Operations Research

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

On the Futility of Dynamics in Robust Mechanism Design

Santiago R. Balseiro, Anthony Kim, Daniel Russo

informs
http://pubsonline.informs.org/journal/opre
OPERATIONS RESEARCH
Articles in Advance, pp. 1–17
ISSN 0030-364X (print), ISSN 1526-5463 (online)

Crosscutting Areas

On the Futility of Dynamics in Robust Mechanism Design

Santiago R. Balseiro,[a] Anthony Kim,[b] Daniel Russo[a]

[a] Columbia University, New York, New York 10027; [b] Amazon, New York, New York 10001

Contact: srb2155@columbia.edu, https://orcid.org/0000-0002-0012-3292 (SRB); tonyekim@gmail.com, https://orcid.org/0000-0001-8908-2660 (AK); djr2174@gsb.columbia.edu, https://orcid.org/0000-0001-5926-8624 (DR)

Received: November 12, 2019
Revised: October 5, 2020
Accepted: December 21, 2020
Published Online in Articles in Advance: October 1, 2021

Subject Classifications: stochastic; games/group decisions
Area of Review: Revenue Management and Market Analytics

https://doi.org/10.1287/opre.2021.2122

Abstract. We consider a principal who repeatedly interacts with a strategic agent holding private information. In each round, the agent observes an idiosyncratic shock drawn independently and identically from a distribution known to the agent but not to the principal. The utilities of the principal and the agent are determined by the values of the shock and outcomes that are chosen by the principal based on reports made by the agent. When the principal commits to a dynamic mechanism, the agent best-responds to maximize his aggregate utility over the whole time horizon. The principal's goal is to design a dynamic mechanism to minimize his worst-case regret, that is, the largest difference possible between the aggregate utility he could obtain if he knew the agent's distribution and the actual aggregate utility he obtains. We identify a broad class of games in which the principal's optimal mechanism is static without any meaningful dynamics. The optimal dynamic mechanism, if it exists, simply repeats an optimal mechanism for a single-round problem in each round. The minimax regret is the number of rounds times the minimax regret in the single-round problem. The class of games includes repeated selling of identical copies of a single good or multiple goods, repeated principal-agent relationships with hidden information, and repeated allocation of a resource without money. Outside this class of games, we construct examples in which a dynamic mechanism provably outperforms any static mechanism.

Funding: The work of A. Kim was supported in part by a Columbia DRO Postdoctoral Award and was done prior to joining Amazon.
Supplemental Material: The online appendices are available at https://doi.org/10.1287/opre.2021.2122.

Keywords: strategic learning • robust mechanism design • minimax regret • dynamic pricing • dynamic contracting

1. Introduction

Individuals increasingly have repeated interactions with the same online platform. Commuters check the same ride-hailing app each morning, freelancers frequently hunt for short-term work on the same online marketplace, and advertisers bid daily on the same ad exchange. These interactions generate data the platform could use to personalize future offerings. Sometimes, the objectives of the platform and users are aligned—such as when an online music service recommends songs tailored to a user's tastes and the user's experience improves as accurate data are gathered. But often, the incentives of the platform and users are misaligned and repeated interactions become more complicated when a user strategically responds to the platform's strategy. Consider a platform that targets discount coupons at users who appear price-sensitive. This incentivizes loyal price-insensitive customers to mimic those who are not, complicating any inference from past data. Similar concerns arise in online ad exchanges—where, due to ad targeting, a meaningful fraction of auctions contain only a single bidder with a significantly high bid and appropriately setting reserve prices is a key driver of revenue—or online freelancing platforms—where a freelancer might reject an otherwise profitable contract to avoid signaling they are open to working for a low wage in the future.

In such environments, the platform could employ a myriad of dynamic strategies under which the offers available to an individual depend on all the past interactions. How much additional benefit can be derived from such dynamic strategies when an individual is strategic? We use the language of robust mechanism design to formalize a stark impossibility result. We identify a broad class of problems in which an optimal dynamic mechanism is static and simply repeats a single-round mechanism over and over. In this sense, the platform cannot benefit by using a more complex mechanism with meaningful dynamics, including any schemes that attempt to infer the private information of an individual and exploit this information using, for example, dynamic schemes (Bakos and Brynjolfsson 1999, Jackson and Sonnenschein 2007) that link together outcomes across periods. Intuitively, dynamic mechanisms that adapt based on previous actions can be manipulated by a strategic individual to induce future outcomes that are beneficial for him at the expense of the platform. Therefore, the platform finds it optimal to commit to implementing a static mechanism that does not exploit the individual's private information beyond what is known at the beginning of their repeated interactions.

For these problems, our results could be interpreted negatively as showing the impossibility of learning and exploiting the private information of a strategic individual. Viewed more positively, these results lead to massive simplification in that static mechanisms are not only robust to strategic manipulations but also optimal, allowing the platform to search over the more tractable space of single-round mechanisms. In addition, these results justify the use of simple mechanisms with substantial practical advantages: static mechanisms are simple to implement and alleviate the need for individuals to engage in complex strategic behavior. Interestingly, for some other problems, it is still possible for the platform to implement a dynamic mechanism and perform strictly better than implementing static mechanisms.

1.1. Contributions

We study a model where a principal and a strategic agent repeatedly play a game over a discrete-time finite horizon of length T. In each round, the agent privately observes an idiosyncratic shock drawn independently and identically from a distribution known to the agent but not to the principal, and the principal and agent interact through the game to realize an outcome and respective utilities. Both parties derive aggregate utility equal to the sum of their utilities across individual rounds. When the principal commits to a dynamic mechanism, the agent is strategic in the sense that he plays a best-response strategy to maximize his aggregate utility. Drawing inspiration from the enormous literature on dynamic learning in nonstrategic environments (see, e.g., Kleinberg and Leighton 2003, Besbes and Zeevi 2009), we measure the performance of a dynamic mechanism through its worst-case regret, that is, the largest difference possible between the aggregate utility he could obtain if he knew the agent's distribution and the actual aggregate utility he obtains. The principal's objective is minimax regret and the principal designs a dynamic mechanism to minimize the worst-case regret.

We provide false-dynamics results for a broad class of games, showing the principal's optimal mechanism is static without any meaningful dynamics. More specifically, we show the minimax regret is T times the minimax regret of a single-round problem and repeating T times a (near) optimal single-round mechanism from the single-round problem is correspondingly (near) optimal in the multiround problem. We prove our results under two assumptions. First, the set of possible distributions for the agent includes all point masses, that is, the principal must guard against the possibility that the agent's preferences are constant over time. Second, the optimal performance achievable by the principal with the knowledge of the agent's distribution should be *extreme-point convex*, that is, for any possible agent's distribution, the optimal performance achievable for that distribution is at most the convex combination of optimal performances corresponding to point masses where the convex combination is determined by the distribution.

Our analysis relies on leveraging point-mass distributions as worst-case distributions. When restricted to point-mass distributions, we obtain a static information structure where the agent's shock is constant and an optimal dynamic mechanism for the principal is static and repeats a single-round mechanism. Under the extreme-point convexity assumption, the optimality of static mechanisms over point masses as worst-case distributions extends to all possible distributions. We explain the extreme-point convexity assumption in terms of two opposing effects of shock uncertainty—information asymmetry and trade across shocks—and provide sufficient conditions for the assumption to hold. To the best of our knowledge, the second effect of trade across shocks is novel and may be of independent interest.

For specific applications of our general false-dynamics results, we consider (1) the dynamic selling mechanism design problem where a seller sells independent units of a single or multiple goods sequentially over time to a buyer and maximizes revenue or welfare, (2) the principal-agent model with hidden costs where a principal has a nonlinear revenue function and repeatedly contracts with an agent to produce at particular output levels, and (3) the repeated resource allocation problem without monetary transfers where a social planner allocates a costly resource in settings where monetary transfers are not allowed. In all these applications, our assumptions hold and an optimal mechanism for the multiperiod problem simply repeats an optimal mechanism for a single-round problem.

When our assumptions do not hold, it is possible that static mechanisms are not optimal and we show specific games in which either assumption does not hold and a dynamic mechanism provably outperforms any static mechanism. Finally, we extend our results in several directions and discuss connections to other related settings: a multiplicative performance guarantee, saddle-point properties, alternative benchmarks, serially correlated shock processes, a stronger notion of regret, and the maximin utility objective.

1.2. Related Work

We discuss connections between our work and several streams of literature including Bayesian dynamic mechanism design, robust mechanism design, and strategic learning.

1.2.1. Bayesian Mechanism Design. This stream of literature studies Bayesian mechanism design problems where the principal and agent share a common known prior over the distribution of shocks. False dynamics is a recurring phenomenon where optimal mechanisms do not display meaningful dynamics in sequential problems with static information (Laffont and Tirole 1993, Börgers et al. 2015). It was first observed by Baron and Besanko (1984) who considered a continuing relationship between a firm who reports cost information and a regulator who grants a license to operate. Similar results hold in many other dynamic allocation models where the agent's shock is constant (e.g., Baron and Besanko 1984) or changing over time (e.g., Bakos and Brynjolfsson 1999, Kakade et al. 2013, Pavan et al. 2014). Although generally dependent on the time horizon, optimal mechanisms may determine the outcomes of all future periods in the first round in these results. To the best of our knowledge, our paper is the first to study false-dynamics in a general class of problems with respect to a minimax regret objective. Although our model is more aligned with sequential screening models (e.g., Courty and Li 2000, Krähmer and Strausz 2015, Bergemann et al. 2017) where the optimal mechanisms are dynamic, we still obtain false-dynamics results. Under more stringent liquidity or participation constraints, false-dynamics disappears and adaptive mechanisms can outperform the optimal static mechanism in Bayesian settings (Krishna et al. 2013, Ashlagi et al. 2016, Balseiro et al. 2018). Our results reveal a stronger collapse of dynamics for a class of problems in that they hold even when these participation constraints are imposed.

1.2.2. Robust Mechanism Design. Although Bayesian models have appealing philosophical foundations, the resulting mechanisms sometimes place impractical requirements on the prior information of the designer. Wilson (1987) argues that mechanisms should not excessively rely on probabilistic assessments on the agents' types. Our work contributes to the robust mechanism design literature that was pioneered by Bergemann and Schlag (2008, 2011). In particular, the authors consider the single-round problem where the principal (i.e., seller) sells a good to an agent (i.e., buyer) to minimize the worst-case regret without the knowledge of the agent's distribution in Bergemann and Schlag (2008). Carrasco et al. (2019) is perhaps the most closely related work to ours. They show a similar false-dynamics result in an auction setting with respect to the maximin utility objective where the principal knows the mean of the unknown distribution and maximizes the worst-case utility over distributions that are potentially correlated over time. They bound the worst-case utility of any dynamic mechanisms by considering a worst-case distribution that is perfectly correlated across time and then invoking standard false-dynamics results from the Bayesian literature. Our approach in the minimax regret setting relies on considering point masses as worst-case distributions, which are not feasible in Carrasco et al. (2019) because of the moment constraints, and applies for a broader class of games. There are many other works in this literature (e.g., Carrasco et al. 2019; Kos and Messner 2015; Carroll 2017; Pınar and Kızılkale 2017; Carrasco et al. 2018a, b; Kocyigit et al. 2018), but they consider single-round problems whereas our problem is a multiround problem.

1.2.3. Strategic Learning. For a special case of our general problem, Amin et al. (2013) has previously shown that regret must grow linearly with the time horizon. Such results show that the principal bears a cost of asymmetric information that does not vanish regardless of the length of the time horizon. This formalizes the common folklore that learning about a strategic agent is fundamentally more difficult than learning about a myopic one. However, such results tend to be only asymptotic in nature and do not speak to the exact, absolute potential benefits of dynamic mechanisms over static ones, which is our main contribution.

Numerous papers in the learning theory literature consider assumptions that enable efficient dynamic learning on the part of the principal. Most notably, positive results are available when the principal repeatedly interacts with a myopic agent who optimizes without internalizing future consequences of his actions or a population of agents who each interact with the principal only once or when the principal simultaneously interacts with multiple agents whose values are drawn i.i.d. from the same distribution (e.g., Kleinberg and Leighton 2003, Kanoria and Nazerzadeh 2020). When the principal interacts with an agent who is forward-looking but less patient, usually modeled through unequal discount factors, the principal can learn and exploit the agent's private information to the extent that their time-preferences differ (e.g., Amin et al. 2013, 2014; Mohri and Munoz 2014, 2015; Golrezaei et al. 2020). Not surprisingly, the performance guarantees obtained in these settings degrade as the difference between the principal's and agent's discount factors becomes small. A distinguishing feature of our model is that the principal and agent are placed on a more equal footing in that they are both forward-looking and equally patient (i.e., the same discount factor). Many online platforms are characterized by short planning horizons and a high frequency of transactions. Although both parties might discount future payoffs in the one-to-one relationship, the difference in the discount factors can be expected to be small as planning horizons span only for weeks or months.

2. Model

We consider games described in terms of a time horizon T and an environment (Ω, Θ, u, v) where Ω is a set of outcomes, Θ is a set of idiosyncratic shocks of the agent, $u : \Theta \times \Omega \to \mathbb{R}$ is the utility function of the principal, and $v : \Theta \times \Omega \to \mathbb{R}$ is the utility function of the agent. A principal and an agent repeatedly interact in the given environment over T rounds, producing outcomes $\omega_1, \ldots, \omega_T \in \Omega$. Independently and identically distributed shocks $\theta_1, \ldots, \theta_T \sim F$ might influence the utility of both the principal and agent. The shock distribution $F \in \Delta(\Theta)$ is private to and learned by the agent in Round 0 before the shocks, and the shock θ_t is privately observed by the agent at the start of Round t. Assume that Ω contains a designated no-interaction outcome denoted by $\emptyset$. Both the principal and agent attain a utility of zero in any round with no interaction.

Example 1 (Selling Problem). The principal (or seller) repeatedly offers identical copies of an item to the agent (or buyer). The outcome $\omega_t = (x_t, p_t)$ realized in Round t consists of a quantity $x_t \geq 0$ of the good received by the agent and a payment $p_t \in \mathbb{R}$. The shock θ_t is the agent's willingness-to-pay or valuation for the good in Round t and F is his private valuation distribution. The agent's utility function is $v(\theta, (x, p)) = \theta x - p$ and the principal's is his revenue $u(\theta, (x, p)) = p$. The no-interaction outcome is one where $x = 0$ and $p = 0$. The private distribution F could be thought of as the agent's type and is persistent over time. Additional randomness in the shock θ_t, beyond F, represents unpredictable external factors that influence the agent's preferences in that round.

As is standard in the mechanism design literature, the principal can commit to implement a dynamic mechanism, which specifies a full protocol of interaction between the principal and agent.[1] Formally, a dynamic mechanism can be written as a tuple $A = (\{\mathcal{M}_t\}_{0:T}, \{\pi_t\}_{1:T}, \{\sigma_t\}_{0:T})$, where we use $a : b$ as shorthand for $a, \ldots, b$. For each t, the set $\mathcal{M}_t$ is the report space and defines the space of possible messages the agent can transmit to the principal in Round t. Let $\mathcal{M} = \{\mathcal{M}_t\}_{0:T}$. A decision rule π_t specifies an outcome $\omega_t = \pi_t(m_t, h_t, z_t) \in \Omega$ in Round t as a function of the report $m_t \in \mathcal{M}_t$, the history of prior interaction $h_t = (m_{0:t-1}, \omega_{1:t-1})$, and a private random variable z_t drawn independently over time from a uniform distribution over $[0, 1]$. Let $\pi = \{\pi_t\}_{1:T}$. The private random variable z_t allows outcomes to be determined on a randomized basis. Assuming it to be uniformly distributed is without loss of generality since more complex random variables can be generated by inverse transform sampling.

A dynamic mechanism A allows the principal to implement per-round decision rules that may be linked across rounds and depend on the history of past interactions. For example, in the selling problem, the principal may post a reserve price that he adjusts dynamically, bundle the current item and future items together, or provide some discount scheme that offers a future item at a low price if the current item is bought at a high price.

Following the convention in the mechanism design literature, the principal's mechanism specifies a recommended agent strategy $\sigma := \{\sigma_t\}_{0:T}$. We later constrain the choice of the recommended strategy such that the mechanism is incentive compatible. In each round t, σ_t specifies the report $m_t = \sigma_t(\theta_t, h_t^+, y_t) \in \mathcal{M}_t$ as a function of the realization of the private shock θ_t, the augmented history h_t^+ containing all information available (i.e., the public history h_t and the agent's private information) to the agent prior to Round t, and a private random variable y_t drawn independently from a uniform distribution over $[0, 1]$. The initial augmented history is $h_0^+ := F$ while $h_t^+ = (F, \theta_{1:t-1}, m_{0:t-1}, \omega_{1:t-1})$ for $t > 0$. Notice that no outcome is determined in Round 0, but the agent's initial message $m_0 = \sigma_0(F, y_0)$ could influence subsequent outcomes.

