"On the Futility of Dynamics in Robust Mechanism Design" Technical Report with Supplementary Materials

Santiago Balseiro
Columbia University
srb2155@columbia.edu

Anthony Kim
Amazon
tonyekim@cs.stanford.edu

Daniel Russo
Columbia University
djr2174@gsb.columbia.edu

January 23, 2021

TR.1 Introduction

In this technical report, we present several results that complement the exposition in the main paper. In Section TR.2, we prove a general minimax theorem and an "asymmetric" minimax theorem using a topological approach for the single-round problem. In Section TR.3, we prove a multi-round saddle-point theorem via an equivalent formulation of the multi-round problem involving distributions over distributions. In Section TR.4, we consider welfare maximization in the single-good case and revenue/welfare maximization in the multiple-goods case of the dynamic selling problem in Section 4.1 of the main paper. In Section TR.5, we apply our results to a repeated resource allocation problem without monetary transfers. Section TR.6 extends our results in the main paper in several directions: alternative benchmarks, serially correlated shock processes, a multiplicative performance guarantee, and a stronger notion of regret. Finally, in Section TR.7, we show equivalence-type connections between the minimax regret and maximin utility objectives for revenue maximization in the dynamic selling mechanism design problem.

Before presenting the results, we provide some preliminaries below.

Single-Round Problem. Using the outcome distribution representation of single-round direct mechanisms and Lemma 1 in the main paper, we can equivalently write the single-round problem (2) as the optimization problem (3) with the alternative objective given in the lemma. For ease of notations, we define $\widehat{\text{Regret}}(S, F) := \int_\Theta \text{OPT}(\delta_\theta, 1) dF(\theta) - \int_\Theta \int_\Omega u(\theta, \omega) dS_\theta(\omega) dF(\theta)$ for a single-round direct mechanism S and distribution F and write the optimization problem as

$$\inf_{\substack{S\in\Delta(\Omega)^\Theta:\\ \text{(IC)},\text{(IR)}}} \sup_{F\in\Delta(\Theta)} \widehat{\text{Regret}}(S,F)\,,$$

where S can be any single-round direct IC/IR mechanism. In this technical report, we interchangeably refer to the above formulation as the single-round problem and to $\int_\Theta \text{OPT}(\delta_\theta, 1)dF(\theta)$ as the single-round benchmark. We let $\widehat{\text{Regret}}$ be the corresponding optimal value, i.e., the minimax regret of the single-round problem. As we will be presenting and leveraging saddle-point results (via minimax duality theory), the above formulation with the inner supremum as a convex program will be frequently used.

Saddle-Point Results. Both multi-round and single-round minimax regret problems can be viewed as sequential-move zero-sum games in which the principal first chooses a mechanism and then nature selects a worst-case distribution to maximize the principal's regret. When appropriate saddle-point theorems hold, these problems are respectively equivalent to ones in which nature chooses first and then the principal optimizes his performance given nature's choice. This provides a framework for establishing the existence of optimal mechanisms and explicitly characterizing them and, also, a direct connection between our robust formulation and a more classical Bayesian formulation in the multi-round problem.

In particular, we can cast the single-round problem as a simultaneous-move zero-sum game between the principal and nature and show it admits an optimal solution via a saddle-point result that says 1) the saddle-point property (or, the duality gap of 0) holds, i.e.,

$$\inf_{\substack{S\in\Delta(\Omega)^\Theta:\\ \text{(IC)},\text{(IR)}}} \sup_{F\in\Delta(\Theta)} \widehat{\text{Regret}}(S,F) = \sup_{F\in\Delta(\Theta)} \inf_{\substack{S\in\Delta(\Omega)^\Theta:\\ \text{(IC)},\text{(IR)}}} \widehat{\text{Regret}}(S,F)\,, \tag{TR-1}$$

and 2) there exists a saddle point (S^*, F^*), which is a single-round direct IC/IR mechanism S^* and a distribution F^* such that $\widehat{\text{Regret}}(S^*, F) \le \widehat{\text{Regret}}(S^*, F^*) \le \widehat{\text{Regret}}(S, F^*)$ for any $S \in \Delta(\Omega)^\Theta$ satisfying the IC/IR constraints and $F \in \Delta(\Theta)$. Note that 2) implies 1) but does not necessarily hold when 1) does. If a saddle point or, equivalently, a Nash equilibrium exists, the minimax regret of the single-round problem is exactly the value of the zero-sum game, i.e., $\widehat{\text{Regret}} = \widehat{\text{Regret}}(S^*, F^*)$, and S^* is an optimal single-round solution.

A saddle point can be shown to exist either through topological arguments (see Section TR.2) or by construction. For several applications treated in the main paper and this technical report, we use the optimality conditions for the inner optimization problems in (TR-1) to explicitly construct a saddle point and, with it, an optimal single-round mechanism. It is possible that saddle points do not exist in the single-round problem; see the resource allocation problem without monetary transfers in Section TR.5.

TR.2 Single-Round Saddle-Point Theorems under Sufficient Conditions

For the single-round problem, we prove a general minimax theorem and an "asymmetric" minimax theorem under general sufficient conditions. Establishing a saddle-point result for the single-round problem in full generality is challenging, and the resource allocation problem without monetary transfers (in Section TR.5) hints at the difficulty, because saddle points do not exist even for this simple game. We can still establish a general saddle-point result under general conditions and the assumption of a finite shock space, and an asymmetric saddle-point result when the game has a continuous shock space. Our approach is topological in nature and involves endowing the space of (randomized) single-round direct mechanisms $\Delta(\Omega)^\Theta$ with the right topology and then leveraging existing minimax theorems from the literature.

Under the assumption to follow, we show a general saddle-point result for the single-round problem:

Assumption TR.1. *The game satisfies:*

(i) The outcome space $\Omega \subset \mathbb{R}^m$ is compact and the shock space Θ is finite.

(ii) The principal's utility function $u(\theta, \omega)$ is upper semi-continuous in ω for all θ.

(iii) The agent's utility function $v(\theta, \omega)$ is continuous in ω for all θ.

Assumption TR.1 guarantees that the single-round benchmark is well-defined. Parts of the assumption may be relaxed as long as the benchmark is well-defined in applications. In most cases including those considered in the main paper and this technical report, this assumption is satisfied, at least, for an approximate result, because the outcome space can be suitably restricted without loss by imposing large bounds (which can be shown to be not binding at an optimal solution) and the shock space can be discretized. Parts (ii) and (iii) are typical in the mechanism design literature. Note Assumption TR.1 holds trivially for finite discrete games and may be dropped. We can show the following:

Theorem TR.1. *Suppose the game satisfies Assumption TR.1. The minimax regret of the single-round problem satisfies:*

$$\widehat{\text{Regret}} = \min_{\substack{S\in\Delta(\Omega)^{\Theta}:\\ \text{(IC),(IR)}}} \max_{F\in\Delta(\Theta)} \widehat{\text{Regret}}(S,F) = \max_{F\in\Delta(\Theta)} \min_{\substack{S\in\Delta(\Omega)^{\Theta}:\\ \text{(IC),(IR)}}} \widehat{\text{Regret}}(S,F)\,.$$

Moreover, there exists a (randomized) single-round direct IC/IR mechanism S^* and a distribution F^* such that $\widehat{\text{Regret}} = \widehat{\text{Regret}}(S^*, F^*)$ and $\widehat{\text{Regret}}(S^*, F) \le \widehat{\text{Regret}}(S^*, F^*) \le \widehat{\text{Regret}}(S, F^*)$ for any $S \in \Delta(\Omega)^{\Theta}$ satisfying the IC/IR constraints and $F \in \Delta(\Theta)$.

Theorem TR.1 shows that the minimax regret is equivalent to the maximin regret for the single-round problem and the minimax regret is achieved by a single-round mechanism. In the maximin regret formulation, nature first chooses the agent's private distribution and then the principal chooses an optimal mechanism based on nature's choice. The result shows there is a distribution that is uniformly challenging for all possible single-round mechanisms. We prove the result in Section TR.2.1 using the well-known von Neumann-Fan minimax theorem (see, e.g., Fan 1953).

The assumption that the shock space is finite appears critical for the existence of a saddle point. In Section TR.5, we exhibit a simple game with a continuous, even compact, shock space that does not admit a worst-case distribution that is uniformly challenging for all mechanisms. When the shock space is arbitrary and the same conditions in Assumption TR.1 otherwise hold, it is possible to show an "asymmetric" minimax theorem, i.e., the saddle-point property holds, the minimax regret problem admits an optimal single-round mechanism, but the maximin regret problem does not necessarily have an optimal worst-case distribution. We prove such an asymmetric minimax result in Section TR.2.2.

TR.2.1 Proof of Theorem TR.1

We prove the result in three steps. First, we endow the space of randomized single-round direct mechanisms with a topology. Second, we show that an optimization problem corresponding to the single-round benchmark $\mathbb{E}_{\theta\sim F}[\text{OPT}(\delta_\theta, 1)]$ admits an optimal solution and the single-round benchmark is upper semi-continuous in F. The single-round benchmark is equivalently the value of the following optimization problem:

$$\int_\Theta \mathrm{OPT}(\delta_\theta,1)\mathrm{d}F(\theta) = \sup_{S\in\Delta(\Omega)^\Theta} \int_\Theta\int_\Omega u(\theta,\omega)\mathrm{d}S_\theta(\omega)\mathrm{d}F(\theta)$$

$$\text{s.t.} \int_\Omega v(\theta,\omega)\mathrm{d}S_\theta(\omega) \geq 0\,, \quad \forall\theta\in\Theta\,.$$

That the above optimization problem admits an optimal solution and the single-round benchmark is the objective value will be useful in showing the upper semi-continuity of the single-round benchmark. Finally, we prove the minimax result using the von Neumann-Fan minimax theorem.

Step 1. We endow the space of distributions over outcomes $\Delta(\Omega)$ with the weak* topology. Namely, for a sequence of probability distributions in $\Delta(\Omega)$, we have $G^k \rightharpoonup G$ if and only if $\int_\Omega \psi(\omega)dG^k(\omega) \to \int_\Omega \psi(\omega)dG(\omega)$ for all ψ in the space of continuous functions $C(\Omega,\mathbb{R})$ from Ω to $\mathbb{R}$ with the sup-norm topology. The space $\Delta(\Omega)$ is weak* compact because Ω is compact (see Aliprantis and Border, 2006, Theorem 15.11 in p. 513). We endow the space of randomized single-round direct mechanisms $\Delta(\Omega)^\Theta$ with pointwise convergence, i.e., for a sequence of mechanisms in $\Delta(\Omega)^\Theta$, we have $S^k \to S$ if and only if $S^k_\theta \rightharpoonup S_\theta$ for all $\theta\in\Theta$. By Tychonoff Product Theorem (see Aliprantis and Border, 2006, Theorem 2.61 in p. 52), the product space $\Delta(\Omega)^\Theta$ is compact because each factor is compact. Because the shock space is finite, the space of distributions over shocks $\Delta(\Theta) = \left\{f\in\mathbb{R}_+^{|\Theta|} : \sum_{\theta\in\Theta} f_\theta = 1\right\}$ is a compact subset of the Euclidean space. We have $F^k \to F$ if and only if $f^k_\theta \to f_\theta$ for all $\theta\in\Theta$.

Step 2. Let $U : \Delta(\Omega)^\Theta \times \Delta(\Omega) \to \mathbb{R}$ be the principal's utility functional which is given by $U(S,F) = \int_\Theta\int_\Omega u(\theta,\omega)\mathrm{d}S_\theta(\omega)\mathrm{d}F(\theta)$ for $S\in\Delta(\Omega)^\Theta$ and $F\in\Delta(\Theta)$. Because Θ is finite, we have $U(S,F) = \sum_{\theta\in\Theta}\int_\Omega u(\theta,\omega)\mathrm{d}S_\theta(\omega)f_\theta$. A similar definition holds for the agent's utility functional $V : \Delta(\Omega)^\Theta \times \Delta(\Omega) \to \mathbb{R}$, that is, $V(S,F) = \int_\Theta\int_\Omega v(\theta,\omega)\mathrm{d}S_\theta(\omega)\mathrm{d}F(\theta)$ for $S\in\Delta(\Omega)^\Theta$ and $F\in\Delta(\Theta)$. We denote by $\mathcal{C}(F) = \left\{S\in\Delta(\Omega)^\Theta : V(S,\delta_\theta)\geq 0, \forall\theta\in\Theta\right\}$ the set of interim individually rational mechanisms when the distribution is F. The single-round benchmark is given by $\mathbb{E}_{\theta\sim F}[\mathrm{OPT}(\delta_\theta,1)] = \sup_{S\in C(F)} U(F,S)$. The following holds:

- The principal's utility functional $U(S,F)$ is jointly upper semi-continuous: Consider sequences $F^k \to F$ and $S^k \to S$. Fix $\theta\in\Theta$. Let $u^k_\theta = \int_\Omega u(\theta,\omega)dS^k_\theta(\omega)$ and $u_\theta = \int_\Omega u(\theta,\omega)dS_\theta(\omega)$. Because $u(\theta,\omega)$ is upper semi-continuous in ω, we have $\limsup_{k\to\infty} u^k_\theta \leq u_\theta$. Because the shock space is finite we obtain, using the multiplication rule for limits, that $\limsup_{k\to\infty} U(S^k,F^k) =$

$\sum_{\theta\in\Theta} \limsup_{k\to\infty} u_\theta^k f_\theta^k \leq \sum_{\theta\in\Theta} u_\theta f_\theta = U(S,F)$.

- The agent's utility functional $V(S,F)$ is jointly continuous: This follows from repeating the same reasoning steps above for the principal's utility functional.

- The feasible set correspondence $\mathcal{C}(F)$ is compact-valued and non-empty: For any fixed distribution F, the set $\{S \in \Delta(\Omega)^\Theta : V(S,F) \geq 0\}$ is closed because upper level sets of upper semi-continuous functions are closed (see Aliprantis and Border, 2006, Corollary 2.60 in p. 52). In particular, $\{S \in \Delta(\Omega)^\Theta : V(S,\delta_\theta) \geq 0\}$ is closed for each $\theta \in \Theta$. Since the intersection of a finite number of closed sets is closed, $\mathcal{C}(F)$ is closed. Because $\Delta(\Omega)^\Theta$ is compact, compactness follows because the intersection of a closed set and a compact set is compact. Non-emptiness follows because the trivial mechanism that always determines the no-interaction outcome satisfies the interim IR constraint and is feasible.

- The feasible set correspondence $\mathcal{C}(F)$ is upper hemi-continuous in F: In fact, the feasible set $\mathcal{C}(F)$ is the same closed set for all distributions $F \in \Delta(\Theta)$ since it is defined in terms of the point-mass distributions δ_θ for $\theta \in \Theta$. Clearly, the feasible set correspondence is upper hemi-continuous.

Berge's Maximum Theorem implies that the optimization problem corresponding to the single-round benchmark admits an optimal solution and the single-round benchmark $\mathbb{E}_{\theta\sim F}[\mathrm{OPT}(\delta_\theta,1)]$ is upper semi-continuous in F (see Aliprantis and Border, 2006, Lemma 17.30 in p. 569).

Step 3. Let $V_{\theta,\theta'} : \Delta(\Omega)^\Theta \to \mathbb{R}$ be the agent's utility functional when his shock is θ and his report is θ', which is given by $V_{\theta,\theta'}(S) = \int_\Omega v(\theta,\omega)\mathrm{d}S_{\theta'}(\omega)$ for $S \in \Delta(\Omega)^\Theta$. We need to show a minimax result for the following problem:

$$\inf_{\substack{S\in\Delta(\Omega)^\Theta: \\ \text{(IC),(IR)}}} \sup_{F\in\Delta(\Omega)} \{\mathbb{E}_{\theta\sim F}[\mathrm{OPT}(\delta_\theta,1)] - U(S,F)\}$$

where (IR) is given by $V_{\theta,\theta}(S) \geq 0$ for all $\theta \in \Theta$ and (IC) is given by $V_{\theta,\theta}(S) \geq V_{\theta,\theta'}(S)$ for all $\theta,\theta' \in \Theta$.

The following hold:

- The space of feasible mechanisms $\hat{\mathcal{S}} := \{S \in \Delta(\Omega)^\Theta : \text{(IC) and (IR)}\}$ is compact: Note that $V_{\theta,\theta}(S) = V(S,\delta_\theta)$ where we denote by δ_θ the point-mass distribution that takes the value θ

with probability one in the expression $V(S,\cdot)$. Joint continuity of V implies that $V(S,\delta_\theta)$ is continuous in S for all θ. The IR constraint is closed because upper level sets of continuous functions are closed (see Aliprantis and Border, 2006, Corollary 2.60 in p. 52). A similar argument follows for the IC constraint by considering the functionals $\tilde{V}_{\theta,\theta'} : \Delta(\Omega)^\Theta \to \mathbb{R}$ given by $\tilde{V}_{\theta,\theta'}(S) = V_{\theta,\theta}(S) - V_{\theta,\theta'}(S)$ for $\theta, \theta' \in \Theta$. The result follows because $\Delta(\Omega)^\Theta$ is compact and the intersection of a finite number of closed sets with a compact set is compact.

- The objective is convex on $\hat{\mathcal{S}} := \{S \in \Delta(\Omega)^\Theta : \text{(IC) and (IR)}\}$, concave on $\Delta(\Theta)$, and lower semi-continuous on $\hat{\mathcal{S}}$, and upper semi-continuous on $\Delta(\Theta)$: The convexity in S and concavity in F follows because the objective is bilinear since we allow for randomized mechanisms and the feasible sets are convex. Lower semi-continuity in S follows because $U(S,F)$ is upper semi-continuous. We next prove the upper semi-continuity in F. From the previous step, we know that the single-round benchmark $\mathbb{E}_{\theta\sim F}[\mathrm{OPT}(\delta_\theta, 1)]$ is upper semi-continuous in F. Moreover, $U(S,F)$ is continuous in F for any fixed S because all linear functionals are continuous in finite dimensional spaces. The result follows because the sum of a continuous function and an upper semi-continuous function is upper semi-continuous.
- The feasible set is non-empty: This follows because the trivial mechanism that always determines the no-interaction outcome, i.e., the outcome $\emptyset$ such that $v(\theta,\emptyset) = 0$ for all θ, satisfies the IC/IR constraints and is feasible.

The theorem then follows from the von Neumann-Fan minimax theorem (Fan, 1953).

TR.2.2 An Asymmetric Saddle-Point Theorem for Arbitrary Shock Spaces

We discuss how to extend our saddle-point result, Theorem TR.1, to arbitrary shock spaces under the following assumption. More specifically, we show an asymmetric saddle-point result where the saddle-point property holds and the minimax regret formulation admits an optimal solution. We adapt the same notations from Section TR.2.1:

Assumption TR.2. *The game satisfies:*

(i) The outcome space Ω is normed and compact, and the shock space Θ is a topological space.

(ii) The principal's utility function $u(\theta,\omega)$ is upper semi-continuous in ω for all θ and uniformly bounded.

(iii) The agent's utility function $v(\theta,\omega)$ is continuous in ω for all θ.

We shall prove that:

$$\widehat{\text{Regret}} = \min_{\substack{S\in\mathcal{S}:\\ \text{(IC),(IR)}}} \sup_{F\in\Delta(\Theta)} \widehat{\text{Regret}}(S,F) = \sup_{F\in\Delta(\Theta)} \min_{\substack{S\in\mathcal{S}:\\ \text{(IC),(IR)}}} \widehat{\text{Regret}}(S,F)\,,$$

where $\mathcal{S} := \left\{S \in \Delta(\Omega)^\Theta : S \text{ is Borel measurable in } \Theta\right\}$. Moreover, there exists a (randomized) single-round direct IC/IR mechanism S^* such that $\widehat{\text{Regret}} = \sup_{F\in\Delta(\Theta)} \widehat{\text{Regret}}(S^*,F)$.

As in the proof of Theorem TR.1, we endow the space of distributions over outcomes $\Delta(\Omega)$ with the weak* topology. We endow the space $\mathcal{S}$ of measurable randomized single-round mechanisms with pointwise convergence. By Tychonoff Product Theorem (see Aliprantis and Border, 2006, Theorem 2.61 in p. 52), the product space $\Delta(\Omega)^\Theta$ is compact and Hausdorff. The space of measurable single-round mechanisms is closed with respect to the pointwise-convergence topology because $\Delta(\Omega)$ is metrizable (see Aliprantis and Border, 2006, Lemma 4.29 in p. 142). In turn, the space $\Delta(\Omega)$ is metrizable because Ω is a compact and normed (see Aliprantis and Border, 2006, Theorem 15.11 in p. 513). Because the intersection of a compact set and a closed set is compact, we obtain that $\mathcal{S}$ is compact.

From Theorem 2 in Fan (1953), it suffices to show that (i) the space of feasible mechanisms $\hat{\mathcal{S}} := \{S \in \mathcal{S} : \text{(IC) and (IR)}\}$ is compact, and (ii) the objective is convex in S, concave in F, and lower semi-continuous in S. For (i), it suffices to show that $V_{\theta,\theta}(S)$ is continuous in S for all θ. Consider a sequence of mechanisms $S^k \to S$, which is equivalent to $S^k_\theta \rightharpoonup S_\theta$ for all θ. Because $V_{\theta,\theta}(S) = \int_\Omega v(\theta,\omega)\mathrm{d}S_\theta(\omega)$, the result follows from weak* convergence in $\Delta(\Omega)$ since $v(\theta,\omega)$ is continuous in ω for all θ. For (ii), it suffices to show that $U(S,F)$ is upper semi-continuous in S. We equivalently write $U(S,F) = \int_\Theta U_\theta(S)\mathrm{d}F(\theta)$ with $U_\theta(S) = \int_\Omega u(\theta,\omega)\mathrm{d}S_\theta(\omega)$. Consider a sequence of mechanisms $S^k \to S$. From weak* convergence in $\Delta(\Omega)$, we obtain that $\limsup_{k\to\infty} U_\theta(S^k) \le U_\theta(S)$ for all θ because $u(\theta,\omega)$ is upper semi-continuous in ω. Because $U_\theta(S)$ is uniformly bounded, the result follows from the reverse Fatou's lemma.

TR.3 A Multi-Round Saddle-Point Theorem and Connections to Bayesian Mechanism Design

In the formulation presented in Section 2 of the main paper, the multi-round problem does not admit a saddle-point result in that there is no worst-case distribution for the agent that is uniformly challenging for all possible dynamic mechanisms. This is because for each possible distribution F, there are mechanisms tailored to the distribution that achieve $\mathrm{OPT}(F, T)$ arbitrarily closely and, hence, incur regret arbitrarily close to 0. Interestingly, a saddle-point result can be recovered if nature is allowed to use mixed strategies and randomize over distributions in $\mathcal{F}$, and this leads to a Bayesian mechanism design interpretation of the multi-round problem in which the principal has a Bayesian prior over the space of possible distributions.

Thus motivated, we introduce the space $\Delta(\mathcal{F})$ of all possible distributions over distributions in $\mathcal{F}$, i.e., Bayesian priors, and the following equivalent formulation of the multi-round problem:

$$\inf_{A \in \mathcal{A}} \sup_{G \in \Delta(\mathcal{F})} \mathbb{E}_{F \sim G} \left[\mathrm{Regret}(A, F, T) \right] .$$

This is equivalent to the original formulation, $\inf_{A \in \mathcal{A}} \sup_{F \in \mathcal{F}} \mathrm{Regret}(A, F, T)$, because nature cannot benefit from randomizing over distributions with respect to the objective $\mathbb{E}_{F \sim G} \left[\mathrm{Regret}(A, F, T) \right]$. The corresponding dual problem is

$$\sup_{G \in \Delta(\mathcal{F})} \inf_{A \in \mathcal{A}} \mathbb{E}_{F \sim G} \left[\mathrm{Regret}(A, F, T) \right] ,$$

where the inner infimum of the maximin regret can be equivalently written as the difference

$$\mathbb{E}_{F \sim G}[\mathrm{OPT}(F, T)] - \sup_{A \in \mathcal{A}} \mathbb{E}_{F \sim G}[\mathrm{PrincipalUtility}(A, F, T)] ,$$

where the first quantity is the optimal performance achievable, when F is known, averaged according to G and the second quantity is the principal utility under an optimal Bayesian mechanism where the principal knows the agent's unknown distribution F is drawn from the known prior G.

In the maximin regret formulation, nature adversarially chooses a Bayesian prior over $\mathcal{F}$ and then the principal responds by choosing a Bayesian optimal mechanism based on the prior. Given any $G \in \Delta(\mathcal{F})$, minimizing regret over mechanisms is equivalent to maximizing the principal utility. This problem – where the agent's distribution F itself is unknown to the principal but is drawn from a known prior G – is known in the literature as sequential screening. When a saddle point exists in the above formulation, our minimax regret problem is equivalent to a sequential screening problem under the least-favorable prior, giving a direct connection to Bayesian mechanism design.

The next result shows connections between the modified formulation of the multi-round problem and the single-round problem in terms of the saddle-point properties and saddle points.

Theorem TR.2. *Suppose Assumptions 1 and 2 of the main paper hold. The saddle-point property holds for the single-round minimax regret problem if and only if the saddle-point property holds for the multi-round minimax regret problem. Moreover, if a saddle point (S^*, F^*) exists in the single-round problem, then the repeated mechanism $(S^*)^{\times T}$ and the distribution over distributions G^* that assigns probability $F^*(\theta)$ to point-mass distribution δ_θ for all $\theta \in \Theta$ form a saddle point $((S^*)^{\times T}, G^*)$ in the multi-round problem. Conversely, if a saddle point exists in the multi-round problem, then a saddle point exists in the single-round problem (see the proof for a construction).*

Note $(\cdot)^{\times T}$ denotes T repetitions of a single-round mechanism. This means that we can similarly cast our multi-round problem as a simultaneous-move zero-sum game between the principal and nature and solve for a Nash equilibrium, assuming a saddle-point result holds for the single-round problem. More importantly, the above theorem shows that the multi-round minimax regret problem reduces to the Bayesian mechanism design problem where the objective is the expected regret and the principal implements a Bayesian optimal mechanism against a worst-case prior, that is, a distribution over distributions. To prove the result, we use our general results from Theorem 1 of the main paper, a classical result in the analysis of online algorithms known as Yao's principle (Yao, 1977), and an extension of the classical false-dynamics results of Baron and Besanko (1984) in a Bayesian setting. We prove it in Section TR.3.1.

TR.3.1 Proof of Theorem TR.2

We first prove that the single-round saddle-point property implies the multi-round saddle-point property and then prove the converse. Then, we prove the statements about the saddle points. Following the discussion in Section TR.3, we write the multi-round saddle-point property in terms of both $\inf_{A\in\mathcal{A}} \sup_{F\in\mathcal{F}} \text{Regret}(A, F, T)$ and $\inf_{A\in\mathcal{A}} \sup_{G\in\Delta(\mathcal{F})} \mathbb{E}_{F\sim G} [\text{Regret}(A, F, T)]$.

