$$\overline{v}(1-x) > \overline{v}(1-c^\star_{\mathrm{cv}}) \ge R^\star_{m,\mathrm{cv}}.$$

Hence (44) implies

$$\min\{\underline{v}x, dx(1-x)\} \le R^\star_{m,\mathrm{cv}}.$$

Equivalently,

$$\min\left\{\underline{v}, d - \frac{R^\star_{m,\mathrm{cv}}}{1-x}\right\} x \le R^\star_{m,\mathrm{cv}}.$$

Now suppose $x > c^\star_{\mathrm{cv}}$ and consider first the case where $p < x$. By (45),

$$(d - (\widetilde{T}^\star)'_-(x))x = R^\star_{m,\mathrm{cv}},$$

so (48) is bounded by $R^\star_{m,\mathrm{cv}}$.

Next, consider $p > x$, since

$$\underline{v}x > \underline{v}c^\star_{\mathrm{cv}} \ge R^\star_{m,\mathrm{cv}},$$

inequality (44) implies

$$\min\{\overline{v}(1-x), dx(1-x)\} \le R^\star_{m,\mathrm{cv}}.$$

Equivalently,

$$\min\left\{\overline{v}, d - \frac{R^\star_{m,\mathrm{cv}}}{x}\right\} (1-x) \le R^\star_{m,\mathrm{cv}}.$$

By (45), this is precisely the bound required in (47).

At $x = c^\star_{\mathrm{cv}}$, the bound for $p < x$ follows from (48), since the inequality established for $x < c^\star_{\mathrm{cv}}$ extends to $x = c^\star_{\mathrm{cv}}$ by continuity of its right-hand side. Similarly, the inequality established for $x > c^\star_{\mathrm{cv}}$ extends to $x = c^\star_{\mathrm{cv}}$ by continuity, so (47) bounds regret when $p > x$. Here we use the left and right derivatives of $\widetilde{T}^\star$, respectively.

Thus every agent-optimal action in every admissible state generates regret at most $R^\star_{m,\mathrm{cv}}$.

Therefore,

$$R_{m,\mathrm{cv}}(T^\star_{m,\mathrm{cv}}) \leq R^\star_{m,\mathrm{cv}}.$$

Together with (42), this gives

$$R_{m,\mathrm{cv}}(T^\star_{m,\mathrm{cv}}) = \inf_{T\in\mathcal{T}} R_{m,\mathrm{cv}}(T) = R^\star_{m,\mathrm{cv}}.$$

A.10. Proof of Lemma 1

Note that, under Assumption 3, $\mathcal{A}_{\mathrm{em}} \subset \mathcal{V} \times \mathcal{B}_m$. Therefore, the upper bound on regret achieved by the kinked tax still holds. In particular, the emissions tax in (13) guarantees regret no greater than $\alpha\beta/(\alpha+\beta)$.

To establish the reverse inequality, we follow a similar strategy to the one used to prove the lower bound in Theorem 2. Start by defining $k_{\mathrm{G}} = \max\{\underline{k}, \underline{h}\}$ and $k_0 = \min\{\overline{k}, \overline{h}\}$. By (11), $k_{\mathrm{G}} < k_0$. Moreover, $\overline{h} - k_{\mathrm{G}} = \alpha$ and $k_0 - \underline{h} = \beta$.

Fix any $\xi \in [0,1]$ and consider the following abatement-cost function:

$$C^{(\xi)}(a) = \min\{k_0 a, K_\xi + k_{\mathrm{G}} a\}, \qquad K_\xi = (k_0 - k_{\mathrm{G}})(1-\xi).$$

Since $k_{\mathrm{G}}, k_0 \in [\underline{k}, \overline{k}]$, the function $C^{(\xi)}$ satisfies Assumption 3. The corresponding private benefit from emissions is

$$C^{(\xi)}(1) - C^{(\xi)}(1-x) = k_{\mathrm{G}} \min\{x, \xi\} + k_0 (x-\xi)_+.$$

Consider two admissible damage functions,

$$H_0(x) = \underline{h} x, \qquad H_1(x) = \overline{h} x.$$

Because the abatement-cost function is identical in both states, any tax induces the same nonempty set of optimal emissions in both.