Given the principal's dynamic mechanism $A = (\mathcal{M}, \pi, \sigma)$, the agent's strategy $\tilde{\sigma} = \{\tilde{\sigma}_t\}_{0:T}$ (which may be different from σ), the distribution F, and the time horizon T, the principal and agent's total expected utilities are defined, respectively, as

$$\text{PrincipalUtility}(A, \tilde{\sigma}, F, T) := \mathbb{E}_{\pi, \tilde{\sigma}}\left[\sum_{t=1}^{T} u(\theta_t, \omega_t)\right]$$

and

$$\text{AgentUtility}(A, \tilde{\sigma}, F, T) := \mathbb{E}_{\pi, \tilde{\sigma}}\left[\sum_{t=1}^{T} v(\theta_t, \omega_t)\right].$$

The expectations above are taken over the realizations of the shocks $(\theta_1, \ldots, \theta_T)$ and the private random variables $\{z_t\}_{1:T}$ and $\{y_t\}_{0:T}$ which are omitted. The subscript indicates that the agent's messages are determined by $\tilde{\sigma}$ and the mechanism's outcomes are determined by π, meaning $m_t = \tilde{\sigma}_t(\theta_t, h_t^+, y_t)$ for $t \in \{0, \ldots, T\}$ and $\omega_t = \pi_t(m_t, h_t, z_t)$ for $t \in \{1, \ldots, T\}$. It is implicitly understood that the agent's augmented histories $h_0^+ := F$ and $h_t^+ = (F, \theta_{1:t-1}, m_{0:t-1}, \omega_{1:t-1})$ for $t \geq 1$ contain information about the same distribution F from which the shocks are drawn. We typically omit the subscripts from the above expectations, as they are clear from the context. We remark that the principal may not directly observe his utility if it depends on the agent's private shocks as, for example, in the case of welfare maximization. Figure 1 summarizes the order of events over the time horizon.

We say that dynamic mechanism $A = (\mathcal{M}, \pi, \sigma)$ is *incentive compatible* (IC) if the inequality

$$\text{AgentUtility}(A, \sigma, F, T) \geq \text{AgentUtility}(A, \tilde{\sigma}, F, T)$$

holds for every probability distribution F over Θ and every feasible agent strategy $\tilde{\sigma}$, that is, the agent is weakly better off following the principal's recommendation. We follow the convention in mechanism design of assuming the agent follows the recommended strategy if the mechanism is incentive compatible.[2] With this convention, we simplify the notation when A is IC and write $\text{PrincipalUtility}(A, F, T) := \mathbb{E}_{\pi,\sigma}\left[\sum_{t=1}^{T} u(\theta_t, \omega_t)\right]$ and $\text{AgentUtility}(A, F, T) := \mathbb{E}_{\pi,\sigma}\left[\sum_{t=1}^{T} v(\theta_t, \omega_t)\right]$ with the understanding that the omitted recommended strategy is the utility-maximizing strategy for the agent chosen by the principal.

We model an *individual rationality* (IR) constraint (equivalently, *participation* constraint) by providing a no-participation option in Round 0 and imposing IC with respect to this option. More specifically, we assume $\mathcal{M}_0$ contains a special report denoted by QUIT. If the agent reports QUIT in Round 0, he does not participate and the outcome is understood to be the no-interaction outcome over the whole time horizon. Because reporting QUIT is a feasible deviation, an incentive compatible mechanism should always provide the agent a nonnegative expected utility. This is commonly known as the ex-ante IR constraint in that the agent decides to participate or not while knowing only his distribution but not future shocks.

Let $\mathcal{A}$ denote the set of all incentive compatible dynamic mechanisms. If the agent has private distribution F, the optimal performance attainable by a principal who knows this distribution is

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

An optimal solution for this Bayesian dynamic mechanism design problem can be characterized recursively using the promised utility framework pioneered by Green (1987), Spear and Srivastava (1987), and Thomas and Worrall (1990). Alternatively, when the T is large, asymptotically optimal mechanisms can sometimes be provided (see, e.g., Fudenberg et al. 1994, Jackson and Sonnenschein 2007). The regret defined as

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

measures the shortfall in the performance of a dynamic mechanism $A \in \mathcal{A}$ against this known-distribution benchmark. The principal's objective is to design an incentive compatible mechanism with minimal worst-case regret where the worst-case regret for mechanism A is defined as

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

where $\mathcal{F} \subseteq \Delta(\Theta)$ is a given set of probability distributions over Θ that are possible for the agent. We can equivalently interpret that nature is selecting a worst-case distribution F against the principal's mechanism and the principal is guarding against all such possibilities in $\mathcal{F}$. The optimal minimax regret in the multiround problem is given by $\text{Regret}(T) := \inf_{A \in \mathcal{A}} \text{Regret}(A, T)$. To ensure this quantity is well defined, we assume throughout that $\sup_{F \in \mathcal{F}} \text{OPT}(F, T) < \infty$.

Remark 1. It may seem natural to instead formulate a maximin utility problem where the principal wants to solve $\sup_{A \in \mathcal{A}} \inf_{F \in \mathcal{F}} \text{PrincipalUtility}(A, F, T)$. The challenge with this formulation is that, in some games such as dynamic selling, nature could select a distribution under which the agent does not value the good at all, leading to a maximin utility of zero. If the distribution F were known, minimizing regret is equivalent to maximizing utility, but worst-case distributions are more natural under a regret objective. Our main result also applies to other robust objectives. We can show that static mechanisms are optimal for the maximin ratio objective where the principal wants to solve $\sup_{A \in \mathcal{A}} \inf_{F \in \mathcal{F}} \text{PrincipalUtility}(A, F, T)/\text{OPT}(F, T)$. In addition, static mechanisms are optimal for a constrained maximin utility objective studied by Carrasco et al. (2019) in the dynamic selling problem. These extensions are discussed in Section 6. Our primary reason for focusing on regret instead of, say, the maximin ratio objective is that minimax regret is a widely studied objective for dynamic learning in nonstrategic environments.

Figure 1. The Order of Events over the Time Horizon

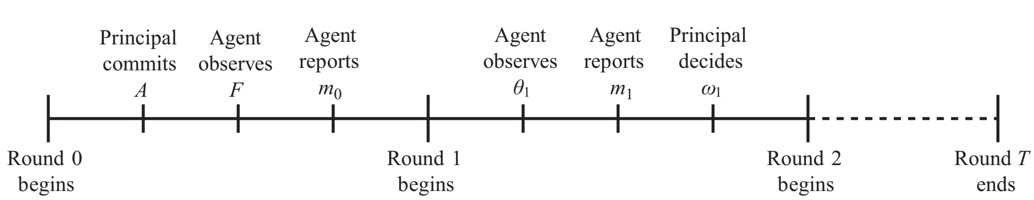

3. Optimality of Direct Static Mechanisms

In this section, we provide our main results. For a general class of games satisfying two sufficient conditions, we determine the minimax regret of the multiround problem to be T times that of a single-round minimax regret problem restricted to point-mass distributions and show an optimal dynamic mechanism is static, that is, it simply repeats a single-round direct mechanism, without any meaningful dynamics.

3.1. Direct Static Mechanisms

Our paper will focus on the special class of direct static mechanisms, denoted by $\mathcal{S}^{\times T} \subset \mathcal{A}$. These mechanisms are direct, meaning that the report spaces $\mathcal{M}_t$ contain the shock space Θ and the recommended strategy for the agent is to truthfully report his shocks to the principal. They satisfy interim IR constraints, meaning that the agent prefers to participate in each round given his information up to that point (including his shock for that round). Finally, they are simple repetitions of single-round mechanisms, meaning that the outcome in a given round is determined solely by the agent's reported shock in that round and does not depend directly on the history or the current round. To be precise, $\mathcal{S}^{\times T}$ is the space of *repeated single-round, direct, IC/IR* mechanisms, but we refer to them as *direct static mechanisms*.

More formally, an IC mechanism $A = (\mathcal{M}, \pi, \sigma)$ is an element of $\mathcal{S}^{\times T}$ if:

1. The message space allows for reporting the agent's shock or a choice to not participate in this round or, equivalently, pass: $\mathcal{M}_t = \Theta \cup \{\mathsf{PASS}\}$ for each $t \in \{1, \ldots, T\}$.
2. The mechanism honors a request to pass: $\pi_t(\mathsf{PASS}, h_t, z_t) = \emptyset$ for all h_t and z_t.
3. The recommended agent strategy is truthful reporting: $\sigma_t(\theta_t, h_t^+, y_t) = \theta_t$ holds for each possible (θ_t, h_t^+, y_t) and for $t \in \{1, \ldots, T\}$.
4. The mechanism's decision rule does not depend on the history or current round: there exists $\tilde{\pi} : \mathcal{M}_1 \times [0,1] \to \Omega$ such that $\pi_t(m_t, h_t, z_t) = \tilde{\pi}(m_t, z_t)$ for each possible (m_t, h_t, z_t) and $t \in \{1, \ldots, T\}$.

A direct static mechanism is uniquely identified by a direct mechanism for a single-round problem, that is, a problem with $T = 1$. For a single-round direct mechanism $S \in \mathcal{S}^{\times 1}$, we let $S^{\times T} \in \mathcal{S}^{\times T}$ denote the direct static mechanism that simply repeats S for T rounds. Note that rounds decouple under direct static mechanisms.

As in the case of the ex-ante IR constraint we considered previously, we implicitly enforce IR constraints by requiring the mechanism to be incentive compatible in the presence of no-participation options. Notice that for static mechanisms, we have not specified a report space $\mathcal{M}_0$ in Round 0. We can take $\mathcal{M}_0 = \{\mathsf{CONTINUE}, \mathsf{QUIT}\}$, where $\mathsf{CONTINUE}$ advances the agent to the first round and, as before, QUIT indicates a choice to not participate in any future rounds. Because the agent has the freedom to choose PASS in each round, but instead prefers to report his shock truthfully, this initial round is redundant and is included only for consistency with the general formulation.

3.2. Main Result

We prove our results for the general class of games satisfying the following two assumptions. We defer most discussion of these until Section 5. For any $\theta \in \Theta$, let δ_θ denote a point mass at θ, that is, a probability distribution with $\delta_\theta(\{\theta\}) = 1$. The first assumption states that the set of possible distributions $\mathcal{F}$ contains all point masses. Interpreted differently, it says that the principal must guard against the possibility that the agent's private shock is some unknown value in Θ and is constant across time.

Assumption 1 (Possibility of Deterministic Shocks). *For every $\theta \in \Theta$, we have $\delta_\theta \in \mathcal{F}$.*

The next assumption imposes a condition on the known-distribution benchmark $\mathrm{OPT}(F, T)$. In words, the assumption means that the principal is better off when the agent's shock is constant across time and publicly known than when shocks are random across time and are privately observed by the agent. As we discuss in Section 5, the condition holds for a broad class of games and has an interesting economic interpretation. Mathematically, it imposes the defining condition of convexity, but only with respect to the point-mass distributions—which are extreme points in the simplex of probability distributions $\Delta(\Theta)$. Note any distribution F is a convex combination of point-mass distributions where the combination is given by the distribution F, that is, F and $\mathbb{E}_{\theta \sim F}[\delta_\theta]$ are equivalent in the distributional sense.

Assumption 2 (Extreme-Point Convexity). *For all $F \in \mathcal{F}$, $\mathrm{OPT}(F, T) \le \mathbb{E}_{\theta \sim F}[\mathrm{OPT}(\delta_\theta, T)]$.*

The following theorem shows a complete reduction from the multiround problem to a single-round problem (with $T = 1$) in terms of their objective values, their (nearly) optimal solutions, and the existence of optimal solutions. Recall that $\mathrm{Regret}(T) = \inf_{A \in \mathcal{A}} \sup_{F \in \mathcal{F}} \mathrm{Regret}(A, F, T)$ is the minimax regret attainable in the multiround problem.

Theorem 1. *Suppose Assumptions 1 and 2 hold. Then,* $\mathrm{Regret}(T) = T \cdot \inf_{S \in \mathcal{S}^{\times 1}} \sup_{\theta \in \Theta} \mathrm{Regret}(S, \delta_\theta, 1)$. *Moreover, for any $\epsilon \ge 0$, if a mechanism $S \in \mathcal{S}^{\times 1}$ satisfies*

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

then,

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

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.*

Our result is easier to interpret when the infimum in (1) is attained, that is, there exists an optimal single-round mechanism. In this case, our main theorem says that an optimal dynamic mechanism can be constructed by first solving for an optimal direct mechanism for the single-round problem $S^* \in \arg\min_{S\in\mathcal{S}^{\times 1}} \sup_{\theta\in\Theta} \text{Regret}(S,\delta_\theta,1)$ and then simply repeating this mechanism over the rounds, that is, implementing $(S^*)^{\times T}$. Such a mechanism is myopic and nonadaptive: it plans over a single round even when there are many and it does not adjust its decision rule based on prior interactions with the same agent. This provides insight into the character of an optimal mechanism and also lets us formulate more tractable optimization problems. Additionally, our result states that the minimax regret in the multiround problem is T times the minimax regret in the single-round problem, that is, $\text{Regret}(T) = T\cdot\text{Regret}(1)$. The final part of the result concerns the existence of direct static mechanisms that are exactly optimal.

Depending on the value of the minimax regret for the single-round problem, different interpretations are possible. If $\text{Regret}(1) = 0$, repeating a single-round mechanism obtains essentially the optimal performance achievable with the knowledge of the agent's private distribution. Then, the distributional information is not necessary and learning/adaptive schemes are not beneficial to begin with and the multiround problem is easy. On the other hand, if $\text{Regret}(1) > 0$, the minimax regret for the multiround problem is linear in the time horizon and repeating a single-round mechanism is still a (near) optimal dynamic mechanism. Even though knowing the agent's private distribution would be valuable, it is impossible to increase the principal's utility by employing an adaptive learning scheme. Welfare maximization in the dynamic selling mechanism design problem (Section 4.1) is of the former kind. Revenue maximization in the same problem, the principal-agent contract model (Section 4.2) and the dynamic resource allocation problem without monetary transfers (Section TR.5 of the technical report Balseiro et al. 2019) are of the latter. We defer further details to respective sections.

3.3. Characterizing Direct Static Mechanisms

Theorem 1 reduces a complex dynamic mechanism design problem to the following single-round optimization problem over direct mechanisms and restricted to point-mass distributions:

$$\inf_{S\in\mathcal{S}^{\times 1}} \sup_{\theta\in\Theta} \text{Regret}(S,\delta_\theta,1). \tag{2}$$

In this subsection, we describe the structure of (2) in more detail.[3] We first describe the space of single-round direct mechanisms in a more explicit form. Recall from Subsection 3.1 that $S\in\mathcal{S}^{\times 1}$ is uniquely determined by decision rule $\pi_1 : \Theta\cup\{\text{PASS}\}\times[0,1]\to\Omega$ that determines the outcome $\omega_1 = \pi_1(\theta_1, z_1)$ as a function of the reported shock θ_1 and the random variable z_1 that enables randomized decision rules. Put differently, S is effectively defined by a rule that associates each possible report $\theta\in\Theta$ with a distribution over outcomes; we ignore the possibility of a report of PASS under the recommended truthful reporting strategy. We make this explicit and, abusing notations, equivalently define a single-round direct mechanism $S\in\Delta(\Omega)^\Theta$ by

$$S_\theta(W) = \mathbb{P}_{z\sim\text{Uniform}[0,1]}(\pi_1(\theta,z)\in W),$$

for every shock $\theta\in\Theta$ and measurable set $W\subseteq\Omega$.

Consider the following optimization problem in terms of the outcome distribution representation:

$$\inf_{S\in\Delta(\Omega)^\Theta} \sup_{\theta\in\Theta} \left\{\text{OPT}(\delta_\theta,1) - \int_\Omega u(\theta,\omega)\mathrm{d}S_\theta(\omega)\right\} \tag{3}$$

$$\text{s.t.} \int_\Omega v(\theta,\omega)\mathrm{d}S_\theta(\omega) \ge \int_\Omega v(\theta,\omega)\mathrm{d}S_{\theta'}(\omega)\ \forall\ \theta,\theta'\in\Theta, \tag{IC}$$

$$\int_\Omega v(\theta,\omega)\mathrm{d}S_\theta(\omega) \ge 0 \qquad \forall\ \theta\in\Theta. \tag{IR}$$

It is simple to conclude that the constraints (IC) and (IR) are a rewriting of the incentive compatibility and (implicit) individual rationality constraints we have placed on the set $\mathcal{S}^{\times 1}$ in Subsection 3.1. The following lemma leverages this representation of single-round direct mechanisms to simplify (2) and shows its equivalence to (3); see Online Appendix A.2 for the proof.

Lemma 1. *The optimization problems* (2) *and* (3) *attain the same objective value. Moreover, a single-round direct mechanism* S^* *with decision rule* $\pi_1 : \Theta\cup\{\text{PASS}\}\times[0,1]\to\Omega$ *is an optimal solution of* (2) *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* (3). *Finally, the objective of* (3) *can be equivalently replaced with* $\sup_{F\in\Delta(\Theta)}\{\int_\Theta \text{OPT}(\delta_\theta,1)dF(\theta) - \int_\Theta\int_\Omega u(\theta,\omega)dS_\theta(\omega)dF(\theta)\}$.

It is easy to see that the known-distribution benchmark in (3) can be equivalently written as

$$\text{OPT}(\delta_\theta,1) = \sup_{G\in\Delta(\Omega)} \int_\Omega u(\theta,\omega)\mathrm{d}G(\omega)$$
$$\text{s.t.} \int_\Omega v(\theta,\omega)\mathrm{d}G(\omega) \ge 0.$$

Therefore, $\text{OPT}(\delta_\theta,1)$ can be thought of as a "first-best" benchmark without IC constraints in which the principal chooses, for the given shock, the best possible distribution over outcomes subject to an ex-ante IR constraint.

When the spaces of shocks and of outcomes are discrete, (3) is a finitely-sized linear program that can be efficiently solved. When the shock space is single-dimensional, an optimal single-round direct mechanism can be sometimes determined analytically using the Myersonian theory (which involves using the envelope theorem to simplify the IC constraints) and minimax duality theory (on the version of (3) with the alternative objective stated in Lemma 1 in which the inner supremum is a convex program). We illustrate this approach in Section 4. We are not aware of a general approach to solve (3) in general, multidimensional mechanism design problems.