Only-If Direction for Saddle-Point Properties. We prove the result by showing first that the minimax regret is at least the maximin regret and then that the minimax regret is at most the maximin regret in the multi-round problem. For any incentive compatible mechanism $A \in \mathcal{A}$ and distribution over distributions $G \in \Delta(\mathcal{F})$, we have

$$\sup_{F \in \mathcal{F}} \text{Regret}(A, F, T) = \mathbb{E}_{F \sim G}\left[\sup_{F' \in \mathcal{F}} \text{Regret}(A, F', T)\right] \geq \mathbb{E}_{F \sim G}\left[\text{Regret}(A, F, T)\right] ,$$

because $\sup_{F' \in \mathcal{F}} \text{Regret}(A, F', T) \geq \text{Regret}(A, F, T)$. Taking the infimum over $A \in \mathcal{A}$ on both sides, we obtain

$$\inf_{A \in \mathcal{A}} \sup_{F \in \mathcal{F}} \text{Regret}(A, F, T) \geq \inf_{A \in \mathcal{A}} \mathbb{E}_{F \sim G}\left[\text{Regret}(A, F, T)\right] .$$

Taking the supremum over $G \in \Delta(\mathcal{F})$ on the right-hand side, we obtain

$$\inf_{A \in \mathcal{A}} \sup_{F \in \mathcal{F}} \text{Regret}(A, F, T) \geq \sup_{G \in \Delta(\mathcal{F})} \inf_{A \in \mathcal{A}} \mathbb{E}_{F \sim G}\left[\text{Regret}(A, F, T)\right] .$$

We next prove that the minimax regret is at most the maximin regret. Fix an arbitrary distribution $F' \in \Delta(\Theta)$ and take G' to be the corresponding distribution over atomic distributions (i.e., point-mass distributions) induced by F' as in the statement of the theorem. That is, for every measurable set $E \subseteq \mathcal{F}$, we have $G'(E) = F'(\{\theta \in \Theta : \delta_\theta \in E\})$ where $\delta_\theta \in \mathcal{F}$ is the point-mass distribution that assigns probability 1 to θ. This is possible under Assumption 1 of the main paper. We have

$$\begin{aligned}
&\sup_{G \in \Delta(\mathcal{F})} \inf_{A \in \mathcal{A}} \mathbb{E}_{F \sim G}\left[\text{Regret}(A, F, T)\right] \\
&\qquad \geq \inf_{A \in \mathcal{A}} \mathbb{E}_{F \sim G'}\left[\text{Regret}(A, F, T)\right] \\
&\qquad = \inf_{A \in \mathcal{A}} \mathbb{E}_{\theta \sim F'}\left[\text{Regret}(A, \delta_\theta, T)\right] \\
&\qquad = \inf_{A \in \mathcal{A}} \mathbb{E}_{\theta \sim F'}\left\{\text{OPT}(\delta_\theta, T) - \text{PrincipalUtility}(A, \delta_\theta, T)\right\} \\
&\qquad = \underbrace{\mathbb{E}_{\theta \sim F'}\left[\text{OPT}(\delta_\theta, T)\right]}_{(I)} - \underbrace{\sup\nolimits_{A \in \mathcal{A}} \mathbb{E}_{\theta \sim F'}\left[\text{PrincipalUtility}(A, \delta_\theta, T)\right]}_{(II)},
\end{aligned}$$

where the first inequality follows because G' is feasible; the first equality from the definition of G'; and the last equality from extracting the constant term (independent of A) from the infimum over $A \in \mathcal{A}$ and flipping the direction of the optimization. For the first term in the last expression, we have

$$(I) = \mathbb{E}_{\theta \sim F'}\left[\text{OPT}(\delta_\theta, T)\right] = T \cdot \mathbb{E}_{\theta \sim F'}\left[\text{OPT}(\delta_\theta, 1)\right] ,$$

where we used Proposition 1 of the main paper. The second term corresponds to a Bayesian mechanism design problem in which the agent's shock is constant throughout the rounds and the realization of the shock is private and drawn according to F'. Note Lemma A.1 of the main paper implies that, when the agent's shock is constant, any incentive compatible dynamic mechanism can be reduced to a single-round direct IC/IR mechanism. Therefore, we obtain

$$\begin{aligned}
(II) &= \sup_{A \in \mathcal{A}} \mathbb{E}_{\theta \sim F'} \left[\text{PrincipalUtility}(A, \delta_\theta, T) \right] \\
&= \sup_{A \in \mathcal{A}} \mathbb{E}_{\theta \sim F'} \left[T \cdot \text{PrincipalUtility}(S(A), \sigma^{\text{TR}}, \delta_\theta, 1) \right] \\
&\leq \sup_{\substack{S \in \Delta(\Omega)^\Theta: \\ \text{(IC),(IR)}}} \mathbb{E}_{\theta \sim F'} \left[T \cdot \text{PrincipalUtility}(S, \sigma^{\text{TR}}, \delta_\theta, 1) \right] \\
&= T \cdot \sup_{\substack{S \in \Delta(\Omega)^\Theta: \\ \text{(IC),(IR)}}} \int_\Theta \int_\Omega u(\theta, \omega) \mathrm{d}S_\theta(\omega) \mathrm{d}F'(\theta) \,,
\end{aligned}$$

where the second equality follows from Lemma A.1 of the main paper and $S(A)$ is the single-round direct IC/IR mechanism corresponding to an incentive compatible dynamic mechanism A; the second-to-last step follows because the set of direct IC/IR mechanisms is a superset of $\{S(A) \mid A \in \mathcal{A}\}$; and the last step follows from explicitly writing the principal's utility and using that S is direct and incentive compatible. Putting these two expressions together, we obtain

$$\begin{aligned}
\sup_{G \in \Delta(\mathcal{F})} \inf_{A \in \mathcal{A}} \mathbb{E}_{F \sim G} \left[\text{Regret}(A, F, T) \right] &\geq (I) - (II) \\
&\geq T \cdot \mathbb{E}_{\theta \sim F'} [\text{OPT}(\delta_\theta, 1)] - T \cdot \sup_{\substack{S \in \Delta(\Omega)^\Theta: \\ \text{(IC),(IR)}}} \int_\Theta \int_\Omega u(\theta, \omega) \mathrm{d}S_\theta(\omega) \mathrm{d}F'(\theta) \\
&= T \cdot \inf_{\substack{S \in \Delta(\Omega)^\Theta: \\ \text{(IC),(IR)}}} \left\{ \mathbb{E}_{\theta \sim F'} [\text{OPT}(\delta_\theta, 1)] - \int_\Theta \int_\Omega u(\theta, \omega) \mathrm{d}S_\theta(\omega) \mathrm{d}F'(\theta) \right\} \\
&= T \cdot \inf_{\substack{S \in \Delta(\Omega)^\Theta: \\ \text{(IC),(IR)}}} \widehat{\text{Regret}}(S, F') \,.
\end{aligned}$$

Note $F' \in \Delta(\Theta)$ was an arbitrary distribution. Taking a supremum over $\Delta(\Theta)$, we conclude that

$$\begin{aligned}\sup_{G\in\Delta(\mathcal{F})}\inf_{A\in\mathcal{A}}\mathbb{E}_{F\sim G}\left[\text{Regret}(A,F,T)\right] &\geq T\cdot \sup_{F\in\Delta(\Theta)}\inf_{\substack{S\in\Delta(\Omega)^{\Theta}:\\ \text{(IC),(IR)}}}\widehat{\text{Regret}}(S,F)\\ &= T\cdot \inf_{\substack{S\in\Delta(\Omega)^{\Theta}:\\ \text{(IC),(IR)}}}\sup_{F\in\Delta(\Theta)}\widehat{\text{Regret}}(S,F)\\ &= \inf_{A\in\mathcal{A}}\sup_{F\in\mathcal{F}}\text{Regret}(A,F,T)\,,\end{aligned}$$

where the first equality follows from the saddle-point property for the single-round problem and the second equality from Theorem 1 of the main paper, which holds since Assumptions 1 and 2 of the main paper hold.

If Direction for Saddle-Point Properties. The max-min inequality implies that the minimax regret is at least the maximin regret in the single-round problem and it remains to show that the minimax regret is at most the maximin regret. To do so, it suffices to show that for every $\epsilon > 0$, there exists some $F \in \Delta(\Theta)$ such that $\widehat{\text{Regret}} \leq \epsilon + \inf_{\substack{S\in\Delta(\Omega)^{\Theta}:\\ \text{(IC),(IR)}}}\widehat{\text{Regret}}(S,F)$ where $\widehat{\text{Regret}} = \inf_{\substack{S\in\Delta(\Omega)^{\Theta}:\\ \text{(IC),(IR)}}}\sup_{F\in\Delta(\Theta)}\widehat{\text{Regret}}(S,F)$.

Fix $\epsilon > 0$. From the multi-round saddle-point property, that is,

$$\text{Regret}(T) = \inf_{A\in\mathcal{A}}\sup_{F\in\mathcal{F}}\text{Regret}(A,F,T) = \sup_{G\in\Delta(\mathcal{F})}\inf_{A\in\mathcal{A}}\mathbb{E}_{F\sim G}\left[\text{Regret}(A,F,T)\right]\,,$$

we know that there exists some $G \in \Delta(\mathcal{F})$ such that

$$\text{Regret}(T) \leq T\cdot\epsilon + \inf_{A\in\mathcal{A}}\mathbb{E}_{F\sim G}\left[\text{Regret}(A,F,T)\right]\,. \tag{TR-2}$$

Consider the distribution $\hat{F} \in \Delta(\Theta)$ obtained from marginalizing over G, i.e., $\hat{F}(E) = \mathbb{E}_{F\sim G}\left[F(E)\right]$ for any measurable set $E \subseteq \Theta$. We have that

$$\mathbb{E}_{F\sim G}\left[\text{OPT}(F,T)\right] \leq \mathbb{E}_{F\sim G,\theta\sim F}[\text{OPT}(\delta_\theta,T)] = \mathbb{E}_{\theta\sim\hat{F}}[\text{OPT}(\delta_\theta,T)] = T\cdot\mathbb{E}_{\theta\sim\hat{F}}[\text{OPT}(\delta_\theta,1)]\,, \tag{TR-3}$$

where the inequality follows from Assumption 2 of the main paper and the law of total expectation; the first equality from the definition of the compounded distribution $\hat{F}$; and the last equality from Proposition 1 of the main paper. For any single-round direct IC/IR mechanism $S \in \Delta(\Omega)^{\Theta}$, consider the direct mechanism $S^{\times T}$ obtained from T repetitions of S. In turn, the principal's utility satisfies

$$\begin{aligned}\mathbb{E}_{F\sim G}\left[\text{PrincipalUtility}(S^{\times T}, F, T)\right] &= T \cdot \mathbb{E}_{F\sim G}\left[\text{PrincipalUtility}(S, \sigma^{\text{TR}}, F, 1)\right] \\ &= T \cdot \text{PrincipalUtility}(S, \sigma^{\text{TR}}, \hat{F}, 1)\,, \end{aligned} \tag{TR-4}$$

where the first step follows from Lemma A.2 of the main paper, and the second step follows from that the mechanism S is static and does not screen the agent for his distribution and from the definition of $\hat{F}$. Theorem 1 of the main paper implies $\text{Regret}(T) = T \cdot \widehat{\text{Regret}}$ and, consequently,

$$\begin{aligned}T \cdot \widehat{\text{Regret}} &\leq T \cdot \epsilon + \inf_{\substack{S \in \Delta(\Omega)^\Theta: \\ \text{(IC),(IR)}}} \mathbb{E}_{F\sim G}\left[\text{Regret}(S^{\times T}, F, T)\right] \\ &= T \cdot \epsilon + \inf_{\substack{S \in \Delta(\Omega)^\Theta: \\ \text{(IC),(IR)}}} \mathbb{E}_{F\sim G}\left[\text{OPT}(F, T) - \text{PrincipalUtility}(S^{\times T}, F, T)\right] \\ &\leq T \cdot \epsilon + T \cdot \inf_{\substack{S \in \Delta(\Omega)^\Theta: \\ \text{(IC),(IR)}}} \left\{\mathbb{E}_{\theta\sim\hat{F}}[\text{OPT}(\delta_\theta, 1)] - \text{PrincipalUtility}(S, \sigma^{\text{TR}}, \hat{F}, 1)\right\} \\ &= T \cdot \epsilon + T \cdot \inf_{\substack{S \in \Delta(\Omega)^\Theta: \\ \text{(IC),(IR)}}} \widehat{\text{Regret}}(S, \hat{F})\,,\end{aligned}$$

where the first inequality follows from (TR-2) and using that the space of static mechanisms repeating a single-round direct IC/IR mechanism is a subset of all possible incentive compatible multi-round mechanisms $\mathcal{A}$; the first equality from the definition of Regret; the second inequality from (TR-3) and (TR-4); and the last equality from the definition of $\widehat{\text{Regret}}$. The result follows from dividing both sides by $T \geq 1$.

First Statement on Saddle Points. Let (S^*, F^*) be a saddle point for the single-round problem such that

$$\widehat{\text{Regret}}(S^*, F) \leq \widehat{\text{Regret}}(S^*, F^*) \leq \widehat{\text{Regret}}(S, F^*)$$

for any $S \in \Delta(\Omega)^\Theta$ satisfying the IC/IR constraints and $F \in \Delta(\Theta)$. Note that the existence of the saddle point implies that $\widehat{\text{Regret}} = \widehat{\text{Regret}}(S^*, F^*)$ and the single-round saddle-point property holds and that S^* and F^* are optimal solutions in respective single-round problems achieving the objective value of $\widehat{\text{Regret}}$. To see this, the existence of the saddle point implies

$$\begin{aligned}
\inf_{\substack{S\in\Delta(\Omega)^\Theta:\\ \text{(IC),(IR)}}} \sup_{F\in\Delta(\Theta)} \widehat{\text{Regret}}(S,F) &\le \sup_{F\in\Delta(\Theta)} \widehat{\text{Regret}}(S^*,F) \\
&\le \widehat{\text{Regret}}(S^*,F^*) \\
&\le \inf_{\substack{S\in\Delta(\Omega)^\Theta:\\ \text{(IC),(IR)}}} \widehat{\text{Regret}}(S,F^*) \le \sup_{F\in\Delta(\Theta)} \inf_{\substack{S\in\Delta(\Omega)^\Theta:\\ \text{(IC),(IR)}}} \widehat{\text{Regret}}(S,F)\,,
\end{aligned}$$

and the max-min inequality implies all the above relations are equalities. By the first statement of the theorem which we proved above, the multi-round saddle-point property holds. By Theorem 1 of the main paper, $(S^*)^{\times T}$ is an optimal solution to the multi-round minimax regret problem since S^* is an optimal solution to the single-round problem; note $(\cdot)^{\times T}$ denotes T repetitions of a single-round mechanism.

To show that $((S^*)^{\times T}, G^*)$, for G^* constructed as in the theorem statement, is a saddle point in the multi-round problem, it suffices to show that G^* is an optimal solution to the multi-round maximin regret problem, i.e., $\sup_{G\in\Delta(\mathcal{F})}\inf_{A\in\mathcal{A}}\mathbb{E}_{F\sim G}\left[\text{Regret}(A,F,T)\right] = \inf_{A\in\mathcal{A}}\mathbb{E}_{F\sim G^*}\left[\text{Regret}(A,F,T)\right]$. Then, we would have

$$\begin{aligned}
\mathbb{E}_{F\sim G^*}\left[\text{Regret}((S^*)^{\times T},F,T)\right] &\ge \inf_{A\in\mathcal{A}} \mathbb{E}_{F\sim G^*}\left[\text{Regret}(A,F,T)\right] \\
&= \sup_{G\in\Delta(\mathcal{F})} \inf_{A\in\mathcal{A}} \mathbb{E}_{F\sim G}\left[\text{Regret}(A,F,T)\right] \\
&= \inf_{A\in\mathcal{A}} \sup_{G\in\Delta(\mathcal{F})} \mathbb{E}_{F\sim G}\left[\text{Regret}(A,F,T)\right] \\
&= \sup_{G\in\Delta(\mathcal{F})} \mathbb{E}_{F\sim G}\left[\text{Regret}((S^*)^{\times T},F,T)\right] \\
&\ge \mathbb{E}_{F\sim G^*}\left[\text{Regret}((S^*)^{\times T},F,T)\right]\,,
\end{aligned}$$

where the first equality would follow if G^* is an optimal solution to the multi-round maximin regret problem; the second equality is from the multi-round saddle-point property; and the last equality is from that $(S^*)^{\times T}$ is an optimal solution to the multi-round minimax regret problem. Since the first and last expressions are the same, all the above relations would be equalities and we would have

$$\sup_{G\in\Delta(\mathcal{F})} \mathbb{E}_{F\sim G}\left[\text{Regret}((S^*)^{\times T},F,T)\right] = \mathbb{E}_{F\sim G^*}\left[\text{Regret}((S^*)^{\times T},F,T)\right] = \inf_{A\in\mathcal{A}} \mathbb{E}_{F\sim G^*}\left[\text{Regret}(A,F,T)\right]\,,$$

which would imply that $((S^*)^{\times T}, G^*)$ is a saddle point in the multi-round problem, as desired.

We now show that G^* is an optimal solution to the multi-round maximin regret problem. Let $A \in \mathcal{A}$ be an arbitrary incentive compatible dynamic mechanism. Then, we have

$$\begin{aligned}
\mathbb{E}_{F\sim G^*}\left[\text{Regret}(A,F,T)\right] &= \mathbb{E}_{\theta\sim F^*}\left[\text{Regret}(A,\delta_\theta,T)\right] \\
&= \mathbb{E}_{\theta\sim F^*}[\text{OPT}(\delta_\theta,T) - \text{PrincipalUtility}(A,\delta_\theta,T)] \\
&= \mathbb{E}_{\theta\sim F^*}[T\cdot\text{OPT}(\delta_\theta,1) - T\cdot\text{PrincipalUtility}(S(A),\sigma^{\text{TR}},\delta_\theta,1)] \\
&= T\cdot\widehat{\text{Regret}}(S(A),F^*) \\
&\geq T\cdot\widehat{\text{Regret}} \\
&= \sup_{G\in\Delta(\mathcal{F})}\inf_{A\in\mathcal{A}}\mathbb{E}_{F\sim G}\left[\text{Regret}(A,F,T)\right]\,,
\end{aligned}$$

where the first equality is by the construction of G^*; the second by the definition of Regret; the third by Proposition 1 and Lemma A.1 of the main paper where $S(A)$ is the single-round direct IC/IR mechanism corresponding to the dynamic mechanism A; the fourth by the definition of $\widehat{\text{Regret}}$; the second-to-last step is from that (S^*,F^*) is a saddle point in the single-round problem and $\widehat{\text{Regret}} = \widehat{\text{Regret}}(S^*,F^*)$; and the last step is because Theorem 1 of the main paper and the multi-round saddle-point property implies

$$T\cdot\widehat{\text{Regret}} = \text{Regret}(T) = \inf_{A\in\mathcal{A}}\sup_{F\in\mathcal{F}}\text{Regret}(A,F,T) = \sup_{G\in\Delta(\mathcal{F})}\inf_{A\in\mathcal{A}}\mathbb{E}_{F\sim G}\left[\text{Regret}(A,F,T)\right]\,.$$

As $A \in \mathcal{A}$ was arbitrary,

$$\inf_{A\in\mathcal{A}}\mathbb{E}_{F\sim G^*}\left[\text{Regret}(A,F,T)\right] \geq \sup_{G\in\Delta(\mathcal{F})}\inf_{A\in\mathcal{A}}\mathbb{E}_{F\sim G}\left[\text{Regret}(A,F,T)\right]\,,$$

and G^* is an optimal solution to the multi-round maximin regret problem. This completes the proof.

Second Statement on Saddle Points. Let the pair (A^*,G^*) of an incentive compatible dynamic mechanism $A^* \in \mathcal{A}$ and a distribution over distributions $G^* \in \Delta(\mathcal{F})$ be a saddle point for the multi-round problem. As in the proof for the first statement on saddle points, the existence of the saddle point implies that the saddle-point property holds in the multi-round problem and, furthermore, that A^* is an optimal dynamic mechanism in the multi-round minimax regret problem and G^* is an optimal solution in the multi-round maximin regret problem, both with the objective value of $\text{Regret}(T)$. By the first part of the theorem, the single-round saddle-point property holds. Given this, it suffices to show that the single-round minimax regret problem admits an optimal solution S^*

and that the single-round maximin regret problem admits an optimal solution F^*, or equivalently, the worst-case distribution. Then, the pair (S^*, F^*) would form a saddle point in the single-round problem, as desired. This is because if S^* and F^* are optimal solutions in respective single-round problems, we would have

$$\begin{aligned}\inf_{\substack{S\in\Delta(\Omega)^\Theta:\\ \text{(IC),(IR)}}} \sup_{F\in\Delta(\Theta)} \widehat{\text{Regret}}(S,F) &= \sup_{F\in\Delta(\Theta)} \widehat{\text{Regret}}(S^*,F)\\ &\geq \widehat{\text{Regret}}(S^*,F^*)\\ &\geq \inf_{\substack{S\in\Delta(\Omega)^\Theta:\\ \text{(IC),(IR)}}} \widehat{\text{Regret}}(S,F^*) = \sup_{F\in\Delta(\Theta)} \inf_{\substack{S\in\Delta(\Omega)^\Theta:\\ \text{(IC),(IR)}}} \widehat{\text{Regret}}(S,F)\,,\end{aligned}$$

where the equalities would follow from the optimality of S^* and F^*. The single-round saddle-point property would imply that all the relations in the above sequence are equalities and $\widehat{\text{Regret}} = \widehat{\text{Regret}}(S^*, F^*)$. In particular,

$$\sup_{F\in\Delta(\Theta)} \widehat{\text{Regret}}(S^*,F) = \widehat{\text{Regret}}(S^*,F^*) = \inf_{\substack{S\in\Delta(\Omega)^\Theta:\\ \text{(IC),(IR)}}} \widehat{\text{Regret}}(S,F^*)\,,$$

and this would imply that (S^*, F^*) is a saddle point.

The existence of an optimal single-round direct IC/IR mechanism S^* in the single-round problem follows from the existence of an optimal dynamic mechanism, in particular, A^*, by Theorem 1 of the main paper. We can construct such an optimal mechanism S^* from A^* as discussed in the proof of Theorem 1 of the main paper.

We now show there exists an optimal distribution F^* in the single-round maximin regret problem. Consider the distribution F^* constructed from G^* such that $F^*(Q) = \Pr(\theta \in Q \mid F \sim G^*, \theta \sim F)$ for any $Q \subseteq \Theta$ and, in particular, $\Pr(\hat{\theta} = \theta \mid \hat{\theta} \sim F^*) = \Pr(\hat{\theta} = \theta \mid F \sim G^*, \hat{\theta} \sim F)$ for any $\theta \in \Theta$. Note

$$\text{Regret}(T) \leq \inf_{A\in\mathcal{A}} \mathbb{E}_{F\sim G^*}[\text{Regret}(A,F,T)] \leq \inf_{\substack{S\in\Delta(\Omega)^\Theta:\\ \text{(IC),(IR)}}} \mathbb{E}_{F\sim G^*}[\text{Regret}(S^{\times T},F,T)]\,, \tag{TR-5}$$

where the first inequality is by the optimality of G^* and the second is because the set of static mechanisms that repeat a single-round direct IC/IR mechanism is a subset of incentive compatible dynamic mechanisms $\mathcal{A}$. Furthermore, note that, for any single-round direct IC/IR mechanism S and distribution $F \in \Delta(\Theta)$,

$$\begin{aligned}
\mathrm{Regret}(S^{\times T}, F, T) &= \mathrm{OPT}(F, T) - \mathrm{PrincipalUtility}(S^{\times T}, \sigma^{\mathrm{TR}}, F, T) \\
&= \mathrm{OPT}(F, T) - T \cdot \mathrm{PrincipalUtility}(S, \sigma^{\mathrm{TR}}, F, 1) \\
&\leq T\left(\mathbb{E}_{\theta\sim F}[\mathrm{OPT}(\delta_\theta, 1)] - \mathrm{PrincipalUtility}(S, \sigma^{\mathrm{TR}}, F, 1)\right) \\
&= T\widehat{\mathrm{Regret}}(S, F)\,,
\end{aligned}$$

where σ^{TR} is the agent's truthful reporting strategy, the second step follows from Lemma A.2 of the main paper, and the second-to-last step follows from Assumption 2 and Proposition 1 of the main paper. Combining with (TR-5), we obtain

$$\mathrm{Regret}(T) \leq T \cdot \inf_{\substack{S\in\Delta(\Omega)^\Theta: \\ \text{(IC),(IR)}}} \mathbb{E}_{F\sim G^*}[\widehat{\mathrm{Regret}}(S, F)]\,.$$

In the last expression, we can equivalently write the expected quantity inside the infimum as, for any single-round direct IC/IR mechanism S,

$$\begin{aligned}
\mathbb{E}_{F\sim G^*}[\widehat{\mathrm{Regret}}(S, F)] &= \mathbb{E}_{F\sim G^*}\left[\int_\Theta \mathrm{OPT}(\delta_\theta, 1)\mathrm{d}F(\theta) - \int_\Theta\int_\Omega u(\theta,\omega)\mathrm{d}S_\theta(\omega)\mathrm{d}F(\theta)\right] \\
&= \mathbb{E}_{F\sim G^*,\theta\sim F}\left[\mathrm{OPT}(\delta_\theta, 1) - \int_\Omega u(\theta,\omega)\mathrm{d}S_\theta(\omega)\right] \\
&= \mathbb{E}_{\theta\sim F^*}\left[\mathrm{OPT}(\delta_\theta, 1) - \int_\Omega u(\theta,\omega)\mathrm{d}S_\theta(\omega)\right] \\
&= \widehat{\mathrm{Regret}}(S, F^*)\,,
\end{aligned}$$

where the first step is by the definition of $\widehat{\mathrm{Regret}}$ notion; the second step is by the total law of expectation; the third step is by the construction of F^*; and the last step is by the definition of $\widehat{\mathrm{Regret}}$ notion. Hence, it follows that

$$\mathrm{Regret}(T) \leq T \cdot \inf_{\substack{S\in\Delta(\Omega)^\Theta: \\ \text{(IC),(IR)}}} \widehat{\mathrm{Regret}}(S, F^*)\,.$$

Since $\mathrm{Regret}(T) = T\cdot\widehat{\mathrm{Regret}}$ by Theorem 1 of the main paper, which holds since Assumptions 1 and 2 of the main paper hold, and the single-round saddle-point property holds, we have

$$\widehat{\mathrm{Regret}} = \sup_{F\in\Delta(\Theta)} \inf_{\substack{S\in\Delta(\Omega)^\Theta: \\ \text{(IC),(IR)}}} \widehat{\mathrm{Regret}}(S, F) \leq \inf_{\substack{S\in\Delta(\Omega)^\Theta: \\ \text{(IC),(IR)}}} \widehat{\mathrm{Regret}}(S, F^*)\,.$$

Then, F^* is an optimal solution in the single-round maximin regret problem. This completes the proof.

TR.4 Additional Materials on Dynamic Selling Mechanisms

We consider other versions of the dynamic selling problem presented in Section 4.1 of the main paper. In Section TR.4.1, we consider welfare maximization in the single-good case. In Section TR.4.2, we consider both revenue and welfare maximization in the multiple-goods case where the principal sells identical copies of n goods, one per good, in each round.

TR.4.1 Welfare Maximization

For welfare maximization in the single-good case, the principal's utility function is $u(\theta, \omega) = \theta \cdot \hat{x}$ for outcome $\omega = (\hat{x}, \hat{p})$. The single-round benchmark is $\mathbb{E}_{\theta \sim F}[\mathrm{OPT}(\delta_\theta, 1)] = \mathbb{E}_{\theta \sim F}[\theta]$ because $\mathrm{OPT}(\delta_\theta, 1) = \theta$ which is achieved by choosing the outcome $(1, 0)$ or, equivalently, always allocating the item at no cost. Since we have $\mathcal{F} = \Delta([0,1])$ as in the revenue maximization version, Assumption 1 of the main paper holds. Assumption 2 also holds because we always have the per-round allocation $\hat{x} \in [0,1]$ and, hence, $\mathrm{OPT}(F, T) \leq T\mathbb{E}_{\theta \sim F}[\theta] = \mathbb{E}_{\theta \sim F}[\mathrm{OPT}(\delta_\theta, T)]$, where we used Proposition 1 of the main paper. Therefore, Theorem 1 of the main paper applies. We show that the minimax regret for the single-round problem for direct IC/IR mechanisms is 0 and, thus, that for the multi-round problem is 0 because $\mathrm{Regret}(T) = T \cdot \widehat{\mathrm{Regret}}$.

We formally state the minimax regret result as follows:

Proposition TR.1. *For welfare maximization in the dynamic selling mechanism design problem with one good, the minimax regret is 0 and an optimal solution is allocating items for free, which is T repetitions of the same strategy.*

Clearly, allocating items for free satisfies the IC/IR constraints when considered as a single-round direct mechanism. This is a degenerate result with a trivial solution but we think it is interesting that our general result captures it. Note the optimal multi-round solution of allocating for free is still optimal if the private shock distribution is known to the principal. In this sense, no dynamic/adaptive strategy with sophisticated learning was necessary to begin with and the distributional information of the agent's private distribution was not needed.

Proof of Proposition TR.1. We solve the single-round minimax regret problem for direct IC/IR mechanisms which is

$$\inf_{\substack{S\in\Delta(\Omega)^{\Theta}:\\ \text{(IC),(IR)}}} \sup_{F\in\Delta(\Theta)} \left\{ \int_{\Theta} \theta \mathrm{d}F(\theta) - \int_{\Theta}\int_{\Omega} u(\theta,\omega)\mathrm{d}S_{\theta}(\omega)\mathrm{d}F(\theta) \right\}.$$

The trivial single-round mechanism S^* that always allocates the item for free, i.e., S^* is the deterministic mechanism π^* for which $\pi^*(\theta) = (1, 0)$ for all $\theta \in \Theta$, satisfies the IC/IR constraints and is an optimal solution because it achieves $\mathbb{E}_{\theta\sim F}[\theta]$ for any agent's distribution F. It is easy to see that we can take any distribution F and (S^*, F) is a saddle point for the single-round problem. Hence, $\widehat{\text{Regret}} = 0$. By Theorem 1 of the main paper, we obtain $\text{Regret}(T) = 0$ and repeating S^* over T rounds is an optimal solution with the minimax regret of $\text{Regret}(T) = 0$. □

TR.4.2 Multiple Goods

Building on Section 4.1 of the main paper and Section TR.4.1, we consider the multiple-goods version of the dynamic selling mechanism design problem where the principal has n goods and sells independent copies (or units) of these goods, one per good, in each round. In particular, this is an application where the agent's private shock is multidimensional. For welfare maximization, it is straightforward to see that the same strategy from the single-good case of giving away for free is an optimal solution and achieves the minimax regret of 0 in both the single-round and multi-round problems. For revenue maximization, we show below that repeatedly selling each good separately according to the randomized posted pricing mechanism given in Proposition 3 of the main paper for the single-good case is optimal.

In each round, the agent sees n goods (that is, copies of) and realizes his value for each good in the range $[0, 1]$ according to a private shock distribution F. The values of the goods can be arbitrarily correlated but are drawn from the joint distribution independently across rounds. The shock space is $\Theta = [0, 1]^n$ and the agent's private shock distribution $F \in \mathcal{F}$ is a distribution over Θ where $\mathcal{F} = \Delta(\Theta)$. The outcome space is $\Omega = \{0, 1\}^n \times \mathbb{R}$. Using superscript i to denote the coordinate corresponding to the i-th good, an outcome $\omega = (\hat{x}, \hat{p}) \in \Omega$ is given by allocations $(\hat{x}^1, \ldots, \hat{x}^n)$ of the goods and payment $\hat{p}$. Given an outcome $\omega = (\hat{x}, \hat{p})$, the agent's utility function is $v(\theta, \omega) = \sum_{i=1}^{n} \theta^i \cdot \hat{x}^i - \hat{p}$. For revenue maximization, the principal's utility function is $u(\theta, \omega) = \hat{p}$, and for welfare maximization, the principal's utility function is $u(\theta,\omega) = \sum_{i=1}^{n} \theta^i \cdot \hat{x}^i$.