In state 0, social welfare is $V_0(x) = C^{(\xi)}(1) - C^{(\xi)}(1-x) - H_0(x)$, which is maximized at

$x = 1$, since $k_G, k_0 \geq \underline{h}$. If the firm chooses $x \leq \xi$, regret is at least

$$\begin{aligned} V_0(1) - V_0(x) &= (k_G - \underline{h})(\xi - x) + (k_0 - \underline{h})(1 - \xi) \\ &\geq \beta(1 - \xi). \end{aligned}$$

In state 1, social welfare is $V_1(x) = C^{(\xi)}(1) - C^{(\xi)}(1 - x) - H_1(x)$, which is maximized at $x = 0$, since $k_G, k_0 \leq \overline{h}$. If the firm chooses $x \geq \xi$, regret is at least

$$\begin{aligned} V_1(0) - V_1(x) &= (\overline{h} - k_G)\xi + (\overline{h} - k_0)(x - \xi) \\ &\geq \alpha\xi. \end{aligned}$$

Thus, independently of the specific action induced by any mechanism, one of the two states generates regret of at least $\min\{\beta(1 - \xi), \alpha\xi\}$. Since both states belong to $\mathcal{A}_{\text{em}}$, every mechanism T satisfies

$$R_{\text{em}}(T) \geq \min\{\beta(1 - \xi), \alpha\xi\}.$$

Since the construction is admissible for every $\xi \in [0, 1]$, we have in particular

$$\begin{aligned} R_{\text{em}}(T) &\geq \max_{\xi \in [0,1]} \min\{\beta(1 - \xi), \alpha\xi\} \\ &= \frac{\alpha\beta}{\alpha + \beta}, \end{aligned}$$

where the maximum is attained at $\xi = c^\star_{\text{em}} = \beta/(\alpha + \beta)$. Together with the upper bound from Theorem 2, this proves the result. □

A.11. Proof of Lemma 2

Fix an admissible pair $(F_A, F_B) \in \mathcal{F}$. Since each p_g is weakly decreasing, P_g is concave. Equation (14) therefore implies that V_F is concave and $V_F(0) = 0$. For $0 \leq x < y \leq 1$,

$$V_F(y) - V_F(x) = \int_{qx}^{qy} p_A(s)\, ds - \int_{q(1-y)}^{q(1-x)} p_B(s)\, ds.$$

Each integral is taken over an interval of length $q(y - x)$, and the integrand lies in $[\underline{w}, \overline{w}]$. Consequently,

$$-L(y - x) \leq V_F(y) - V_F(x) \leq L(y - x).$$

Hence, $V_F \in \mathcal{V}_{\text{cv}}(L, L)$.

Next, we show the reverse inclusion, i.e., given any concave function $V \in \mathcal{V}_{\text{cv}}(L, L)$ we can find distributions (F_A, F_B) such that $V_F(x) = V(x)$ for all $x \in [0, 1]$, where V_F is defined as in Equation 14. Fix an arbitrary $V \in \mathcal{V}_{\text{cv}}(L, L)$ and let $\psi(x) = V'_+(x)$ denote its right derivative on $[0, 1)$, extended arbitrarily at $x = 1$. Concavity implies that ψ is (weakly) decreasing, while the Lipschitz bound gives $-L \leq \psi(x) \leq L$. We have to construct functions p_A and p_B such that

$$q[p_A(qx) - p_B(q(1-x))] = \psi(x)$$

A natural choice is to let $\psi^+(x) = \max\{\psi(x), 0\}$ and $\psi^-(x) = \max\{-\psi(x), 0\}$ and define on $[0, q]$

$$p_A(z) = \underline{w} + \frac{\psi^+(z/q)}{q}, \qquad p_B(z) = \underline{w} + \frac{\psi^-(1 - z/q)}{q}.$$

Both functions are (weakly) decreasing and take values in $[\underline{w}, \overline{w}]$, since $L = q(\overline{w} - \underline{w})$. Extend them to $(q, \mu_g]$ by setting $p_g(z) = \underline{w}$. Any such function is induced, up to values on a set of measure zero, by the upper quantiles of some probability distribution F_g supported on $[\underline{w}, \overline{w}]$. Hence the constructed pair belongs to $\mathcal{F}$. For this pair, Equation (14) yields

$$V_F(x) = q \int_0^x \big[p_A(qt) - p_B(q(1-t))\big] dt = \int_0^x \big[\psi^+(t) - \psi^-(t)\big] dt = V(x),$$

where the last equality uses absolute continuity and $V(0) = 0$. This proves the reverse inclusion.