3.4. Proof Sketch for Theorem 1

To prove our main theorem, we show a lower bound and an upper bound on the minimax regret of the multiround problem in terms of direct static mechanisms. Essentially, the multiround problem reduces to a static problem when the principal restricts to point-mass distributions that are possible candidates for the agent's distribution under Assumption 1 and this restriction is without loss for the principal under Assumption 2. Both the lower bound and upper bound arguments rely crucially on how the single-round benchmark $\mathrm{OPT}(\delta_\theta, 1)$ relates to the optimal performance achievable $\mathrm{OPT}(F, T)$ in T rounds. As already mentioned, point-mass distributions will be important and we have the following result on the benchmark. When the distribution F is a point-mass and is known to the principal, the agent holds no private information and the principal can (nearly) attain $\mathrm{OPT}(\delta_\theta, T)$ by simply repeating a (near) optimal mechanism that (nearly) attains $\mathrm{OPT}(\delta_\theta, 1)$ over T rounds.

Proposition 1. *For any* $\theta \in \Theta$, $\mathrm{OPT}(\delta_\theta, T) = T \cdot \mathrm{OPT}(\delta_\theta, 1)$.

The following result lower bounds the regret of any incentive compatible dynamic mechanism in terms of the regret of direct static mechanisms restricted to point-mass distributions and it directly implies a lower bound on the minimax regret in the multiround problem.

Lemma 2 (Lower Bound). *Suppose Assumption 1 holds. For any 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*

$$\sup_{F \in \mathcal{F}} \mathrm{Regret}(A, F, T) \geq T \cdot \sup_{\theta \in \Theta} \mathrm{Regret}(S, \delta_\theta, 1). \tag{4}$$

To see the lower bound $\mathrm{Regret}(T) \geq T \cdot \inf_{S \in \mathcal{S}^{\times 1}} \sup_{\theta \in \Theta} \mathrm{Regret}(S, \delta_\theta, 1)$, first take the infimum over all single-round direct IC/IR mechanisms S on the right-hand side and then take the infimum over all incentive compatible dynamic mechanisms A on the left-hand side. To prove the lemma, we use a revelation-principle-type argument to reduce the multiround problem to the single-round problem. The main idea involves using Assumption 1 to focus on point-mass distributions and then imposing structural constraints (the IC/IR constraints) as we effectively shrink the time horizon; this idea also appears in Amin et al. (2013). More specifically, we can construct a single-round direct mechanism S from any incentive compatible dynamic mechanism A by letting S_θ for $\theta \in \Theta$ to be the time-averaged distribution of outcomes when the agent's distribution is the point-mass distribution δ_θ and the agent plays the recommended strategy (as given in A). By construction, truthfully reporting a shock θ under S gives the same utilities to both parties, when scaled by T, as implementing the recommended strategy under A. Then, the resulting single-round mechanism S is incentive compatible and individually rational because the recommended strategy is utility-maximizing for the agent and guarantees the agent utility of at least 0 for the point-mass distributions in the multiround problem.

The next result upper bounds the regret of direct static mechanisms in the multiround problem under Assumption 2. For any single-round direct IC/IR mechanism, the regret incurred by repeating it T times is no greater than T times its regret in the single-round problem when restricted to point-mass distributions.

Lemma 3 (Upper Bound). *Suppose Assumption 2 holds. For every single-round direct IC/IR mechanism $S \in \mathcal{S}^{\times 1}$,*

$$\sup_{F \in \mathcal{F}} \mathrm{Regret}\big(S^{\times T}, F, T\big) \leq T \cdot \sup_{\theta \in \Theta} \mathrm{Regret}(S, \delta_\theta, 1). \tag{5}$$

The lemma implies the upper bound $\mathrm{Regret}(T) \leq T \cdot \inf_{S \in \mathcal{S}^{\times 1}} \sup_{\theta \in \Theta} \mathrm{Regret}(S, \delta_\theta, 1)$ because we can take the infimum over all single-round direct IC/IR mechanisms on both sides in the stated inequality and note repetitions of single-round direct IC/IR mechanisms are a subset of all incentive compatible dynamic mechanisms $\mathcal{A}$. For a sketch of the proof of the lemma, we note that when the principal implements $S^{\times T}$ for a single-round direct IC/IR mechanism S and the agent reports truthfully (as recommended for direct mechanisms), the individual rounds correspondingly decouple and $\mathrm{PrincipalUtility}(S^{\times T}, F, T) = T \cdot \mathrm{PrincipalUtility}(S, F, 1)$ for any distribution F. We then note Assumption 2 guarantees that the worst-case regret against point-mass distributions extends to that against any distributions in $\mathcal{F}$ in the single-round problem. That is, by the extreme-point convexity assumption, the principal can control his regret by protecting against all point-mass distributions.

Combining Lemmas 2 and 3, we can prove Theorem 1. From the above discussion, we already have $\text{Regret}(T) = T \cdot \inf_{S \in \mathcal{S}^{\times 1}} \sup_{\theta \in \Theta} \text{Regret}(S, \delta_\theta, 1)$. Similarly, we can prove the second part about the (near) optimality of direct static mechanisms and the third part about the existence of optimal mechanisms from these lemmas. For a complete proof, we refer to Online Appendix A.1. We refer to Online Appendix A.3 for proofs of Lemmas 2 and 3 and to Online Appendix A.3.3. for that of Proposition 1.

To conclude, we note restricting to point-mass distributions is reasonable in hindsight because point-mass distributions happen to be the right class of "worst-case" distributions in that for any dynamic mechanism, there exists a point-mass distribution for the agent against which the dynamic mechanism is forced to obtain a regret at least the minimax regret. See the following proposition; its proof is provided in Online Appendix A.3.3.

Proposition 2. *Suppose Assumptions 1 and 2 hold. For any incentive compatible dynamic mechanism $A \in \mathcal{A}$ and $\epsilon > 0$, there exists a point-mass distribution δ_θ such that* $\text{Regret}(A, \delta_\theta, T) \geq \text{Regret}(T) - \epsilon$.

4. Applications

In this section, we apply our results to a dynamic selling mechanism design problem where a seller sells independent goods sequentially over time and a principal-agent model with hidden costs in which a principal repeatedly contracts with an agent to produce at particular output levels. A third application to a repeated resource allocation problem without monetary transfers is presented, due to space considerations, in Section TR.5 of the technical report Balseiro et al. (2019).

4.1. Dynamic Selling Mechanism

Consider a repeated setting where the principal (i.e., seller) sells independent and identical items to a strategic agent (i.e., buyer) over T rounds and seeks to maximize the revenue; this is an extension of the seminal single-round model from Bergemann and Schlag (2008) to our robust, dynamic mechanism design framework. The items are being sold one by one sequentially and, in each round, the agent realizes his value (equivalently, his willingness to pay) for the current item. The agent's values are drawn from an underlying private distribution known only to him. The principal does not know the agent's private value distribution except that the agent's value is in the range $[0, 1]$.

In the language of the general model, the agent's shock is his private value for the item and the shock space is $\Theta = [0, 1]$. We assume that $\mathcal{F} = \Delta([0, 1])$, which implies that Assumption 1 holds. The outcome space is $\Omega = \{0, 1\} \times \mathbb{R}$ and an outcome is $\omega = (\hat{x}, \hat{p}) \in \Omega$ where $\hat{x}$ is the allocation and $\hat{p}$ is the payment, that is, whether the item is allocated to the agent and the payment the agent makes to the principal. Given an outcome $\omega = (\hat{x}, \hat{p})$, the agent's utility function is $v(\theta, \omega) = \theta \cdot \hat{x} - \hat{p}$ and the principal's utility function is $u(\theta, \omega) = \hat{p}$. Abusing notations, for single-round direct mechanisms, we use $x : \Theta \to [0, 1]$ and $p : \Theta \to \mathbb{R}$ to denote the interim rules mapping reported shocks to expected allocations and payments, respectively, and the pair (x, p) to represent a single-round direct mechanism when convenient.

Note that $\text{OPT}(\delta_\theta, T) = T\theta$ for all $\theta \in \Theta$ because the principal can extract the full surplus of the agent and satisfy the IR constraint by charging the agent's value when his shock is constant. Because of the agent's participation constraint, the principal's revenue is at most the agent's surplus and, thus, $\text{OPT}(F, T) \leq T\mathbb{E}_{\theta \sim F}[\theta] = \mathbb{E}_{\theta \sim F}[\text{OPT}(\delta_\theta, T)]$ for every distribution $F \in \mathcal{F}$. Therefore, Assumption 2 holds and Theorem 1 applies. We next show the minimax regret is T/e and an optimal dynamic mechanism is T repetitions of a randomized posted pricing mechanism to be specified below.

Proposition 3. *For revenue maximization in the dynamic selling mechanism design problem with one good, the minimax regret is T/e and an optimal solution is T repetitions of the randomized posted pricing mechanism S^* with price distribution Φ^* given by*

$$\Phi^*(p) = \begin{cases} 0, & \text{if } p \in [0, 1/e) \\ 1 + \ln p, & \text{if } p \in [1/e, 1] \end{cases},$$

such that the interim allocation and payment rules are $x^*(\theta) = \Phi^*(\theta)$ *and* $p^*(\theta) = [\theta - 1/e]_+$ *for* $\theta \in [0, 1]$.

For comparison, Amin et al. (2013) showed a lower bound of $T/12$ for the restricted class of dynamic posted pricing mechanisms; they considered a slightly different benchmark, but their results still hold in our setting (see section TR.6.1 of the technical report Balseiro et al. 2019). Their results would imply that static posted pricing mechanisms are asymptotically optimal in the restricted class. We consider more general dynamic mechanisms and determine an exactly optimal dynamic mechanism in this larger class with the performance that matches the improved lower bound of T/e exactly. To prove Proposition 3, it suffices to show that the randomized posted pricing mechanism S^* is a solution to the single-round minimax regret problem (3) and the corresponding minimax regret is $1/e$ via a saddle-point result, by Theorem 1. Its proof is deferred to Online Appendix B. An analogous result restricted to randomized posted pricing strategies exists due to Bergemann and Schlag (2008). Our single-round saddle-point result is for the slightly more general class of single-round direct IC/IR mechanisms and is still obtained with the same optimality structure and minimax regret value.[4]

Our results extend to multiple-goods settings (selling n goods in each round to an agent with additive, multidimensional valuations) and welfare maximization. In the case of welfare maximization, the minimax regret is zero because, using the Vickrey-Clarke-Groves mechanism (which reduces to allocating items for free for one buyer), it is possible to allocate efficiently without any prior beliefs on the agent's private information. In the case of revenue maximization with multiple goods, a robust optimal solution involves selling each item independently using the above mechanism S^* for a total regret of $(n/e)T$. Due to space considerations, we defer the analysis of these models to Sections TR.4.1 and TR.4.2 of the technical report Balseiro et al. (2019).

4.2. Principal-Agent Model with Hidden Costs

We consider a repeated principal-agent problem that captures various applications such as retail franchising, labor contracts, and procurement contracts. Similar to revenue maximization in the dynamic selling problem, we show that the minimax regret is linear in T and an optimal mechanism is static and repeats a single-round mechanism. Due to nonlinearity in the problem, our analysis is more involved.

More formally, the principal repeatedly contracts with the agent to produce output on his behalf and obtains revenue $R(\hat{q})$ when the agent produces $\hat{q}$ units of output, which is publicly observable. A contract specifies a payment $\hat{p}$ from the principal to the agent as a function of the number of output units $\hat{q}$. The agent has a private marginal production cost $\theta \in [\underline{\theta}, \bar{\theta}]$ where $0 < \underline{\theta} < \bar{\theta} < \infty$ which is assumed to be independently and identically distributed according to distribution F across the rounds. The agent observes his private cost and then decides on the production level $\hat{q}$ in each round. When he produces $\hat{q}$ units of output and receives a payment $\hat{p}$, his utility is $\hat{p} - \theta \cdot \hat{q}$ where θ is his marginal cost for that round. The principal does not know the agent's private distribution F but only that realized costs are in the range $[\underline{\theta}, \bar{\theta}]$. In particular, $\mathcal{F} = \Delta([\underline{\theta}, \bar{\theta}])$ and Assumption 1 holds. We assume $R(x)$ is a strictly increasing, strictly concave function that is twice continuously differentiable on $(0, \infty)$ with $R(0) = 0$ and $\lim_{x\to 0} R'(x) = \infty$; for example, $R(x) = \sqrt{x}$.

In terms of our general model, the agent's shock is his marginal cost of production and $\Theta = [\underline{\theta}, \bar{\theta}]$. The outcome space is $\Omega = \mathbb{R}_+ \times \mathbb{R}$ and an outcome $\omega = (\hat{q}, \hat{p}) \in \Omega$ is a pair of the production level $\hat{q}$ and the payment $\hat{p}$. When the outcome is $\omega = (\hat{q}, \hat{p})$ in a round, the agent's utility function is $v(\theta, \omega) = \hat{p} - \theta \cdot \hat{q}$ and the principal utility function is $u(\theta, \omega) = R(\hat{q}) - \hat{p}$. Abusing notations, for single-round direct mechanisms, we use $q : \Theta \to \mathbb{R}_+$ and $p : \Theta \to \mathbb{R}$ to denote the interim rules mapping reported shocks to production levels and payments, respectively, and the pair (q, p) to represent a single-round direct mechanism when convenient.

We now discuss the first-best mechanism that the principal can implement when he knows the agent's shock in a round. Because monetary transfers are allowed, the principal would set payments so that the IR constraint of the agent binds. Denote by $\bar{q}(\theta) = \operatorname{argmax}_{x\geq 0}\{R(x) - \theta \cdot x\}$ the optimal production level when the agent's shock is known. The first-best mechanism involves requesting the agent to produce $\bar{q}(\theta)$ units and paying the agent the minimum amount $\theta \cdot \bar{q}(\theta)$ that makes him indifferent between participating or not in the contract (see, e.g., Laffont and Martimort 2001). Let $\bar{R}(\theta) = \max_{x\geq 0}\{R(x) - \theta \cdot x\}$ be the first-best utility of the principal when the shock is known to be θ. Because of the assumptions on $R(\cdot)$, $\bar{q}(\theta)$ is uniquely defined such that $R'(\bar{q}(\theta)) = \theta$ and $\bar{R}(\theta)$ is a strictly decreasing convex function.

A similar reasoning to the dynamic selling problem yields that $\mathrm{OPT}(\delta_\theta, T) = T\bar{R}(\theta)$ because, when the agent's shocks are constant and equal to θ, the principal can implement the above first-best solution. The known-distribution benchmark for distribution F can be bounded as follows:

$$\begin{aligned}\mathrm{OPT}(F, T) &\leq \sup_{A\in\mathcal{A}} \mathbb{E}_{\pi,\sigma}\left[\sum_{t=1}^{T} R(\hat{q}_t) - \theta_t \hat{q}_t\right] \\ &\leq \sum_{t=1}^{T} \mathbb{E}_{\theta_t}\left[\bar{R}(\theta_t)\right] = \mathbb{E}_{\theta\sim F}\left[\mathrm{OPT}(\delta_\theta, T)\right],\end{aligned}$$

where the first inequality follows from the agent's participation constraint, which implies $\mathbb{E}_{\pi,\sigma}[\sum_{t=1}^{T} \hat{p}_t - \theta_t \hat{q}_t] \geq 0$, the second from relaxing the IC constraint and optimizing pointwise over the shocks, and the equality because shocks are identically distributed. Therefore, Assumption 2 holds and Theorem 1 applies. We formally state the main result of the subsection as follows:

Proposition 4. *For the principal-agent model with hidden costs, the minimax regret of the multiround problem is cT for some constant $c > 0$ and an optimal solution is T repetitions of offering the menu of deterministic contracts $\{(q^*(\theta), p^*(\theta))\}_{\theta\in\Theta}$, which is a single-round direct IC/IR mechanism. The allocation rule is continuous and satisfies the differential equation characterization*

$$(q^*)'(\theta) = -\frac{\bar{q}(\theta)}{R'(q^*(\theta)) - \theta}, \text{for } \theta \in (\underline{\theta}, \kappa),$$

with boundary conditions $q^*(\underline{\theta}) = \bar{q}(\underline{\theta})$ and $q^*(\theta) = 0$ for $\theta \in [\kappa, \bar{\theta}]$, where κ is the smallest cost for which q^* equals *to 0 and is assumed to be $\bar{\theta}$ if q^* is positive over $[\underline{\theta}, \bar{\theta}]$, and the payment rule is given by*

$$p^*(\theta) = \theta \cdot q^*(\theta) + \int_{\theta}^{\bar{\theta}} q^*(x) \mathrm{d}x.$$

For Proposition 4, we prove the stated single-round direct IC/IR mechanism is an optimal solution to the single-round minimax regret problem (3) via a saddle-point result. Because the revenue function R is concave, no randomization is needed and we can restrict our search of an optimal solution to those single-round direct IC/IR mechanisms that can be described in terms of a menu of deterministic contracts $(q(\theta), p(\theta))$ for $\theta \in [\underline{\theta}, \bar{\theta}]$ where the contract terms are all deterministic without randomization. We refer to Online Appendix C.1 for further details on the single-round problem including the IC/IR constraints and to Online Appendix C.2 for the proof of Proposition 4.

5. Interpretation and Necessity of Assumptions 1 and 2

In this section, we provide counterexamples showing that direct static mechanisms are not necessarily optimal in the absence of Assumptions 1 or 2. We also include a richer discussion of the extreme-point convexity assumption; we provide an economic interpretation, establish general sufficient conditions under which it holds, and provide an approximate form of Theorem 1 under an approximate extreme-point convexity condition.

5.1. Suboptimality of Static Mechanisms Without Assumption 1

Assumption 1 requires that the principal must guard against point-mass distributions because these are feasible choices for nature for the agent's distribution. To explore the necessity of this assumption, we show a game for which Assumption 1 does not hold and in which repeating a single-round direct mechanism leads to a suboptimal regret. The game will involve a class of shock distributions for which the knowledge of the mean immediately constrains the variance.