The main result is as follows.

Proposition TR.2. *For revenue maximization in the dynamic selling mechanism design problem with n goods, the minimax regret is $\frac{n}{e}T$ and an optimal solution is T repetitions of selling each good separately via the randomized posted pricing mechanism S^* given in Proposition 3 of the main paper. For welfare maximization, the minimax regret is 0 and an optimal solution is allocating goods for free, which is T repetitions of the same strategy.*

The proof for the welfare maximization part is almost identical to the single-good case in Section TR.4.1 and thus omitted. Henceforth, we discuss the revenue maximization part. In each round, we implement the single-good solution (x^*, p^*) for each good to determine $\hat{x}^i = x^*(\theta^i)$ and $\hat{p}^i = p^*(\theta^i)$. Then, the overall allocation and payment are $(\hat{x}^1, \ldots, \hat{x}^n)$ and $\hat{p} = \sum_i \hat{p}^i$. Note $\mathrm{OPT}(\delta_\theta, 1) = \sum_{i=1}^{n} \theta^i$ because the principal can extract the full surplus of the agent by allocating all the goods and charging $\sum_{i=1}^{n} \theta^i$. Hence, the single-round benchmark is $\mathbb{E}_{\theta\sim F}[\mathrm{OPT}(\delta_\theta, 1)] = \mathbb{E}_{\theta\sim F}[\sum_{i=1}^{n} \theta^i]$. Since $\mathcal{F} = \Delta(\Theta)$, Assumption 1 of the main paper holds. For Assumption 2 of the main paper, we follow the same reasoning steps as in the single-good case in the main paper.

As in the single-good case, the general result (Theorem 1 of the main paper) applies. To show the claimed dynamic mechanism is an optimal solution to the multi-round problem, we show that the single-round solution of selling each good separately is an optimal solution to the single-round minimax regret problem, via a saddle-point formulation. Note Kocyigit et al. (2018) have recently shown the same single-round result. We prove it for completeness below.

Proposition TR.3. *Let $S^{*,n}$ denote the single-round direct IC/IR mechanism that separately implements $S^* = (x^*, p^*)$ for each good and $F^{*,n}$ denote the agent's distribution with $F^{*,n}(\theta) = F^*(\min_i \theta^i)$ for any $\theta \in \Theta$ (i.e., perfectly correlated values) where S^* and F^* are as given in Propositions 3 and B.1 of the main paper. Then, $\widehat{\mathrm{Regret}} = \widehat{\mathrm{Regret}}(S^{*,n}, F^{*,n}) = \frac{n}{e}$ and*

$$\widehat{\mathrm{Regret}}(S^{*,n}, F) \leq \widehat{\mathrm{Regret}}(S^{*,n}, F^{*,n}) \leq \widehat{\mathrm{Regret}}(S, F^{*,n}),$$

for any $S \in \Delta(\Omega)^\Theta$ satisfying the IC/IR constraints and $F \in \Delta(\Theta)$.

Proof. First, we show that $F^{*,n}$ is a solution to

$$\max_{F \in \Delta(\Theta)} \widehat{\text{Regret}}(S^{*,n}, F)\,.$$

The above expression is equivalent to

$$\max_{F \in \Delta(\Theta)} \mathbb{E}_{\theta \sim F}\left[\sum_{i=1}^{n} (\theta^i - p^*(\theta^i))\right] = \max_{F \in \Delta(\Theta)} \sum_{i=1}^{n} \mathbb{E}_{\theta^i \sim F^i}\left[\theta^i - p^*(\theta^i)\right]\,,$$

where F^i is the marginal distribution of F for the i-th good. Interchanging the maximum and summation, the last expression is upper bounded by

$$\sum_{i=1}^{n} \max_{F^i \sim \Delta(\Theta^i)} \mathbb{E}_{\theta^i \in F^i}\left[\theta^i - p^*(\theta^i)\right]\,,$$

where we can independently choose distributions F_i over $\Theta^i := [0,1]$ as one-dimensional distributions. By Proposition B.1 of the main paper, an optimal one-dimensional distribution F^i in each summand is F^* which yields the value of $\frac{1}{e}$ and the upper bound evaluates to $\frac{n}{e}$. Since $F^{*,n}$ is a distribution with marginal distributions equal to F^* for each good, we have $\mathbb{E}_{\theta \sim F}[\sum_i (\theta^i - p^*(\theta^i))] = \frac{n}{e}$ and, hence, the distribution is a solution to the original maximization we started out with. Furthermore, it follows that $\widehat{\text{Regret}}(S^{*,n}, F^{*,n}) = \frac{n}{e}$.

Second, we show that $S^{*,n}$ is a solution to

$$\min_{\substack{S \in \Delta(\Omega)^{\Theta}: \\ \text{(IC),(IR)}}} \widehat{\text{Regret}}(S, F^{*,n})\,.$$

Substituting in $F^{*,n}$, the above optimization problem reduces to

$$\min_{\substack{S=(x,p) \in \Delta(\Omega)^{\Theta}: \\ \text{(IC),(IR)}}} \mathbb{E}_{\phi \sim F^*}[n\phi - p(\phi, \ldots, \phi)]$$

where (x, p) are the interim allocation and payment rules. We relax the IC/IR constraints by restricting to the one-dimensional space where the agent's per-good shocks are the same. Recall:

$$\theta^{\mathsf{T}} \cdot x(\theta) - p(\theta) \geq \theta^{\mathsf{T}} \cdot x(\theta') - p(\theta')\,, \quad \forall \theta, \theta' \in [0,1]^n\,. \tag{IC}$$

$$\theta^{\mathsf{T}} \cdot x(\theta) - p(\theta) \geq 0\,, \quad \forall \theta \in [0,1]^n\,. \tag{IR}$$

The relaxed IC/IR constraints are:

$$\phi \cdot \sum_{i=1}^{n} x^i(\phi \cdot \vec{\mathbf{1}}) - p(\phi \cdot \vec{\mathbf{1}}) \geq \phi \cdot \sum_{i=1}^{n} x^i(\phi' \cdot \vec{\mathbf{1}}) - p(\phi' \cdot \vec{\mathbf{1}})\,, \quad \forall \phi, \phi' \in [0,1]\,. \tag{IC'}$$

$$\phi \cdot \sum_{i=1}^{n} x^i(\phi \cdot \vec{\mathbf{1}}) - p(\phi \cdot \vec{\mathbf{1}}) \geq 0\,, \quad \forall \phi \in [0,1]\,. \tag{IR'}$$

Note $\vec{\mathbf{1}}$ is the all-ones vector $(1, \ldots, 1)$. Then, the value of the optimization problem is lower bounded by

$$\min_{\substack{S=(x,p)\in\Delta(\Omega)^{\Theta}:\\ \text{(IC'),(IR')}}} \mathbb{E}_{\phi\sim F^*}\left[n\phi - p(\phi, \ldots, \phi)\right]. \tag{TR-6}$$

We can transform a single-round mechanism $S = (x, p)$ for multiple goods to $\tilde{S} = (\tilde{x}, \tilde{p})$ for a single good by taking $\tilde{x}(\phi) = \frac{1}{n}\sum_{i=1}^{n} x^i(\phi \cdot \vec{\mathbf{1}})$ and $\tilde{p}(\phi) = \frac{1}{n}p(\phi \cdot \vec{\mathbf{1}})$ for report $\phi \in [0,1]$. For any multi-good single-round mechanism S satisfying the relaxed IC/IR constraints, the transformation yields a single-good single-round mechanism satisfying the original IC/IR constraints for the single-good case. Furthermore, note that the multi-good mechanism implementing a single-good single-round direct IC/IR mechanism separately for each good satisfies the relaxed IC/IR constraints and yields the same single-good mechanism via the transformation. Hence, (TR-6) is equal to the following in terms of the objective value:

$$\min_{\substack{\tilde{S}=(\tilde{x},\tilde{p})\in\tilde{\mathcal{S}}:\\ \text{(IC"),(IR")}}} n \cdot \mathbb{E}_{\phi\sim F^*}\left[\phi - \tilde{p}(\phi)\right],$$

where $\tilde{\mathcal{S}}$ denotes the single-round direct mechanisms for the single-good case and the IC/IR constraints are for the single-good case. This is equal to $\frac{n}{e}$ from the single-round minimax regret determined for the single-good case in Section 4.1 of the main paper.

To complete, we note $S^{*,n}$ is an optimal solution to (TR-6) and, hence, the original optimization problem because it is a feasible solution and achieves the lower bound $\frac{n}{e}$. More specifically, $\widehat{\text{Regret}}(S^{*,n}, F^{*,n})$ is equal to

$$n \cdot \mathbb{E}_{\phi\sim F^*}\left[\phi - p^*(\phi)\right] = n \cdot \widetilde{\text{Regret}} = \frac{n}{e}\,,$$

where $\widetilde{\text{Regret}}$ is the single-good single-round minimax regret (so, it is equal to $\widehat{\text{Regret}} = \frac{1}{e}$ in Section 4.1 of the main paper). □

TR.5 Resource Allocation without Monetary Transfers

In this section, we consider a dynamic resource allocation problem without monetary transfers. In particular, a social planner is repeatedly allocating a costly resource in settings where monetary transfers may not be practical for legal, ethical and various other reasons. Real-life applications include an organization or government allocating an internal resource and a nurse attending a patient. This problem has been recently studied by Guo and Hörner (2015), Balseiro et al. (2019), and Gorokh et al. (2019). These papers assume that the principal has access to samples or knows the agents' distributions. In comparison, here we consider a setting in which the principal does not know the agent's distribution. We show that the minimax regret of the multi-round problem is linear in T and an optimal solution simply repeats a single-round mechanism.

Our model is closely related to Guo and Hörner (2015). The principal repeatedly allocates an independent and identical unit of a resource in each round over the time horizon. In each round, the agent privately observes his current value for the resource which is drawn independently and identically from an underlying distribution known to him. The principal does not know the agent's distribution nor per-round values but knows that the per-round values are in the range $[0, 1]$. The outcome is a singleton $\omega = \hat{x}$ where $\hat{x}$ is the allocation, i.e., whether or not the resource is allocated to the agent. The principal incurs an opportunity cost $c \in (0, 1)$ when allocating the resource.[1] Equivalently, the principal wants to allocate the resource only when the value exceeds the cost, but the agent wants to be allocated always.

In the formal language of Section 2 of the main paper, the agent's shock is his value for the resource and $\Theta = [0, 1]$. The agent's distribution $F \in \mathcal{F}$ can be any distribution over Θ, i.e., $\mathcal{F} = \Delta(\Theta)$. The outcome space is $\Omega = \{0, 1\}$ and an outcome $\omega = \hat{x} \in \Omega$ is the allocation $\hat{x}$. When the outcome is $\omega = \hat{x}$ in a round, the agent's utility function is $v(\theta, \hat{x}) = \theta \cdot \hat{x}$ and the principal's utility function is $u(\theta, \hat{x}) = (\theta - c) \cdot \hat{x}$ where c is the fixed opportunity cost. For notational convenience, we represent a decision rule in terms of the corresponding allocation rule for a single-round direct mechanism and represent the interim allocation rule with $x : \Theta \to [0, 1]$, with the understanding that when the agent reports θ, the probability of allocation is $x(\theta)$.

Since $\mathcal{F} = \Delta(\Theta)$, Assumption 1 of the main paper holds. Since the principal would want to allocate

[1] Instead of a fixed opportunity cost c, we can alternatively think the cost per allocation to be a random Bernoulli variable with the average of c and the principal only observes the cost after an allocation decision and reasons in terms of the average cost c. Furthermore, we focus on the cases where $c \in (0, 1)$ because there exists a trivial optimal solution when $c = 0$ or $c = 1$. If $c = 0$, always allocating is optimal. If $c = 1$, not allocating is optimal.

only if $\theta \geq c$ when the agent's shock is θ, $\text{OPT}(\delta_\theta, 1) = \max\{\theta - c, 0\}$ for any $\theta \in \Theta$. The known-distribution benchmark for distribution F can be bounded as follows:

$$\text{OPT}(F, T) = \sup_{A \in \mathcal{A}} \mathbb{E}_{\pi,\sigma}\left[\sum_{t=1}^{T} (\theta_t - c)\hat{x}_t\right] \leq \sum_{t=1}^{T} \mathbb{E}_{\theta_t}[\max\{\theta_t - c, 0\}] = T\mathbb{E}_{\theta\sim F}\left[\text{OPT}(\delta_\theta, 1)\right] ,$$

where the second step follows from relaxing the incentive compatibility constraint and noting $\hat{x}_t \in [0, 1]$, and the last step follows because shocks are identically distributed. The last expression is equal to $\mathbb{E}_{\theta\sim F}\left[\text{OPT}(\delta_\theta, T)\right]$ by Proposition 1 of the main paper and, hence, Assumption 2 of the main paper holds. Our general results from Section 3 of the main paper apply and we obtain the following:

Proposition TR.4. *For the dynamic resource allocation problem without monetary transfers, the minimax regret of the multi-round problem is $c(1-c)T$ and an optimal solution is T repetitions of the probabilistic allocation rule x^* where $x^*(\theta) = 1 - c$ for $\theta \in (0, 1]$ and $x^*(0)$ can be any probability in the range $[0, 1-c]$.*

We can show that the probability allocation rule x^* is an optimal solution to the single-round problem, but not by finding a saddle point. Interestingly, saddle points do not exist for the corresponding single-round minimax regret problem, in contrast to other applications considered in the main paper and the previous section of this technical report. We instead show an asymmetric saddle-point result. See the proof and details in Appendix A.

TR.6 Extensions

In this section, we extend our results in the main paper in several directions. First, we show our results still hold for other alternative benchmarks that are considered in the learning literature. Second, we consider serially correlated shock processes and show our results still apply. Third, we consider multiplicative performance guarantees and prove analogous results connecting the multi-round and single-round problems. Fourth, we explore a stronger notion of regret in which the agent plays a utility-maximizing strategy that is the least favorable for the principal.

TR.6.1 Two Alternative Benchmarks

Instead of the optimal performance achievable $\mathrm{OPT}(F,T)$, we consider two different alternative benchmarks and show our results in the main paper still hold for the minimax regret defined with respect to these benchmarks. The first one is $T\cdot\bar{u}(F)$ (with $\bar{u}(F)$ as defined in Section 5.2 of the main paper) which can be thought of as a stronger benchmark than $\mathrm{OPT}(F,T)$ since $T\cdot\bar{u}(F)\geq\mathrm{OPT}(F,T)$ by Proposition E.1 of the main paper. It is equivalently the first-best performance that the principal can achieve in the full-information version of the multi-round problem. The second one is $T\cdot\tilde{u}(F)$ where $\tilde{u}(F)$ is the performance of an optimal single-round incentive compatible and individually rational mechanism in the single-round problem. This benchmark has been studied previously in the literature and can be thought of as a weaker benchmark as we will show $T\cdot\tilde{u}(F)\leq\mathrm{OPT}(F,T)$. Interestingly, the same general results from Section 3 of the main paper hold for these benchmarks under a stronger assumption than Assumption 2 of the main paper. The main observation is that the alternative benchmarks and $\mathrm{OPT}(F,T)$ coincide for point-mass distributions which form the worst cases in so far as determining the minimax regret.

For the agent's distribution F, we define

$$\tilde{u}(F) := \sup_{\substack{S\in\Delta(\Omega)^{\Theta}:\\ \text{(IC),(IR)}}} \int_{\Theta}\int_{\Omega} u(\theta,\omega)dS_{\theta}(\omega)dF(\theta)\,,$$

which can be thought of as the "second-best" benchmark in the single-round problem when F is known to the principal. We have $\tilde{u}(F)\leq\bar{u}(F)$, because we can relax the IC constraint in the definition of $\tilde{u}(F)$ and solve for the best outcome distribution for each shock separately. This leads to an upper bound of $\mathbb{E}_{\theta\sim F}[\bar{u}(\delta_\theta)]$ which is at most $\bar{u}(F)$, by Proposition E.1 of the main paper. The multi-round benchmark $T\cdot\tilde{u}(F)$ is a weaker benchmark in the sense that $T\cdot\tilde{u}(F)\leq\mathrm{OPT}(F,T)$, because the principal can repeat the single-round solution (or approximately optimal) to the optimization problem defining $\tilde{u}(F)$ and realize the performance of $T\cdot\tilde{u}(F)$ since a utility-maximizing strategy for the agent is to participate and then truthfully report in each round. In the case of revenue maximization in the dynamic selling mechanism design problem with one good, $\bar{u}(F)$ evaluates to the average shock, equivalently, the full surplus of the agent, and $\tilde{u}(F)$ evaluates to the optimal revenue achievable, say, by posting the Myerson's price, i.e., $\max_{x\geq 0} x\cdot\Pr_{\theta\sim F}(\theta\geq x)$ (Myerson, 1981).

To distinguish the regret notions with respect to different benchmarks for an incentive compatible mechanism A, we use $\text{Regret}^{\mathsf{OPT}}$ to denote

$$\text{Regret}^{\mathsf{OPT}}(A, F, T) := \text{OPT}(F, T) - \text{PrincipalUtility}(A, F, T)$$

which is the original regret notion as defined in Section 2 of the main paper and $\text{Regret}^{\mathsf{FB}}$ (FB for "first-best") and $\text{Regret}^{\mathsf{SB}}$ (SB for "second-best") to denote, respectively,

$$\begin{aligned}&\text{Regret}^{\mathsf{FB}}(A, F, T) := T \cdot \bar{u}(F) - \text{PrincipalUtility}(A, F, T)\,, \text{ and}\\&\text{Regret}^{\mathsf{SB}}(A, F, T) := T \cdot \tilde{u}(F) - \text{PrincipalUtility}(A, F, T)\,.\end{aligned}$$

We distinguish $\text{Regret}^{\mathsf{OPT}}(T)$, $\text{Regret}^{\mathsf{FB}}(T)$ and $\text{Regret}^{\mathsf{SB}}(T)$ similarly.

Of the three variants, the $\text{Regret}^{\mathsf{SB}}(T)$ notion is closely related to the standard regret notion in the learning literature that considers the best fixed "action" in hindsight, which naturally corresponds to the best fixed single-round mechanism in our setting, which is repeated across the time horizon. Since a dynamic mechanism can potentially do better, this regret can be negative sometimes. In particular, Amin et al. (2013) and subsequent works studied the $\text{Regret}^{\mathsf{SB}}(T)$ notion (what they call "strategic regret") for the restricted class of dynamic posted pricing strategies for the problem of repeatedly selling a single good.

As the following theorem shows, we obtain identical results as Theorem 1 of the main paper with respect to above alternative benchmarks with the same minimax regret and structural characterization of an optimal dynamic mechanism. We still have the same single-round minimax regret problem as defined in Section 2 of the main paper. Instead of Assumption 2 of the main paper, we assume a stronger assumption in terms of $\bar{u}(F)$ for the stronger benchmark $T \cdot \bar{u}(F)$; this is stronger by Proposition 7 of the main paper. See Appendix B.1 for its proof and further details.

Theorem TR.3. *Suppose Assumption 1 of the main paper holds and $\bar{u}(F) = \mathbb{E}_{\theta \sim F}[\bar{u}(\delta_\theta)]$ for all $F \in \mathcal{F}$. Using $\text{Regret}^{(\cdot)}$ to denote both $\text{Regret}^{\mathsf{FB}}$ and $\text{Regret}^{\mathsf{SB}}$ notions, the following statements hold with respect to both notions: $\text{Regret}^{(\cdot)}(T) = T \cdot \inf_{S \in \mathcal{S}^{\times 1}} \sup_{\theta \in \Theta} \text{Regret}^{\mathsf{OPT}}(S, \delta_\theta, 1)$. For any $\epsilon \geq 0$, if a single-round direct mechanism $S \in \mathcal{S}^{\times 1}$ satisfies*

$$\sup_{\theta \in \Theta} \text{Regret}^{\mathsf{OPT}}(S, \delta_\theta, 1) \leq \inf_{S' \in \mathcal{S}^{\times 1}} \sup_{\theta \in \Theta} \text{Regret}^{\mathsf{OPT}}(S', \delta_\theta, 1) + \frac{\epsilon}{T}\,,$$

then,

$$\sup_{F \in \mathcal{F}} \text{Regret}^{(\cdot)}(S^{\times T}, F, T) \leq \text{Regret}^{(\cdot)}(T) + \epsilon \,.$$

In addition, $\arg\min_{A \in \mathcal{A}} \sup_{F \in \mathcal{F}} \text{Regret}^{(\cdot)}(A, F, T)$ *is empty if and only if* $\arg\min_{S \in \mathcal{S}^{\times 1}} \sup_{\theta \in \Theta} \text{Regret}^{\mathsf{OPT}}(S, \delta_\theta, 1)$ *is empty.*

TR.6.2 Arbitrary Shock Processes

We consider more general arbitrary shock processes under which the agent's private shocks may be correlated across rounds. For example, in the dynamic selling mechanism design problem, the agent's private value may be given by permanent and transitory components. His permanent component v_0 is drawn from a privately known distribution and transitory components ϵ_t are drawn from separate privately known distributions over the rounds such that his private shock in Round t is $\theta_t = v_0 + \epsilon_t$. For comparison, Carrasco et al. (2015) considered arbitrary shock processes for the maximin utility objective; we further discuss Carrasco et al. (2015) in Section TR.7. The repeated i.i.d. setting described in Section 2 of the main paper and considered throughout the paper is a special case where the shocks are independently and identically drawn from a fixed underlying distribution. In this section, we show that the same general results from Section 3 of the main paper still hold in the general shock process setting.

To distinguish the shock processes, we use superscript T. We use $\mathcal{F}^T$ to denote the set of possible T-round shock processes for the agent with the support of each per-round marginal distribution contained in Θ, i.e., $\mathcal{F}^T \subseteq \Delta(\Theta^T)$, and F^T to denote a particular T-round shock process. The minimax regret value of the multi-round problem is, using the same Regret notation,

$$\text{Regret}(T) = \inf_{A \in \mathcal{A}} \sup_{F^T \in \mathcal{F}^T} \text{Regret}(A, F^T, T) \,, \tag{TR-7}$$

where function Regret explicitly takes a T-round shock process. Generalizing $\text{OPT}(F, T)$, the multi-round benchmark is $\text{OPT}(F^T, T)$ which is the optimal performance achievable when the principal knows the agent's private shock process F^T:

$$\text{OPT}(F^T, T) = \sup_{A \in \mathcal{A}} \text{PrincipalUtility}(A, F^T, T) \,.$$

Hence, the regret for an incentive compatible dynamic mechanism $A \in \mathcal{A}$ when the agent's T-round shock process is F^T is the difference between the optimal performance achievable and the actual performance achieved:

$$\mathrm{Regret}(A, F^T, T) = \mathrm{OPT}(F^T, T) - \mathrm{PrincipalUtility}(A, F^T, T)\,.$$

As in the repeated i.i.d. setting, the agent plays the utility-maximizing strategy that is recommended under A. The corresponding single-round problem in the general shock process setting is still the same problem as in the repeated i.i.d. setting (in Section 2 of the main paper).

We use the same notations from the main paper to refer to the repeated i.i.d. setting such that F without superscript T denotes a single-round distribution and $\mathrm{Regret}(A, F, T)$ is the regret of an incentive compatible dynamic mechanism A when the agent's private distribution is F in the repeated i.i.d. setting. We use both $F^{\times T}$ and $(F)^{\times T}$ to denote the shock process in which the per-round shocks are drawn i.i.d. from F in the general shock process setting. Whether the Regret notation refers to the repeated i.i.d. setting or the general shock process setting will be clear from the context and the parameters; in particular, $\mathrm{Regret}(A, F, T) = \mathrm{Regret}(A, F^{\times T}, T)$ for $A \in \mathcal{A}$ and $F \in \Delta(\Theta)$.

We have the following result:

Theorem TR.4. *Suppose that $\delta_\theta^{\times T} \in \mathcal{F}^T$ for all $\theta \in \Theta$ and that $\mathrm{OPT}(F^T, T) \le \mathbb{E}_{t\sim[T],\theta\sim(F^T)_t}[\mathrm{OPT}(\delta_\theta^{\times T}, T)]$ for all $F^T \in \mathcal{F}^T$, where $t \sim [T]$ means we draw a round uniformly at random and $\theta \sim (F^T)_t$ means we draw a shock from the marginal shock distribution of F^T in Round t. Then, the following statements hold in the general shock process setting: $\mathrm{Regret}(T) = T \cdot \inf_{S\in\mathcal{S}^{\times 1}} \sup_{\theta\in\Theta} \mathrm{Regret}(S, \delta_\theta, 1)$. For any $\epsilon \ge 0$, if a single-round direct mechanism $S \in \mathcal{S}^{\times 1}$ satisfies*

$$\sup_{\theta\in\Theta} \mathrm{Regret}(S, \delta_\theta, 1) \le \inf_{S'\in\mathcal{S}^{\times 1}} \sup_{\theta\in\Theta} \mathrm{Regret}(S', \delta_\theta, 1) + \frac{\epsilon}{T}\,,$$

then,

$$\sup_{F^T\in\mathcal{F}^T} \mathrm{Regret}(S^{\times T}, F^T, T) \le \mathrm{Regret}(T) + \epsilon\,.$$

In addition, $\arg\min_{A\in\mathcal{A}} \sup_{F^T\in\mathcal{F}^T} \mathrm{Regret}(A, F^T, T)$ is empty if and only if $\arg\min_{S\in\mathcal{S}^{\times 1}} \sup_{\theta\in\Theta} \mathrm{Regret}(S, \delta_\theta, 1)$ is empty.

Instead of Assumptions 1 and 2 of the main paper, we have analogous assumptions in terms of $\mathrm{OPT}(F^T, T)$. The right-hand side of the second assumption can be equivalently written as $\sum_{t=1}^{T} \mathbb{E}_{\theta\sim(F^T)_t}[\mathrm{OPT}(\delta_\theta, 1)]$ by Proposition 1 of the main paper. We defer the proof of Theorem TR.4 and further details to Appendix B.2. In the same appendix, we also prove analogous results to Propositions 7, 8 and E.1 of the main paper in terms of a generalization of $\bar{u}(F)$ in the general shock process setting. Therefore, Theorem TR.4 holds for all games with payments that enter linearly into the utility functions of the principal and agent or with a nonnegative utility function for the agent. In particular, it holds for all applications considered in Section 4 of the main paper and in Sections TR.4 and TR.5. Similar to the repeated i.i.d. setting, the class of constant shock processes, which are equivalently point-mass distributions in the repeated i.i.d. setting, form the worst cases in the general shock process setting; we refer to Proposition 2 of the main paper and note the same proof still works.

TR.6.3 Multiplicative Guarantees

We show that our results and analyses of the main paper extend to analogous results for a multiplicative performance metric. In particular, the best multiplicative guarantee for the multi-round problem is equal to the best multiplicative guarantee for the corresponding single-round problem and the principal can achieve the best multiplicative guarantee arbitrarily closely by repeating a single-round mechanism and exactly by repeating, if it exists, an optimal single-round mechanism to the single-round problem.

The minimax regret and regret are standard performance metrics in the sequential learning literature. If a learning algorithm has the minimax regret of $o(T)$, the worst-case regret when averaged over the rounds diminishes towards 0 as the time horizon T increases and the algorithm achieves the optimal performance (or some suitable benchmark) asymptotically. Alternatively, we can consider multiplicative guarantees in terms of a ratio, similar to the approximation and competitive ratios that are common in the theoretical computer science literature. This is a reasonable performance metric that is scale-free and may be more interpretable than a difference in absolute terms; at least 50% of the benchmark versus at most \$50 less than the benchmark.

Suppose there exist some constants $0 < L < U < \infty$ such that $\mathrm{OPT}(\delta_\theta, 1) \in [L, U]$ for all shocks $\theta \in \Theta$ and $\mathrm{OPT}(F, T) \in [LT, UT]$ for all distributions $F \in \mathcal{F}$, i.e., bounded from above and away from 0. We define the multi-round multiplicative guarantee $\mathrm{Ratio}(T)$ and the corresponding multi-round problem as

$$\mathrm{Ratio}(T) := \sup_{A \in \mathcal{A}} \inf_{F \in \mathcal{F}} \mathrm{Ratio}(A, F, T)\,, \tag{TR-8}$$

where the multiplicative ratio is defined as $\mathrm{Ratio}(A, F, T) := \frac{\mathrm{PrincipalUtility}(A,F,T)}{\mathrm{OPT}(F,T)}$ for any incentive compatible mechanism A and distribution F. As for the regret objective, we assume the agent follows the recommended strategy that is given as part of the incentive compatible mechanism A.

Note $\text{Ratio}(T) \leq 1$ since the realized performance of an incentive compatible dynamic mechanism is upper bounded by the optimal performance achievable with the knowledge of the agent's distribution. It is also at least 0 because the principal can guarantee the total utility of 0 via the trivial mechanism that always forces the no-interaction outcome.