We consider a dynamic selling problem as in Section 4.1 with one good when the shock space is $\Theta = \mathbb{R}_+$ and the class of distributions $\mathcal{F}$ is a *scaled family* (see, e.g., Casella and Berger 2002, Definition 3.5.4). For a fixed cumulative distribution function G over $\mathbb{R}_+$, which is assumed to be known by the principal, we consider the parametric class of distributions $\mathcal{F}_G = \{F(\theta; \tau) = G(\theta/\tau) \text{ for some } \tau \in [0,1]\}$ where the scale parameter τ is known by the agent but not by the principal. Denoting by γ_t a random variable distributed according to G, we have that shocks are multiplicatively separable in the sense that $\theta_t \overset{(d)}{=} \tau\gamma_t$ (i.e., equivalent in the distributional sense). The effect of the scale parameter τ is to contract the distribution G while maintaining the same shape. For example, when G is an exponential distribution with mean 1, $\mathcal{F}_G$ is the class of all exponential distributions with mean $\tau \in [0,1]$. Alternatively, when G is a lognormal distribution with mean 0 and standard deviation σ, $\mathcal{F}_G$ is the class of all lognormal distributions with mean $\log(\tau) \in (-\infty, 0]$ and standard deviation σ.

Note we have an ex-ante participation constraint. Because monetary payments are allowed, Proposition 8 (in Subsection 5.2) implies that Assumption 2 holds. Assumption 1, however, does not hold because not all point-mass distributions are feasible distributions in $\mathcal{F}_G$.

We characterize an optimal mechanism using the relax and verify approach of Kakade et al. (2013):

Proposition 5. *Let G be an arbitrary distribution over $\mathbb{R}_+$ with mean $\mathbb{E}[\gamma] > 0$. For revenue maximization in the dynamic selling problem with one good and the agent's distribution restricted to $\mathcal{F}_G$, the minimax regret is $(\mathbb{E}[\gamma]/e)T$ and an optimal mechanism initially screens the agent by the parameter τ in Round 0, and then in each round, allocates each item with probability $x^*(\tau)$ and charges $p^*(\tau)$ with (x^*, p^*) as stated in Proposition 3.*

Recall that we can decompose shocks as $\theta_t = \tau\gamma_t$ with γ_t drawn i.i.d. from G. We prove the above result by considering a relaxed environment in which γ_t are public and observable by the principal—the agent's only private information is the parameter τ. Leveraging our ex-ante participation constraint, we show the relaxed problem can be reduced to a single-round problem that can solved to optimality as in the proof of Proposition 3. We conclude by showing that the resulting mechanism is feasible and, thus, optimal for the original problem. We remark that the mechanism presented in the above result is ex-ante individually rational but not ex-post individually rational, that is, the agent's total utility might be negative with a positive probability. We conjecture that the minimax regret of $(\mathbb{E}[\gamma]/e)T$ can also be asymptotically attained with an ex-post individually rational mechanism.

Let $\text{Regret}^{\mathcal{S}}(T) := \inf_{S \in \mathcal{S}^{\times 1}} \sup_{F \in \mathcal{F}_G} \text{Regret}(S^{\times T}, F, T)$ be the optimal regret for the class of direct static mechanisms. Because $\mathcal{S}^{\times T} \subset \mathcal{A}$, we have that $\text{Regret}(T) \le \text{Regret}^{\mathcal{S}}(T)$. For a specific game where direct static mechanisms without screening are strictly suboptimal or, equivalently, $\text{Regret}(T) < \text{Regret}^{\mathcal{S}}(T)$, we let G be the exponential distribution with mean 1. In this case, we have that $\text{Regret}(T) = T/e$ from the above proposition. On the other hand, the following proposition shows that $\text{Regret}^{\mathcal{S}}(T) = (1 - 1/e)T$.

Proposition 6. *Let G be the exponential distribution with mean 1. For revenue maximization in the dynamic selling problem with one good and the agent's distribution restricted to $\mathcal{F}_G$, the minimax regret for direct static mechanisms* *is* $(1-1/e)T$ *and repeatedly offering a posted price of 1 is an optimal direct static mechanism.*

Therefore, when Assumption 1 does not hold, there can be a separation between the minimax regrets achievable by incentive compatible dynamic mechanisms and by direct static mechanisms. The proofs of the above propositions are provided in Online Appendix D.

5.2 Sufficient Conditions for Assumption 2

This subsection provides sufficient conditions for the extreme-point convexity requirement in Assumption 2. Its economic interpretation is deferred until the next subsection. All proofs are deferred to Online Appendix E.

It is helpful to replace the known-distribution benchmark $\mathrm{OPT}(F,T)$ with a more tractable object. The first-best (equivalently, full information) benchmark $\bar{u}(F)$, defined below, is the optimal performance attainable when incentive compatibility constraints are dropped and the mechanism is subject only to an ex-ante individual rationality constraint. In other words, this is the optimal principal utility in a surrogate single-round problem in which the agent's distribution F is commonly known, shocks are publicly observable, and the agent commits to participate before observing the shock. Relatedly, we also consider the measure $\mathbb{E}_{\theta\sim F}[\bar{u}(\delta_\theta)]$ which differs from $\bar{u}(F)$ because it contains an interim IR constraint (on the per-shock basis) rather than an ex-ante IR constraint. It is the optimal principal utility in the single-round problem with full information as in $\bar{u}(F)$, but where the agent can choose not to participate (or report PASS) after the shock is realized. For any distribution F, we formally define $\bar{u}(F)$ and the (derived) measure $\mathbb{E}_{\theta\sim F}[\bar{u}(\delta_\theta)]$ as follows:

$$\begin{aligned}\bar{u}(F) := \sup_{S\in\mathcal{S}^{\times 1}} &\int_\Theta\int_\Omega u(\theta,\omega)\mathrm{d}S_\theta(\omega)\mathrm{d}F(\theta)\\ \text{s.t.} &\int_\Theta\int_\Omega v(\theta,\omega)\mathrm{d}S_\theta(\omega)\mathrm{d}F(\theta)\geq 0\end{aligned}$$

$$\begin{aligned}\mathbb{E}_{\theta\sim F}[\bar{u}(\delta_\theta)] = \sup_{S\in\mathcal{S}^{\times 1}} &\int_\Theta\int_\Omega u(\theta,\omega)\mathrm{d}S_\theta(\omega)\mathrm{d}F(\theta)\\ \text{s.t.} &\int_\Omega v(\theta,\omega)\mathrm{d}S_\theta(\omega)\geq 0,\ \forall\,\theta\in\Theta.\end{aligned}$$

Because the agent's distribution F and the shock θ in Round 1 are publicly known in the full information setting, it suffices that the principal designs a direct mechanism that assigns an outcome distribution for each possible shock value. When scaled by T, $T\bar{u}(F)$ and $T\mathbb{E}_{\theta\sim F}[\bar{u}(\delta_\theta)]$ are optimal performances achievable by the principal in the multiround problem with full information, respectively, under the ex-ante IR constraint and under the per-round interim IR constraint.

The next proposition relates these two quantities to Assumption 2:

Proposition 7. *For any* $F\in\mathcal{F}$, *if* $\bar{u}(F)=\mathbb{E}_{\theta\sim F}[\bar{u}(\delta_\theta)]$ *then* $\mathrm{OPT}(F,T)\leq\mathbb{E}_{\theta\sim F}[\mathrm{OPT}(\delta_\theta,T)]$.

Building on this broader sufficient condition, the next proposition gives more specific sufficient conditions that cover all applications in Section 4 and the technical report Balseiro et al. (2019). First, Assumption 2 holds for games with monetary payments that enter linearly into the utility functions of the principal and agent. This is because, given a mechanism satisfies the ex-ante IR constraint, the principal can use monetary transfers to satisfy the interim IR constraints without changing the expected overall utilities. The second part observes that Assumption 2 holds when agent utilities are always nonnegative, because both the ex-ante and interim IR constraints are satisfied automatically. When monetary transfers are allowed, and the principal wants to implement outcomes that lead to negative utility to the agent, our result may require negative payments to make the agent participate.

Proposition 8. *Assume the game is such that either (1) the outcome space factorizes as* $\Omega=\Omega^0\times\mathbb{R}$ *and the utility functions can be written* $u(\theta,(\omega^0,p))=u^0(\theta,\omega^0)+\alpha p$ *and* $v(\theta,(\omega^0,p))=v^0(\theta,\omega^0)-\beta p$ *for all outcomes* $(\omega^0,p)\in\Omega^0\times\mathbb{R}$ *for some functions* $u^0:\Theta\times\Omega^0\to\mathbb{R}$ *and* $v^0:\Theta\times\Omega^0\to\mathbb{R}$ *and scalars* $\alpha\geq 0$ *and* $\beta>0$, *or (2)* $v(\theta,\omega)\geq 0$ *for all* $\theta\in\Theta$ *and* $\omega\in\Omega$. *Then,* $\bar{u}(F)=\mathbb{E}_{\theta\sim F}[\bar{u}(\delta_\theta)]$ *for all* $F\in\mathcal{F}$ *and, therefore, Assumption 2 holds.*

5.3. Economic Intuition for Assumption 2 and Trade Across Shocks

Whether $\mathrm{OPT}(F,T)$ is extreme-point convex depends on two opposing effects of shock uncertainty. First, greater uncertainty in the shock distribution leads to greater information asymmetry between the principal and agent. To elicit the agent's private information, the principal needs to concede information rents, leading to lower principal utility. This suggests shock uncertainty can be undesirable from the principal's point of view. The second, more subtle effect, which we call *trade across shocks*, sometimes allows the principal to implement a broader set of outcomes when there is greater shock uncertainty. Subject to the agent's participation, the principal can implement outcomes that are beneficial to him but unfavorable for the agent under some realizations of the shock by, in return, offering other more favorable outcomes for the agent under other realizations of the shock. This kind of trading off is possible when there is shock uncertainty and, in this sense, shock uncertainty can be beneficial for the principal. Whether Assumption 2 holds depends on which of these effects dominates.

To see these effects directly, we can equivalently write the inequality in Assumption 2 as

$$T\bar{u}(F) - T\mathbb{E}_{\theta\sim F}[\bar{u}(\delta_\theta)] \le T\bar{u}(F) - \mathrm{OPT}(F,T),$$

by noting $\mathrm{OPT}(\delta_\theta, T) = T\bar{u}(\delta_\theta)$ (see Proposition E.1 in Online Appendix E).[5] The right-hand side is the effect of information asymmetry, which is commonly recognized to be the difference between the first-best (without IC) and "second-best" (with IC) performances in the Bayesian version of the multiround problem where F is known to the principal but not the realized shocks. The left-hand side is the effect of trade across shocks, which is the difference between optimal performances under different participation options, that is, the ex-ante IR versus per-round interim IR constraints, in the full information version of the problem where both F and shocks are known to the principal.

In many games where the principal knows F but not the shocks, the principal may link decisions across rounds subject to the agent's participation constraint so that $\mathrm{OPT}(F,T)/T \to \bar{u}(F)$ as $T \to \infty$ (see, e.g., Fudenberg et al. 1994, Jackson and Sonnenschein 2007). Therefore, the information asymmetry effect is negligible when the number of time periods is large and the trade across shocks effect ends up being the determining factor of whether the extreme-point convexity holds. In fact, Propositions 7 and 8 are statements about when the left-hand side is exactly 0.

To further explain trade across shocks, we focus on the full information setting, that is, F and θ_t are observed publicly. Consider the game in Table 1. There are two shocks or states of the world, reflecting the volume of rain in a given farming season. The principal cultivates a crop that requires heavy rain, whereas the agent cultivates a crop that grows only when the rain is light. There are two possible outcomes: they do not interact and each earns a utility of 0, or they share, which yields a utility of −1 for whoever must share his season-appropriate crop and a utility of 2 for the recipient (with the magnitude difference reflecting diminishing marginal returns).

For simplicity, assume each state θ^i of the world is equally likely under F. Then, a simple calculation shows $\bar{u}(F) = 3/4$ by sharing in θ^1 and sharing with probability $1/2$ in θ^2. On the other hand, $\bar{u}(\delta_{\theta^1}) = 0$ and $\bar{u}(\delta_{\theta^2}) = 0$, because the principal prefers not to share in θ^1 and the agent prefers not to in θ^2. In particular, $\bar{u}(F) > \mathbb{E}_{\theta\sim F}[\bar{u}(\delta_\theta)]$ because the principal can implement a broader set of outcomes when there is shock uncertainty. This gap reflects that sharing is individually rational for the agent ex-ante, but may not be after the state of the world is observed, and that the principal strictly does better under a sharing scheme of his choice that the agent agrees to before observing shocks.

Table 1. A Game (Ω, Θ, u, v) with Outcome Space $\Omega = \{\emptyset, \omega^1\}$, Shock Space $\Theta = \{\theta^1, \theta^2\}$, and Utility Functions u and v of the Principal and Agent in Matrix Representation

$u(\cdot,\cdot)$, $v(\cdot,\cdot)$	$\emptyset$	ω^1 = SHARE
θ^1 = LIGHT RAIN	0,0	2, −1
θ^2 = HEAVY RAIN	0,0	−1, 2

Note. The no-interaction outcome is denoted by $\emptyset$.

To tie back to Theorem 1 where the principal does not know F, Assumption 1 implies the principal has to account for the possibility that F is a point-mass distribution and, when shocks are constant over time, direct static mechanisms that do away with trade across shocks and satisfy the more stringent interim individual rationality constraint are optimal. When Assumption 2 holds, by protecting against point-mass distributions, the principal can hedge against any other distribution as regret is the largest for point masses. In other words, because there is no benefit from trading across shocks, direct static mechanisms are optimal when the principal does not know the distribution F.

Now suppose that the extreme-point convexity does not hold and a nondegenerate distribution is worst-case optimal. Because the trade across shock effect dominates, direct static mechanisms might be suboptimal for two reasons. First, as discussed in the farming example, the interim individual rationality constraint limits the set of outcomes implementable by the principal. Second, even if we relax the participation constraint to ex-ante individual rationality, a dynamic mechanism might be necessary to attain low regret. This follows because, to implement outcomes that are individually rational for the agent in expectation over his shocks, the principal needs a dynamic scheme to infer the agent's distribution of shocks. In Section 5.4, we exhibit a game where direct static mechanisms incur linear regret and provide a dynamic mechanism that attains sublinear regret by inferring the agent's distribution of shocks.

Finally, imagine introducing money to the game above. From Proposition 8, we know $\bar{u}(F) = \mathbb{E}_{\theta\sim F}[\bar{u}(\delta_\theta)]$, because the principal could satisfy the more stringent interim IR constraint by paying the agent to share if necessary. Put more simply, the principal could buy crops from the agent in light rain seasons and sell his own crops in heavy rain seasons. Money allows for economic interactions that are individually rational to the agent in every state of the world. When monetary transfers are infeasible, the principal can try to mimic them by guaranteeing the agent beneficial outcomes in other states of the world. Trade across shocks is this phenomenon where one mimics money by transferring utilities across different states of the world to incentivize the agent's participation.

5.4. On the Regret of Static Mechanisms Without Assumption 2

The previous subsection provides intuition that dynamics may be beneficial when Assumption 2 fails due to the possibility of trade across shocks. But because our discussion was focused on a full information setting, it remains to show that dynamics can be beneficial in games with partial information. Due to space considerations, all proofs of the following results are deferred to Online Appendix F.

On the positive side, we show that direct static mechanisms are still near-optimal when Assumption 2 nearly holds, in the sense that the gap $\mathrm{OPT}(F,T) - \mathbb{E}_{\theta\sim F}[\mathrm{OPT}(\delta_\theta,T)]$ is small for all $F \in \mathcal{F}$.

Theorem 2. *Suppose Assumption 1 holds. For any $\epsilon \geq 0$, if a mechanism $S \in \mathcal{S}^{\times 1}$ satisfies (1), then*

$$\mathrm{Regret}(S^{\times T},T) \leq \mathrm{Regret}(T) + \epsilon + \sup_{F\in\mathcal{F}} \{\mathrm{OPT}(F,T) - \mathbb{E}_{\theta\sim F}[\mathrm{OPT}(\delta_\theta,T)]\}.$$

On the negative side, we next show a game for which Assumption 2 does not hold and direct static mechanisms are suboptimal. For this game, we prove that direct static mechanisms suffer linear minimax regret but that implementing a dynamic mechanism leads to a sublinear minimax regret. The performance gap shows that dynamic or adaptive schemes can effectively take advantage of the history of past outcomes and reports by the agent when Assumption 2 does not hold.

Consider the game in Table 2. It is not difficult to show that Assumption 2 fails. Let $F = (f_1, f_2)$ be the agent's private distribution over Θ where the shock is θ^i with probability f_i for $i = 1, 2$. We assume $\mathcal{F} = \{(f_1,f_2) : f_1 + f_2 = 1, f_i \geq 0\}$ so that Assumption 1 holds. The next result characterizes the optimal performance achievable.

Proposition 9. *For the game in Table 2, we have $\mathrm{OPT}(F,T) = T \cdot \bar{u}(F)$ and $\bar{u}(F) = \min\{f_1, f_2\}$.*

For point-mass distributions, we have $\bar{u}(\delta_{\theta^1}) = \bar{u}(\delta_{\theta^2}) = 0$. For example, when $F = (1/2, 1/2)$, we have $\bar{u}(F) = 1/2$ but $\mathbb{E}_{\theta\sim F}[\bar{u}(\delta_\theta)] = f_1 \cdot \bar{u}(\delta_{\theta^1}) + f_2 \cdot \bar{u}(\delta_{\theta^2}) = 0$; that is, $\mathrm{OPT}(F,T) > \mathbb{E}_{\theta\sim F}[\mathrm{OPT}(\delta_\theta,T)]$.

Like the game in Table 1, in this game, the principal would like to sometimes implement outcomes that are worse for the agent than not participating. Because the game does not allow for monetary transfers, the principal can only do this by making desirable expected outcomes in future rounds contingent on the agent's participation—implementing what we have called trade across shocks. The main difficulty in establishing the next proposition is in showing that the principal can implement a form of trade across shocks without knowing the shock distribution or observing the realized shocks.

Proposition 10. *For the game in Table 2, a separation exists:*

1. *For every T, the minimax regret of direct static mechanisms that repeat a single-round direct IC/IR mechanism is $\inf_{S\in\mathcal{S}^{\times 1}} \sup_{F\in\mathcal{F}} \mathrm{Regret}(S^{\times T}, F, T) = T/2$.*
2. *For every T, there exists an incentive compatible dynamic mechanism A such that $\sup_{F\in\mathcal{F}} \mathrm{Regret}(A,F,T) = O((\ln T)^{1/2} T^{2/3})$.*

The first part of the proposition shows that any repetition of a single-round direct IC/IR mechanism necessarily incurs a linear regret. Because of the structure of the game, single-round direct IC/IR mechanisms are severely limited in generating utilities for the principal. For the second part, we design a dynamic mechanism A with two phases. In the first phase, the principal implements some default mechanism that places positive probabilities on outcomes ω^1 and ω^2 when the reported shocks are θ^1 and θ^2, respectively, and induces the agent to truthfully report his per-round shocks, as the present disutility from misreporting overwhelms any potential future gains. Then in the second phase, the principal estimates the agent's distribution F from the reports in the first phase and implements a better-tuned mechanism, which is a version of the optimal ex-ante mechanism that is perturbed to account for the statistical errors introduced by estimating F with limited samples. Using standard concentration inequalities, we can balance the loss of offering suboptimal mechanisms in the first and second phases. By choosing the number of rounds in the first phase to grow sublinearly relative to the time horizon, we can show that the dynamic mechanism incurs a sublinear regret.

Although we do not provide details due to the space consideration, (1) we can still show a separation between direct static mechanisms and dynamic mechanisms when the utility functions of the principal and agent are bounded, that is, bounded entries in Table 2; and (2) we can show the class of sequential screening mechanisms, which ask the agent to report his distribution and then implement a static mechanism based on the screened information, performs better than the class of (naive) direct static mechanisms considered above but still obtains the minimax regret that is linear in the time horizon (approximately, $0.217812 \cdot T$).

It is worth mentioning that, from Proposition 8, Assumption 2 holds in the modified version of the

Table 2. A Game (Ω, Θ, u, v) with Outcome Space $\Omega = \{\emptyset, \omega^1, \omega^2\}$, Shock Space $\Theta = \{\theta^1, \theta^2\}$, and Utility Functions u and v of the Principal and Agent in Matrix Representation

$u(\cdot,\cdot)$, $v(\cdot,\cdot)$	$\emptyset$	ω^1	ω^2
θ^1	0, 0	1, −1	0, −∞
θ^2	0, 0	1, −∞	0, 1

Note. The no-interaction outcome is denoted by $\emptyset$.

above game that allows for monetary transfers. In that case, trade across shocks is not needed to incentivize participation and Theorem 1 shows direct static mechanisms are optimal.

6. Extensions and Discussion

Our results can be extended in several directions. Due to space considerations, we describe these briefly below and provide details in the technical report Balseiro et al. (2019).

• **Multiplicative guarantees.** We prove analogous results in terms of multiplicative performance guarantees instead of regrets. Similar to the approximation and competitive ratios in the theoretical computer science literature, we consider the multiplicative performance guarantee that is the ratio of the principal utility and the optimal performance achievable, as in $\text{PrincipalUtility}(A, F, T)/\text{OPT}(F, T)$, and the principal's goal is to maximize the worst-case ratio. Under Assumptions 1 and 2, we can show that the multiround multiplicative guarantee is equal to an appropriately defined single-round multiplicative guarantee and an optimal mechanism, if it exists, is static in that it repeats a single-round mechanism.

• **Dual Perspective and Saddle-Point Theorems.** The multiround 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. We show that, under certain conditions, these problems are respectively equivalent to ones in which nature chooses first and then the principal optimizes his performance given nature's choice. These results provide 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 multiround problem.

• **Alternative benchmarks.** We show our results still hold for other alternative multiround benchmarks that are considered in the learning literature. Instead of the optimal performance achievable $\text{OPT}(F, T)$, we consider $T \cdot \bar{u}(F)$, which is a stronger benchmark by Proposition E.1 in Online Appendix E, and a weaker benchmark which naturally corresponds to the performance achievable by repeating the best fixed single-round direct IC/IR mechanism (i.e., the best fixed "action" in hindsight). The latter has been considered by Amin et al. (2013) and subsequent works.

• **Arbitrary shock processes.** Our results also apply in the general shock process setting where the agent's shocks can be serially correlated according to a stochastic process that is known to the agent but not to the principal. This is a natural generalization of the repeated i.i.d. setting considered thus far where the per-round shocks are drawn independently and identically from an underlying distribution. As the set of shock processes is more general, the multiround minimax regret problem is more challenging for the principal. Not surprisingly, the constant shock processes where the agent's shock is fixed over the whole time horizon are the corresponding counterparts of point-mass distributions, which are worst cases in the repeated i.i.d. setting.

• **Principal pessimism.** We consider a stronger notion of regret in which the agent plays a utility-maximizing strategy that is least favorable for the principal. Under this alternative tie-breaking possibility, the worst-case uncertainty that the principal faces is in both the agent's distribution and his utility-maximizing strategy. We can show our general result (Theorem 1) and those in Section 4 still hold with respect to this more robust notion of minimax regret.

• **Connections to maximin utility objective.** We discuss some connections to the maximin utility objective for revenue maximization in the dynamic selling problem with a single good (Section 4.1). Despite differences in the settings and objectives, our results with respect to the minimax regret objective and those in Carrasco et al. (2019) with respect to the maximin utility objective have similar analyses and solution structures. We show this is because both papers rely on essentially the same single-round saddle-point problem involving direct IC/IR mechanisms and show equivalence-type connections using saddle-point results.

• **More stringent participation constraints.** Our model as stated assumes an ex-ante participation constraint in the sense that the agent determines whether to participate or not in Round 0 while knowing his distribution but not the realization of future shocks. This participation constraint is standard in the dynamic screening literature (see, e.g., Courty and Li 2000). Our results extend to other, more stringent participation constraints that have been considered in the literature such as the dynamic individual rationality constraint (Kakade et al. 2013, Pavan et al. 2014) and the per-period individual rationality constraint (Krishna et al. 2013, Ashlagi et al. 2016, Balseiro et al. 2018). This is because the minimax regret can be achieved by repeating a single-round direct IC/IR mechanism, which naturally satisfies the latter participation constraints.

• **Discounting.** We assume no discounting in Theorem 1, but when the principal and agent discount future payoffs using the same discount factor $\gamma \in (0, 1)$, the same results would still hold with minimal changes and the minimax regret would be linear in the effective time horizon $T_\gamma := 1 + \gamma + \ldots + \gamma^{T-1}$.

7. Conclusion

In this paper, we proved false-dynamics results for a finite horizon setting where the principal and agent repeatedly play a game. Our results hold whenever the set of possible distributions for the agent includes all point masses and when the optimal performance achievable is extreme-point convex. The latter condition is satisfied by all games with linear dependence on monetary transfers or in which the agent utility function is always nonnegative. In particular, this includes the dynamic selling problem, the principal-agent model with hidden costs, and resource allocation without monetary transfers, and we determined the minimax regret and characterized an optimal dynamic mechanism that simply repeats a single-round mechanism in these applications. When either assumption does not hold, it is possible that a dynamic mechanism can outperform static mechanisms and we showed a separation in terms of performance between dynamic and static mechanisms for specific games. Furthermore, we showed our techniques extend in several directions.

For future research, it would be interesting to better understand the extreme-point convexity assumption and find a more general class of games where false-dynamics-type results hold. On the other hand, it would be also interesting to further explore where false-dynamics-type results do not hold and identify the class of games in which there is a separation between dynamic and static mechanisms, that is, dynamics strictly helps.

Other possible research directions include restricting the space of distributions and considering multiple agents. Point-mass distributions happen to be the right class of worst-case distributions in our analysis, and we have shown that false-dynamics-type results do not necessarily hold when we rule out point-mass distributions. We considered one strategic agent who is forward-looking and responds to the principal's mechanism. When there are multiple forward-looking agents, equilibrium considerations become important as the outcome may differ depending on whether the agents know each other's distribution or not, which, in turn, may affect the principal's optimal dynamic mechanism. A final interesting research direction is to study the design of robust dynamic mechanisms when the principal has access to samples from the agent's distribution of shocks.

Acknowledgments

The authors thank the area editor Prof. Johari, an anonymous associate editor, and three anonymous reviewers for providing a number of valuable comments that greatly improved the paper. The authors also thank Dirk Bergemann, Ying-Ju Chen, and the participants at the INFORMS Revenue Management & Pricing Conference, ACM EC Workshop on Learning in Presence of Strategic Behavior, and Cornell ORIE Young Researchers Workshop for valuable feedback.

Endnotes

[1] Commitment can be sustained, for example, by writing a contract that is enforced by a court of law, by reputation, or when planning horizons are short. Lack of commitment can be shown to lead to worse performance for the principal.

[2] It is possible to apply the revelation principle and restrict attention to direct mechanisms in which the agent reports his distribution and the realized shocks. Because a direct mechanism asks the agent to report his private information and nothing more, the mechanism should only be defined for shocks that have a positive probability under the reported distribution. This creates some complexities, which we sidestep by proving a weaker version of the revelation principle: under some additional assumptions, there exists an optimal direct static mechanism.

[3] Note that, in general, (2) is not equivalent to $\text{Regret}(1) := \inf_{A \in \mathcal{A}} \text{Regret}(A, 1)$ because direct static mechanisms have the more stringent interim participation constraint and they do not allow the principal to screen the agent in Round 0 based on his distribution. Under our assumptions, however, these two problems are equivalent.

[4] To see the connection, note that for any single-round direct IC/IR mechanism S, there exists a randomized posted pricing strategy with nearly matching interim allocation and payment rules, that is, over $[0,1]$ except a set of measure 0. For example, we can interpret a suitable extension and modification of the interim allocation rule of S as the cumulative distribution function from which posted prices are randomly drawn.

[5] Recall $T\bar{u}(F)$ and $T\mathbb{E}_{\theta\sim F}[\bar{u}(\delta_\theta)]$ are optimal performances achievable by the principal in the multiround problem with full information, respectively, under the ex-ante IR constraint and under the per-round interim IR constraint.

References

Amin K, Rostamizadeh A, Syed U (2013) Learning prices for repeated auctions with strategic buyers. Burges CJC, Bottou L, Welling M, Ghahramani Z, Weinberger KQ, eds. *Advances in Neural Information Processing Systems*, vol. 26 (Curran Associates, Inc., Red Hook, NY), 1169–1177.

Amin K, Rostamizadeh A, Syed U (2014) Repeated contextual auctions with strategic buyers. Ghahramani Z, Welling M, Cortes C, Lawrence ND, Weinberger KQ, eds. *Advances in Neural Information Processing Systems*, vol. 27 (Curran Associates, Inc., Red Hook, NY), 622–630.

Ashlagi I, Daskalakis C, Haghpanah N (2016) Sequential mechanisms with ex-post participation guarantees. *Proc. 2016 ACM Conf. Econom. Computat., EC '16* (Association for Computing Machinery, Maastricht, The Netherlands), 213–214.

Bakos Y, Brynjolfsson E (1999) Bundling information goods: Pricing, profits, and efficiency. *Management Sci.* 45(12):1613–1630.

Balseiro S, Kim A, Russo D (2019) On the futility of dynamics in robust mechanism design–technical report. Technical report, Columbia University. http://www.columbia.edu/˜srb2155/papers/futility_technical_report.pdf.

Balseiro SR, Mirrokni VS, Leme RP (2018) Dynamic mechanisms with martingale utilities. *Management Sci.* 64(11):5062–5082.

Baron DP, Besanko D (1984) Regulation and information in a continuing relationship. Inform. Econom. Policy 1(3):267–302.

Bergemann D, Schlag KH (2008) Pricing without priors. *J. Eur. Econom. Assoc.* 6(2-3):560–569.

Bergemann D, Schlag K (2011) Robust monopoly pricing. J. Econom. Theory 146(6):2527–2543.

Bergemann D, Castro F, Weintraub G (2017) The scope of sequential screening with ex post participation constraints. *Proc. 2017 ACM Conf. Econom. Computat. EC '17* (ACM, New York), 163–164.

Besbes O, Zeevi A (2009) Dynamic pricing without knowing the demand function: risk bounds and near optimal algorithms. *Oper. Res.* 57:1407–1420.

Börgers T, Krahmer D, Strausz R (2015) *An Introduction to the Theory of Mechanism Design* (Oxford University Press, Oxford, UK).

Carrasco V, Luz VF, Monteiro P, Moreira H (2019) Robust mechanisms: The curvature case. *J. Econom. Theory* 68(1):203–222.

Carrasco V, Luz VF, Monteiro PK, Moreira H (2018b) Robust mechanisms: the curvature case. *Econom. Theory.*

Carrasco V, Luz VF, Kos N, Messner M, Monteiro P, Moreira H (2018a) Optimal selling mechanisms under moment conditions. *J. Econom. Theory* 177:245–279.

Carroll G (2017) Robustness and separation in multidimensional screening. *Econometrica* 85(2):453–488.

Casella G, Berger RL (2002) *Statistical Inference*, vol. 2 (Duxbury Pacific Grove, CA).

Courty P, Li H (2000) Sequential screening. *Rev. Econom. Stud.* 67(4):697–717.

Fudenberg D, Levine D, Maskin E (1994) The folk theorem with imperfect public information. *Econometrica* 62(5):997–1039.

Golrezaei N, Javanmard A, Mirrokni V (2020) Dynamic incentive-aware learning: Robust pricing in contextual auctions. Preprint, submitted February 25, https://arxiv.org/abs/2002.11137.

Green E (1987) Lending and the smoothing of uninsurable income. *Contractual Arrange Intertemp Trade* 1:3–25.

Jackson MO, Sonnenschein HF (2007) Overcoming incentive constraints by linking decisions1. *Econometrica* 75(1):241–257.

Kakade S, Lobel I, Nazerzadeh H (2013) Optimal dynamic mechanism design and the virtual pivot mechanism. *Oper. Res.* 61(3):837–854.

Kanoria Y, Nazerzadeh H (2020) Dynamic reserve prices for repeated auctions: Learning from bids. Preprint, submitted February 18, https://arxiv.org/abs/2002.07331.

Kleinberg R, Leighton T (2003) The value of knowing a demand curve: Bounds on regret for online posted-price auctions. *Proc. 44th Annual IEEE Sympos. Foundations Comput. Sci. FOCS '03* (IEEE Computer Society, Washington, DC), 594.

Kocyigit C, Rujeerapaiboon N, Kuhn D (2018) Robust multidimensional pricing: Separation without regret. Preprint, submitted August 9, https://dx.doi.org/10.2139/ssrn.3219680.

Kos N, Messner M (2015) Selling to the mean. Preprint, submitted July 18, https://dx.doi.org/10.2139/ssrn.2632014.

Krähmer D, Strausz R (2015) Optimal sales contracts with withdrawal rights. *Rev. Econom. Stud.* 82(2):762.

Krishna RV, Lopomo G, Taylor CR (2013) Stairway to heaven or highway to hell: Liquidity, sweat equity, and the uncertain path to ownership. *RAND J. Econom.* 44(1):104–127.

Laffont J-J, Martimort D (2001) *The Theory of Incentives: The Principal-Agent Model* (Princeton University Press, Princeton, NJ).

Laffont J-J, Tirole J (1993) *A Theory of Incentives in Procurement and Regulation*, vol. 1, 1st ed. (The MIT Press, Cambridge, MA).

Mohri M, Munoz A (2014) Optimal regret minimization in posted-price auctions with strategic buyers. Ghahramani Z, Welling M, Cortes C, Lawrence ND, Weinberger KQ, eds. *Advances in Neural Information Processing Systems*, vol. 27 (Curran Associates, Inc., Red Hook, NY), 1871–1879.

Mohri M, Munoz A (2015) Revenue optimization against strategic buyers. Cortes C, Lawrence ND, Lee DD, Sugiyama M, Garnett R, eds. *Advances in Neural Information Processing Systems*, vol. 28 (Curran Associates, Inc., Red Hook, NY), 2530–2538.

Pavan A, Segal I, Toikka J (2014) Dynamic mechanism design: A Myersonian approach. *Econometrica* 82(2):601–653.

Pınar MÇ, Kızılkale C (2017) Robust screening under ambiguity. Math. Programming 163(1):273–299.

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

Thomas J, Worrall T (1990) Income fluctuation and asymmetric information: An example of a repeated principal-agent problem. *J. Econom. Theory* 51(2):367–390.

Wilson R (1987) Game-theoretic analysis of trading processes. Bewley T, ed. Adv. Econom. Theory *Fifth World Congress* (Cambridge University Press, Cambridge, UK), 33–70.

Santiago R. Balseiro is the Stanton Associate Professor of Business at the Graduate School of Business, Columbia University. His primary research interests are in the area of dynamic optimization, stochastic systems, and game theory with applications in revenue management and internet advertising.