Similarly, we define the corresponding single-round problem for direct IC/IR mechanisms as

$$\sup_{S \in \mathcal{S}^{\times 1}} \inf_{\theta \in \Theta} \text{Ratio}(S, \delta_\theta, 1)\,, \tag{TR-9}$$

with the value (i.e., the single-round multiplicative guarantee) in the interval $[0, 1]$. The upper bound of 1 and lower bound of 0 follow by the same argument as above.

Similar to the general minimax regret result, Theorem 1 of the main paper, we can show the following. We defer its proof to Appendix B.3.

Theorem TR.5. *Suppose Assumptions 1 and 2 of the main paper hold. Suppose there exist some constants* $0 < L < U < \infty$ *such that* $\text{OPT}(\delta_\theta, 1) \in [L, U]$ *for all* $\theta \in \Theta$ *and* $\text{OPT}(F, T) \in [LT, UT]$ *for all* $F \in \mathcal{F}$*. Then,* $\text{Ratio}(T) = \sup_{S \in \mathcal{S}^{\times 1}} \inf_{\theta \in \Theta} \text{Ratio}(S, \delta_\theta, 1)$*. Moreover, for any* $\epsilon \geq 0$*, if a mechanism* $S \in \mathcal{S}^{\times 1}$ *satisfies*

$$\inf_{\theta \in \Theta} \text{Ratio}(S, \delta_\theta, 1) \geq \sup_{S' \in \mathcal{S}^{\times 1}} \inf_{\theta \in \Theta} \text{Ratio}(S', \delta_\theta, 1) - \frac{\epsilon L}{U}\,,$$

then,

$$\inf_{F \in \mathcal{F}} \text{Ratio}(S^{\times T}, F, T) \geq \text{Ratio}(T) - \epsilon\,.$$

Finally, $\arg\max_{A \in \mathcal{A}} \inf_{F \in \mathcal{F}} \text{Ratio}(A, F, T)$ *is empty if and only if* $\arg\max_{S \in \mathcal{S}^{\times 1}} \inf_{\theta \in \Theta} \text{Ratio}(S, \delta_\theta, 1)$ *is empty.*

We remark that for the same game, the minimax regret and multiplicative guarantee are, in general, different. For example, in the case of revenue maximization in the dynamic selling problem with one good (Section 4.1 of the main paper), the minimax regret is $\frac{T}{e}$ while the multiplicative guarantee is unbounded whenever zero is the lowest value for the shock.

TR.6.4 Principal Pessimism

We consider a stronger notion of minimax regret under which the agent can choose to play any utility-maximizing strategy. We analyze in terms of what we call the principal pessimism constraint and show most of our results of the main paper continue to hold with respect to the stronger notion of minimax regret. Throughout the main paper, we assumed the agent plays the recommended strategy that is given as part of the principal's incentive compatible mechanism. Essentially, the agent plays a best-response strategy chosen in the principal's favor and the uncertainty that the principal faces is in terms of the agent's distribution only. Alternatively, we can allow for a different kind of "tie-breaking" possibility in which the agent plays any utility-maximizing strategy and, in particular, one that minimizes the principal utility among such utility-maximizing strategies. This leads to a stronger and more robust notion of minimax regret under which guarantees hold for any shock distribution and any best-response strategy for the agent.

Going back to the main paper where an optimal dynamic mechanism repeats a single-round direct IC/IR mechanism, truthful reporting is one utility-maximizing strategy and the regret guarantees hold for this strategy and, also, for any utility-maximizing strategy chosen in the principal's favor. As far as the regret notion is concerned, if there are multiple utility-maximizing strategies for the agent, the agent plays one that also maximizes the principal utility. But it is possible that different utility-maximizing strategies lead to different total expected utilities for the principal and, hence, that the regret is greater for other utility-maximizing strategies.[2]

Departing from the main paper's exposition, we drop the recommended strategy from dynamic mechanisms' specification and use $\mathcal{A}$ to denote the set of dynamic mechanisms given by tuple $(\mathcal{M}, \pi)$. For a dynamic mechanism $A = (\mathcal{M}, \pi)$, let $\mathcal{B}(A, T)$ be the set of all utility-maximizing strategies for the agent in the sense that if $\sigma \in \mathcal{B}(A, T)$, then $\text{AgentUtility}(A, \sigma, F, T) \geq \text{AgentUtility}(A, \tilde{\sigma}, F, T)$ holds for every probability distribution F over Θ and every feasible agent strategy $\tilde{\sigma}$. The minimax regret of the main paper that corresponds to the agent playing a utility-maximizing strategy that

[2]Different selections of the agent's utility-maximizing strategies lead to different principal utilities and, hence, regrets. Consider a principal selling a single good to an agent where the agent's value for the good is $\theta = \frac{1}{2}$. Assume the principal's mechanism is as follows: given the agent's report $\hat{\theta}$, 1) charge $\frac{1}{2}$ and allocate with probability 1 if $\hat{\theta} > \frac{1}{2}$; 2) charge $\frac{1}{4}$ and allocate with probability $\frac{1}{2}$ if $\hat{\theta} = \frac{1}{2}$; and 3) do not allocate if $\hat{\theta} < \frac{1}{2}$. This mechanism satisfies the IC/IR constraints. For the agent, truthful reporting is a utility-maximizing strategy and leads to the agent utility of 0 and the principal utility of $\frac{1}{4}$. But other reporting strategies are also utility-maximizing and can lead to a different principal utility. Reporting some $\hat{\theta} < \frac{1}{2}$ leads to a smaller principal utility of 0 and reporting some $\hat{\theta} > \frac{1}{2}$ leads to a greater principal utility of $\frac{1}{2}$, while both lead to the same agent utility of 0 as truthful reporting.

also maximizes the principal utility can be equivalently written as

$$\mathrm{Regret}(T) = \inf_{A\in\mathcal{A}} \inf_{\tilde{\sigma}\in\mathcal{B}(A,T)} \sup_{F\in\mathcal{F}} \mathrm{Regret}(A, \tilde{\sigma}, F, T)\,,$$

where $\mathrm{Regret}(A, \tilde{\sigma}, F, T) = \mathrm{OPT}(F, T) - \mathrm{PrincipalUtility}(A, \tilde{\sigma}, F, T)$. For the stronger minimax regret notion, the inner infimum becomes a supremum as in

$$\inf_{A\in\mathcal{A}} \sup_{\tilde{\sigma}\in\mathcal{B}(A,T)} \sup_{F\in\mathcal{F}} \mathrm{Regret}(A, \tilde{\sigma}, F, T)\,,$$

and the uncertainty that the principal faces is both in terms of the agent's distribution and best-response strategy in the worst-case sense.

In what follows, for any principal's mechanism A and time horizon T, let $\sigma^*(A, T)$ be a utility-maximizing strategy that, if multiple ones exist, minimizes the principal utility among such utility-maximizing strategies for any distribution $F \in \Delta(\Theta)$. That is, $\mathrm{PrincipalUtility}(A, \sigma^*(A, T), F, T) \leq \mathrm{PrincipalUtility}(A, \tilde{\sigma}, F, T)$ holds for every $F \in \Delta(\Theta)$ and $\tilde{\sigma} \in \mathcal{B}(A, T)$. We refer to such strategy as a *principal-pessimistic utility-maximizing strategy*. The multi-round problem with respect to the stronger regret notion is equivalently

$$\inf_{A\in\mathcal{A}} \sup_{F\in\mathcal{F}} \mathrm{Regret}(A, \sigma^*(A, T), F, T)\,.$$

When clear from the context, we suppress the principal-pessimistic utility-maximizing strategy and use the same notations as introduced in Section 2 of the main paper for the stronger notions of regret and minimax regret. Note we keep the same optimal performance achievable $\mathrm{OPT}(F, T)$ from the main paper. This is what the principal can achieve when he knows the agent's distribution F and the agent plays a utility-maximizing strategy in the principal's favor; this is the best-case scenario from the principal's perspective.[3]

We use $\mathcal{S}^{\times T} \subset \mathcal{A}$ to denote repeated single-round direct mechanisms for which truthful reporting (i.e., reports $\mathsf{CONTINUE}$ in Round 0 and then truthfully reports shocks) is a principal-pessimistic utility-maximizing strategy. More specifically, these are IC/IR mechanisms because truthful reporting is a utility-maximizing strategy for the agent and they satisfy the additional constraint, which we

[3]Note the set $\mathcal{B}(A, T)$ and the strategy $\sigma^*(A, T)$ may not be well-defined. For ease of presentation, we assume these are well-defined for any dynamic mechanism A and time horizon T. Without the assumption, we can instead reason with sequences of "approximately" utility-maximizing strategies and define AgentUtility, PrincipalUtility and Regret using limit superior and inferior.

call the *principal pessimism constraint* (PP), that truthful reporting is a principal-pessimistic utility-maximizing strategy. In the same manner for IC/IR constraints, we say a direct IC/IR mechanism is *principal-pessimistic* (PP) if it satisfies the PP constraint. The mechanisms in $\mathcal{S}^{\times T}$ are uniquely identified by single-round direct mechanisms $\mathcal{S}^{\times 1}$ that they repeat and we let $S^{\times T}$ denote the mechanism that repeats S over T rounds for a single-round direct IC/IR/PP mechanism $S \in \mathcal{S}^{\times 1}$.

The corresponding single-round problem is

$$\inf_{S \in \mathcal{S}^{\times 1}} \sup_{\theta \in \Theta} \mathrm{Regret}(S, \sigma^*(S,1), \delta_\theta, 1)\,, \tag{TR-10}$$

where we take $\sigma^*(S,1)$ to be the truthful reporting strategy σ^{TR} without loss in terms of the agent utility, principal utility and regret. In the outcome distribution representation form, the analogue of (3) of the main paper is

$$\inf_{S \in \Delta(\Omega)^\Theta} \sup_{\theta \in \Theta} \left\{ \mathrm{OPT}(\delta_\theta, 1) - \int_\Omega u(\theta,\omega) dS_\theta(\omega) \right\} \quad \text{s.t.} \quad \text{(IC)}, \text{(IR)}, \text{(PP)}\,, \tag{TR-11}$$

with the same IC/IR constraints and where PP is formulated as:

$$\begin{aligned} &\int_\Omega u(\theta,\omega)\mathrm{d}S_\theta(\omega) \leq \int_\Omega u(\theta,\omega)\mathrm{d}S_{\theta'}(\omega)\,, \quad \forall \theta \in \Theta, \theta' \in \mathcal{B}^*(S,\theta) \\ &\int_\Omega u(\theta,\omega)\mathrm{d}S_\theta(\omega) \leq 0\,, \quad \forall \theta \in \Theta_0 \end{aligned} \tag{PP}$$

where $\mathcal{B}^*(S,\theta) = \{\theta' \in \Theta \mid \int_\Omega v(\theta,\omega)\mathrm{d}S_{\theta'}(\omega) = \int_\Omega v(\theta,\omega)\mathrm{d}S_\theta(\omega)\}$ is the set of utility-maximizing reports for the agent when his shock is θ given the IC constraint holds, and $\Theta_0 = \{\theta \in \Theta \mid \int_\Omega v(\theta,\omega)\mathrm{d}S_\theta(\omega) = 0\}$ is the set of shocks for which truthful reporting leads to the agent utility of 0. The first part of PP constraint, given the IC constraint holds, stipulates that truthful reporting leads to the lowest principal utility among all best responses for the agent. The second part stipulates that the principal utility is at most 0 if truthful reporting leads to the agent utility of 0. Note if S satisfies the IC/IR constraints, we have $\theta \in \mathcal{B}^*(S,\theta)$ for all $\theta \in \Theta$. The IC/IR/PP constraints together imply that truthful reporting is an interim principal-pessimistic utility-maximizing strategy.

Our main general result (Theorem 1 of the main paper) still holds with suitable changes in the proof (in choosing a utility-maximizing strategy) to show that the multi-round minimax regret is T times the single-round minimax regret and we can achieve this minimax regret arbitrarily closely by repeating a single-round mechanism. We provide the formal statement below and refer to Appendix B.4.1 for proof details.

Theorem TR.6. *Suppose Assumptions 1 and 2 of the main paper hold. The following statements hold with respect to the stronger notion of regret and minimax regret where the agent plays a principal-pessimistic utility-maximizing strategy:* $\text{Regret}(T) = T \cdot \inf_{S \in \mathcal{S}^{\times 1}} \sup_{\theta \in \Theta} \text{Regret}(S, \delta_\theta, 1)$. *For any* $\epsilon \geq 0$*, if a single-round direct IC/IR/PP mechanism* $S \in \mathcal{S}^{\times 1}$ *satisfies*

$$\sup_{\theta \in \Theta} \text{Regret}(S, \delta_\theta, 1) \leq \inf_{S' \in \mathcal{S}^{\times 1}} \sup_{\theta \in \Theta} \text{Regret}(S', \delta_\theta, 1) + \frac{\epsilon}{T},$$

then,

$$\sup_{F \in \mathcal{F}} \text{Regret}(S^{\times T}, F, T) \leq \text{Regret}(T) + \epsilon .$$

Finally, $\arg\min_{A \in \mathcal{A}} \sup_{F \in \mathcal{F}} \text{Regret}(A, F, T)$ *is empty if and only if* $\arg\min_{S \in \mathcal{S}^{\times 1}} \sup_{\theta \in \Theta} \text{Regret}(S, \delta_\theta, 1)$ *is empty.*

For the single-round problem, we also prove the analogue of Lemma 1 of the main paper in Appendix B.4.2. By the analogue, we can equivalently solve the following version for single-round direct IC/IR/PP mechanisms:

$$\inf_{\substack{S \in \Delta(\Omega)^\Theta: \\ \text{(IC),(IR),(PP)}}} \sup_{F \in \Delta(\Theta)} \widehat{\text{Regret}}(S, F),$$

where $\widehat{\text{Regret}}(S, F) := \int_\Theta \text{OPT}(\delta_\theta, 1) dF(\theta) - \int_\Theta \int_\Omega u(\theta, \omega) dS_\theta(\omega) dF(\theta)$ for a single-round direct mechanism S and distribution F.

One way to find an optimal solution to the above single-round problem is to solve the relaxed version without the PP constraint and show the solution to the relaxed version satisfies the PP constraint and, thus, is also a solution to the original version. Note we already have optimal solutions to the relaxed versions for the applications in Section 4 of the main paper and Sections TR.4–TR.5. We can characterize the PP constraint in these applications and show the optimal single-round direct IC/IR mechanisms found already satisfy the PP constraint. For instance, for the dynamic selling mechanism design problem with a single good, the if-and-only-if condition for the PP constraint is roughly that any flat part of the interim allocation rule, if exists, is closed on the right side. The optimal single-round direct IC/IR mechanism in Proposition 3 of the main paper has an interim allocation rule that is continuous and indeed satisfies the PP constraint. Consequently, Propositions 3 and 4 of the main paper and Propositions TR.1, TR.2, and TR.4 (when we let $x^*(0) = 1 - c$) of this technical report still hold the same with respect to the stronger notion of minimax regret. We defer further details to Appendix B.4.3.

In the most general setting, it is not clear how to mathematically formalize the PP constraint in a tractable way and we do not have an explicit characterization. We believe other results still hold with respect to the stronger notions of regret and minimax regret.

TR.7 Connections to Maximin Utility Objective

We show equivalence-type connections between the minimax regret and maximin utility objectives for revenue maximization in the single-good case of the dynamic selling mechanism design problem (considered in Section 4.1 of the main paper). Similar connections hold more generally for other robust mechanism design problems as long as a saddle-point result exists for the minimax regret objective and a corresponding saddle-point result exists for the maximin utility objective (with the known mean). Carrasco et al. (2015) recently showed a similar false-dynamics result for the multi-round dynamic selling problem with respect to the maximin utility objective where the principal maximizes the minimum utility achieved in the worst-case sense and only the mean of the agent's shock distribution is known a priori. More specifically, we show the minimax regret criterion and maximin utility criterion with a known mean share the same single-round saddle-point problem involving direct IC/IR mechanisms. For the principal-agent model with hidden costs, we are not aware of a multi-round false-dynamics result but our analysis of the single-round minimax regret problem is, again, similar to Carrasco et al. (2018) that considered a single-round utility-maximization problem subject to a nonlinear quality cost function with respect to the other robust objective.

In contrast to the minimax regret objective, the maximin utility objective without any additional distributional information (other than the shock space Θ) leads to a trivial answer in the case of revenue maximization in the dynamic selling mechanism design problem. The worst case is when the agent has the lowest possible value and the principal will accordingly price the good at this value always. If the lowest possible value is 0, then the maximin utility would be 0. The maximin utility objective leads to more meaningful solutions with additional fixed-moment-type distributional information. In particular, Carrasco et al. (2015) considered a dynamic setting where the agent's values within $[0, 1]$ follow an unknown arbitrary value process over T rounds with the property that each per-round marginal distribution has the same mean known a priori to the principal. They showed a saddle-point result where the optimal dynamic mechanism is T repetitions of a single-round mechanism and the worst-case value process is one where a value is drawn once from a specified distribution and is fixed across all rounds. Interestingly, the worst-case value process is such that T

repetitions of a single-round mechanism is still optimal even when the value process is known to the principal.

In Section 4.1 of the main paper, we considered the same multi-round problem with respect to the minimax regret objective without any additional distributional information (other than the shock space Θ) and value processes were such that per-round value distributions are identical and independent. Our version of the multi-round problem does not allow the above worst-case value process in which values are not drawn independently from the same marginal shock distribution. Furthermore, there is no saddle-point result for the multi-round problem in our setting.

Despite these differences, our results and Carrasco et al. (2015) have similar analyses and solution structures. This is because both papers rely on essentially the same single-round saddle-point problem involving direct IC/IR mechanisms. Strictly speaking, the single-round problem with the maximin utility objective and a known mean is "finer-grained" in the sense that we can use the solutions to this problem to find a solution to the single-round problem with the minimax regret objective. This is assuming we have saddle-point results with respect to both objectives which were indeed proved, independently, in Carrasco et al. (2015) and in Section 4.1 of the main paper. The saddle-point results allow us to change the order of the infimum and supremum in the single-round problems.

In the remainder, we describe the connection via single-round problems. Recall $\mathcal{F} = \Delta(\Theta)$ in the dynamic selling mechanism design problem. We start with the saddle-point result for the single-round problem for direct IC/IR mechanisms with respect to the minimax regret objective using the formulation given in Section TR.1. Recall the single-round benchmark $\mathbb{E}_{\theta\sim F}[\mathrm{OPT}(\delta_\theta, 1)]$ is equal to $\mathbb{E}_{\theta\sim F}[\theta]$ from Section 4.1 of the main paper. Denoting by $\hat{\mathcal{S}} = \{S \in \Delta(\Omega)^\Theta : (\mathrm{IC}), (\mathrm{IR})\}$ the set of single-round direct IC/IR mechanisms, we have

$$\begin{aligned}\widehat{\mathrm{Regret}} &= \inf_{S\in\hat{\mathcal{S}}} \sup_{F\in\Delta(\Theta)} \widehat{\mathrm{Regret}}(S,F) = \sup_{F\in\Delta(\Theta)} \inf_{S\in\hat{\mathcal{S}}} \widehat{\mathrm{Regret}}(S,F) \\ &= \sup_{F\in\Delta(\Theta)} \inf_{S\in\hat{\mathcal{S}}} \left\{\mathbb{E}_{\theta\sim F}[\theta] - \mathrm{PrincipalUtility}(S, \sigma^{\mathrm{TR}}, F, 1)\right\} ,\end{aligned}$$

where σ^{TR} denotes the truthful reporting strategy. Let $\mathcal{F}_\mu = \{F \in \Delta(\Theta) \mid \mathbb{E}_{\theta\sim F}[\theta] = \mu\}$, i.e., distributions with the mean equal to μ. The optimization problem can be equivalently written as

$$\begin{aligned}
\widehat{\text{Regret}} &= \sup_{\mu\in[0,1]} \sup_{F\in\mathcal{F}_\mu} \inf_{S\in\hat{\mathcal{S}}} \{\mu - \text{PrincipalUtility}(S, \sigma^{\text{TR}}, F, 1)\} \\
&= \sup_{\mu\in[0,1]} \left\{\mu - \inf_{F\in\mathcal{F}_\mu} \sup_{S\in\hat{\mathcal{S}}} \text{PrincipalUtility}(S, \sigma^{\text{TR}}, F, 1)\right\} \\
&= \sup_{\mu\in[0,1]} \left\{\mu - \sup_{S\in\hat{\mathcal{S}}} \inf_{F\in\mathcal{F}_\mu} \text{PrincipalUtility}(S, \sigma^{\text{TR}}, F, 1)\right\} ,
\end{aligned}$$

where the second equation follows by extracting the mean μ from the objective and accounting for the negative sign in front of the principal utility expression, and the last by the saddle-point result for the maximin utility objective with a known mean from Carrasco et al. (2015). Now, assume we have saddle-point solutions (S_μ^*, F_μ^*) for this problem for all possible $\mu \in [0,1]$. We arrive at a saddle-point solution for the single-round problem with respect to the minimax regret objective by choosing (S_μ^*, F_μ^*) for μ that maximizes the quantity

$$\mu - \text{PrincipalUtility}(S_\mu^*, \sigma^{\text{TR}}, F_\mu^*, 1) .$$

Note the connection holds when the saddle-point results exist with respect to both objectives, which we indeed have.

References

Charalambos D Aliprantis and Kim Border. *Infinite dimensional analysis: a hitchhiker's guide.* Springer Science & Business Media, 2006.

Kareem Amin, Afshin Rostamizadeh, and Umar Syed. Learning prices for repeated auctions with strategic buyers. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 26*, pages 1169–1177. Curran Associates, Inc., 2013.

Santiago R. Balseiro, Huseyin Gurkan, and Peng Sun. Multiagent mechanism design without money. *Operations Research*, 67(5):1417–1436, 2019.

David P. Baron and David Besanko. Regulation and information in a continuing relationship. *Information Economics and Policy*, 1(3):267 – 302, 1984. ISSN 0167-6245.

Vinicius Carrasco, Vitor Farinha Luz, Paulo Monteiro, and Humberto Moreira. Robust selling mechanisms. Textos para discussão 641, Pontifícia Universidade Católica do Rio de Janeiro (PUC-Rio), Departamento de Economia, Rio de Janeiro, 2015. URL `http://hdl.handle.net/10419/176124`.

Vinicius Carrasco, Vitor Farinha Luz, Paulo K. Monteiro, and Humberto Moreira. Robust mechanisms: the curvature case. *Economic Theory*, April 2018.

Ky Fan. Minimax theorems. *Proceedings of the National Academy of Sciences of the United States of America*, 39(1):42–47, 1953.

Artur Gorokh, Siddhartha Banerjee, and Krishnamurthy Iyer. From monetary to non-monetary mechanism design via artificial currencies. *Mathematics of Operations Research (forthcoming)*, 2019.

Yingni Guo and Johannes Hörner. Dynamic allocation without money. Available at SSRN: https://ssrn.com/abstract=2563005, February 2015.

Cagil Kocyigit, Napat Rujeerapaiboon, and Daniel Kuhn. Robust multidimensional pricing: Separation without regret. Available at SSRN: https://ssrn.com/abstract=3219680, July 2018.

Roger B. Myerson. Optimal auction design. *Math. Oper. Res.*, 6(1):58–73, February 1981. ISSN 0364-765X. doi: 10.1287/moor.6.1.58. URL `http://dx.doi.org/10.1287/moor.6.1.58`.

A. C. C. Yao. Probabilistic computations: Toward a unified measure of complexity. *Foundations of Computer Science (FOCS), 18th IEEE Symposium on*, pages 222–227, 1977.

A Missing Proofs from Section TR.5

We prove Proposition TR.4. By Theorem 1 of the main paper, it suffices to show the following:

Proposition TR.5. *Let S^* be the single-round direct IC/IR mechanism with the allocation rule given in Proposition TR.4. Then, $\widehat{\text{Regret}} = c \cdot (1-c)$ and S^* is an optimal solution to the single-round minimax problem, i.e.,*

$$\widehat{\text{Regret}} = \sup_{F \in \Delta(\Theta)} \widehat{\text{Regret}}(S^*, F) .$$

Furthermore, there exist no saddle points but the following asymmetric saddle-point result holds:

$$\sup_{F \in \Delta(\Theta)} \widehat{\text{Regret}}(S^*, F) = \sup_{F \in \Delta(\Theta)} \inf_{\substack{S \in \Delta(\Omega)^\Theta: \\ \text{(IC),(IR)}}} \widehat{\text{Regret}}(S, F) .$$

Before proving Proposition TR.5, we will argue that the IC/IR constraints restrict the allocation rule x to be constant on $(0,1]$ and $x(0)$ to be at most the constant value. This reduces the single-round problem to optimizing over two decision variables and Proposition TR.5 would follow from a case-by-case analysis. For intuition on the non-existence of a worst-case distribution in the supremum-infimum problem $\sup_{F \in \Delta(\Theta)} \inf_{\substack{S \in \Delta(\Omega)^\Theta: \\ \text{(IC),(IR)}}} \widehat{\text{Regret}}(S, F)$, we note that the IC constraint allows for allocation probabilities that are discontinuous at $\theta = 0$. In order to achieve a regret of $\widehat{\text{Regret}} = c(1-c)$, nature would like the agent's distribution F to be $\theta = 1$ with probability c and $\theta = \epsilon$ with probability $1-c$ for some small value $\epsilon \in (0, c)$. In this case, the single-round benchmark would coincide with $c(1-c)$ and the principal's maximum utility would be $\mathbb{E}_{\theta \sim F}[\theta - c] = (1-c)c + (\epsilon - c)(1-c) = \epsilon(1-c)$, which is achieved by always allocating, and the corresponding regret would be $(c-\epsilon)(1-c)$. Letting $\epsilon \downarrow 0$

would yield the optimal regret $\widehat{\text{Regret}}$. The regret, however, has a discontinuity at $\epsilon = 0$ because in this case, the principal can respond by not allocating when $\theta = 0$ to obtain an expected utility of $\mathbb{E}[(\theta - c)\mathbf{1}\{\theta > 0\}] = c(1-c)$ and, hence, a regret of 0 at $\epsilon = 0$.

Recall we represent single-round direct mechanisms with the interim allocation rule $x : \Theta \to [0,1]$ with the understanding that when the agent reports θ, the probability of allocation is $x(\theta)$. Note $\text{OPT}(\delta_\theta, 1) = \begin{cases} \theta - c & \text{, if } \theta \geq c \\ 0 & \text{, otherwise} \end{cases}$ for any $\theta \in \Theta$ and for any distribution F, the single-round benchmark is

$$\mathbb{E}_{\theta\sim F}[\text{OPT}(\delta_\theta, 1)] = \mathbb{E}_{\theta\sim F}[\mathbf{1}\{\theta \geq c\} \cdot (\theta - c)]\,,$$

where $\mathbf{1}\{\theta \geq c\}$ is the indicator that is 1 if $\theta \geq c$, and 0 otherwise. Then, the single-round minimax regret problem is:

$$\widehat{\text{Regret}} = \inf_{x:(\text{IC}),(\text{IR})} \sup_{F\in\Delta(\Theta)} \left\{\mathbb{E}_{\theta\sim F}[\mathbf{1}\{\theta \geq c\} \cdot (\theta - c)] - \mathbb{E}_{\theta\sim F}[(\theta - c)\cdot x(\theta)]\right\}\,,$$

and the IC/IR constraints for single-round direct mechanisms are:

$$\theta \cdot x(\theta) \geq \theta \cdot x(\theta')\,, \quad \forall \theta, \theta' \in \Theta \tag{IC}$$

$$\theta \cdot x(\theta) \geq 0\,, \quad \forall \theta \in \Theta\,. \tag{IR}$$

The IR constraint is always satisfied. From the IC constraint, we show x is constant on $(0,1]$ and $x(0)$ is at most the constant value. For θ and θ' arbitrarily chosen in $(0,1]$, the IC constraint implies $\theta \cdot x(\theta) \geq \theta \cdot x(\theta')$ and, dividing θ on both sides, $x(\theta) \geq x(\theta')$. Changing the roles of θ and θ', we also have $x(\theta) \leq x(\theta')$. Hence, x is constant on $(0,1]$. Assume $\theta \in (0,1]$ and $\theta' = 0$. Then, the IC constraint implies $\theta \cdot x(\theta) \geq \theta \cdot x(0)$ and, consequently, $x(\theta) \geq x(0)$. It follows that $x(0)$ is at most the constant value. Note the IC constraint is always satisfied when $\theta = 0$. Hence, in what follows, we parametrize single-round direct IC/IR mechanisms in terms of $0 \leq x_0 \leq x_1 \leq 1$ such that $x(0) = x_0$ and $x(\theta) = x_1$ for $\theta \in (0,1]$.

We now prove Proposition TR.5. For ease of presentation, we present the proof for the first part below and the second part on the nonexistence of saddle points and the third part on the asymmetric saddle-point result in Appendices A.1 and A.2, respectively.

Proof of the First Part of Proposition TR.5. Let $F = F^- + F(0)$ where $F(0)$ is the point-mass, which may have the zero mass, at $\theta = 0$ and F^- is the rest of the distribution over $(0,1]$. Then, the objective of the single-round minimax regret problem is

$$\begin{aligned}\widehat{\text{Regret}}((x_0,x_1),F) &= \mathbb{E}_{\theta\sim F^-}[\mathbf{1}\{\theta \geq c\} \cdot (\theta - c)] - \mathbb{E}_{\theta\sim F^-}[(\theta - c)\cdot x_1] + c \cdot x_0 \cdot F(0) \\ &= \mathbb{E}_{\theta\sim F^-}[\mathbf{1}\{\theta \geq c\} \cdot (\theta - c) - \theta \cdot x_1] + c \cdot x_1 \cdot (1 - F(0)) + c \cdot x_0 \cdot F(0)\,.\end{aligned}$$

Let $g(\theta) = \mathbf{1}\{\theta \geq c\} \cdot (\theta - c) - \theta \cdot x_1$ be the expression inside the expectation with c and x_1 fixed. For $\theta \in (0,c)$, $g(\theta) = -\theta \cdot x_1$. For $\theta \in [c,1]$, $g(\theta) = \theta \cdot (1 - x_1) - c$.

For each possible pair (x_0, x_1) such that $0 \leq x_0 \leq x_1 \leq 1$, we compute $\sup_{F\in\Delta(\Theta)} \widehat{\text{Regret}}((x_0,x_1),F)$ and find the corresponding worst-case distributions.

Case 1) $1 - x_1 - c \geq 0$ (Equivalently, $1 - c \geq x_1$.)

Whether $x_1 = 0$ or $x_1 > 0$, $g(1) > g(\theta)$ for $\theta \in (0, 1)$. That is, $g(\theta)$ achieves the unique maximum at $\theta = 1$ over the interval $(0, 1]$. If $x_1 = 0$, $g(\theta) = 0$ for $\theta \in (0, c)$ and $g(\theta) = \theta - c$ for $\theta \in [c, 1]$. If $x_1 > 0$, $g(\theta) = -\theta \cdot x_1 < 0$ for $\theta \in (0, c)$ and $g(\theta) = \theta \cdot (1 - x_1) - c$ for $\theta \in [c, 1]$, which is maximized at $\theta = 1$ and at least 0 at that point.

Then, the worst-case partial distribution F^- over $(0, 1]$ given $F(0)$ is the probability mass of $1 - F(0)$ at $\theta = 1$. It follows that

$$\begin{aligned}
&\sup_{F \in \Delta(\Theta)} \widehat{\text{Regret}}((x_0, x_1), F) \\
&\qquad = \sup_{F(0) \in [0,1]} \left\{ (1 - x_1 - c) \cdot (1 - F(0)) + c \cdot x_1 \cdot (1 - F(0)) + c \cdot x_0 \cdot F(0) \right\} \\
&\qquad = \sup_{F(0) \in [0,1]} \left\{ (1 - x_1 - c + c \cdot x_1) \cdot (1 - F(0)) + c \cdot x_0 \cdot F(0) \right\} .
\end{aligned}$$

We further analyze the simplified expression in the following cases. In all cases, we show

$$\sup_{F \in \Delta(\Theta)} \widehat{\text{Regret}}((x_0, x_1), F) = 1 - c - x_1 + c \cdot x_1.$$

Case a) $1 - x_1 - c > 0$

Note $1 - x_1 - c + c \cdot x_1 > c \cdot x_0$ since $c \cdot x_1 \geq c \cdot x_0$. The worst-case distribution F is uniquely determined to be the point-mass of 1 at $\theta = 1$, i.e., $F(0) = 0$. Then,

$$\sup_{F \in \Delta(\Theta)} \widehat{\text{Regret}}((x_0, x_1), F) = 1 - c - (1 - c) \cdot x_1 .$$

Case b) $1 - x_1 - c = 0$ and $x_0 < x_1$

Since $1 - x_1 - c = 0$, the maximization problem further simplifies to

$$\sup_{F(0) \in [0,1]} \left\{ c \cdot x_1 \cdot (1 - F(0)) + c \cdot x_0 \cdot F(0) \right\} .$$

The worst-case distribution F is uniquely determined to be the point-mass of 1 at $\theta = 1$, i.e., $F(0) = 0$. Then,

$$\sup_{F \in \Delta(\Theta)} \widehat{\text{Regret}}((x_0, x_1), F) = c \cdot x_1 .$$

Case c) $1 - x_1 - c = 0$ and $x_0 = x_1$

As in Case 1b), the maximization problem becomes

$$\sup_{F(0) \in [0,1]} \left\{ c \cdot x_1 \cdot (1 - F(0)) + c \cdot x_0 \cdot F(0) \right\} .$$

The worst-case distribution F is not uniquely determined. Any split between $\theta = 0$ and $\theta = 1$ is a worst-case, i.e., any $F(0) \in [0, 1]$. Then,

$$\sup_{F \in \Delta(\Theta)} \widehat{\text{Regret}}((x_0, x_1), F) = c \cdot x_1 = c \cdot x_0 .$$

Case 2) $1 - x_1 - c < 0$ (Equivalently, $1 - c < x_1$.)

Whether $x_1 = 1$ or $x_1 < 1$, $g(\theta) < 0$ for all $\theta \in (0, 1]$. For $\theta \in (0, c)$, $g(\theta) = -\theta \cdot x_1 < -\theta \cdot (1-c) < 0$. For $\theta \in [c, 1]$, $g(\theta) = \theta \cdot (1 - x_1) - c \leq 1 - x_1 - c < 0$. Note $\widehat{\text{Regret}}((x_0, x_1), F) \leq c \cdot x_1 \cdot (1 - F(0)) + c \cdot x_0 \cdot F(0)$ for any distribution F. Then,

$$\begin{aligned}\sup_{F \in \Delta(\Theta)} \widehat{\text{Regret}}((x_0, x_1), F) &\leq \sup_{F(0) \in [0,1]} \{c \cdot x_1 \cdot (1 - F(0)) + c \cdot x_0 \cdot F(0)\} \\ &= c \cdot x_1 .\end{aligned}$$

In fact, $\sup_{F \in \Delta(\Theta)} \widehat{\text{Regret}}((x_0, x_1), F) = c \cdot x_1$. For any arbitrarily small $\epsilon \in (0, c)$, let F^ϵ be the point-mass distribution such that $\theta = \epsilon$ with probability 1. Then, $\widehat{\text{Regret}}((x_0, x_1), F_\epsilon) = (c - \epsilon) \cdot x_1$. As ϵ was arbitrary, we indeed have

$$\sup_{F \in \Delta(\Theta)} \widehat{\text{Regret}}((x_0, x_1), F) = c \cdot x_1 .$$

In the following cases, we determine corresponding worst-case distributions.

Case a) $x_0 < x_1$

There exists no worst-case distribution that achieves the supremum exactly, because while the expression $c \cdot x_1 \cdot (1 - F(0)) + c \cdot x_0 \cdot F(0)$ in the regret objective is maximized by putting probability mass over $(0, 1]$, the remainder $\mathbb{E}_{\theta \sim F^-}[\mathbf{1}\{\theta \geq c\} \cdot (\theta - c) - \theta \cdot x_1]$ is strictly negative. We can achieve the supremum arbitrarily closely by point-mass distribution F_ϵ for which $\theta = \epsilon$ with probability 1 for arbitrarily small $\epsilon \in (0, c)$.

Case b) $x_0 = x_1$

The worst-case distribution F is uniquely determined to be the point-mass of 1 at $\theta = 0$, i.e., $F(0) = 1$. To see this, we note the regret objective reduces to

$$\widehat{\text{Regret}}((x_0, x_1), F) = \mathbb{E}_{\theta \sim F^-}[\mathbf{1}\{\theta \geq c\} \cdot (\theta - c) - \theta \cdot x_1] + c \cdot x_0 .$$

Since $\mathbb{E}_{\theta \sim F^-}[\mathbf{1}\{\theta \geq c\} \cdot (\theta - c) - \theta \cdot x_1]$ is strictly negative for any probability mass placed over $(0, 1]$, the maximum regret is realized for $F(0) = 1$.

Then, we choose $x_1^* = 1 - c$ such that $1 - x_1 - c + c \cdot x_1 = c \cdot x_1$ and let x_0^* be any number in $[0, x_1^*]$. Any mechanism of this form is optimal and achieves the minimax regret of $\widehat{\text{Regret}} = c \cdot (1 - c)$. Hence, the first part of the proposition follows. Note x_1^* is uniquely determined. If $x_1 < x_1^*$, we have

$$\sup_{F \in \Delta(\Theta)} \widehat{\text{Regret}}((x_0, x_1), F) = 1 - c - (1 - c) \cdot x_1 > 1 - c - (1 - c) \cdot x_1^* = c \cdot (1 - c) .$$

If $x > x_1^*$, then

$$\sup_{F \in \Delta(\Theta)} \widehat{\text{Regret}}((x_0, x_1), F) = c \cdot x_1 > c \cdot x_1^* = c \cdot (1 - c) . \qquad \square$$

A.1 Second Part of Proposition TR.5: Nonexistence of Saddle Points

First, we consider the implications from the existence of a saddle point. Note it is always true that

$$\inf_{\substack{S\in\Delta(\Omega)^\Theta:\\ \text{(IC),(IR)}}} \sup_{F\in\Delta(\Theta)} \widehat{\text{Regret}}(S,F) \geq \sup_{F\in\Delta(\Theta)} \inf_{\substack{S\in\Delta(\Omega)^\Theta:\\ \text{(IC),(IR)}}} \widehat{\text{Regret}}(S,F)\,,$$

which is the max-min inequality. If there exists a saddle point (S^*, F^*) such that

$$\widehat{\text{Regret}}(S^*,F) \leq \widehat{\text{Regret}}(S^*,F^*) \leq \widehat{\text{Regret}}(S,F^*)$$

for any single-round direct IC/IR mechanism S and distribution F, it follows that

$$\begin{aligned}
\inf_{\substack{S\in\Delta(\Omega)^\Theta:\\ \text{(IC),(IR)}}} \sup_{F\in\Delta(\Theta)} \widehat{\text{Regret}}(S,F) &\leq \sup_{F\in\Delta(\Theta)} \widehat{\text{Regret}}(S^*,F) \\
&\leq \widehat{\text{Regret}}(S^*,F^*) \\
&\leq \inf_{\substack{S\in\Delta(\Omega)^\Theta:\\ \text{(IC),(IR)}}} \widehat{\text{Regret}}(S,F^*) \\
&\leq \sup_{F\in\Delta(\Theta)} \inf_{\substack{S\in\Delta(\Omega)^\Theta:\\ \text{(IC),(IR)}}} \widehat{\text{Regret}}(S,F)\,.
\end{aligned}$$

Combining with the max-min inequality, it follows that all the inequalities are actually equalities. In particular, the existence of a saddle point (S^*, F^*) implies that 1) S^* is an optimal solution to the infimum-supremum problem (i.e., $\inf_{\substack{S\in\Delta(\Omega)^\Theta:\\ \text{(IC),(IR)}}} \sup_{F\in\Delta(\Theta)} \widehat{\text{Regret}}(S,F)$) and achieves the objective value of $\widehat{\text{Regret}}(S^*,F^*)$ and F^* is a worst-case distribution for S^*; and 2) F^* is an optimal solution to the supremum-infimum problem (i.e., $\sup_{F\in\Delta(\Theta)} \inf_{\substack{S\in\Delta(\Omega)^\Theta:\\ \text{(IC),(IR)}}} \widehat{\text{Regret}}(S,F)$) and achieves the objective value of $\widehat{\text{Regret}}(S^*,F^*)$ and S^* is an optimal mechanism against F^*.

Given the above discussion, it suffices we show that there exists no optimal solution F^* to the infimum-supremum problem that achieves the objective value of $\widehat{\text{Regret}} = c(1-c)$, or, mathematically,

$$\inf_{\substack{S\in\Delta(\Omega)^\Theta:\\ \text{(IC),(IR)}}} \widehat{\text{Regret}}(S,F^*) = \widehat{\text{Regret}} = c(1-c)\,. \tag{TR-12}$$

We prove it by contradiction. Let F^* be such distribution that satisfies (TR-12). We have

$$\begin{aligned}
\inf_{\substack{S\in\Delta(\Omega)^\Theta:\\ \text{(IC),(IR)}}} \widehat{\text{Regret}}(S,F^*) &= \inf_{x:\text{(IC),(IR)}} \{\mathbb{E}_{\theta\sim F^*}[\text{OPT}(\delta_\theta,1)] - \mathbb{E}_{\theta\sim F^*}[(\theta-c)\cdot x(\theta)]\} \\
&= \mathbb{E}_{\theta\sim F^*}[\text{OPT}(\delta_\theta,1)] - \sup_{x:\text{(IC),(IR)}} \mathbb{E}_{\theta\sim F^*}[(\theta-c)\cdot x(\theta)] \\
&= \mathbb{E}_{\theta\sim F^*}[\text{OPT}(\delta_\theta,1)] - \sup_{0\le x_0\le x_1\le 1} \{x_1\mathbb{E}_{\theta\sim F^*}[(\theta-c)\mathbf{1}\{\theta>0\}] - cx_0F^*(0)\} \\
&= \mathbb{E}_{\theta\sim F^*}[\text{OPT}(\delta_\theta,1)] - \max\{\underbrace{\mathbb{E}_{\theta\sim F^*}[(\theta-c)\mathbf{1}\{\theta>0\}]}_{(*)},0\}\,, \qquad \text{(TR-13)}
\end{aligned}$$

where the second equality follows from extracting the constant term $\mathbb{E}_{\theta\sim F^*}[\text{OPT}(\delta_\theta,1)]$ and flipping the direction of the infimum; the third equality because single-round direct IC/IR mechanisms can be parametrized in terms of $0\le x_0\le x_1\le 1$ such that $x(0)=x_0$ and $x(\theta)=x_1$ for $\theta\in(0,1]$; and the last equality because it is optimal to set $x_0=0$ and $x_1=1$ if $\mathbb{E}_{\theta\sim F^*}[(\theta-c)\mathbf{1}\{\theta>0\}]\ge 0$ and $x_1=0$ otherwise.

We claim that F^* satisfies $(*)=0$. First, note that the single-round benchmark satisfies

$$\begin{aligned}
\mathbb{E}_{\theta\sim F^*}[\text{OPT}(\delta_\theta,1)] &= \mathbb{E}_{\theta\sim F^*}[\max\{\theta-c,0\}] \\
&= \mathbb{E}_{\theta\sim F^*}[\max\{\theta-c,0\}\cdot\mathbf{1}(\theta>0)] \\
&\le (1-c)\mathbb{E}_{\theta\sim F^*}[\theta\mathbf{1}\{\theta>0\}]\,, \qquad \text{(TR-14)}
\end{aligned}$$

where the second equality follows because $c\in(0,1)$, and the inequality because $\max\{\theta-c,0\}\le(1-c)\theta$ since $\theta\in[0,1]$ and $c\in(0,1)$. Suppose $(*)<0$. Combining (TR-12), (TR-13) and (TR-14), we obtain

$$c(1-c)\le(1-c)\mathbb{E}_{\theta\sim F^*}[\theta\mathbf{1}\{\theta>0\}] < c(1-c)(1-F^*(0)) \le c(1-c)\,,$$

where the strict inequality follows because $c\in(0,1)$ and $(*)<0$ implies $\mathbb{E}_{\theta\sim F^*}[\theta\mathbf{1}\{\theta>0\}] < \mathbb{E}_{\theta\sim F^*}[c\mathbf{1}\{\theta>0\}]=c(1-F^*(0))$, and the last inequality because $F^*(0)\in[0,1]$. This is a contradiction.

Similarly, suppose $(*)>0$. Combining (TR-12), (TR-13), and (TR-14), we obtain

$$c(1-c)\le c(1-F^*(0)) - c\mathbb{E}_{\theta\sim F^*}[\theta\mathbf{1}(\theta>0)] < c(1-c)(1-F^*(0)) \le c(1-c)\,,$$

where the strict inequality follows because $c\in(0,1)$ and $(*)>0$ implies $\mathbb{E}_{\theta\sim F^*}[\theta\mathbf{1}\{\theta>0\}] > \mathbb{E}_{\theta\sim F^*}[c\mathbf{1}\{\theta>0\}]=c(1-F^*(0))$, and the last inequality because $F^*(0)\in[0,1]$. Again, a contradiction. Hence, we have $(*)=0$.

We now argue that $(*)=0$ implies that $F^*(0)=1$. Combining (TR-12), (TR-13), and (TR-14) together with $(*)=0$ implies that $c(1-c)\le c(1-c)(1-F^*(0))\le c(1-c)$; to see this, we follow the same argument above under the assumption $(*)<0$ where the strict inequality becomes an equality. That is, $c(1-c)(1-F^*(0))=c(1-c)$. Because $c\in(0,1)$, we can divide both sides by $c(1-c)$ and obtain that $F^*(0)=1$. Hence, the only possible candidate distribution F^* is the point-mass distribution under which the shock is 0 with probability 1. For this particular distribution F^*, $\mathbb{E}_{\theta\sim F^*}[\text{OPT}(\delta_\theta,1)]=0$ and the IC/IR mechanism that always does not allocate (i.e., $x_0=x_1=0$) is such that $\mathbb{E}_{\theta\sim F^*}[(\theta-c)\cdot x(\theta)]=0$. This means $\inf_{\substack{S\in\Delta(\Omega)^\Theta:\\ \text{(IC),(IR)}}}\widehat{\text{Regret}}(S,F^*)$ can be at most 0 and cannot be equal to $\widehat{\text{Regret}} = c(1-c) > 0$. We, thus, conclude there exists no such distribution F^* satisfying (TR-12).

A.2 Third Part of Proposition TR.5: an Asymmetric Result

While saddle points do not exist, we can still show an asymmetric saddle-point result, that is, the saddle-point property holds and the single-round minimax regret problem admits an optimal solution. In Part 1, we showed that S^* is an optimal solution to the single-round minimax regret problem and achieves $\widehat{\text{Regret}}$. In what follows, we show the saddle-point property holds:

$$\inf_{\substack{S\in\Delta(\Omega)^{\Theta}:\\ \text{(IC),(IR)}}} \sup_{F\in\Delta(\Theta)} \widehat{\text{Regret}}(S,F) = \sup_{F\in\Delta(\Theta)} \inf_{\substack{S\in\Delta(\Omega)^{\Theta}:\\ \text{(IC),(IR)}}} \widehat{\text{Regret}}(S,F)\,.$$

From Part 1, we know the left-hand side is equal to $\widehat{\text{Regret}} = c(1-c)$.

For an arbitrary $\epsilon \in (0,c)$, consider the distribution F_ϵ under which $\theta = \epsilon$ with probability $1-c$ and $\theta = 1$ with probability c. Then,

$$\begin{aligned}
\widehat{\text{Regret}}((x_0,x_1),F) &= \mathbb{E}_{\theta\sim F^-}[\mathbf{1}\{\theta \geq c\}\cdot(\theta-c) - \theta\cdot x_1] + c\cdot x_1\cdot(1-F(0)) + c\cdot x_0\cdot F(0)\\
&= -\epsilon\cdot x_1\cdot(1-c) + (1-c-x_1)\cdot c + c\cdot x_1\\
&= -\epsilon\cdot(1-c)\cdot x_1 + c\cdot(1-c)\,.
\end{aligned}$$

It follows that

$$\inf_{x:\text{(IC),(IR)}} \widehat{\text{Regret}}((x_0,x_1),F_\epsilon) = -\epsilon\cdot(1-c) + c\cdot(1-c)\,,$$

where the infimum is achieved by the single-round mechanism with $x_1 = 1$ and $x_0 \in [0,1]$. As ϵ was arbitrary,

$$\sup_{F\in\Delta(\Theta)} \inf_{x:\text{(IC),(IR)}} \widehat{\text{Regret}}((x_0,x_1),F_\epsilon) = c\cdot(1-c)\,,$$

as desired.

B Additional Materials for Section TR.6

B.1 Two Alternative Benchmarks

Before proving Theorem TR.3, we prove the following proposition relating $\bar{u}(F)$ and $\tilde{u}(F)$ which will be used in proving Theorem TR.3.

Proposition TR.6. *We have the following relations:*

1. *For any distribution $F \in \Delta(\Theta)$, $\tilde{u}(F) \leq \bar{u}(F)$.*
2. *For any point-mass distribution $\theta \in \Theta$, $\tilde{u}(\delta_\theta) = \bar{u}(\delta_\theta)$.*

Proof. We prove the first part. Let F be any arbitrary distribution. We equivalently write $\tilde{u}(F)$ as

$$\begin{aligned} \tilde{u}(F) = \sup_{S \in \Delta(\Omega)^{\Theta}} & \int_{\Theta} \int_{\Omega} u(\theta, \omega) \mathrm{d}S_{\theta}(\omega) \mathrm{d}F(\theta) \\ \text{s.t.} & \int_{\Omega} v(\theta, \omega) \mathrm{d}S_{\theta}(\omega) \geq \int_{\Omega} v(\theta, \omega) \mathrm{d}S_{\theta'}(\omega)\,, \quad \forall \theta, \theta' \in \Theta \\ & \int_{\Omega} v(\theta, \omega) \mathrm{d}S_{\theta}(\omega) \geq 0\,, \quad \forall \theta \in \Theta\,. \end{aligned} \tag{TR-15}$$

Relaxing the IC constraint, we obtain the optimization problem in Section 5.2 of the main paper that defines $\mathbb{E}_{\theta \sim F}[\bar{u}(\delta_\theta)]$ (also (E-4) in Appendix E of the main paper). Hence, $\tilde{u}(F) \leq \mathbb{E}_{\theta \sim F}[\bar{u}(\delta_\theta)]$. By Part 3 of Proposition E.1 in Appendix E of the main paper, it follows that $\tilde{u}(F) \leq \bar{u}(F)$.

For the second part, we now show that $\tilde{u}(\delta_\theta)$ is equal to $\bar{u}(\delta_\theta)$ for any $\theta \in \Theta$. Fix an arbitrary θ. Let G_θ be an arbitrary outcome distribution over Ω that is feasible for the following equivalent form of the optimization problem defining $\bar{u}(\delta_\theta)$:

$$\begin{aligned} \bar{u}(\delta_\theta) = \sup_{G \in \Delta(\Omega)} & \int_{\Omega} u(\theta, \omega) \mathrm{d}G(\omega) \\ \text{s.t.} & \int_{\Omega} v(\theta, \omega) \mathrm{d}G(\omega) \geq 0\,. \end{aligned}$$

Consider the single-round direct IC/IR mechanism S that given report θ' returns a random outcome drawn from G_θ if $\mathbb{E}_{\omega \sim G_\theta}[v(\theta', \omega)] \geq 0$, and the no-interaction outcome otherwise. Mechanism S provides two possibilities and chooses the better possible outcome (or a distribution of outcomes) for each θ' if the agent's private shock was θ'. Hence, by construction, S satisfies the IC/IR constraints and is a feasible solution in the optimization problem defining $\tilde{u}(\delta_\theta)$. Furthermore, $\mathbb{E}_{\omega \sim S_\theta}[u(\theta, \omega)] = \mathbb{E}_{\omega \sim G_\theta}[u(\theta, \omega)]$ and G_θ and S obtain the same objectives in respective optimization problems. As G_θ was arbitrary, it follows that $\tilde{u}(\delta_\theta) \geq \bar{u}(\delta_\theta)$. Combined with the observation that $\tilde{u}(F) \leq \bar{u}(F)$ for any distribution $F \in \Delta(\Theta)$, it follows that $\tilde{u}(\delta_\theta) = \bar{u}(\delta_\theta)$. □

We now prove Theorem TR.3 below:

Proof of Theorem TR.3. We show that Lemmas 2 and 3 of the main paper hold for the $\text{Regret}^{\mathsf{FB}}$ and $\text{Regret}^{\mathsf{SB}}$ notions in the multi-round problem (and with the same $\text{Regret}^{\mathsf{OPT}}$ notion in the single-round problem). Then, the theorem statements would follow for these notions by the same reasoning steps used in the proof of Theorem 1 of the main paper for the $\text{Regret}^{\mathsf{OPT}}$ notion in Appendix A.1 of the main paper. The assumptions of the theorem are similar to those of Theorem 1 of the main paper. The second part, the linearity assumption of $\bar{u}(F)$, is more stringent and needed for our false-dynamics results.

We start by showing Lemma 2 of the main paper holds with respect to the alternative benchmarks. We use the same reasoning used for when we start with the $\text{Regret}^{\mathsf{OPT}}$ notion for the multi-round problem. Fix an arbitrary incentive compatible dynamic mechanism $A \in \mathcal{A}$. Then,

$$\begin{aligned}
\sup_{F\in\mathcal{F}} \text{Regret}^{\mathsf{FB}}(A, F, T) &\geq \sup_{\theta\in\Theta} \text{Regret}^{\mathsf{FB}}(A, \delta_\theta, T) \\
&= \sup_{\theta\in\Theta} \{T \cdot \bar{u}(\delta_\theta) - \text{PrincipalUtility}(A, \sigma, \delta_\theta, T)\} \\
&= T \cdot \sup_{\theta\in\Theta} \{\bar{u}(\delta_\theta) - \text{PrincipalUtility}(S(A), \sigma^{\text{TR}}, \delta_\theta, 1)\} \\
&= T \cdot \sup_{\theta\in\Theta} \{\text{OPT}(\delta_\theta, 1) - \text{PrincipalUtility}(S(A), \sigma^{\text{TR}}, \delta_\theta, 1)\} \\
&= T \cdot \sup_{\theta\in\Theta} \text{Regret}^{\mathsf{OPT}}(S(A), \delta_\theta, 1)\,,
\end{aligned}$$

where the first step is because point-mass distributions are a subset of $\mathcal{F}$ under Assumption 1 of the main paper; the second step is by the definition of the $\text{Regret}^{\mathsf{FB}}$ notion and where σ is the recommended agent strategy under A; the third step is by Lemma A.1 of the main paper and $S(A)$ is the single-round direct IC/IR mechanism derived from A as described in the proof of the lemma; the second-to-last step is because $\text{OPT}(\delta_\theta, 1) = \bar{u}(\delta_\theta)$ which is by Proposition E.1 of the main paper; and the last step is by the definition of the $\text{Regret}^{\mathsf{OPT}}$ notion. Hence, Lemma 2 holds for the $\text{Regret}^{\mathsf{FB}}$ notion.

For the $\text{Regret}^{\mathsf{SB}}$ notion, we follow the same reasoning with $\text{Regret}^{\mathsf{SB}}$ and $\tilde{u}(F)$ in places of $\text{Regret}^{\mathsf{FB}}$ and $\bar{u}(F)$, respectively. The second-to-last step in the above sequence still follows from Proposition E.1 of the main paper and Proposition TR.6 which imply that $\text{OPT}(\delta_\theta, 1) = \bar{u}(\delta_\theta) = \tilde{u}(\delta_\theta)$ for $\theta \in \Theta$.

We now show Lemma 3 of the main paper holds with respect to the alternative benchmarks. Let S be an arbitrary single-round direct IC/IR mechanism and consider the direct static mechanism $S^{\times T}$ which is T repetitions of S. Note $S^{\times T}$ is incentive compatible by Lemma A.2 of the main paper. Then,

$$\begin{aligned}
\sup_{F\in\mathcal{F}} \text{Regret}^{\mathsf{FB}}(S^{\times T}, F, T) &= \sup_{F\in\mathcal{F}} \left\{T \cdot \bar{u}(F) - \text{PrincipalUtility}(S^{\times T}, \sigma^{\text{TR}}, F, T)\right\} \\
&= \sup_{F\in\mathcal{F}} \left\{T \cdot \mathbb{E}_{\theta\sim F}[\text{OPT}(\delta_\theta, 1)] - \text{PrincipalUtility}(S^{\times T}, \sigma^{\text{TR}}, F, T)\right\} \\
&= T \cdot \sup_{F\in\mathcal{F}} \{\mathbb{E}_{\theta\sim F}[\text{OPT}(\delta_\theta, 1)] - \text{PrincipalUtility}(S, \sigma^{\text{TR}}, F, 1)\} \\
&\leq T \cdot \sup_{\theta\in\Theta} \text{Regret}^{\mathsf{OPT}}(S, \sigma^{\text{TR}}, \delta_\theta, 1)\,,
\end{aligned}$$

where the first step is by the definition of the $\text{Regret}^{\mathsf{FB}}$ notion and the truthful reporting strategy σ^{TR} (i.e., the agent reports $\mathsf{CONTINUE}$ in Round 0 and truthfully reports his shocks in future rounds) which is the recommended strategy for direct mechanisms; the second step is because for any distribution $F \in \mathcal{F}$,

$$T \cdot \bar{u}(F) = T \cdot \mathbb{E}_{\theta\sim F}[\bar{u}(\delta_\theta)] = T \cdot \mathbb{E}_{\theta\sim F}[\text{OPT}(\delta_\theta, 1)]\,,$$

by the linearity assumption on $\bar{u}(F)$ and Proposition E.1 of the main paper; the second-to-last step is by Lemma A.2 of the main paper; and the last step is by Lemma A.3 of the main paper. Hence, Lemma 3 holds for the $\text{Regret}^{\mathsf{FB}}$ notion.