Anthony Kim is an applied scientist at Amazon working on Sponsored Products. His research interests include revenue management, mechanism design, algorithmic game theory and online algorithms/learning. He received his PhD in Computer Science from Stanford University. Prior to Amazon, he was a postdoctoral researcher in the Decision, Risk, and Operations division at the Graduate School of Business, Columbia University.

Daniel Russo is an assistant professor in the Decision, Risk, and Operations Division of Columbia Business School. His research lies at the intersection of statistical machine learning and sequential decision-making, and he contributes to the fields of online optimization and reinforcement learning.

Electronic Companion:

On the Futility of Dynamics in Robust Mechanism Design

Santiago Balseiro, Anthony Kim, Daniel Russo

January 4, 2021

A Missing Proofs from Section 3

A.1 Proof of Theorem 1

As explained in Section 3.4, we use Lemmas 2 and 3.

Part 1): We use Lemma 2. Taking the infimum over all single-round direct IC/IR mechanisms S on the right-hand side of (4), we obtain for any incentive compatible dynamic mechanism $A \in \mathcal{A}$,

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

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

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

Now, it remains to show that $\mathrm{Regret}(T) \leq T \cdot \inf_{S \in \mathcal{S}^{\times 1}} \sup_{\theta \in \Theta} \mathrm{Regret}(S, \delta_\theta, 1)$. Fix an arbitrary $\epsilon > 0$. By the definition of infimum, there exists a single-round direct IC/IR mechanism S satisfying

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

Then,

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

where the first inequality is by Lemma 3 and the second by the choice of S. Since $S^{\times T}$ is a particular incentive compatible dynamic mechanism for T rounds (see the proof of Lemma 3), it follows that

$$\begin{aligned}\mathrm{Regret}(T) = \inf_{A \in \mathcal{A}} \sup_{F \in \mathcal{F}} \mathrm{Regret}(A, F, T) &\leq \sup_{F \in \mathcal{F}} \mathrm{Regret}(S^{\times T}, F, T) \\ &\leq T \cdot \inf_{S' \in \mathcal{S}^{\times 1}} \sup_{\theta \in \Theta} \mathrm{Regret}(S', \delta_\theta, 1) + \epsilon \,.\end{aligned}$$

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

$$\mathrm{Regret}(T) \leq T \cdot \inf_{S' \in \mathcal{S}^{\times 1}} \sup_{\theta \in \Theta} \mathrm{Regret}(S', \delta_\theta, 1) \,.$$

Part 2): For any $\epsilon \geq 0$, assume a single-round direct IC/IR 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,

$$\begin{aligned} \sup_{F \in \mathcal{F}} \text{Regret}(S^{\times T}, F, T) &\leq T \cdot \sup_{\theta \in \Theta} \text{Regret}(S, \delta_\theta, 1) \\ &\leq T \cdot \inf_{S' \in \mathcal{S}^{\times 1}} \sup_{\theta \in \Theta} \text{Regret}(S', \delta_\theta, 1) + \epsilon = \text{Regret}(T) + \epsilon \,, \end{aligned}$$

where the first inequality is by Lemma 3, the second inequality is by the property of S and the last equality is by Part 1.

Part 3): Equivalently, we show that $\arg\min_{A \in \mathcal{A}} \sup_{F \in \mathcal{F}} \text{Regret}(A, F, T)$ is non-empty if and only if $\arg\min_{S \in \mathcal{S}^{\times 1}} \sup_{\theta \in \Theta} \text{Regret}(S, \delta_\theta, 1)$ is non-empty. The if direction follows directly from Part 2. If there exists an optimal single-round direct IC/IR mechanism S^* to the single-round problem, the optimal solution S^* satisfies

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

By Part 2 (with $\epsilon = 0$),

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

It follows that the direct static mechanism that repeats S^* T times is an optimal incentive compatible dynamic mechanism in the multi-round problem and, hence, there exists an optimal dynamic mechanism in the multi-round problem. Note the incentive compatibility of the direct static mechanism $(S^*)^{\times T}$ follows from Lemma 2.

For the only-if direction, assume there exists an optimal incentive compatible dynamic mechanism A^* such that $\text{Regret}(T) = \sup_{F \in \mathcal{F}} \text{Regret}(A^*, F, T)$. By Lemma 2, there exists a single-round direct IC/IR mechanism S such that

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

By Part 1 that $\text{Regret}(T) = T \cdot \inf_{S' \in \mathcal{S}^{\times 1}} \sup_{\theta \in \Theta} \text{Regret}(S', \delta_\theta, 1)$, it follows that

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

The above implies that S is an optimal single-round mechanism because it achieves the single-round minimax regret. Hence, there exists an optimal single-round mechanism in the single-round problem. In particular, the single-round direct IC/IR mechanism constructed from A^*, $S(A^*)$, as in the proof of Lemma 1 is one such single-round mechanism that satisfies the statement of Lemma 2 and, hence, is an optimal solution to the single-round problem.

A.2 Proof of Lemma 1

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 $\text{Regret}(S, \delta_\theta, 1) = \text{OPT}(\delta_\theta, 1) - \int_\Omega u(\theta, \omega) dS_\theta(\omega)$. To prove the first and second statements, it suffices to show that a single-round direct mechanism S 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 as formulated in (3).

For the only-if direction, we proceed as follows. Assume an arbitrary incentive compatible mechanism $S \in \mathcal{S}^{\times 1}$. Recall that S being incentive compatible means $\text{AgentUtility}(S, \sigma^{\text{TR}}, F, 1) \geq \text{AgentUtility}(S, \tilde{\sigma}, F, 1)$ for any distribution F and any feasible agent strategy $\tilde{\sigma}$. For the point-mass distribution δ_θ for $\theta \in \Theta$ and the agent strategy $\tilde{\sigma}$ that deterministically reports $\mathsf{CONTINUE}$ in Round 0 and then shock $\theta' \in \Theta$, the inequality reduces to

$$\int_\Omega v(\theta, \omega) \mathrm{d}S_\theta(\omega) \geq \int_\Omega v(\theta, \omega) \mathrm{d}S_{\theta'}(\omega)$$

which is the IC constraint of (3), in terms of the outcome distributions of S. Now, for the point-mass distribution δ_θ for $\theta \in \Theta$ and the agent strategy $\tilde{\sigma}$ that deterministically reports QUIT in Round 0 and does not participate, the inequality $\text{AgentUtility}(S, \sigma^{\text{TR}}, F, 1) \geq \text{AgentUtility}(S, \tilde{\sigma}, F, 1)$ reduces to

$$\int_\Omega v(\theta, \omega) \mathrm{d}S_\theta(\omega) \geq 0$$

which is the IR constraint of (3). Therefore, the outcome distributions of S satisfy the IC and IR constraints.

We now show the if direction. Assume an arbitrary single-round direct mechanism S (with decision rule π_1) with its outcome distributions satisfying the IC and IR constraints in (3). Let σ^{TR} be the recommended strategy of truthfully reporting for the agent. We consider the following three cases depending on how an arbitrary alternative 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 $\mathsf{CONTINUE}$ in Round 0

The IC and IR constraints in (3) imply that for each possible shock $\theta_1 \in \Theta$ in Round 1, truthfully reporting the shock θ_1 is weakly better than deterministically reporting some other shock or PASS for the agent. That is,

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

for any $\theta \in \Theta$ and $\hat{m} \in \Theta \cup \{\mathsf{PASS}\}$. Note the left-hand side is equal to $\int_\Omega v(\theta, \omega) \mathrm{d}S_\theta(\omega)$ when $\theta_1 = \theta$, and the right-hand side is equal to $\int_\Omega v(\theta, \omega) \mathrm{d}S_{\theta'}(\omega)$ when $\theta_1 = \theta$ and $\hat{m} = \theta'$ and is equal to 0 when $\hat{m} = \mathsf{PASS}$. Hence, truthfully reporting in Round 1 is weakly better than any reporting strategy that is potentially randomized:

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

for any $\theta \in \Theta$, which follows by averaging the above inequality over possible $\hat{m}$ values under $\tilde{\sigma}$.

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

$$\text{AgentUtility}(S, \sigma^{\text{TR}}, F, 1) = \mathbb{E}_{\pi_1, \sigma^{\text{TR}}}[v(\theta_1, \omega_1)] \geq \mathbb{E}_{\pi_1, \tilde{\sigma}}[v(\theta_1, \omega_1)] = \text{AgentUtility}(S, \tilde{\sigma}, F, 1)\,.$$

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

The IR constraint in (3) implies that for each possible shock $\theta_1 \in \Theta$ in Round 1, truthfully reporting is weakly better for the agent than reporting PASS which yields the utility of 0. Then, for all $\theta \in \Theta$,

$$\mathbb{E}_{\pi_1, \sigma^{\text{TR}}}[v(\theta_1, \pi_1(\theta_1, h_1, z_1))|\theta_1 = \theta] \geq 0\,,$$

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

Since $\tilde{\sigma}$ reports QUIT in Round 0, the agent does not participate at all and $\text{AgentUtility}(S, \tilde{\sigma}, F, 1) = 0$. Clearly, $\text{AgentUtility}(S, \sigma^{\text{TR}}, F, 1) \geq \text{AgentUtility}(S, \tilde{\sigma}, F, 1)$.

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, we have

$$\begin{aligned}
&\text{AgentUtility}(S, \sigma^{\text{TR}}, F, 1) \geq \mathbb{E}_{\pi_1, \tilde{\sigma}}[v(\theta_1, \omega_1)|m_0 = \mathsf{CONTINUE}] \text{ and} \\
&\text{AgentUtility}(S, \sigma^{\text{TR}}, F, 1) \geq \mathbb{E}_{\pi_1, \tilde{\sigma}}[v(\theta_1, \omega_1)|m_0 = \mathsf{QUIT}]\,.
\end{aligned}$$

Hence,

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

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

As F was arbitrary, $\text{AgentUtility}(S, \sigma^{\text{TR}}, F, 1) \geq \text{AgentUtility}(S, \tilde{\sigma}, F, 1)$ in all cases for any distribution F. The above cases cover all possibilities for the alternative strategy $\tilde{\sigma}$ and it follows that S is incentive compatible, i.e., $\text{AgentUtility}(S, \sigma^{\text{TR}}, F, 1) \geq \text{AgentUtility}(S, \tilde{\sigma}, F, 1)$ for any distribution F and any feasible agent strategy $\tilde{\sigma}$.

For the last part of the lemma, we show the original objective of (3) and the alternative objective lead to the same value. Fix an arbitrary single-round direct mechanism $S \in \Delta(\Omega)^\Theta$. We have

$$\begin{aligned}
&\sup_{\theta' \in \Theta} \left\{ \text{OPT}(\delta_{\theta'}, 1) - \int_\Omega u(\theta', \omega) dS_{\theta'}(\omega) \right\} \\
&\qquad = \sup_{\theta' \in \Theta} \left\{ \int_\Theta \text{OPT}(\delta_\theta, 1) d\delta_{\theta'}(\theta) - \int_\Theta \int_\Omega u(\theta, \omega) dS_\theta(\omega) d\delta_{\theta'}(\theta) \right\} \\
&\qquad \leq \sup_{F \in \Delta(\Theta)} \left\{ \int_\Theta \text{OPT}(\delta_\theta, 1) dF(\theta) - \int_\Theta \int_\Omega u(\theta, \omega) dS_\theta(\omega) dF(\theta) \right\}\,,
\end{aligned}$$

where the first step is by rewriting the inner expressions 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

$$\begin{aligned}
\sup_{F\in\Delta(\Theta)} & \left\{\int_\Theta \mathrm{OPT}(\delta_\theta,1)\mathrm{d}F(\theta) - \int_\Theta\int_\Omega u(\theta,\omega)\mathrm{d}S_\theta(\omega)\mathrm{d}F(\theta)\right\} \\
&= \sup_{F\in\Delta(\Theta)} \int_\Theta \left(\mathrm{OPT}(\delta_\theta,1) - \int_\Omega u(\theta,\omega)\mathrm{d}S_\theta(\omega)\right)\mathrm{d}F(\theta) \\
&\le \sup_{F\in\Delta(\Theta)}\sup_{\theta\in\Theta}\left\{\mathrm{OPT}(\delta_\theta,1) - \int_\Omega u(\theta,\omega)dS_\theta(\omega)\right\} \\
&= \sup_{\theta\in\Theta}\left\{\mathrm{OPT}(\delta_\theta,1) - \int_\Omega u(\theta,\omega)dS_\theta(\omega)\right\}.
\end{aligned}$$

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 (3) and the version with the alternative objective are equivalent in terms of the optimal value and optimal solutions.

A.3 Additional Materials for Section 3.4

We prove Lemmas 2 and 3 in Appendices A.3.1 and A.3.2, respectively. We use Proposition 1 which is proved in Appendix A.3.3.

A.3.1 Proof of Lemma 2

We need to relate the two regret objectives of the multi-round and single-round problems and reduce the multi-round problem to the single-round problem for direct IC/IR mechanisms. The following lemma is useful. It follows from a revelation-principle-type argument and shows that if the agent's distribution is restricted to point-mass distributions, the principal's dynamic mechanism effectively reduces to a single-round direct mechanism with IC/IR properties and we can assume the agent's recommended strategy is the truthful reporting strategy σ^{TR}. See Appendix A.3.4 for the proof of the lemma.

Lemma 1. *For any incentive compatible dynamic mechanism A with a recommended agent strategy σ, there exists a single-round direct IC/IR mechanism, denoted $S(A)$, such that for any $\theta \in \Theta$,*

$$\mathrm{PrincipalUtility}(A,\sigma,\delta_\theta,T) = T\cdot \mathrm{PrincipalUtility}(S(A),\sigma^{TR},\delta_\theta,1)\,,$$

where σ^{TR} is the agent's truthful reporting strategy under which the agent participates (i.e., reports CONTINUE *in Round 0) and truthfully reports his shock.*

Using the above lemma, we prove Lemma 2 as follows:

Proof of Lemma 2. Fix an arbitrary incentive compatible dynamic mechanism $A \in \mathcal{A}$. Note that

$$\sup_{F\in\mathcal{F}} \mathrm{Regret}(A,F,T) \ge \sup_{\theta\in\Theta}\mathrm{Regret}(A,\delta_\theta,T)\,,$$

since point-mass distributions are a subset of general distributions $\mathcal{F}$ by Assumption 1. We can equivalently write the last expression as

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

where σ is the recommended agent strategy as part of the mechanism A in the first step, $S(A)$ in the second step is the single-round direct IC/IR mechanism derived from A as described in the proof of Lemma 1, and the second step follows from the same lemma and Proposition 1.

Putting the above together, for the single-round direct IC/IR mechanism $S(A)$, we have

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

A.3.2 Proof of Lemma 3

To prove Lemma 3, we need the following lemmas. The first one is about direct static mechanisms that are simply T repetitions of a single-round direct IC/IR mechanism. The second one is a technical step that involves a variant of the regret notion with a different benchmark other than $\text{OPT}(F,T)$. We prove these lemmas in Appendix A.3.4.

Lemma 2. *Let $S^{\times T}$ denote the direct static mechanism that repeats single-round direct IC/IR mechanism S for T rounds. For any single-round direct IC/IR mechanism $S\in\mathcal{S}^{\times 1}$, $S^{\times T}$ is incentive compatible and*

$$\text{PrincipalUtility}(S^{\times T},\sigma^{TR},F,T)=T\cdot\text{PrincipalUtility}(S,\sigma^{TR},F,1)$$

for any agent's distribution F, where σ^{TR} is the agent's truthful reporting strategy under which the agent participates (i.e., reports CONTINUE *in Round 0) and truthfully reports his shock(s).*

Lemma 3. *For any single-round direct IC/IR mechanism $S\in\mathcal{S}^{\times 1}$,*

$$\sup_{F\in\mathcal{F}}\{\mathbb{E}_{\theta\sim F}[\text{OPT}(\delta_\theta,1)]-\text{PrincipalUtility}(S,F,1)\}\leq\sup_{\theta\in\Theta}\text{Regret}(S,\delta_\theta,1)\,.$$

We now have:

Proof of Lemma 3. Let $S\in\mathcal{S}^{\times 1}$ be any 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 2. By the definition of Regret notion,

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

where σ^{TR} is the agent's truthful reporting strategy (i.e., the agent reports CONTINUE in Round 0 and truthfully reports his shocks in future rounds) that is the recommended strategy for direct mechanisms. Note that for any distribution $F \in \mathcal{F}$,

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

where the inequality is by Assumption 2 and the equality is by Proposition 1. Then,

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

where the last step is by Lemma 2. By Lemma 3, note the optimization problem in the last expression can be upper bounded as follows:

$$\sup_{F \in \mathcal{F}} \left\{\mathbb{E}_{\theta \sim F}[\mathrm{OPT}(\delta_\theta, 1)] - \mathrm{PrincipalUtility}(S, \sigma^{\mathrm{TR}}, F, 1)\right\} \le \sup_{\theta \in \Theta} \mathrm{Regret}(S, \sigma^{\mathrm{TR}}, \delta_\theta, 1)\,.$$

It follows that

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

A.3.3 Proofs of Propositions 1 and 2

Proof of Proposition 1. The proof follows straightforwardly from Proposition 3 in Appendix E. For any $\theta \in \Theta$,

$$\mathrm{OPT}(\delta_\theta, T) = T \cdot \bar{u}(\delta_\theta) = T \cdot \mathrm{OPT}(\delta_\theta, 1)\,,$$

by the second part of Proposition 3. Alternatively, we can prove the proposition directly using the same ideas in the proof of Proposition 3. We keep this presentation to avoid repeating proofs. $\square$

Proof of Proposition 2. In what follows, let $\widehat{\mathrm{Regret}} := \inf_{S \in \mathcal{S}^{\times 1}} \sup_{\theta \in \Theta} \mathrm{Regret}(S, \delta_\theta, 1)$. Fix an arbitrary incentive compatible dynamic mechanism A. Following the same line of reasoning in the beginning of the proof of Lemma 2, we have for any $\theta \in \Theta$,

$$\mathrm{Regret}(A, \delta_\theta, T) = T \cdot \mathrm{Regret}(S(A), \delta_\theta, 1)\,,$$

where $S(A)$ is the single-round direct IC/IR mechanism derived from A as described in the proof of Lemma 1. Note $\sup_{\theta \in \Theta} \mathrm{Regret}(S(A), \delta_\theta, 1) \ge \widehat{\mathrm{Regret}}$. By the definition of supremum, for any $\epsilon > 0$, there exists a point-mass distribution δ_{θ^*} for some $\theta^* \in \Theta$ such that

$$\mathrm{Regret}(S(A), \delta_{\theta^*}, 1) \ge \widehat{\mathrm{Regret}} - \frac{\epsilon}{T}\,.$$

Combining with the above observation, we then have

$$\mathrm{Regret}(A, \delta_{\theta^*}, T) \ge T \cdot \widehat{\mathrm{Regret}} - \epsilon\,.$$

Since Assumptions 1 and 2 hold, we have that $\mathrm{Regret}(T) = T \cdot \widehat{\mathrm{Regret}}$ by Theorem 1 and, hence, that

$$\mathrm{Regret}(A, \delta_{\theta^*}, T) \ge \mathrm{Regret}(T) - \epsilon\,. \qquad \square$$

A.3.4 Remaining Proofs from Appendix A.3

Proof of Lemma 1. First, we show a construction of a single-round direct IC/IR mechanism which will be our choice of $S(A)$ and then prove the claimed statements.