Similarly, we follow the same reasoning for the $\text{Regret}^{\mathsf{SB}}$ notion in terms of $\tilde{u}(F)$. The second step in the above sequence will be an inequality and it follows because

$$T \cdot \tilde{u}(F) \leq T \cdot \bar{u}(F) = T \cdot \mathbb{E}_{\theta \sim F}[\bar{u}(\delta_\theta)] = T \cdot \mathbb{E}_{\theta \sim F}[\mathrm{OPT}(\delta_\theta, 1)]\,,$$

where we used Proposition TR.6, the linearity assumption on $\bar{u}(F)$ and Proposition E.1 of the main paper, in that order. □

B.2 Arbitrary Shock Processes

We first prove Theorem TR.4. We then introduce a generalization of $\bar{u}(F)$ and prove analogous results to Propositions 7, 8 and E.1 of the main paper in the general shock process setting.

Proof of Theorem TR.4. Note the two assumptions of the theorem are the analogues of Assumptions 1 and 2 of the main paper, respectively, in the general shock process setting. We prove the analogues of Lemmas 2 and 3 of the main paper for the general shock process setting. Then, the theorem statements would follow directly from the analogues via the same reasoning steps in the proof of Theorem 1 of the main paper.

For an analogue of Lemma 2 of the main paper, we proceed as follows. Note that for any incentive compatible dynamic mechanism $A \in \mathcal{A}$, we have

$$\begin{aligned}
\sup_{F^T \in \mathcal{F}^T} \mathrm{Regret}(A, F^T, T) &\geq \sup_{\theta \in \Theta} \mathrm{Regret}(A, \delta_\theta^{\times T}, T) \\
&= \sup_{\theta \in \Theta} \{\mathrm{OPT}(\delta_\theta, T) - \mathrm{PrincipalUtility}(A, \sigma, \delta_\theta, T)\} \\
&= \sup_{\theta \in \Theta} \{\mathrm{OPT}(\delta_\theta, T) - T \cdot \mathrm{PrincipalUtility}(S(A), \sigma^{\mathrm{TR}}, \delta_\theta, 1)\} \\
&= T \cdot \sup_{\theta \in \Theta} \{\mathrm{OPT}(\delta_\theta, 1) - \mathrm{PrincipalUtility}(S(A), \sigma^{\mathrm{TR}}, \delta_\theta, 1)\} \\
&= T \cdot \sup_{\theta \in \Theta} \mathrm{Regret}(S(A), \delta_\theta, 1)\,,
\end{aligned}$$

where the first step is because the set of shock processes where per-round shocks are some fixed constant is a subset of $\mathcal{F}^T$ under the first assumption; the second step is by the definition of the Regret notion and where σ is the recommended agent strategy under A; the third step is by Lemma A.1 of the main paper and $S(A)$ is the single-round direct IC/IR mechanism derived from A as described in the proof of the lemma; the second-to-last step is by Proposition 1 of the main paper; and the last step is by the definition of the Regret notion. This is an analogous lower bound to Lemma 2 of the main paper for the general shock process setting.

We now show an analogous upper bound to Lemma 3 of the main paper. Let $S \in \mathcal{S}^{\times 1}$ be an arbitrary single-round direct IC/IR mechanism. Let $S^{\times T}$ be the direct static mechanism that is T repetitions of S, which is incentive compatible by Lemma A.2 of the main paper. For any $F^T \in \mathcal{F}^T$, we have

$$\mathrm{Regret}(S^{\times T}, F^T, F) = \mathrm{OPT}(F^T, T) - \mathrm{PrincipalUtility}(S^{\times T}, \sigma^{\mathrm{TR}}, F^T, T)\,,$$

where the truthful reporting strategy σ^{TR} (i.e., the agent reports $\mathsf{CONTINUE}$ in Round 0 and truthfully reports his shocks in future rounds) is the recommended strategy for direct mechanisms. We can upper bound $\mathrm{OPT}(F^T, T)$ as follows:

$$\begin{aligned}\mathrm{OPT}(F^T, T) &\le \mathbb{E}_{t\sim[T],\theta\sim(F^T)_t}[\mathrm{OPT}(\delta_\theta^{\times T}, T)] = \frac{1}{T}\sum_{t=1}^{T}\mathbb{E}_{\theta\sim(F^T)_t}[\mathrm{OPT}(\delta_\theta, T)]\\ &= \sum_{t=1}^{T}\mathbb{E}_{\theta\sim(F^T)_t}[\mathrm{OPT}(\delta_\theta, 1)]\,,\end{aligned}$$

where the first step is by the second assumption of the theorem and the last step is by Proposition 1 of the main paper. Then,

$$\begin{aligned}\mathrm{Regret}(S^{\times T}, F^T, T) &\le \sum_{t=1}^{T}\mathbb{E}_{\theta\sim(F^T)_t}[\mathrm{OPT}(\delta_\theta, 1)] - \mathrm{PrincipalUtility}(S^{\times T}, \sigma^{\mathrm{TR}}, F^T, T)\\ &= \sum_{t=1}^{T}\Big(\mathbb{E}_{\theta\sim(F^T)_t}[\mathrm{OPT}(\delta_\theta, 1)] - \mathrm{PrincipalUtility}(S, \sigma^{\mathrm{TR}}, (F^T)_t, 1)\Big)\\ &\le \sum_{t=1}^{T}\sup_{F'\in\Delta(\Theta)}\{\mathbb{E}_{\theta\sim F'}[\mathrm{OPT}(\delta_\theta, 1)] - \mathrm{PrincipalUtility}(S, \sigma^{\mathrm{TR}}, F', 1)\}\\ &\le T\cdot\sup_{\theta\in\Theta}\mathrm{Regret}(S, \delta_\theta, 1)\,,\end{aligned}$$

where the second step follows from that since the principal repeats S and the agent repeats the same strategy of truthful reporting, the rounds become independent and the principal utility in each round can be written in terms of the marginal shock distributions; the third step is by taking the supremum over $\Delta(\Theta)$ for the marginal distribution in each summand; and the last step is by upper bounding each summand via Lemma A.3 of the main paper (which also holds for $\mathcal{F} = \Delta(\Theta)$). As F^T was arbitrary, it follows that

$$\sup_{F^T\in\mathcal{F}^T}\mathrm{Regret}(S^{\times T}, F^T, T) \le T\cdot\sup_{\theta\in\Theta}\mathrm{Regret}(S, \delta_\theta, 1)\,,$$

which is the upper bound analogue of Lemma 3 of the main paper. □

We now define a generalization of $\bar{u}(F)$ for T-round shock processes and prove analogous results to Propositions 7, 8 and E.1 of the main paper. We still use the same notation and extend it to $F^T \in \Delta(\Theta^T)$. For T-round shock processes $F^T \in \Delta(\Theta^T)$, we define $\bar{u}(F^T)$ as follows:

$$\begin{aligned}\bar{u}(F^T) := \sup_{S_1,\ldots,S_T}\ &\mathbb{E}_{\theta\sim F^T,\omega\sim\prod_{t=1}^{T}S_{t,\theta_t}}\left[\sum_{t=1}^{T}u(\theta_t,\omega_t)\right]\\ \text{s.t. } &\mathbb{E}_{\theta\sim F^T,\omega\sim\prod_{t=1}^{T}S_{t,\theta_t}}\left[\sum_{t=1}^{T}v(\theta_t,\omega_t)\right] \ge 0\,,\end{aligned} \tag{TR-16}$$

where the supremum is over sequences of single-round direct mechanisms $\{S_t\}_{t\in[T]}$ where each mechanism S_t is a collection of outcome distributions $\{S_{t,\theta}\}_{\theta\in\Theta}$ and the expectation is with respect to the shocks $\theta = (\theta_1, \ldots, \theta_T)$ determined according to the T-round shock process F^T and the outcomes $\omega = (\omega_1, \ldots, \omega_T)$ where the t-th round outcome ω_t is determined according to the t-th single-round mechanism S_t and the t-th shock θ_t, i.e., $\omega_t \sim S_{t,\theta_t}$, in the product notation. Similar to $\bar{u}(F)$ in the repeated i.i.d. setting, we can think of $\bar{u}(F^T)$ as the first-best performance that the principal can achieve subject to the ex-ante IR constraint when the shock process and per-round shocks are known. Correspondingly, the linearity property for $\bar{u}(F^T)$ can be stated as $\bar{u}(F^T) = \sum_{t=1}^{T} \mathbb{E}_{\theta \sim (F^T)_t}[\bar{u}(\delta_\theta)]$ for $F^T \in \mathcal{F}^T$. This is a generalization of the linearity property for $\bar{u}(F)$ considered in the main paper in the following sense:

Proposition TR.7. *For any $F \in \mathcal{F}$, if $\bar{u}(F^{\times T}) = \sum_{t=1}^{T} \mathbb{E}_{\theta \sim (F^{\times T})_t}[\bar{u}(\delta_\theta)]$ then $\bar{u}(F) = \mathbb{E}_{\theta \sim F}[\bar{u}(\delta_\theta)]$.*

Proof. Fix any $F \in \mathcal{F}$. Note $\bar{u}(F^{\times T}) \geq T \cdot \bar{u}(F)$ because the right-hand side is the supremum value of (TR-16) for $F^{\times T}$ when the IR constraint is imposed in each round instead of once over all rounds. We also have $\bar{u}(F^{\times T}) \leq T \cdot \bar{u}(F)$ because any feasible solution $\{S_t\}$ to (TR-16) for $F^{\times T}$ can be aggregated by uniformly randomizing over $\{S_t\}$ into a feasible solution in the optimization problem defining $\bar{u}(F)$ that achieves $\frac{1}{T}$ times the objective value achieved by $\{S_t\}$ in (TR-16). Then, we have $\bar{u}(F^{\times T}) = T \cdot \bar{u}(F)$. By the linearity assumption of $\bar{u}(F^{\times T})$,

$$\bar{u}(F^{\times T}) = \sum_{t=1}^{T} \mathbb{E}_{\theta \sim F}[\bar{u}(\delta_\theta)] = T \cdot \mathbb{E}_{\theta \sim F}[\bar{u}(\delta_\theta)] \,.$$

Combining the above two observations, it follows that $\bar{u}(F) = \mathbb{E}_{\theta \sim F}[\bar{u}(\delta_\theta)]$. $\square$

The following proposition shows analogous results similar to those in Propositions 7 and E.1 of the main paper. It relates $\bar{u}(F^T)$ and the optimal performance achievable $\mathrm{OPT}(F^T, T)$.

Proposition TR.8. *The following hold:*

1. *For any $F^T \in \Delta(\Theta^T)$, $\mathrm{OPT}(F^T, T) \leq \bar{u}(F^T)$.*
2. *For any $F^T \in \Delta(\Theta^T)$, $\bar{u}(F^T) \geq \sum_{t=1}^{T} \mathbb{E}_{\theta \sim (F^T)_t}[\bar{u}(\delta_\theta)]$.*
3. *For any $F^T \in \mathcal{F}^T$, if $\bar{u}(F^T) = \sum_{t=1}^{T} \mathbb{E}_{\theta \sim (F^T)_t}[\bar{u}(\delta_\theta)]$ then $\mathrm{OPT}(F^T, T) \leq \mathbb{E}_{t \sim [T], \theta \sim (F^T)_t}[\mathrm{OPT}(\delta_\theta^{\times T}, T)]$.*

Proof. Part 1): Fix an arbitrary T-round shock process $F^T \in \Delta(\Theta^T)$. Assume the principal commits to any arbitrary incentive compatible dynamic mechanism A and the agent plays the recommended strategy σ (given as part of A). Let $\{\omega_t\}_{t=1}^{T}$ be the resulting random sequence of realized outcomes. If the agent reports QUIT in Round 0 and does not participate, the sequence is the sequence of the no-interaction outcome. For each $t \in [T]$ and $\theta \in \Theta$, we define measure $\mu_{t,\theta}(Q) = \Pr(\omega_t \in Q \mid \theta_t = \theta)$ for any $Q \subseteq \Omega$ and let $S_{t,\theta}$ be the corresponding distribution over Ω such that $\omega \sim S_{t,\theta}$ means an outcome ω is realized with probability $\mu_{t,\theta}(\omega)$.

For each $t \in [T]$, we define a single-round direct mechanism $S_t = \{S_{t,\theta}\}_{\theta \in \Theta}$ that given a report θ returns an outcome $\omega \sim S_{t,\theta}$. Consider the sequence of single-round direct mechanisms $\{S_t\}_{t \in [T]}$ thus constructed. We show that $\{S_t\}$ is a feasible solution to (TR-16) and achieves the objective value that is equal to $\mathrm{PrincipalUtility}(A, \sigma, F^T, T)$. As the incentive compatible dynamic mechanism A was arbitrarily chosen, the proposition statement would follow.

First, we note that

$$
\begin{aligned}
\text{AgentUtility}(A, \sigma, F^T, T) &= \mathbb{E}\left[\sum_{t=1}^{T} v(\theta_t, \omega_t)\right] \\
&= \sum_{t=1}^{T} \mathbb{E}_{\theta_t \sim (F^T)_t}\left[\mathbb{E}[v(\theta_t, \omega_t)|\theta_t]\right] \\
&= \sum_{t=1}^{T} \mathbb{E}_{\theta_t \sim (F^T)_t}\left[\mathbb{E}_{\omega_t \sim S_{t,\theta_t}}[v(\theta_t, \omega_t)]\right] \\
&= \mathbb{E}_{\theta \sim F^T, \omega \sim \prod_{t=1}^{T} S_{t,\theta_t}}\left[\sum_{t=1}^{T} v(\theta_t, \omega_t)\right],
\end{aligned}
$$

where the second equality follows from the linearity of expectations and the tower rule, the third from that the t-th round idiosyncratic shock can be thought to be drawn independently from the marginal distribution $(F^T)_t$, and the last from the construction of $\{S_t\}_{t \in [T]}$. Since the agent's recommended strategy σ is a utility-maximizing strategy and the agent can achieve the aggregate utility of 0 by not participating, it must be that $\text{AgentUtility}(A, \sigma, F^T, T) \geq 0$. Hence, $\{S_t\}$ is a feasible solution to (TR-16).

Similarly, we have

$$
\text{PrincipalUtility}(A, \sigma, F^T, T) = \mathbb{E}_{\theta \sim F^T, \omega \sim \prod_{t=1}^{T} S_{t,\theta_t}}\left[\sum_{t=1}^{T} u(\theta_t, \omega_t)\right].
$$

It follows that $\{S_t\}$ is a feasible solution to (TR-16) and achieves the objective value of $\text{PrincipalUtility}(A, \sigma, F^T, T)$. This completes the proof.

Part 2): This is because the right-hand side is the supremum value of (TR-16) with the IR constraint for each shock in each round which is a more constrained version of (TR-16).

Part 3): Let $F^T \in \mathcal{F}^T$ be an arbitrary T-round shock process. Then,

$$
\text{OPT}(F^T, T) \leq \bar{u}(F^T) = \sum_{t=1}^{T} \mathbb{E}_{\theta \sim (F^T)_t}[\bar{u}(\delta_\theta)] = \sum_{t=1}^{T} \mathbb{E}_{\theta \sim (F^T)_t}[\text{OPT}(\delta_\theta, 1)],
$$

where the first step is by the first part, the second step is by the linearity assumption on $\bar{u}(F^T)$ and the last step is by Proposition E.1 of the main paper. The last expression is equivalently

$$
\sum_{t=1}^{T} \mathbb{E}_{\theta \sim (F^T)_t}[\text{OPT}(\delta_\theta, 1)] = T \cdot \mathbb{E}_{t \sim [T], \theta \sim (F^T)_t}[\text{OPT}(\delta_\theta, 1)] = \mathbb{E}_{t \sim [T], \theta \sim (F^T)_t}[\text{OPT}(\delta_\theta, T)],
$$

where we used Proposition 1 of the main paper. □

The following is the analogue of Proposition 8 of the main paper in the general shock process setting. It holds for games with payments that enter linearly into the utility functions of the principal and agent or with a nonnegative utility function for the agent:

Proposition TR.9. *Assume the game satisfies either conditions of Proposition 8 of the main paper. Then,* $\bar{u}(F^T) = \sum_{t=1}^T \mathbb{E}_{\theta\sim(F^T)_t}[\bar{u}(\delta_\theta)]$ *for all* $F^T \in \mathcal{F}^T$.

Proof. By the second part of Proposition TR.8, it remains to show $\bar{u}(F^T) \le \sum_{t=1}^T \mathbb{E}_{\theta\sim(F^T)_t}[\bar{u}(\delta_\theta)]$. It suffices to show that $\bar{u}(F^T) \le \sum_{t=1}^T \bar{u}((F^T)_t)$. It then would follow that

$$\bar{u}(F^T) \le \sum_{t=1}^T \bar{u}((F^T)_t) = \sum_{t=1}^T \mathbb{E}_{\theta\sim(F^T)_t}[\bar{u}(\delta_\theta)]\,,$$

where the equality is by the linearity condition in the repeated i.i.d. setting which holds by Proposition 8 of the main paper.

Recall that the optimization problem (TR-16) is

$$\begin{aligned}\sup_{S_1,\ldots,S_T}\ & \mathbb{E}_{\theta\sim F^T,\omega\sim\prod_{t=1}^T S_{t,\theta_t}}\left[\sum_{t=1}^T u(\theta_t,\omega_t)\right]\\ \text{s.t. } & \mathbb{E}_{\theta\sim F^T,\omega\sim\prod_{t=1}^T S_{t,\theta_t}}\left[\sum_{t=1}^T v(\theta_t,\omega_t)\right] \ge 0\,.\end{aligned}$$

Note $\sum_{t=1}^T \bar{u}((F^T)_t)$ in the value of the following optimization problem which is a modified version of (TR-16) with the nonnegativity constraint for each round:

$$\begin{aligned}\sup_{S_1,\ldots,S_T}\ & \mathbb{E}_{\theta\sim F^T,\omega\sim\prod_{t=1}^T S_{t,\theta_t}}\left[\sum_{t=1}^T u(\theta_t,\omega_t)\right]\\ \text{s.t. } & \mathbb{E}_{\hat{\theta}\sim(F^T)_t,\hat{\omega}\sim S_{t,\hat{\theta}}}\left[v(\hat{\theta},\hat{\omega})\right] \ge 0\,,\quad \forall t\in[T]\,.\end{aligned} \tag{TR-17}$$

Part 1: We use the same representation used in the statement of Proposition 8 of the main paper for separating out the payment from the outcome, outcome space and utility functions of the principal and agent. We refer to the statement of Proposition 8 of the main paper for further details of the representation.

Fix an arbitrary T-round shock process $F^T \in \mathcal{F}^T$. Let $\{S_t\}$ be an arbitrary sequence of single-round direct mechanisms that is a feasible solution to (TR-16). For $t \in [T]$ and $\hat{\theta} \in \Theta$, let the payment offset be defined as

$$q_{t,\hat{\theta}} = \frac{1}{\beta}\left(\mathbb{E}_{\hat{\omega}\sim S_{t,\hat{\theta}}}[v(\hat{\theta},\hat{\omega})] - \frac{1}{T}\cdot\mathbb{E}_{\theta\sim F^T,\omega\sim\prod_{t=1}^T S_{t,\theta_t}}\left[\sum_{t=1}^T v(\theta_t,\omega_t)\right]\right)\,.$$

Note the second term in the payment offset is a constant that does not depend on t or $\hat{\theta}$. By construction, $\mathbb{E}_{\theta\sim F^T}[\sum_{t=1}^T q_{t,\theta_t}] = 0$.

Consider the sequence of single-round direct mechanisms $\{S'_t\}$ constructed as follows. For each $t \in [T]$, S'_t is the single-round direct mechanism where for each $\hat{\theta} \in \Theta$, $S'_{t,\hat{\theta}}$ is the outcome distribution $S_{t,\hat{\theta}}$

modified with the fixed offset $q_{t,\hat{\theta}}$ such that to realize an outcome $\hat{\omega} \sim S'_{t,\hat{\theta}}$, we draw $(\hat{\omega}^0, \hat{p}) \sim S_{t,\hat{\theta}}$ and set $\hat{\omega} = (\hat{\omega}^0, \hat{p} + q_{t,\hat{\theta}})$. Note $q_{t,\hat{\theta}}$ is the same constant adjustment for every realized outcome.

Note, for each $t \in [T]$,

$$\begin{aligned}
\mathbb{E}_{\substack{\hat{\theta} \sim (F^T)_t \\ \hat{\omega} \sim S'_{t,\hat{\theta}}}} \left[v(\hat{\theta}, \hat{\omega}) \right] &= \mathbb{E}_{\substack{\hat{\theta} \sim (F^T)_t \\ (\hat{\omega}^0, \hat{p}) \sim S_{t,\hat{\theta}}}} \left[v(\hat{\theta}, (\hat{\omega}^0, \hat{p} + q_{t,\hat{\theta}})) \right] \\
&= \mathbb{E}_{\substack{\hat{\theta} \sim (F^T)_t \\ \hat{\omega} \sim S_{t,\hat{\theta}}}} \left[v(\hat{\theta}, \hat{\omega}) \right] - \beta \cdot \mathbb{E}_{\hat{\theta} \sim (F^T)_t} [q_{t,\hat{\theta}}] \\
&= \mathbb{E}_{\substack{\hat{\theta} \sim (F^T)_t \\ \hat{\omega} \sim S_{t,\hat{\theta}}}} \left[v(\hat{\theta}, \hat{\omega}) \right] - \left(\mathbb{E}_{\substack{\hat{\theta} \sim (F^T)_t \\ \hat{\omega} \sim S_{t,\hat{\theta}}}} \left[v(\hat{\theta}, \hat{\omega}) \right] - \frac{1}{T} \cdot \mathbb{E}_{\substack{\theta \sim F^T \\ \omega \sim \prod_{t=1}^T S_{t,\theta_t}}} \left[\sum_{t=1}^T v(\theta_t, \omega_t) \right] \right) \\
&= \frac{1}{T} \cdot \mathbb{E}_{\substack{\theta \sim F^T \\ \omega \sim \prod_{t=1}^T S_{t,\theta_t}}} \left[\sum_{t=1}^T v(\theta_t, \omega_t) \right] \\
&\geq 0\,,
\end{aligned}$$

where the second step follows because

$$v(\hat{\theta}, (\hat{\omega}^0, \hat{p} + q_{t,\hat{\theta}})) = v^0(\hat{\theta}, \hat{\omega}^0) - \beta(\hat{p} + q_{t,\hat{\theta}}) = v(\hat{\theta}, (\hat{\omega}^0, \hat{p})) - \beta q_{t,\hat{\theta}}$$

and the payment offset can be separated out, the third step follows by substituting in the payment offsets, and the last step is because $\{S_t\}$ is a feasible solution to (TR-16). Hence, $\{S'_t\}$ is a feasible solution to (TR-17).

Furthermore, note that

$$\begin{aligned}
\mathbb{E}_{\substack{\theta \sim F^T \\ \omega \sim \prod_{t=1}^T S'_{t,\theta_t}}} \left[\sum_{t=1}^T u(\theta_t, \omega_t) \right] &= \sum_{t=1}^T \mathbb{E}_{\substack{\hat{\theta} \sim (F^T)_t \\ \hat{\omega} \sim S'_{t,\hat{\theta}}}} \left[u(\hat{\theta}, \hat{\omega}) \right] \\
&= \sum_{t=1}^T \mathbb{E}_{\substack{\hat{\theta} \sim (F^T)_t \\ (\hat{\omega}^0, \hat{p}) \sim S_{t,\hat{\theta}}}} \left[u(\hat{\theta}, (\hat{\omega}^0, \hat{p} + q_{t,\hat{\theta}})) \right] \\
&= \sum_{t=1}^T \mathbb{E}_{\substack{\hat{\theta} \sim (F^T)_t \\ (\hat{\omega}^0, \hat{p}) \sim S_{t,\hat{\theta}}}} \left[u(\hat{\theta}, (\hat{\omega}^0, \hat{p})) \right] + \alpha \cdot \sum_{t=1}^T \mathbb{E}_{\substack{\hat{\theta} \sim (F^T)_t \\ (\hat{\omega}^0, \hat{p}) \sim S_{t,\hat{\theta}}}} [q_{t,\hat{\theta}}] \\
&= \sum_{t=1}^T \mathbb{E}_{\substack{\hat{\theta} \sim (F^T)_t \\ \hat{\omega} \sim S_{t,\hat{\theta}}}} \left[u(\hat{\theta}, \hat{\omega}) \right] + \alpha \cdot \sum_{t=1}^T \mathbb{E}_{\hat{\theta} \sim (F^T)_t} [q_{t,\hat{\theta}}] \\
&= \mathbb{E}_{\substack{\theta \sim F^T \\ \omega \sim \prod_{t=1}^T S_{t,\theta_t}}} \left[\sum_{t=1}^T u(\theta_t, \omega_t) \right] + \alpha \cdot \mathbb{E}_{\theta \sim F^T} \left[\sum_t q_{t,\theta_t} \right] \\
&= \mathbb{E}_{\substack{\theta \sim F^T \\ \omega \sim \prod_{t=1}^T S_{t,\theta_t}}} \left[\sum_{t=1}^T u(\theta_t, \omega_t) \right]\,,
\end{aligned}$$

where the third step follows by separating out the payment offset and the last step follows because $\mathbb{E}_{\theta \sim F^T} \left[\sum_t q_{t,\theta_t} \right] = 0$. It follows that $\{S'_t\}$ is a feasible solution to (TR-17) that achieves the same objective value as $\{S_t\}$ in (TR-16).

As $\{S_t\}$ was arbitrarily chosen, it follows that (TR-16) $\leq \sum_{t=1}^{T} \bar{u}((F^T)_t)$. Since F^T was arbitrary, the proposition follows.

Part 2: Assume the utility function of the agent is always nonnegative. Let $F^T \in \mathcal{F}^T$ be an arbitrary T-round shock process. Any feasible solution to (TR-16) satisfies the nonnegativity constraint in each round, because the agent's utility function is nonnegative, and, hence, is a feasible solution to (TR-17). In addition, note that the objective functions of (TR-16) and (TR-17) are identical. Therefore, (TR-16) $\leq$ (TR-17) for any $F^T \in \mathcal{F}^T$ and the proposition follows. □

B.3 Multiplicative Guarantees

Suppose there exist some constants $0 < L < U < \infty$ such that $\mathrm{OPT}(\delta_\theta, 1) \in [L, U]$ for all $\theta \in \Theta$ and $\mathrm{OPT}(F, T) \in [LT, UT]$ for all distributions $F \in \mathcal{F}$. Recall that $\mathrm{Ratio}(T) \in [0, 1]$ and $\sup_{S \in \mathcal{S}^{\times 1}} \inf_{\theta \in \Theta} \mathrm{Ratio}(S, \delta_\theta, 1) \in [0, 1]$. We introduce a parametrized notion of regret suitable for our analysis of the multiplicative guarantee. For any $\lambda \in [0, 1]$, we define

$$\mathrm{Regret}(T, \lambda) := \inf_{A \in \mathcal{A}} \sup_{F \in \mathcal{F}} \mathrm{Regret}(A, F, T, \lambda)\,,$$

where

$$\mathrm{Regret}(A, F, T, \lambda) := \lambda \cdot \mathrm{OPT}(F, T) - \mathrm{PrincipalUtility}(A, F, T)\,,$$

for an incentive compatible mechanism A. Note $\mathrm{Regret}(T, \lambda)$ is monotonically increasing in λ.

We first prove the following proposition relating the multiplicative guarantees and the parametrized regret notion.

Proposition TR.10. *We have the following relations:*

1. $\mathrm{Ratio}(T) = \sup\{\lambda \in [0, 1] \mid \inf_{A \in \mathcal{A}} \sup_{F \in \mathcal{F}} \mathrm{Regret}(A, F, T, \lambda) \leq 0\}$.
2. $\sup_{S \in \mathcal{S}^{\times 1}} \inf_{\theta \in \Theta} \mathrm{Ratio}(S, \delta_\theta, 1) = \sup\{\lambda \in [0, 1] \mid \inf_{S \in \mathcal{S}^{\times 1}} \sup_{\theta \in \Theta} \mathrm{Regret}(S, \delta_\theta, 1, \lambda) \leq 0\}$.

Proof. We prove the first relation. The second relation follows similarly and we omit the proof.

Note that by the definition of infimum, for any $\epsilon > 0$, there exists an incentive compatible dynamic mechanism A_ϵ such that $\inf_{F \in \mathcal{F}} \mathrm{Ratio}(A_\epsilon, F, T) \geq \mathrm{Ratio(T)} - \epsilon$, or

$$\frac{\mathrm{PrincipalUtility}(A_\epsilon, F, T)}{\mathrm{OPT}(F, T)} \geq \mathrm{Ratio(T)} - \epsilon\,,$$

for all distributions $F \in \mathcal{F}$. The above expression can be rearranged:

$$\epsilon \cdot \mathrm{OPT}(F, T) \geq \mathrm{Ratio}(T) \cdot \mathrm{OPT}(F, T) - \mathrm{PrincipalUtility}(A_\epsilon, F, T)\,.$$

Since $\text{OPT}(F,T) \le UT$, we then have

$$\begin{aligned}\text{Regret}(T,\text{Ratio}(T)) &\le \sup_{F\in\mathcal{F}} \{\text{Ratio}(T)\cdot\text{OPT}(F,T) - \text{PrincipalUtility}(A_\epsilon,F,T)\}\\ &\le \epsilon UT\,.\end{aligned}$$

As $\epsilon > 0$ was arbitrary, it follows that $\text{Regret}(T,\text{Ratio}(T)) \le 0$ and $\text{Ratio}(T)$ is in the set $\{\lambda \in [0,1] \mid \text{Regret}(T,\lambda) \le 0\}$.

We now show there cannot be a $\lambda > \text{Ratio}(T)$ in the set because reversing the above reasoning leads to a contradiction. Assume there exists $\lambda > \text{Ratio}(T)$ for which

$$\inf_{A\in\mathcal{A}} \sup_{F\in\mathcal{F}} \{\lambda\cdot\text{OPT}(F,T) - \text{PrincipalUtility}(A,F,T)\} \le 0\,.$$

Then, for any $\epsilon > 0$, there exists an incentive compatible dynamic mechanism A_ϵ for which the inner supremum is at most ϵLT, or

$$\epsilon LT \ge \lambda\cdot\text{OPT}(F,T) - \text{PrincipalUtility}(A_\epsilon,F,T)\,,$$

for any distribution $F \in \mathcal{F}$. After rearranging and using $\text{OPT}(F,T) \ge LT$, we have

$$\text{Ratio}(A_\epsilon,F,T) = \frac{\text{PrincipalUtility}(A_\epsilon,F,T)}{\text{OPT}(F,T)} \ge \lambda - \frac{\epsilon LT}{\text{OPT}(F,T)} \ge \lambda - \epsilon\,,$$

for any distribution $F \in \mathcal{F}$. Taking the infimum over $F \in \mathcal{F}$ on the leftmost expression, we obtain

$$\inf_{F\in\mathcal{F}} \text{Ratio}(A_\epsilon,F,T) \ge \lambda - \epsilon\,,$$

and as ϵ was arbitrary, this would mean $\text{Ratio}(T) \ge \lambda$, a contradiction. Hence, the first relation follows. □

We now prove Theorem TR.5:

Proof of Theorem TR.5. Similar to the proof of Theorem 1 of the main paper, we prove using analogues of Lemmas 2 and 3 of the main paper in terms of the parametrized regret notions. Since λ is a scalar multiplier in front of the benchmarks, all the propositions and lemmas used to prove these analogues still apply with the same proofs and the analogues of Lemmas 2 and 3 of the main paper follow with little changes in their proofs but with the multiplier λ. Omitting the proofs, we state and use the following analogues of Lemmas 2 and 3 of the main paper. For the lower bound, for any $\lambda \in [0,1]$ and incentive compatible dynamic mechanism $A \in \mathcal{A}$, there exists a single-round direct IC/IR mechanism $S \in \mathcal{S}^{\times 1}$ such that

$$\sup_{F\in\mathcal{F}} \text{Regret}(A,F,T,\lambda) \ge T\cdot\sup_{\theta\in\Theta}\text{Regret}(S,\delta_\theta,1,\lambda)\,. \tag{TR-18}$$

For the upper bound, for any $\lambda \in [0,1]$ and single-round direct IC/IR mechanism $S \in \mathcal{S}^{\times 1}$,

$$\sup_{F\in\mathcal{F}} \text{Regret}(S^{\times T},F,T,\lambda) \le T\cdot\sup_{\theta\in\Theta}\text{Regret}(S,\delta_\theta,1,\lambda)\,. \tag{TR-19}$$

(First Part): Fix an arbitrary $\lambda \in [0,1]$. Taking the infimum over all single-round direct IC/IR mechanisms S on the right-hand side of the lower bound (TR-18), we have for any incentive compatible dynamic mechanism A,

$$\sup_{F\in\mathcal{F}} \text{Regret}(A, F, T, \lambda) \geq T \cdot \inf_{S\in\mathcal{S}^{\times 1}} \sup_{\theta\in\Theta} \text{Regret}(S, \delta_\theta, 1, \lambda)\,.$$

Taking the infimum over all incentive compatible dynamic mechanisms A on the left-hand side of the above, we obtain

$$\text{Regret}(T, \lambda) = \inf_{A\in\mathcal{A}} \sup_{F\in\mathcal{F}} \text{Regret}(A, F, T, \lambda) \geq T \cdot \inf_{S\in\mathcal{S}^{\times 1}} \sup_{\theta\in\Theta} \text{Regret}(S, \delta_\theta, 1, \lambda)\,.$$

Then, for any $\lambda \in [0,1]$,

$$\text{Regret}(T, \lambda) \geq T \cdot \inf_{S\in\mathcal{S}^{\times 1}} \sup_{\theta\in\Theta} \text{Regret}(S, \delta_\theta, 1, \lambda)\,.$$

This implies whenever $\text{Regret}(T, \lambda) \leq 0$, we have $\inf_{S\in\mathcal{S}^{\times 1}} \sup_{\theta\in\Theta} \text{Regret}(S, \delta_\theta, 1, \lambda) \leq 0$. By Proposition TR.10, it follows that $\text{Ratio}(T) \leq \sup_{S\in\mathcal{S}^{\times 1}} \inf_{\theta\in\Theta} \text{Ratio}(S, \delta_\theta, 1)$.

Again, fix an arbitrary $\lambda \in [0,1]$. Let $\epsilon > 0$ be arbitrary and S be a single-round direct IC/IR mechanism satisfying

$$\sup_{\theta\in\Theta} \text{Regret}(S, \delta_\theta, 1, \lambda) \leq \inf_{S\in\mathcal{S}^{\times 1}} \sup_{\theta\in\Theta} \text{Regret}(S, \delta_\theta, 1, \lambda) + \frac{\epsilon}{T}\,.$$

Note that such mechanism S exists by the definition of infimum. By the upper bound (TR-19) and the property of S,

$$\sup_{F\in\mathcal{F}} \text{Regret}(S^{\times T}, F, T, \lambda) \leq T \cdot \sup_{\theta\in\Theta} \text{Regret}(S, \delta_\theta, 1, \lambda) \leq T \cdot \inf_{S\in\mathcal{S}^{\times 1}} \sup_{\theta\in\Theta} \text{Regret}(S, \delta_\theta, 1, \lambda) + \epsilon\,.$$

Since $S^{\times T}$ is incentive compatible by Lemma A.2 of the main paper, it follows that

$$\text{Regret}(T, \lambda) \leq \sup_{F\in\mathcal{F}} \text{Regret}(S^{\times T}, F, T, \lambda) \leq T \cdot \inf_{S\in\mathcal{S}^{\times 1}} \sup_{\theta\in\Theta} \text{Regret}(S, \delta_\theta, 1, \lambda) + \epsilon\,.$$

As $\epsilon > 0$ was arbitrary and can be made arbitrarily small, it follows that

$$\text{Regret}(T, \lambda) \leq T \cdot \inf_{S\in\mathcal{S}^{\times 1}} \sup_{\theta\in\Theta} \text{Regret}(S, \delta_\theta, 1, \lambda)\,.$$

As $\lambda \in [0,1]$ was arbitrary, the above holds for any $\lambda \in [0,1]$. In particular, $\inf_{S\in\mathcal{S}^{\times 1}} \sup_{\theta\in\Theta} \text{Regret}(S, \delta_\theta, 1, \lambda) \leq 0$ implies $\text{Regret}(T, \lambda) \leq 0$. By Proposition TR.10, we have $\text{Ratio}(T) \geq \sup_{S\in\mathcal{S}^{\times 1}} \inf_{\theta\in\Theta} \text{Ratio}(S, \delta_\theta, 1)$. Combining with the earlier observation that $\text{Ratio}(T) \leq \sup_{S\in\mathcal{S}^{\times 1}} \inf_{\theta\in\Theta} \text{Ratio}(S, \delta_\theta, 1)$, we have the first part.

(Second Part): For any $\epsilon \geq 0$, let S be a single-round direct IC/IR mechanism satisfying

$$\inf_{\theta\in\Theta} \text{Ratio}(S, \delta_\theta, 1) \geq \sup_{S'\in\mathcal{S}^{\times 1}} \inf_{\theta'\in\Theta} \text{Ratio}(S', \delta_{\theta'}, 1) - \frac{\epsilon L}{U}\,.$$

Then, for any $\theta \in \Theta$, we have $\frac{\text{PrincipalUtility}(S,\delta_\theta,1)}{\text{OPT}(\delta_\theta,1)} \geq \sup_{S'\in\mathcal{S}^{\times 1}} \inf_{\theta'\in\Theta} \text{Ratio}(S', \delta_{\theta'}, 1) - \frac{\epsilon L}{U}$. After rearranging terms and upper bounding $\mathrm{OPT}(\delta_\theta, 1)$ by U, we obtain that for any $\theta \in \Theta$,

$$\sup_{S' \in \mathcal{S}^{\times 1}} \inf_{\theta' \in \Theta} \mathrm{Ratio}(S', \delta_{\theta'}, 1) \cdot \mathrm{OPT}(\delta_\theta, 1) - \mathrm{PrincipalUtility}(S, \delta_\theta, 1) \leq \frac{\epsilon L}{U} \cdot \mathrm{OPT}(\delta_\theta, 1) \leq \epsilon L\,.$$

Then, $\sup_{\theta \in \Theta} \mathrm{Regret}(S, \delta_\theta, 1, \sup_{S' \in \mathcal{S}^{\times 1}} \inf_{\theta' \in \Theta} \mathrm{Ratio}(S', \delta_{\theta'}, 1)) \leq \epsilon L$ and by the upper bound (TR-19) for $\lambda = \sup_{S' \in \mathcal{S}^{\times 1}} \inf_{\theta' \in \Theta} \mathrm{Ratio}(S', \delta_{\theta'}, 1)$,

$$\sup_{F \in \mathcal{F}} \mathrm{Regret}(S^{\times T}, F, T, \lambda) \leq T \cdot \sup_{\theta \in \Theta} \mathrm{Regret}(S, \delta_\theta, 1, \lambda) \leq \epsilon L T\,.$$

Substituting $\mathrm{Ratio}(T) = \sup_{S' \in \mathcal{S}^{\times 1}} \inf_{\theta' \in \Theta} \mathrm{Ratio}(S', \delta_{\theta'}, 1)$ which is by Part 1 into the leftmost expression, we obtain for any $F \in \mathcal{F}$,

$$\mathrm{Ratio}(T) \cdot \mathrm{OPT}(F, T) - \mathrm{PrincipalUtility}(S^{\times T}, F, T) \leq \epsilon L T\,.$$

After rearranging terms and lower bounding $\mathrm{OPT}(F, T)$ by LT, we have that for any $F \in \mathcal{F}$,

$$\frac{\mathrm{PrincipalUtility}(S^{\times T}, F, T)}{\mathrm{OPT}(F, T)} \geq \mathrm{Ratio}(T) - \frac{\epsilon L T}{\mathrm{OPT}(F, T)} \geq \mathrm{Ratio}(T) - \epsilon\,.$$

Taking the infimum of the leftmost expression over $F \in \mathcal{F}$, we obtain

$$\inf_{F \in \mathcal{F}} \mathrm{Ratio}(S^{\times T}, F, T) \geq \mathrm{Ratio}(T) - \epsilon\,.$$

(Third Part): Equivalently, we show that $\arg\max_{A \in \mathcal{A}} \inf_{F \in \mathcal{F}} \mathrm{Ratio}(A, F, T)$ is non-empty if and only if $\arg\max_{S \in \mathcal{S}^{\times 1}} \inf_{\theta \in \Theta} \mathrm{Ratio}(S, \delta_\theta, 1)$ is non-empty. For the if direction, assume there exists an optimal single-round direct IC/IR mechanism S^* in the single-round problem. The optimal mechanism S^* satisfies

$$\inf_{\theta \in \Theta} \mathrm{Ratio}(S^*, \delta_\theta, 1) \geq \sup_{S \in \mathcal{S}^{\times 1}} \inf_{\theta \in \Theta} \mathrm{Ratio}(S, \delta_\theta, 1)\,.$$

Then, by Part 2 with $\epsilon = 0$, it follows that

$$\inf_{F \in \mathcal{F}} \mathrm{Ratio}((S^*)^{\times T}, F, T) \geq \mathrm{Ratio}(T)\,.$$

This implies that the static mechanism $(S^*)^{\times T}$ that repeats S^* is optimal in the multi-round problem. Hence, an optimal dynamic mechanism exists in the multi-round problem.

For the only-if direction, assume there exists an optimal incentive compatible dynamic mechanism A^* in the multi-round problem. In particular, $\inf_{F \in \mathcal{F}} \mathrm{Ratio}(A^*, F, T) \geq \mathrm{Ratio}(T)$. Then, for any distribution $F \in \mathcal{F}$, we have $\frac{\mathrm{PrincipalUtility}(A^*, F, T)}{\mathrm{OPT}(F, T)} \geq \mathrm{Ratio}(T)$, or

$$\mathrm{Ratio}(T) \cdot \mathrm{OPT}(F, T) - \mathrm{PrincipalUtility}(A^*, F, T) \leq 0\,.$$

Taking the supremum over $F \in \mathcal{F}$ on the left-hand side, we have

$$\sup_{F \in \mathcal{F}} \mathrm{Regret}(A^*, F, T, \mathrm{Ratio}(T)) \leq 0\,.$$

By the lower bound (TR-18) for $\lambda = \mathrm{Ratio}(T)$, there exists a single-round direct IC/IR mechanism

S such that

$$\sup_{\theta \in \Theta} \text{Regret}(S, \delta_\theta, 1, \text{Ratio}(T)) \leq \frac{1}{T} \sup_{F \in \mathcal{F}} \text{Regret}(A^*, F, T, \text{Ratio}(T)) \leq 0\,.$$

Since $\text{Ratio}(T) = \sup_{S' \in \mathcal{S}^{\times 1}} \inf_{\theta' \in \Theta} \text{Ratio}(S', \delta_{\theta'}, 1)$ by Part 1,

$$\sup_{\theta \in \Theta} \text{Regret}(S, \delta_\theta, 1, \sup_{S' \in \mathcal{S}^{\times 1}} \inf_{\theta' \in \Theta} \text{Ratio}(S', \delta_{\theta'}, 1)) \leq 0\,.$$

This implies that for any $\theta \in \Theta$,

$$\sup_{S' \in \mathcal{S}^{\times 1}} \inf_{\theta' \in \Theta} \text{Ratio}(S', \delta_{\theta'}, 1) \cdot \text{OPT}(\delta_\theta, 1) - \text{PrincipalUtility}(S, \delta_\theta, 1) \leq 0\,.$$

Rearranging terms, we obtain

$$\sup_{S' \in \mathcal{S}^{\times 1}} \inf_{\theta' \in \Theta} \text{Ratio}(S', \delta_{\theta'}, 1) \leq \text{Ratio}(S, \delta_\theta, 1)\,.$$

Taking the infimum over $\theta \in \Theta$ on the right-hand side, we obtain

$$\sup_{S' \in \mathcal{S}^{\times 1}} \inf_{\theta' \in \Theta} \text{Ratio}(S', \delta_{\theta'}, 1) \leq \inf_{\theta \in \Theta} \text{Ratio}(S, \delta_\theta, 1)\,.$$

Then, S is an optimal single-round direct IC/IR mechanism in the single-round problem and the result follows. For S, we can take the single-round direct IC/IR mechanism constructed from A^* via the method described in the proof of Lemma A.1 of the main paper. The same single-round mechanism satisfies the lower bound (TR-18). □

We prove the analogue of Lemma 1 of the main paper below. Note the following is the analogue of (3) of the main paper for the multiplicative guarantee:

$$\sup_{S \in \Delta(\Omega)^\Theta} \inf_{\theta \in \Theta} \frac{\int_\Omega u(\theta, \omega) dS_\theta(\omega)}{\text{OPT}(\delta_\theta, 1)} \quad \text{s.t.} \quad (\text{IC}), (\text{IR})\,, \tag{TR-20}$$

with the same incentive compatibility and individual rationality constraints.

Lemma TR.1. *The optimization problems* (TR-9) *and* (TR-20) *attain the same objective value. Moreover, a single-round direct mechanism* S^* *with decision rule* $\pi_1 : \Theta \cup \{\mathsf{PASS}\} \times [0,1] \to \Omega$ *is an optimal solution of* (TR-9) *if and only if its outcome distributions* $S^*_\theta(W) := \mathbb{P}_{z \sim \text{Uniform}(0,1)}(\pi_1(\theta, z) \in W)$ *for* $\theta \in \Theta, W \subseteq \Omega$ *are an optimal solution of* (TR-20)*. Finally, the objective of* (TR-20) *can be equivalently replaced with* $\inf_{F \in \Delta(\Theta)} \frac{\int_\Theta \int_\Omega u(\theta,\omega) dS_\theta(\omega) dF(\theta)}{\int_\Theta \text{OPT}(\delta_\theta, 1) dF(\theta)}$.

Proof. Note for any single-round direct mechanism S (and the recommended strategy of truthful reporting for the agent) and any point-mass distribution δ_θ, we have $\text{Ratio}(S, \delta_\theta, 1) = \frac{\int_\Omega u(\theta,\omega) dS_\theta(\omega)}{\text{OPT}(\delta_\theta, 1)}$. To prove the first and second statements, it suffices to show that a single-round direct mechanism S is incentive compatible (i.e., in the set $\mathcal{S}^{\times 1}$) if and only if its outcome distributions S_θ satisfy the IC and IR constraints. The same reasoning from the proof of Lemma 1 of the main paper works.

For the third part of the lemma, we show the original objective of (TR-20) and the alternative objective lead to the same value. Fix an arbitrary single-round direct mechanism $S \in \Delta(\Omega)^\Theta$. We have

$$\inf_{\theta'\in\Theta} \frac{\int_\Omega u(\theta',\omega)dS_{\theta'}(\omega)}{\mathrm{OPT}(\delta_{\theta'},1)} = \inf_{\theta'\in\Theta} \frac{\int_\Theta\int_\Omega u(\theta,\omega)dS_\theta(\omega)d\delta_{\theta'}(\theta)}{\int_\Theta \mathrm{OPT}(\delta_\theta,1)d\delta_{\theta'}(\theta)}$$
$$\le \inf_{F\in\Delta(\Theta)} \frac{\int_\Theta\int_\Omega u(\theta,\omega)dS_\theta(\omega)dF(\theta)}{\int_\Theta \mathrm{OPT}(\delta_\theta,1)dF(\theta)},$$

where the first step is by rewriting the denominator and numerator and the second step is because point-mass distributions are a subset of all probability distributions supported on Θ, $\Delta(\Theta)$. For the other direction, we note that

$$\inf_{F\in\Delta(\Theta)} \frac{\int_\Theta\int_\Omega u(\theta,\omega)\mathrm{d}S_\theta(\omega)\mathrm{d}F(\theta)}{\int_\Theta \mathrm{OPT}(\delta_\theta,1)\mathrm{d}F(\theta)} = \inf_{F\in\Delta(\Theta)} \frac{\int_\Theta \frac{\int_\Theta u(\theta,\omega)\mathrm{d}S_\theta(\omega)}{\mathrm{OPT}(\delta_\theta,1)}\cdot \mathrm{OPT}(\delta_\theta,1)\mathrm{d}F(\theta)}{\int_\Theta \mathrm{OPT}(\delta_{\theta'},1)\mathrm{d}F(\theta')}$$
$$\ge \inf_{F\in\Delta(\Theta)} \frac{\int_\Theta \inf_{\theta''\in\Theta}\frac{\int_\Omega u(\theta'',\omega)\mathrm{d}S_{\theta''}(\omega)}{\mathrm{OPT}(\delta_{\theta''},1)}\cdot \mathrm{OPT}(\delta_\theta,1)\mathrm{d}F(\theta)}{\int_\Theta \mathrm{OPT}(\delta_{\theta'},1)\mathrm{d}F(\theta')}$$
$$= \inf_{\theta''\in\Theta} \frac{\int_\Omega u(\theta'',\omega)\mathrm{d}S_{\theta''}(\omega)}{\mathrm{OPT}(\delta_{\theta''},1)},$$

where the first step is by rewriting the numerator, the second step is by lower bounding the fraction inside the integral in the numerator, and the last step follows regardless of F.

As S was arbitrary, it follows that the original objective and alternative objective achieve the same value for $S \in \Delta(\Omega)^\Theta$. Therefore, the optimization problem (TR-20) and the version with the alternative objective are equivalent in terms of the optimal value and optimal solutions. □

B.4 Principal Pessimism

B.4.1 Proof of Theorem TR.6

We can prove Theorem TR.6 by the same reasoning as in the proof of Theorem 1 of the main paper using Lemmas 2 and 3 of the main paper with respect to the stronger notions of regret and minimax regret and single-round direct IC/IR/PP mechanisms. In particular, Lemma 2 of the main paper will show the existence of a single-round direct IC/IR/PP mechanism for any (not necessarily incentive compatible) dynamic mechanism and Lemma 3 of the main paper will hold for any single-round direct IC/IR/PP mechanism. These lemmas hold with almost identical proofs with modifications to account for that the agent plays a utility-maximizing strategy that minimizes the principal utility among utility-maximizing strategies. Below, we sketch the proofs of Lemmas 2 and 3 of the main paper with respect to the stronger notions and single-round direct IC/IR/PP mechanisms.

Analogue of Lemma 2 of the Main Paper. In the proof of Lemma 2 of the main paper with respect to the original notions of regret and minimax regret, we use Lemma A.1 of the main paper. Instead of single-round direct IC/IR mechanisms, we can prove the lemma for single-round direct IC/IR/PP mechanisms. Then, Lemma 2 of the main paper would follow in terms of single-round direct IC/IR/PP mechanisms with respect to the stronger notions of regret and minimax regret by the same proof using the modified version of the lemma.

For Lemma A.1 of the main paper, the same construction provided in its proof yields a single-round direct IC/IR/PP mechanism S satisfying the lemma statement for any given dynamic mechanism $A \in \mathcal{A}$ (not necessarily incentive compatible). That is, we construct $S = \{S_\theta\}_{\theta \in \Theta}$ from the random sequences of outcomes when the principal commits to a dynamic mechanism $A \in \mathcal{A}$, the agent plays $\sigma^*(A, T)$, and the agent's distribution is a point-mass distribution δ_θ for $\theta \in \Theta$. As in the proof of Lemma A.1 of the main paper, the IC/IR constraints follow from that $\sigma^*(A, T)$ is a utility-maximizing strategy for the agent when his shock distribution is δ_θ and guarantees the agent utility of at least 0.

For the PP constraint, we prove by contradiction. For the first part, suppose there exist θ and θ' such that $\mathbb{E}_{\omega \sim S_\theta}[v(\theta, \omega)] = \mathbb{E}_{\omega \sim S_{\theta'}}[v(\theta, \omega)]$ and $\mathbb{E}_{\omega \sim S_\theta}[u(\theta, \omega)] > \mathbb{E}_{\omega \sim S_{\theta'}}[u(\theta, \omega)]$. Then, the agent can implement the strategy $\sigma^*(A, T)$ as if his distribution is $\delta_{\theta'}$ when his actual distribution is δ_θ in the multi-round problem. Let σ' denote this modified strategy that coincides with $\sigma^*(A, T)$ for other distributions than δ_θ. By construction, σ' is a utility-maximizing strategy for the agent. Then, for point-mass distribution δ_θ, the agent can obtain the utility of

$$\begin{aligned}
\text{AgentUtility}(A, \sigma', \delta_\theta, T) &= T \cdot \mathbb{E}_{\omega \sim S_{\theta'}}[v(\theta, \omega)] \\
&= T \cdot \mathbb{E}_{\omega \sim S_\theta}[v(\theta, \omega)] \\
&= \text{AgentUtility}(A, \sigma^*(A, T), \delta_\theta, T)\,,
\end{aligned}$$

which is the same utility as under $\sigma^*(A, T)$. At the same time, the principal utility will be strictly lower:

$$\begin{aligned}
\text{PrincipalUtility}(A, \sigma', \delta_\theta, T) &= T \cdot \mathbb{E}_{\omega \sim S_{\theta'}}[u(\theta, \omega)] \\
&< T \cdot \mathbb{E}_{\omega \sim S_\theta}[u(\theta, \omega)] \\
&= \text{PrincipalUtility}(A, \sigma^*(A, T), \delta_\theta, T)\,.
\end{aligned}$$

This contradicts that $\sigma^*(A, T)$ is a principal-pessimistic utility-maximizing strategy when the agent's distribution is δ_θ.

For the second part of the PP constraint, suppose there exists θ such that $\mathbb{E}_{\omega \sim S_\theta}[v(\theta, \omega)] = 0$ and $\mathbb{E}_{\omega \sim S_\theta}[u(\theta, \omega)] > 0$. When his distribution is δ_θ, the agent can choose to not participate and obtain the same agent utility of 0 as under $\sigma^*(A, T)$, and the principal utility will be 0 which is strictly worse than that achieved under $\sigma^*(A, T)$. As above, we can construct an alternative utility-maximizing strategy based on this observation and contradict the choice of $\sigma^*(A, T)$ as a principal-pessimistic utility-maximizing strategy. Hence, the singe-round direct mechanism constructed is IC/IR/PP and the rest of the proof of Lemma A.1 of the main paper follows.

Analogue of Lemma 3 of the Main Paper. For Lemma 3 of the main paper with respect to the original notions of regret and minimax regret, we use Lemmas A.2 and A.3 of the main paper in its proof. Lemma A.3 of the main paper with the same proof still holds with respect to the stronger notions of regret and minimax regret. For Lemma A.2 of the main paper, we adapt the same proof but account for that a best-response strategy for the agent is a utility-maximizing strategy that also minimizes the principal utility among utility-maximizing strategies. We provide more details below. With these lemmas, Lemma 3 of the main paper would then follow in terms of single-round direct IC/IR/PP mechanisms with respect to the stronger notions of regret and minimax regret.

We now provide more details for proving Lemma A.2 of the main paper for the stronger regret notion.

Let $S \in \mathcal{S}^{\times 1}$ be any single-round direct IC/IR/PP mechanism with decision rule $\tilde{\pi}$ and $S^{\times T}$ be the mechanism that repeats S over T rounds. The statement that $\text{PrincipalUtility}(S^{\times T}, \sigma^{\text{TR}}, F, T) = T \cdot \text{PrincipalUtility}(S, \sigma^{\text{TR}}, F, 1)$ follows from the same reasoning in the proof of Lemma A.2 of the main paper. We argue that $S^{\times T}$ is indeed in the set $\mathcal{S}^{\times T}$, i.e., that truthful reporting is a principal-pessimistic utility-maximizing strategy. This means we can take $\sigma^*(S^{\times T}, T)$ to be truthful reporting without loss in terms of the principal utility, agent utility and regret. In particular, we can use

$$\sup_{F \in \mathcal{F}} \text{Regret}(S^{\times T}, F, T) = \sup_{F \in \mathcal{F}} \left\{ \text{OPT}(F, T) - \text{PrincipalUtility}(S^{\times T}, \sigma^{\text{TR}}, F, T) \right\}$$

in proving Lemma 3 of the main paper for the stronger regret notion.

To see $S^{\times T} \in \mathcal{S}^{\times T}$, we note that truthful reporting is a utility-maximizing strategy for the agent, because S satisfies the IC/IR constraints and Lemma A.2 of the main paper applies. For the sake of contradiction, assume there exists a distribution F' and a utility-maximizing strategy σ' for the agent such that $\text{PrincipalUtility}(S^{\times T}, \sigma^{\text{TR}}, F', T) > \text{PrincipalUtility}(S^{\times T}, \sigma', F', T)$. We define per-round expected agent utility V_t and principal utility U_t when the principal implements $S^{\times T}$ and the agent plays σ' as

$$V_t = \mathbb{E}[v(\theta_t, \tilde{\pi}(\sigma'_t(\theta_t, h_t^+)))] \quad \text{and} \quad U_t = \mathbb{E}[u(\theta_t, \tilde{\pi}(\sigma'_t(\theta_t, h_t^+)))]$$

for Rounds $t \in [T]$. Note the principal's mechanism has no dependence on histories while the agent's strategy may depend on the augmented history h_t^+. Since σ' is a utility-maximizing strategy, V_t is equal to the agent utility achieved under the single-round mechanism S and σ^{TR} for all t. If not, there must exist a round in which V_t is greater (say, where the maximum is achieved) and, by the same claim in the proof of Lemma A.2 of the main paper, there exists a single-round strategy that the agent can implement against S (when his distribution is F') and obtain a greater overall utility than the truthful reporting strategy. This would contradict that truthful reporting is a utility-maximizing strategy against S.

Furthermore, since σ' leads to a lower principal utility than σ^{TR}, there must exist a particular round t' in which $U_{t'}$ is strictly less than the principal utility achieved under the single-round mechanism S and σ^{TR}. By the same claim in the proof of Lemma A.2 of the main paper, there exists an alternative agent strategy against the single-round mechanism S that yields the expected agent utility equal to $V_{t'}$ and the expected principal utility equal to $U_{t'}$. Note $V_{t'} = \text{AgentUtility}(S, \sigma^{\text{TR}}, F', 1)$ and $U_{t'} < \text{PrincipalUtility}(S, \sigma^{\text{TR}}, F', 1)$. This contradicts that σ^{TR} is a principal-pessimistic utility-maximizing strategy for the agent against S.