Let $\{\omega_{\theta,t}\}_{t=1}^{T}$ be a sequence of outcomes realized when the agent plays the recommended strategy σ against the principal's mechanism A when his distribution is δ_θ. When the agent reports QUIT in Round 0 and does not participate, the sequence is simply the no-interaction outcome in all rounds. Consider the following single-round direct mechanism S which is a collection of distributions S_θ on Ω indexed by $\theta \in \Theta$. For any $\theta \in \Theta$ and measurable set $W \subset \Omega$, we define

$$S_\theta(W) := \frac{1}{T}\sum_{t=1}^{T} \mathbb{P}_{\pi,\sigma}\left(\omega_{\theta,t} \in W | F = \delta_\theta\right) ,$$

where the expectation is taken over the randomness of π and σ (here, the shocks are deterministic and equal to θ). We can interpret S_θ as the time-averaged distribution of outcomes when the agent's distribution is the point-mass distribution δ_θ and the agent plays the recommended strategy σ.

Using the representation of S in terms of outcome distributions (as described in Section 3.3), we can show that S satisfies both IC and IR constraints as formulated in the optimization problem (3). For any $\theta, \theta' \in \Theta$,

$$\begin{aligned}
\mathbb{E}_{\omega\sim S_{\theta'}}[v(\theta,\omega)] &= \frac{1}{T}\mathbb{E}\left[\sum_{t=1}^{T} v(\theta,\omega_{\theta',t})\right] \\
&= \frac{1}{T}\,\mathrm{AgentUtility}(A,\sigma',\delta_\theta,T) \\
&\le \frac{1}{T}\,\mathrm{AgentUtility}(A,\sigma,\delta_\theta,T) \\
&= \frac{1}{T}\mathbb{E}\left[\sum_{t=1}^{T} v(\theta,\omega_{\theta,t})\right] \\
&= \mathbb{E}_{\omega\sim S_\theta}[v(\theta,\omega)]\,,
\end{aligned}$$

where σ' is an alternative strategy under which the agent, in particular, reports according the recommended strategy σ as if his distribution is $\delta_{\theta'}$ when his actual distribution is δ_θ and the inequality follows from that A is incentive compatible and σ is a utility-maximizing strategy for the agent when, in particular, his distribution is δ_θ. Similarly, for any $\theta \in \Theta$,

$$\begin{aligned}
\mathbb{E}_{\omega\sim S_\theta}[v(\theta,\omega)] &= \frac{1}{T}\mathbb{E}\left[\sum_{t=1}^{T} v(\theta,\omega_{\theta,t})\right] \\
&= \frac{1}{T}\,\mathrm{AgentUtility}(A,\sigma,\delta_\theta,T) \\
&\ge 0\,,
\end{aligned}$$

where the inequality follows because a utility-maximizing agent can guarantee the total utility of at least 0 by not participating. Since A is incentive compatible, the recommended strategy σ ensures the agent obtains a non-negative utility. Hence, S constructed above is a single-round direct IC/IR mechanism.

By construction, we have for any θ,

$$\begin{aligned}\text{PrincipalUtility}(A, \sigma, \delta_\theta, T) &= T \cdot \frac{1}{T}\mathbb{E}\left[\sum_{t=1}^{T} u(\theta, \omega_{\theta,t})\right] \\ &= T \cdot \mathbb{E}_{\omega \sim S_\theta}[u(\theta, \omega)] \\ &= T \cdot \text{PrincipalUtility}(S, \sigma^{\text{TR}}, \delta_\theta, 1)\,. \qquad \square\end{aligned}$$

Proof of Lemma 2. Fix an arbitrary single-round direct IC/IR mechanism $S \in \mathcal{S}^{\times 1}$ with decision rule $\tilde{\pi}$ and let $S^{\times T}$ be the direct static mechanism that repeats S. Abusing notations, we use σ^{TR} to denote the truthful reporting strategy (that reports CONTINUE in Round 0 and truthfully reports the shocks) to be the associated recommended strategy for the agent for both the single-round and multi-round direct mechanisms. In particular, we have that

$$\text{AgentUtility}(S, \sigma^{\text{TR}}, F, 1) \geq \text{AgentUtility}(S, \tilde{\sigma}, F, 1)\,,$$

for every probability distribution F over Θ and every feasible agent strategy $\tilde{\sigma}$.

First, we argue that $\text{PrincipalUtility}(S^{\times T}, \sigma^{\text{TR}}, F, T) = T \cdot \text{PrincipalUtility}(S, \sigma^{\text{TR}}, F, 1)$ for every distribution F. Since the same decision rule $\tilde{\pi}$ is used under $S^{\times T}$ by the principal and the same truthful reporting strategy is used by the agent in each round, the realized distribution of outcomes associated with each shock under F is identical across rounds and, hence, the principal utility restricted to each round is identical across rounds. The principal utility restricted to each round is exactly the principal utility under S (and the agent truthfully reports). Therefore, the principal utility under $S^{\times T}$ is T times the principal utility under S. Similarly, we have that the agent utility restricted to each round under $S^{\times T}$ is exactly the agent utility under S and, hence, that $\text{AgentUtility}(S^{\times T}, \sigma^{\text{TR}}, F, T) = T \cdot \text{AgentUtility}(S, \sigma^{\text{TR}}, F, 1)$ for every distribution F.

Now, we argue that $S^{\times T}$ is incentive compatible. For the sake of contradiction, assume there exists a distribution F' and an agent strategy $\sigma' = \{\sigma'_t\}_{0:T}$ such that $\text{AgentUtility}(S^{\times T}, \sigma^{\text{TR}}, F', T) < \text{AgentUtility}(S^{\times T}, \sigma', F', T)$. 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' outperforms σ^{TR}, there is a particular round t in which V_t is strictly greater than the agent utility achieved under the single-round mechanism S and σ^{TR}. We use the following claim:

Claim 1. *For any t, there is an agent strategy against S that achieves the expected agent utility and principal utility equal to V_t and U_t, respectively.*

Proof. Fix arbitrary $t \in [T]$. The agent can implement the t-th round strategy σ'_t as a standalone agent strategy against S by implementing σ'_0 in Round 0 and then σ'_t in Round 1. To implement σ'_t, the agent internally chooses randomness z_t and y_t and simulates the history $h_t^+ = (F', \theta_{1:t-1}, m_{0:t-1}, \omega_{1:t-1})$ that is needed for σ'_t. This is possible from the knowledge of S and $\{\sigma'_{t'}\}_{0:t-1}$. By construction, when the agent implements the above strategy against S, the expected agent utility and principal utility equal V_t and U_t, respectively. $\square$

By the above claim, there exists an alternative agent strategy against the single-round mechanism S that yields the expected agent utility equal to V_t which is strictly greater than $\text{AgentUtility}(S, \sigma^{\text{TR}}, F', 1)$. This contradicts that S is incentive compatible with respect to the recommended strategy σ^{TR}. $\square$

Proof of Lemma 3. Let $S \in \mathcal{S}^{\times 1}$ be an arbitrary single-round direct IC/IR mechanism and σ^{TR} be the recommended truthful reporting strategy for the agent. Using the representation of S in terms of outcome distributions as described in Section 3.3, we then have

$$\begin{aligned}
&\sup_{F \in \mathcal{F}} \{\mathbb{E}_{\theta \sim F}[\text{OPT}(\delta_\theta, 1)] - \text{PrincipalUtility}(S, \sigma^{\text{TR}}, F, 1)\} \\
&\qquad = \sup_{F \in \mathcal{F}} \left\{ \int_\Theta \text{OPT}(\delta_\theta, 1) \mathrm{d}F(\theta) - \int_\Theta \int_\Omega u(\theta, \omega) \mathrm{d}S_\theta(\omega) \mathrm{d}F(\theta) \right\} \\
&\qquad = \sup_{F \in \mathcal{F}} \int_\Theta \left(\text{OPT}(\delta_\theta, 1) - \int_\Omega u(\theta, \omega) \mathrm{d}S_\theta(\omega) \right) \mathrm{d}F(\theta) \\
&\qquad = \sup_{F \in \mathcal{F}} \int_\Theta \text{Regret}(S, \sigma^{\text{TR}}, \delta_\theta, 1) \mathrm{d}F(\theta) \\
&\qquad \le \sup_{\theta \in \Theta} \text{Regret}(S, \sigma^{\text{TR}}, \delta_\theta, 1) \,.
\end{aligned}$$

$\square$

B Additional Materials for Section 4.1

B.1 Proof of Proposition 3

As discussed already, Theorem 1 applies and Proposition 3 follows if we show that the single-round direct IC/IR mechanism S^* is an optimal solution to the single-round problem (2) and achieves the value of $\frac{1}{e}$. By Lemma 1, we can equivalently solve the optimization problem (3) with the alternative objective $\sup_{F \in \Delta(\Theta)} \left\{ \int_\Theta \text{OPT}(\delta_\theta, 1) dF(\theta) - \int_\Theta \int_\Omega u(\theta, \omega) dS_\theta(\omega) dF(\theta) \right\}$. Using the notation $\widehat{\text{Regret}}(S, F) := \int_\Theta \text{OPT}(\delta_\theta, 1) dF(\theta) - \int_\Theta \int_\Omega u(\theta, \omega) dS_\theta(\omega) dF(\theta)$, this optimization problem can be written as

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

where S can be any single-round direct IC/IR mechanism. Let $\widehat{\text{Regret}}$ be the corresponding optimal value.

We now solve the above optimization problem by proving the following saddle-point result, which is closely related to a similar result due to Bergemann and Schlag (2008).

Proposition 1. *Let S^* be the randomized posted pricing strategy given in Proposition 3, which is a single-round direct IC/IR mechanism, and the agent's distribution F^* be given by*

$$F^*(\theta) = \begin{cases} 0, & \text{if } \theta \in [0, \frac{1}{e}) \\ 1 - \frac{1}{e\theta}, & \text{if } \theta \in [\frac{1}{e}, 1) \\ 1, & \text{if } \theta = 1 \end{cases} .$$

Then, $\widehat{\text{Regret}} = \widehat{\text{Regret}}(S^*, F^*) = \frac{1}{e}$ *and*

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

for any single-round direct IC/IR mechanism $S \in \mathcal{S}^{\times 1}$ *and distribution* $F \in \Delta(\Theta)$*.*

Proof. When the principal implements a randomized posted pricing strategy, it is best for the agent to truthfully respond, that is, buy the item if the price is lower than his value, and, therefore, the randomized posted pricing strategy satisfies the IC and IR constraints. Furthermore, a randomized posted pricing mechanism has a direct implementation in that the principal can internally draw a random posted price and determine the allocation and payment given the agent's report θ.

First, we prove $\widehat{\text{Regret}}(S^*, F) \le \widehat{\text{Regret}}(S^*, F^*)$ for any agent's distribution F. Note we can represent any single-round direct IC/IR mechanism S with the corresponding interim rules (x, p) where, by standard arguments in mechanism design, the allocation rule x is monotonically non-decreasing and the payment rule p satisfies

$$p(\theta) = p(0) + x(\theta) \cdot \theta - \int_0^\theta x(t)dt\,, \quad \forall \theta \in [0, 1]$$

and $p(0) \le 0$. In particular, let (x^*, p^*) be the interim allocation and payment rules of S^* where $x^*(\theta) = 0$ for $\theta \in [0, \frac{1}{e})$ and $x^*(\theta) = 1 + \ln\theta$ for $\theta \in [\frac{1}{e}, 1]$ and $p^*(\theta) = x^*(\theta) \cdot \theta - \int_0^\theta x^*(t)dt$ for all θ.

It suffices to show that F^* is a solution to the following optimization problem:

$$\max_{F \in \Delta(\Theta)} \left\{ \int_\Theta \theta - p^*(\theta) dF(\theta) \right\}\,,$$

where F can be any distribution over Θ. Given x^*, we can simplify p^* as $p^*(\theta) = 0$ for $\theta \in [0, \frac{1}{e})$ and

$$p^*(\theta) = (1 + \ln\theta)\theta - \int_{\frac{1}{e}}^\theta 1 + \ln t dt = \theta - \frac{1}{e}\,,$$

for $\theta \in [\frac{1}{e}, 1]$. Then, the integrand in the objective function is, equivalently, $\theta\mathbf{1}\{\theta < \frac{1}{e}\} + \frac{1}{e}\mathbf{1}\{\theta \ge \frac{1}{e}\}$ and the optimization problem becomes

$$\max_{F \in \Delta(\Theta)} \left\{ \Pr_{\theta\sim F}\left(\theta < \frac{1}{e}\right) \cdot \mathbb{E}_{\theta\sim F}\left[\theta \mid \theta < \frac{1}{e}\right] + \Pr_{\theta\sim F}\left(\theta \ge \frac{1}{e}\right) \cdot \frac{1}{e} \right\}\,.$$

It follows that any distribution with its support contained in $[\frac{1}{e}, 1]$ is an optimal solution and F^* is one such distribution. Furthermore, we see that the optimization problem has the value of $\frac{1}{e}$ and, so, $\widehat{\text{Regret}} = \frac{1}{e}$.

Next, we show $\widehat{\text{Regret}}(S^*, F^*) \le \widehat{\text{Regret}}(S, F^*)$ for any single-round direct IC/IR mechanism S. Similar to the above argument, we show that S^* is a solution to:

$$\min_{(x,p)} \left\{ \int_\Theta \theta - p(\theta) dF^*(\theta) \ \text{s.t. (IC), (IR)} \right\}\,,$$

where (x, p) are over all possible interim rules satisfying the IC/IR constraints. By standard arguments, the payment rule p satisfies $p(\theta) = p(0) + x(\theta) \cdot \theta - \int_0^\theta x(t)dt$ for $\theta \in [0, 1]$ and $p(0) \leq 0$ and the allocation rule x is monotonically non-decreasing. Then, the above optimization problem becomes

$$\min_{\text{non-decreasing } x, p(0) \leq 0} \left\{ -p(0) + \int_0^1 \left(\theta - x(\theta) \cdot \theta + \int_0^\theta x(t)dt \right) f^*(\theta)d\theta \right\} ,$$

where f^* is the probability density function for F^* with $f^*(\theta) = 0$ for $\theta \in [0, \frac{1}{e})$, $f^*(\theta) = \frac{1}{e\theta^2}$ for $\theta \in [\frac{1}{e}, 1)$ and a point-mass of $\frac{1}{e}$ at $\theta = 1$. By changing the ordering of the integrals, $\int_0^1 \int_0^\theta x(t) f^*(\theta)dtd\theta = \int_0^1 \int_t^1 x(t) f^*(\theta)d\theta dt = \int_0^1 (1 - F^*(t))x(t)dt$, and the optimization problem is equivalently

$$\min_{\text{non-decreasing } x, p(0) \leq 0} \left\{ -p(0) + \mathbb{E}_{\theta \sim F^*}[\theta] + \int_0^1 \left(-\theta \cdot f^*(\theta) + (1 - F^*(\theta)) \right) x(\theta)d\theta \right\} .$$

In the integral, the expression $\phi(\theta) := -\theta \cdot f^*(\theta) + (1 - F^*(\theta))$ can be further simplified as 1 if $\theta \in [0, \frac{1}{e})$, 0 if $\theta \in [\frac{1}{e}, 1)$ and a point-mass of $-\frac{1}{e}$ if $\theta = 1$. Then, the objective function is equal to

$$-p(0) + \mathbb{E}_{\theta \sim F^*}[\theta] + \int_0^{\frac{1}{e}} 1 \cdot x(\theta)d\theta + \int_{\frac{1}{e}}^1 0 \cdot x(\theta)d\theta - \frac{1}{e} \cdot x(1) .$$

It follows that an optimal solution has $p(0) = 0$ and a non-decreasing $x(\cdot)$ such that $x(\theta) = 0$ for $\theta \in [0, \frac{1}{e})$, $x(\theta) \geq 0$ for $\theta \in [\frac{1}{e}, 1)$ and $x(\theta) = 1$ for $\theta = 1$ (in the almost everywhere sense for $\theta < 1$). Clearly, S^* satisfies these conditions and is, therefore, an optimal solution. $\square$

C Additional Materials for Section 4.2

C.1 Single-Round Problem

We provide further details on the single-round direct IC/IR mechanisms and a justification for the restriction to those with deterministic contracts for the single-round problem. Instead of (2), we equivalently consider the optimization problem (3) with the alternative objective $\sup_{F \in \Delta(\Theta)} \{ \int_\Theta \text{OPT}(\delta_\theta, 1)dF(\theta) - \int_\Theta \int_\Omega u(\theta, \omega)dS_\theta(\omega)dF(\theta) \}$, by Lemma 1. Introducing the notation $\widehat{\text{Regret}}(S, F) := \int_\Theta \text{OPT}(\delta_\theta, 1)dF(\theta) - \int_\Theta \int_\Omega u(\theta, \omega)dS_\theta(\omega)dF(\theta)$, this optimization problem can be written as

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

where S can be any single-round direct IC/IR mechanism. Let $\widehat{\text{Regret}}$ be the corresponding optimal value.