B.4.2 Analogue of Lemma 1 of the Main Paper

In what follows, we state and prove the analogue of Lemma 1 of the main paper for the stronger notion of regret:

Lemma TR.2. *The optimization problems* (TR-10) *and* (TR-11) *attain the same objective value. Moreover, a single-round direct mechanism S^* with decision rule $\pi_1 : \Theta \cup \{\mathsf{PASS}\} \times [0,1] \to \Omega$ is an optimal solution of* (TR-10) *if and only if its outcome distributions $S^*_\theta(W) := \mathbb{P}_{z \sim \text{Uniform}(0,1)}(\pi_1(\theta, z) \in W)$ for $\theta \in \Theta, W \subseteq \Omega$ are an optimal solution of* (TR-11). *Finally, the objective of* (TR-11) *can be equivalently replaced with $\sup_{F \in \Delta(\Theta)} \left\{ \int_\Theta \text{OPT}(\delta_\theta, 1) dF(\theta) - \int_\Theta \int_\Omega u(\theta, \omega) dS_\theta(\omega) dF(\theta) \right\}$.*

Similar to how we analyzed specific applications in the main paper and this technical report, instead of (TR-10), we can apply the above lemma and equivalently consider

$$\inf_{\substack{S\in\Delta(\Omega)^{\Theta}:\\ \text{(IC),(IR),(PP)}}} \sup_{F\in\Delta(\Theta)} \widehat{\text{Regret}}(S,F)\,,$$

where S can be any single-round direct IC/IR/PP mechanism and $\widehat{\text{Regret}}(S,F) := \int_{\Theta} \text{OPT}(\delta_\theta,1)dF(\theta) - \int_{\Theta}\int_{\Omega} u(\theta,\omega)dS_\theta(\omega)dF(\theta)$.

We now prove the above lemma:

Proof of Lemma TR.2. Note for any single-round direct mechanism S and any point-mass distribution δ_θ, we have $\text{Regret}(S,\sigma^{\text{TR}},\delta_\theta,1) = \text{OPT}(\delta_\theta,1) - \int_{\Omega} u(\theta,\omega)dS_\theta(\omega)$. To prove the first and second statements, it suffices to show that a single-round direct mechanism S has truthful reporting as a principal-pessimistic utility-maximizing strategy (i.e., in the set $\mathcal{S}^{\times 1}$) if and only if its outcome distributions S_θ satisfy the IC/IR/PP constraints as formulated in (TR-11). Lemma 1 of the main paper already shows that a single-round direct mechanism S has truthful reporting as a utility-maximizing strategy if and only if its outcome distributions S_θ satisfy the IC/IR constraints. Hence, it remains to show that a single-round direct IC/IR mechanism S has truthful reporting as a principal-pessimistic utility-maximizing strategy if and only if its outcome distributions S_θ satisfy the PP constraint.

For the only-if direction, we proceed as follows. Assume an arbitrary single-round direct IC/IR mechanism S against which truthful reporting σ^{TR} is a principal-pessimistic utility-maximizing strategy for the agent. Recall, under σ^{TR}, the agent reports CONTINUE in Round 0 and then truthfully reports shocks in future rounds. Then, $\text{PrincipalUtility}(S,\sigma^{\text{TR}},F,1) \le \text{PrincipalUtility}(A,\tilde{\sigma},F,1)$ holds for every $F \in \Delta(\Theta)$ and $\tilde{\sigma} \in \mathcal{B}(S,1)$. In particular, the inequality holds for point-mass distribution δ_θ for any $\theta \in \Theta$ and the alternative strategy $\tilde{\sigma}$ which reports shock $\theta' \in \mathcal{B}^*(S,\theta)$ (if it is not empty) when the agent's distribution is δ_θ but truthfully reports for other distributions. Clearly, $\tilde{\sigma}$ is a utility-maximizing strategy for the agent by construction. Then, the inequality reduces to $\int_{\Omega} u(\theta,\omega)\text{d}S_\theta(\omega) \le \int_{\Omega} u(\theta,\omega)\text{d}S_{\theta'}(\omega)$ and the first part of the PP constraint follows. If $\mathcal{B}^*(S,\theta)$ is empty for some θ, the first part holds trivially for θ.

The same inequality $\text{PrincipalUtility}(S,\sigma^{\text{TR}},F,1) \le \text{PrincipalUtility}(A,\tilde{\sigma},F,1)$ holds for point-mass distribution δ_θ for any $\theta \in \Theta_0$ (if it is not empty) and the alternative strategy $\tilde{\sigma}$ which reports QUIT in Round 0 when the agent's distribution is δ_θ but reports CONTINUE in Round 0 and shocks truthfully for other distributions. Note that $\tilde{\sigma} \in \mathcal{B}(S,1)$ by construction. Then, we obtain $\int_{\Omega} u(\theta,\omega)\text{d}S_\theta(\omega) \le 0$ and the second part of the PP constraint follows. If Θ_0 is empty, then the second part holds trivially. Therefore the outcome distributions of S satisfy the PP constraint.

We now show the if direction. Assume an arbitrary single-round direct IC/IR mechanism S (with decision rule π_1) with its outcome distributions satisfying the PP constraint in (TR-11). Let σ^{TR} be the strategy of reporting CONTINUE in Round 0 and truthfully reporting in Round 1 for the agent. We consider the following three cases depending on how an arbitrary alternative utility-maximizing strategy $\tilde{\sigma}$ reports in Round 0 for each possible distribution $F \in \Delta(\Theta)$. Fix an arbitrary distribution $F \in \Delta(\Theta)$.

Case 1) $\tilde{\sigma}$ deterministically reports CONTINUE in Round 0

For any θ, we define

$$\tilde{\mathcal{M}}(\theta) = \{\hat{m} \in \Theta \cup \{\mathsf{PASS}\} \mid \mathbb{E}_{\pi_1,\sigma^{\mathrm{TR}}}[v(\theta_1, \pi_1(\theta_1, h_1, z_1))|\theta_1 = \theta] = \mathbb{E}_{\pi_1}[v(\theta_1, \pi_1(m_1, h_1, z_1))|\theta_1 = \theta, m_1 = \hat{m}]\} .$$

Note that $\tilde{\mathcal{M}}(\theta) \subset \mathcal{B}^*(S, \theta) \cup \{\mathsf{PASS}\}$. It is possible that $\mathsf{PASS} \in \tilde{\mathcal{M}}(\theta)$. Since $\tilde{\sigma}$ is a utility-maximizing strategy, its probability distribution of messages for shock θ has a measure of 1 over $\tilde{\mathcal{M}}(\theta)$.

The first part of the PP constraint implies

$$\mathbb{E}_{\pi_1,\sigma^{\mathrm{TR}}}[u(\theta_1, \pi_1(\theta_1, h_1, z_1))|\theta_1 = \theta] \le \mathbb{E}_{\pi_1}[u(\theta_1, \pi_1(m_1, h_1, z_1))|\theta_1 = \theta, m_1 = \hat{m}] ,$$

for any $\theta \in \Theta$ and $\hat{m} \in \mathcal{B}^*(S, \theta)$. If $\mathsf{PASS} \in \tilde{\mathcal{M}}(\theta)$, the second part of the PP constraint implies the same inequality above with the right-hand side equal to 0. Averaging the above inequality over $\hat{m}$ according to $\tilde{\sigma}$ which is potentially randomized, we obtain

$$\mathbb{E}_{\pi_1,\sigma^{\mathrm{TR}}}[u(\theta_1, \pi_1(\theta_1, h_1, z_1))|\theta_1 = \theta] \le \mathbb{E}_{\pi_1,\tilde{\sigma}}[u(\theta_1, \pi_1(m_1, h_1, z_1))|\theta_1 = \theta] ,$$

Since the last inequality holds for each possible value of θ_1, we average it over $\theta_1 \sim F$ and obtain

$$\mathrm{PrincipalUtility}(S, \sigma^{\mathrm{TR}}, F, 1) = \mathbb{E}_{\pi_1,\sigma^{\mathrm{TR}}}[u(\theta_1, \omega_1)] \le \mathbb{E}_{\pi_1,\tilde{\sigma}}[u(\theta_1, \omega_1)] = \mathrm{PrincipalUtility}(S, \tilde{\sigma}, F, 1) .$$

Case 2) $\tilde{\sigma}$ deterministically reports QUIT in Round 0

Since $\tilde{\sigma}$ is a utility-maximizing strategy for the agent, the maximal agent utility is 0 and we have $\mathrm{AgentUtility}(S, \sigma^{\mathrm{TR}}, F, 1) = 0$ because σ^{TR} is also a utility-maximizing strategy. From $\mathrm{AgentUtility}(S, \sigma^{\mathrm{TR}}, F, 1) = \int_\Theta \int_\Omega v(\theta, \omega) dS_\theta(\omega) dF(\theta) = 0$, it follows that the agent's distribution F has a measure of 1 over Θ_0. By the second part of the PP constraint, for all $\theta \in \Theta_0$,

$$\mathbb{E}_{\pi_1,\sigma^{\mathrm{TR}}}[u(\theta_1, \pi_1(\theta_1, h_1, z_1))|\theta_1 = \theta] \le 0 ,$$

where the left-hand side is equal to $\int_\Omega u(\theta, \omega) \mathrm{d}S_\theta(\omega)$. Averaging the above over $\theta_1 \sim F$, we obtain $\mathrm{PrincipalUtility}(S, \sigma^{\mathrm{TR}}, F, 1) = \mathbb{E}_{\pi_1,\sigma^{\mathrm{TR}}}[u(\theta_1, \omega_1)] \le 0$.

Since $\tilde{\sigma}$ reports QUIT in Round 0, there is no interaction between the principal and agent at all and $\mathrm{PrincipalUtility}(S, \tilde{\sigma}, F, 1) = 0$. Clearly, $\mathrm{PrincipalUtility}(S, \sigma^{\mathrm{TR}}, F, 1) \le \mathrm{PrincipalUtility}(S, \tilde{\sigma}, F, 1)$.

Case 3) $\tilde{\sigma}$ probabilistically reports $\mathsf{CONTINUE}$ or QUIT in Round 0

Let $\tilde{\sigma} = \{\tilde{\sigma}_t\}_{0:1}$ where $m_0 = \tilde{\sigma}_0(h_0^+, y_0)$ can be $\mathsf{CONTINUE}$ or QUIT. From the above cases,

$$\mathrm{PrincipalUtility}(S, \sigma^{\mathrm{TR}}, F, 1) \le \mathbb{E}_{\pi_1,\tilde{\sigma}}[u(\theta_1, \omega_1)|m_0 = \mathsf{CONTINUE}] \text{ and}$$
$$\mathrm{PrincipalUtility}(S, \sigma^{\mathrm{TR}}, F, 1) \le \mathbb{E}_{\pi_1,\tilde{\sigma}}[u(\theta_1, \omega_1)|m_0 = \mathsf{QUIT}] .$$

We then have

$$\begin{aligned}\mathrm{PrincipalUtility}(S, \sigma^{\mathrm{TR}}, F, 1) &\le \mathbb{E}_{\pi_1,\tilde{\sigma}}[u(\theta_1, \omega_1)|m_0 = \mathsf{CONT}] \cdot \mathbb{P}(m_0 = \mathsf{CONT}) \\ &\quad + \mathbb{E}_{\pi_1,\tilde{\sigma}}[u(\theta_1, \omega_1)|m_0 = \mathsf{QUIT}] \cdot \mathbb{P}(m_0 = \mathsf{QUIT}) \\ &= \mathrm{PrincipalUtility}(S, \tilde{\sigma}, F, 1) ,\end{aligned}$$

where CONT stands for $\mathsf{CONTINUE}$.

As F was arbitrary, $\text{PrincipalUtility}(S, \sigma^{\text{TR}}, F, 1) \leq \text{PrincipalUtility}(S, \tilde{\sigma}, F, 1)$ in all cases for any distribution F. The above cases cover all possibilities for any alternative utility-maximizing strategy $\tilde{\sigma} \in \mathcal{B}(S, 1)$ and it follows that truthful reporting is a principal-pessimistic utility-maximizing strategy for the agent.

For the last part of the lemma, we follow the same reasoning steps in the proof of Lemma 1 of the main paper. □

B.4.3 Characterizations in Specific Settings

For the applications considered in Section 4 of the main paper and Sections TR.4–TR.5, we show the following characterizations of the PP constraint and that the optimal single-round direct IC/IR mechanisms found for these applications satisfy the PP constraint:

Proposition TR.11. *For revenue maximization in the dynamic selling mechanism design problem with one good, a single-round direct IC/IR mechanism (x, p) satisfies the PP constraint if and only if 1) there exists no $\theta' < \theta \in [0, 1]$ such that $x(\theta') < x(\theta)$ and $x(\hat{\theta}) = x(\theta')$ for all $\hat{\theta} \in [\theta', \theta)$, and 2) $p(\theta) \leq 0$ for all θ such that $\theta \cdot x(\theta) - p(\theta) = 0$.*

Proposition TR.12. *For the principal-agent model with hidden costs, a (deterministic) single-round direct IC/IR mechanism (q, p) satisfies the PP constraint if and only if 1) there exists no interval (θ', θ) where $q(\hat{\theta}) = q_0$ for $\hat{\theta} \in (\theta', \theta)$ for some q_0 and at least one of following sets of conditions holds:*

1. *$q(\theta) < q_0$ and $R(q_0) - \theta \cdot q_0 < R(q(\theta)) - \theta \cdot q(\theta)$*
2. *$q(\theta') > q_0$ and $R(q_0) - \theta \cdot q_0 < R(q(\theta')) - \theta \cdot q(\theta')$,*

and 2) $R(q(\theta)) - p(\theta) \leq 0$ for all θ such that $p(\theta) - \theta \cdot q(\theta) = 0$.

Proposition TR.13. *For the dynamic resource allocation problem without monetary transfers, a single-round direct IC/IR mechanism x satisfies the PP constraint if and only if the probabilistic allocation rule x is constant.*

Except for the last one, the if-and-only-if conditions for PP consist of two parts for the two parts of the PP constraint. The second part of the conditions follows directly from the second part of the PP constraint and is written in terms of the corresponding interim rules in respective problems. The first part of the conditions is more complicated and, intuitively, it says that any flat part of the interim allocation rule $x(\cdot)$, if exists, is closed on the right side in the dynamic selling problem or satisfies some similar condition on both sides in the principal-agent model with hidden costs. In particular, if the interim allocation rule $x(\cdot)$ is continuous, the first part of the conditions will be satisfied in these problems.

For revenue maximization in the dynamic selling problem with a single good, the optimal single-round direct IC/IR mechanism in Proposition 3 of the main paper has a continuous interim allocation rule and the interim rule is such that the principal utility is 0 whenever the agent utility is 0. By Proposition TR.11, the optimal single-round solution is PP. The optimal single-round solution $S^{*,n}$ in Proposition TR.2 in the multiple-goods case satisfies the PP constraint because it uses S^* in

Proposition 3 of the main paper on each good and inherits the same properties from the single-good case.

For welfare maximization in both the single-good and multiple-goods cases, the principal's utility function $u(\cdot, \cdot)$ coincides with the agent's utility function $v(\cdot, \cdot)$ when there is no payment and the realized principal utility is the same for all utility-maximizing strategies for the agent. That is, tie-breaking among utility-maximizing strategies is meaningless and all utility-maximizing strategies are also principal-pessimistic utility-maximizing strategies by default. Hence, the optimal mechanism of allocating items for free in Proposition TR.1 and the same mechanism in the multiple-goods case in Proposition TR.2 satisfy the PP constraint.

Similarly, in the principal-agent model, the optimal single-round direct IC/IR mechanism in Proposition 4 of the main paper is deterministic and has a continuous interim allocation rule. It also satisfies the second part of the if-and-only-if condition. By Proposition TR.12, it satisfies the PP constraint.

For the dynamic resource allocation problem, there are multiple optimal single-round solutions in Proposition TR.4 and, in particular, the constant probabilistic allocation rule $x^*(\theta) = 1 - c$ for all $\theta \in [0, 1]$ is optimal. By Proposition TR.13, it satisfies the PP constraint.

We prove the above propositions below.

Revenue Maximization in the Dynamic Selling Problem with One Good We primarily work with the interim allocation rule $x(\cdot)$ and payment rule $p(\cdot)$. The IC/IR/PP constraints can be equivalently written as

$$\theta \cdot x(\theta) - p(\theta) \geq \theta \cdot x(\theta') - p(\theta'), \quad \forall \theta, \theta' \in [0, 1] \tag{IC}$$

$$\theta \cdot x(\theta) - p(\theta) \geq 0, \quad \forall \theta \in [0, 1] \tag{IR}$$

$$\begin{aligned} &p(\theta) \leq p(\theta'), \quad \forall \theta \in [0, 1], \theta' \in \mathcal{B}^*(\theta) \\ &p(\theta) \leq 0, \quad \forall \theta \in \Theta_0 \end{aligned} \tag{PP}$$

where $\mathcal{B}^*(\theta) = \{\theta' \in \Theta \mid \theta \cdot x(\theta') - p(\theta') = \theta \cdot x(\theta) - p(\theta)\}$ and $\Theta_0 = \{\theta \in \Theta \mid \theta \cdot x(\theta) - p(\theta) = 0\}$. Note if a single-round direct mechanism satisfies the IC/IR constraints, x is non-decreasing and p is given by the payment equivalence formula and, similarly, non-decreasing. Let $V(\theta) = \theta \cdot x(\theta) - p(\theta)$. Note $V(\theta) = \max_{\theta'}\{\theta \cdot x(\theta') - p(\theta')\}$. Equivalently, the utility curve $V(\cdot)$ is supported by lines with slope-intercept pairs $(x(\theta'), -p(\theta'))$ for $\theta' \in [0, 1]$ and, in particular, the line with slope-intercept pair $(x(\theta), -p(\theta))$ goes through the point $(\theta, V(\theta))$.

Proof of Proposition TR.11. It is straightforward to see that the second part of the if-and-only-if condition is equivalent to the second part of the PP constraint in terms of the interim rules. It remains to prove that the first part corresponds to the first part of the PP constraint.

(If part): We prove the contrapositive. Assume the first part of the PP constraint is not satisfied at θ. Note $\theta \neq 0$. Given that IC/IR constraints are satisfied, the first part of the PP constraint not holding at θ means that truthful reporting is not a principal-pessimistic utility-maximizing strategy when the realized value is θ. There exists another utility-maximizing report θ' leading to outcome $(x(\theta'), p(\theta'))$ with a strictly worse principal utility, i.e., $p(\theta') < p(\theta)$. Since the payment rule is non-decreasing, it

must be that $\theta' < \theta$. Since

$$V(\theta) = \theta \cdot x(\theta) - p(\theta) = \theta \cdot x(\theta') - p(\theta'),$$

and $\theta \neq 0$, it follows that $x(\theta') < x(\theta)$. Note the line with slope-intercept pair $(x(\theta'), -p(\theta'))$ goes through $(\theta, V(\theta))$. For this to happen, x has to be constant over the interval $[\theta', \theta)$. If there exists some $\hat{\theta} \in [\theta', \theta)$ at which $x(\hat{\theta}) > x(\theta')$ (this is the only possibility since x is non-decreasing), the line with slope-intercept pair $(x(\hat{\theta}), -p(\hat{\theta}))$ leads to a strictly higher utility for the agent at the realized value θ than the line with slope-intercept pair $(x(\theta'), -p(\theta'))$. This is because the point $(\hat{\theta}, V(\hat{\theta}))$ is above or on the line with slope-intercept pair $(x(\theta'), -p(\theta'))$ and the line with slope-intercept pair $(x(\hat{\theta}), -p(\hat{\theta}))$ goes through the point and has a greater slope. This would contradict that reporting θ' achieves $V(\theta)$ when the realized value is θ. For θ and θ', we have $x(\theta') < x(\theta)$ and x constant over $[\theta', \theta)$.

(Only if part): We similarly prove the contrapositive. Assume there exists a pair $\theta' < \theta$ such that $x(\theta') < x(\theta)$ and $x(\hat{\theta}) = x(\theta')$ for $\hat{\theta} \in [\theta', \theta)$. Since the utility curve $V(\cdot)$ is convex and absolutely continuous, the line with slope-intercept pair $(x(\theta'), -p(\theta'))$ which is tangent to V on $[\theta', \theta)$ contains the point $(\theta, V(\theta))$. This implies

$$V(\theta) = \theta \cdot x(\theta) - p(\theta) = \theta \cdot x(\theta') - p(\theta'),$$

and reporting θ' and truthfully reporting θ lead to the same agent utility when the realized value is θ. By the payment equivalence formula, $p(\hat{\theta}) = p(\theta')$ for $\hat{\theta} \in [\theta', \theta)$ and, in particular, $p(\theta') < p(\theta)$. Then, reporting θ' leads to a strictly worse principal utility. The first part of the PP constraint is not satisfied at θ. □

Principal-Agent Model with Hidden Costs This is after the without-loss restriction of the single-round direct IC/IR mechanisms to those that can be described as a menu of deterministic contracts $\{(q(\theta), p(\theta))\}_{\theta \in [\underline{\theta}, \bar{\theta}]}$. By standard arguments, $q(\cdot)$ is non-increasing and $p(\cdot)$ is given by the payment equivalence formula. See Appendix C.1 of the main paper for details. In this setting, the IC/IR/PP constraints in consideration are

$$p(\theta) - \theta \cdot q(\theta) \geq p(\theta') - \theta \cdot q(\theta'), \quad \forall \theta, \theta' \in [\underline{\theta}, \bar{\theta}] \tag{IC}$$

$$p(\theta) - \theta \cdot q(\theta) \geq 0, \quad \forall \theta \in [\underline{\theta}, \bar{\theta}] \tag{IR}$$

$$\begin{aligned} &R(q(\theta)) - p(\theta) \leq R(q(\theta')) - p(\theta'), \quad \forall \theta \in [\underline{\theta}, \bar{\theta}], \theta' \in \mathcal{B}^*(\theta) \\ &R(q(\theta)) - p(\theta) \leq 0, \quad \forall \theta \in \Theta_0 \end{aligned} \tag{PP}$$

where $\mathcal{B}^*(\theta) = \{\theta' \in \Theta \mid p(\theta') - \theta \cdot q(\theta') = p(\theta) - \theta \cdot q(\theta)\}$ and $\Theta_0 = \{\theta \in \Theta \mid p(\theta) - \theta \cdot q(\theta) = 0\}$. In what follows, let $V(\theta) = p(\theta) - \theta \cdot q(\theta)$ for $\theta \in [\underline{\theta}, \bar{\theta}]$. As in the dynamic selling problem, we can describe the utility curve $V(\cdot)$ as the upper envelope of lines with slope-intercept pairs $(-q(\theta), p(\theta))$ for $\theta \in [\underline{\theta}, \bar{\theta}]$.

Proof of Proposition TR.12. The second part of the if-and-only-if condition is exactly the second part of the PP constraint written in terms of the interim rules. We prove the first part is equivalent to the first part of the PP constraint.

(If part): We prove the contrapositive. Assume the first part of the PP constraint is not satisfied and

reporting some θ' yields the same agent utility as truthful reporting but less principal utility when the realized cost is θ, that is:

$$p(\theta') - \theta \cdot q(\theta') = p(\theta) - \theta \cdot q(\theta) \quad \text{and}$$
$$R(q(\theta')) - p(\theta') < R(q(\theta)) - p(\theta)\,.$$

Combining the two relations, we obtain

$$R(q(\theta')) - \theta \cdot q(\theta') < R(q(\theta)) - \theta \cdot q(\theta)\,. \tag{TR-21}$$

In particular, this implies $q(\theta) \neq \bar{q}(\theta) = \arg\max_{q\geq 0}\{R(q) - \theta \cdot q\}$. There are two cases depending on how $q(\theta)$ and $\bar{q}(\theta)$ compare.

If $q(\theta) < \bar{q}(\theta)$, it must be that $q(\theta') > q(\theta)$ for (TR-21) to hold. As q is non-increasing, it follows that $\theta' < \theta$. We claim q is constant over $[\theta', \theta)$. If otherwise, $q(\hat{\theta}) < q(\theta')$ for some $\hat{\theta} \in [\theta', \theta)$ and this would mean the line with slope-intercept pair $(-q(\hat{\theta}), p(\hat{\theta}))$ leads to a higher utility at realized cost θ than the line with slope-intercept pair $(-q(\theta'), p(\theta'))$. Note the line with slope-intercept pair $(-q(\hat{\theta}), p(\hat{\theta}))$ goes through $(\hat{\theta}, V(\hat{\theta}))$ which is above or on the line with slope-intercept pair $(q(\theta'), p(\theta'))$ and has a greater slope. This would contradict the choice of θ'. For the interval (θ', θ), we have q constant over the interval and the first set of conditions hold.

If $q(\theta) > \bar{q}(\theta)$, then $q(\theta') < q(\theta)$ and $\theta' > \theta$. By a similar argument as above, it follows that q is constant over $(\theta, \theta']$. With the roles of θ and θ' reversed, we have that the second set of conditions hold.

(Only if part): We prove the contrapositive. Assume there exists an interval (θ', θ) within which q is constant, say, equal to q_0. Assume the first set of conditions hold. The argument is similar when the second set of conditions hold instead. We have $q(\theta) < q_0$ and $R(q_0) - \theta \cdot q_0 < R(q(\theta)) - \theta \cdot q(\theta)$. We show that truthful reporting is not a principal-pessimistic utility-maximizing strategy when the realized cost is θ because reporting some $\hat{\theta} \in (\theta', \theta)$ is a utility-maximizing strategy that gives a worse principal utility. Fix an arbitrary $\hat{\theta} \in (\theta', \theta)$. In the interval $[\hat{\theta}, \theta)$, q is equal to q_0 and p is also constant, say p_0. Then, the utility curve is on the line with slope-intercept pair $(-q_0, p_0)$ on interval $[\hat{\theta}, \theta)$. Since it is absolutely continuous, $(\theta, V(\theta))$ is also on the same line. Hence,

$$V(\theta) = p(\theta) - \theta \cdot q(\theta) = p(\hat{\theta}) - \theta \cdot q(\hat{\theta})\,,$$

and reporting $\hat{\theta}$ is a utility-maximizing strategy when the realized cost is θ. Combining the above with the assumption that $R(q_0) - \theta \cdot q_0 < R(q(\theta)) - \theta \cdot q(\theta)$, we obtain

$$R(q(\hat{\theta})) - p(\hat{\theta}) < R(q(\theta)) - p(\theta)\,.$$

This shows reporting $\hat{\theta}$ leads to a strictly worse principal utility. It follows that the first part of the PP constraint does not hold when realized cost is θ. □

Resource Allocation Problem without Monetary Transfers We represent a single-round direct mechanism with its interim allocation rule $x : [0,1] \to [0,1]$. The IC/IR/PP constraints are

$$\theta \cdot x(\theta) \geq \theta \cdot x(\theta')\,, \quad \forall \theta, \theta' \in [0,1] \tag{IC}$$

$$\theta \cdot x(\theta) \geq 0\,, \quad \forall \theta \in [0,1] \tag{IR}$$

$$\begin{aligned} &(\theta - c) \cdot x(\theta) \leq (\theta - c) \cdot x(\theta')\,, \quad \forall \theta \in [0,1], \theta' \in \mathcal{B}^*(\theta) \\ &(\theta - c) \cdot x(\theta) \leq 0\,, \quad \forall \theta \in \Theta_0 \end{aligned} \tag{PP}$$

where $\mathcal{B}^*(\theta) = \{\theta' \in [0,1] \mid \theta \cdot x(\theta') = \theta \cdot x(\theta)\}$ and $\Theta_0 = \{\theta \in [0,1] \mid \theta \cdot x(\theta) = 0\}$. Recall that $c \in (0,1)$.

Proof of Proposition TR.13. We exhaustively go through the following cases and show that the proposition statement holds. As discussed in Appendix A, we parametrize single-round direct IC/IR mechanisms in terms of $0 \leq x_0 \leq x_1 \leq 1$ such that $x(0) = x_0$ and $x(\theta) = x_1$ for $\theta \in (0,1]$.

Case 1) $x_0 = x_1$

In this case, the probabilistic allocation rule x is constant. That is, any report θ leads to the same probability of allocation. Then, any report leads to the same principal and agent utilities and the first part of the PP constraint holds trivially. For the second part, we divide into two subcases depending on whether $x_0 = 0$ or $x_0 > 0$. If $x_0 = 0$, the corresponding single-round mechanism does not allocate at all and the principal and agent utilities will be 0. Hence, the second part of the PP constraint holds. If $x_0 > 0$, then Θ_0 consists of exactly $\theta = 0$ for which the principal utility evaluates to $-c \cdot x_0 < 0$. Again, the second part of the PP constraint holds. It follows that when the allocation rule is constant, the single-round direct IC/IR mechanism satisfies the PP constraint.

Case 2) $x_0 < x_1$

For $\theta \in (0,1]$, any reports $\theta' \in [0,1]$ lead to either the same agent utility or a strictly smaller agent utility compared to truthful reporting. Reporting $\theta' \in (0,1]$ leads to the same allocation probability and, hence, the same principal and agent utilities. Reporting $\theta' = 0$ leads to the agent utility of $\theta \cdot x_0$ which is less than that of $\theta \cdot x_1$ under truthful reporting. When $\theta = 0$, all reports $\theta' \in [0,1]$ lead to the agent utility of 0, but can lead to different principal utilities. Reporting $\theta' = 0$ leads to the principal utility of $-c \cdot x_0$ whereas reporting any $\theta' \in (0,1]$ leads to the principal utility of $-c \cdot x_1$. Since $c > 0$ and $x_0 < x_1$, reporting some $\theta' \in (0,1]$ leads to the same agent utility but a strictly smaller principal utility compared to truthful reporting. Hence, the PP constraint does not hold.

It follows that only the single-round direct IC/IR mechanisms with $x_0 = x_1$, i.e., with a constant probability allocation rule, satisfy the PP constraint. □