We can think of a single-round direct IC/IR mechanism as a collection of distributions S_θ on $\mathbb{R}^+ \times \mathbb{R}$ indexed by $\theta \in [\underline{\theta}, \bar{\theta}]$ such that when the agent reports θ, the outcome is determined by drawing from S_θ, i.e., $(\hat{q}, \hat{p}) \sim S_\theta$ for production level $\hat{q}$ and payment $\hat{p}$. Abusing notations, let $q(\theta) = \mathbb{E}_{(\hat{q}, \hat{p}) \sim S_\theta}[\hat{q}]$ and $p(\theta) = \mathbb{E}_{(\hat{q}, \hat{p}) \sim S_\theta}[\hat{p}]$ be the interim allocation and payment rules, respectively. Then, the IC and

IR constraints can be expressed as follows:

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

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

Fix an arbitrary single-round direct IC/IR mechanism. Let $V(\theta) = p(\theta) - \theta \cdot q(\theta)$ for $\theta \in [\underline{\theta}, \bar{\theta}]$. Note $V(\theta)$ is convex and by standard arguments (similar to the auction case, e.g., in Chapter 5 in Krishna (2009)), $q(\theta)$ is non-increasing and V is absolutely continuous and $V'(\theta) = -q(\theta)$ where the derivative exists. As $q(\theta)$ is nonnegative, $V(\theta)$ is non-increasing. Furthermore, we can write

$$p(\theta) = V(\bar{\theta}) + \theta \cdot q(\theta) + \int_{\theta}^{\bar{\theta}} q(x)dx\,, \text{ for } \theta \in [\underline{\theta}, \bar{\theta}]\,.$$

Using the notation $\widehat{\text{Regret}}$ and noting $\text{OPT}(\delta_\theta, 1) = \bar{R}(\theta)$, we have

$$\widehat{\text{Regret}}(S, F) = \mathbb{E}_{\theta \sim F}[\bar{R}(\theta)] - \mathbb{E}_{\theta \sim F, (\hat{q}, \hat{p}) \sim S_\theta}[R(\hat{q}) - \hat{p}]\,,$$

and

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

Since the revenue function $R(x)$ is concave, $\mathbb{E}_{(\hat{q},\hat{p}) \sim S_\theta}[R(\hat{q})] \leq R(\mathbb{E}_{(\hat{q},\hat{p}) \sim S_\theta}[\hat{q}])$ for any θ. Given any single-round direct IC/IR mechanism S, we can potentially improve (but not hurt) its performance by modifying S_θ to always return a deterministic production level $\hat{q}$ that is the average $q(\theta)$:

$$\mathbb{E}_{(\hat{q},\hat{p}) \sim S_\theta}[R(\hat{q}) - \hat{p}] \leq R(\mathbb{E}_{(\hat{q},\hat{p}) \sim S_\theta}[\hat{q}]) - \mathbb{E}_{(\hat{q},\hat{p}) \sim S_\theta}[\hat{p}] = R(q(\theta)) - p(\theta)\,.$$

Note the IC/IR constraints are still satisfied. Without loss in the minimax regret objective, we can restrict to those single-round IC/IR mechanisms that can be described in terms of a menu of deterministic contracts $(q(\theta), p(\theta))$ for $\theta \in [\underline{\theta}, \bar{\theta}]$. Using Ω^Θ to denote this restricted set of single-round direct mechanisms, the minimax regret for the single-round problem is equal to:

$$\widehat{\text{Regret}} = \inf_{\substack{(q,p) \in \Omega^\Theta: \\ \text{(IC),(IR)}}} \sup_{F \in \Delta(\Theta)} \int_\Theta \bar{R}(\theta) - (R(q(\theta)) - p(\theta))\, dF(\theta)\,,$$

and for $S = (q, p) \in \Omega^\Theta$ and $F \in \Delta(\Theta)$,

$$\widehat{\text{Regret}}(S, F) = \int_\Theta \bar{R}(\theta) - (R(q(\theta)) - p(\theta))\, dF(\theta)\,.$$

C.2 Proof of Proposition 4

By Theorem 1, to prove Proposition 4, it suffices that we prove the stated single-round direct IC/IR mechanism is an optimal solution to the single-round problem. For the single-round problem, we consider the equivalent optimization problem (3) with the alternative objective given in Lemma 1 and restrict attention to those single-round direct mechanisms that can be described in terms of a menu of deterministic contracts, denoted Ω^{Θ}; see details in Appendix C.1. More specifically, we consider the following optimization problem

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

where Ω^{Θ} denotes the restricted class of single-round direct mechanisms and $\widehat{\text{Regret}}(S,F) = \int_{\Theta} \bar{R}(\theta) - (R(q(\theta)) - p(\theta))\, dF(\theta)$ for $S = (q,p) \in \Omega^{\Theta}$ and $F \in \Delta(\Theta)$. Let $\widehat{\text{Regret}}$ be the corresponding optimal value of the optimization problem.

We show the single-round direct IC/IR mechanism given in Proposition 4 is an optimal solution to the above optimization problem via the following saddle-point result.

Proposition 2. *Let S^* be the single-round direct IC/IR mechanism corresponding to the menu of deterministic contracts $\{(q^*(\theta), p^*(\theta)\}_{\theta\in\Theta}$ given in Proposition 4 and the agent's distribution be given by a point-mass of $F^*(\underline{\theta})$, which can be 0, at $\theta = \underline{\theta}$ and a density $\frac{d}{d\theta}F^*(\theta)$ for $\theta \in (\underline{\theta}, \kappa]$ where*

$$F^*(\theta) = e^{-\int_{\theta}^{\kappa} \frac{1}{R'(q^*(x))-x}dx}\,,$$

for the same κ in the definition of S^*. Then, S^* and F^* are well-defined and the minimax regret is $\widehat{\text{Regret}} = \widehat{\text{Regret}}(S^*, F^*) = \int_{\underline{\theta}}^{\bar{\theta}} q^*(x)dx$ and is strictly positive, and

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

for any single-round direct IC/IR mechanism $S \in \mathcal{S}^{\times 1}$ and distribution $F \in \Delta(\Theta)$.

We first show that single-round direct IC/IR mechanism S^* and distribution F^* in the statement of Proposition 2 are well-defined with the claimed characterizations in the following lemma and then prove Proposition 2.

Lemma 4. *The single-round direct IC/IR mechanism S^* and distribution F^* in Propositions 4 and 2 are well-defined. Furthermore, q^* is continuous over the shock space and, in particular, strictly decreases over $[\underline{\theta}, \kappa]$.*

Proof. We proceed in two steps showing S^* and then F^* are well-defined.

(Single-Round Direct IC/IR Mechanism S^*): It suffices to show that a strictly decreasing solution q^* exists to the ordinary differential equation

$$\begin{aligned} &\frac{dq}{d\theta}(\theta) = \frac{-\bar{q}(\theta)}{R'(q(\theta)) - \theta}\,, \text{ for } \theta \in (\underline{\theta}, \kappa) \\ &\lim_{\theta\to\underline{\theta}^+} q(\theta) = \bar{q}(\underline{\theta})\,, \end{aligned} \tag{C-1}$$

for some κ to be determined and then complete it for $\theta = \underline{\theta}$ and $\theta > \kappa$ accordingly. We follow similar reasoning steps as in the proof of Lemma 2 in Carrasco et al. (2018). We equivalently solve the following differential system, with the roles of q and θ interchanged:

$$\begin{aligned} &\frac{d\theta}{dq}(q) = \frac{R'(q) - \theta(q)}{-\bar{q}(\theta(q))}\,, \text{ for } q \le \bar{q}(\underline{\theta}) \\ &\theta(\bar{q}(\underline{\theta})) = \underline{\theta}\,. \end{aligned} \tag{C-2}$$

As we show, this system has a solution θ^* that is strictly decreasing over a suitable interval and we can invert the relationship between θ and q to obtain a solution q^* to the original differential system with a well-defined κ.

Recall a solution to an ordinary differential equation (ODE) is a continuously differentiable function defined on some interval satisfying the specified relations. Let $\psi(q,\theta) := \frac{R'(q)-\theta}{-\bar{q}(\theta)}$ defined on the domain $D = (0, \bar{q}(\underline{\theta}) + \epsilon] \times [\underline{\theta} - \epsilon, \bar{\theta} + \epsilon]$ for arbitrarily small $\epsilon > 0$; ϵ is there to make the domain an open set. By the assumptions on R, ψ is continuous on the domain. Furthermore, it is continuously differentiable on any closed set of the domain and, hence, locally Lipschitz with respect to θ. For any initial value point in D, there exists a unique solution to the differential equation $\frac{d\theta}{dq}(q) = \psi(q,\theta)$ in a neighborhood of the initial value point (e.g., Theorem 3.1 in Hale (1969)). In particular, the above system of differential equation (C-2) has a unique solution θ^* in a neighborhood of the point $(\bar{q}(\underline{\theta}), \underline{\theta})$.

Let $(\underline{q}, \bar{q}(\underline{\theta})]$ be the left maximal interval of definition of the ordinary differential equation (C-2). We show θ^* is strictly decreasing in this interval. Note if a solution $\theta(\cdot)$ has $\theta'(q) = 0$, then

$$\frac{d^2\theta}{dq^2}(q) = \psi_1(q,\theta) + \psi_2(q,\theta) \cdot \frac{d\theta}{dq} = -\frac{R''(q)}{\bar{q}(\theta)} > 0,$$

where ψ_i denotes the partial derivative with respect to the i-th parameter. Since $(\theta^*)'(\bar{q}(\underline{\theta})) = 0$, θ^* is strictly convex at $q = \bar{q}(\underline{\theta})$ and decreases over $[\bar{q}(\underline{\theta}) - \epsilon, \bar{q}(\underline{\theta})]$ for sufficiently small $\epsilon > 0$. Fix an arbitrary $q \in (\underline{q}, \bar{q}(\underline{\theta}))$ and assume θ^* achieves the maximum at some $x \in [q, \bar{q}(\underline{\theta})]$. Note x cannot be in the interior because the first-order condition $(\theta^*)'(x) = 0$ is satisfied and it would mean θ^* is strictly convex and is increasing to the left or right of x. By the above observation, x cannot be $\bar{q}(\underline{\theta})$. Hence, the maximum is achieved at the left-end $x = q$. As q was arbitrary, the argument extends and it implies θ^* is strictly decreasing over $(\underline{q}, \bar{q}(\underline{\theta})]$.

Now, we invert θ^* to obtain q^* that is a solution to the original differential system (C-1) that we want to solve. Let $\theta^*(\underline{q}) = \sup_{q \in (\underline{q}, \bar{q}(\underline{\theta})]} \theta^*(q)$ which may be ∞. If $\theta^*(\underline{q}) < \bar{\theta}$, then $(\theta^*)'(\underline{q}) = \lim_{q \to \underline{q}^+} \psi(q, \theta^*(q))$ would be equal to $\frac{R'(\underline{q}) - \theta^*(\underline{q})}{-\bar{q}(\theta^*(\underline{q}))}$ which is not defined, more specially, $R'(\underline{q})$ in the numerator, if $\underline{q} = 0$ but defined if $\underline{q} > 0$. Since we chose the left maximal interval of definition of the ODE, it must be that $\underline{q} = 0$. Then, we let $\kappa = \theta^*(\underline{q})$ and truncate the solution θ^* so that its range is exactly $[\underline{\theta}, \kappa)$. We let q^* be the inverted curve of the truncated solution on $[\underline{\theta}, \kappa)$ which strictly decreases and converges to 0 over the interval and extend $q^*(\theta) = 0$ for $\theta \in [\kappa, \bar{\theta}]$.

In the other case when $\theta^*(\underline{q}) \ge \bar{\theta}$, we truncate the solution θ^* such that its range is exactly $[\underline{\theta}, \bar{\theta}]$ and consider q^* to be the corresponding inverted solution over the interval $[\underline{\theta}, \bar{\theta}]$. By construction, q^* satisfies the desired differential system and stays positive over the whole interval. We choose $\kappa = \bar{\theta}$.

In both cases, since θ^* is continuous at $q = \bar{q}(\underline{\theta})$ with $\theta^*(\bar{q}(\underline{\theta})) = \underline{\theta}$, we have $q^*(\underline{\theta}) = \bar{q}(\underline{\theta})$ and $\lim_{\theta \to \underline{\theta}^+} q^*(\theta) = \bar{q}(\underline{\theta})$. Also, by our choice of κ, $\lim_{\theta \to \kappa^-} q^*(\theta) = q^*(\kappa)$. That is, q^* is continuous over the whole interval $[\underline{\theta}, \bar{\theta}]$.

(Distribution F^*): Given that we have a solution q^* that is strictly decreasing over $[\underline{\theta}, \kappa]$ and continuously differentiable over $(\underline{\theta}, \kappa)$, the fraction $\frac{1}{R'(q^*(\theta))-\theta}$ is well-defined and positive for $\theta \in (\underline{\theta}, \kappa)$. We argue that the integral $\int_\theta^\kappa \frac{1}{R'(q^*(x))-x}dx$ exists over the same interval. If $\kappa = \bar{\theta}$ and q^* stays positive, the integrand is well-defined and continuous over the compact set. Hence, the integral exists. If $\kappa < \bar{\theta}$, then the integrand goes to 0 as x approaches κ and thus bounded. In this case, again, the integral exists.

As θ approaches $\underline{\theta}$, the integral $\int_\theta^\kappa \frac{1}{R'(q^*(x))-x}dx$ can potentially grow to ∞. But, $F^*(\theta) = e^{-\int_\theta^\kappa \frac{1}{R'(q^*(x))-x}dx}$ is absolutely continuous over $(\underline{\theta}, \kappa]$ and the distribution F^* can be described with a point-mass of $\lim_{\theta' \to \underline{\theta}^+} F^*(\theta')$ at $\theta = \underline{\theta}$, which can be 0, and the absolute continuous part with density $f^*(\theta) = \frac{d}{d\theta}F^*(\theta) = F^*(\theta) \cdot \frac{1}{R'(q^*(\theta))-\theta}$. □

Proof of Proposition 2. We restrict, without loss, to those single-round direct mechanisms that can be described in terms of a menu of deterministic contracts $(q(\theta), p(\theta))$ for $\theta \in [\underline{\theta}, \bar{\theta}]$ in our analysis; we use Ω^Θ to denote this class of mechanisms. By well-definedness, we mean both S^* and F^* exist with the stated characterizations. In particular, it would mean that q^* is a continuous monotone function over $[\underline{\theta}, \bar{\theta}]$ and is integrable. The well-definedness of S^* and F^* has been proved in Lemma 4.

For the first part of the saddle-point result, we show that F^* is an optimal solution to $\max_F \widehat{\text{Regret}}(S^*, F)$ which is equivalent to:

$$\max_{F \in \Delta(\Theta)} \left\{ \int_\Theta \bar{R}(\theta) - (R(q^*(\theta)) - p^*(\theta))\, dF(\theta) \right\}.$$

By the definition of p^*, the optimization problem is equivalent to

$$\max_{F \in \Delta(\Theta)} \left\{ \int_\Theta \left(\bar{R}(\theta) - R(q^*(\theta)) + \theta \cdot q^*(\theta) + \int_\theta^{\bar{\theta}} q^*(x)dx \right) dF(\theta) \right\}.$$

The integrand is continuous and its derivative with respect to θ is

$$-\bar{q}(\theta) - (R'(q^*(\theta)) - \theta) \cdot (q^*)'(\theta),$$

where we used $\bar{R}'(\theta) = -\bar{q}(\theta)$. Since $(q^*)'(\theta) = \frac{-\bar{q}(\theta)}{R'(q^*(\theta))-\theta}$ for $\theta \in (\underline{\theta}, \kappa)$, the derivative is equal to 0 over the same interval. For $\theta \in [\kappa, \bar{\theta}]$, the integrand is equal to $\bar{R}(\theta)$ and the derivative is equal to $-\bar{q}(\theta)$, which is negative. Since q^* is continuous, it follows that the integrand stays constant for $\theta \in [\underline{\theta}, \kappa]$ and then decreases for $\theta \in [\kappa, \bar{\theta}]$. Since F^* has support equal to exactly $[\underline{\theta}, \kappa]$, it maximizes the objective and is an optimal solution, as desired.

Similarly, for the second part, we show that S^* is an optimal solution to

$$\min_{\substack{S \in \Omega^\Theta: \\ \text{(IC),(IR)}}} \left\{ \int_\Theta \bar{R}(\theta) - (R(q(\theta)) - p(\theta))\, dF^*(\theta) \right\}.$$

By the standard arguments (see Appendix C.1), it suffices to show that S^* is an optimal solution to the following equivalent problem:

$$\min_{\text{non-increasing } q, V(\bar{\theta})\geq 0} \left\{ \int_{\Theta} \left(\bar{R}(\theta) - R(q(\theta)) + V(\bar{\theta}) + \theta \cdot q(\theta) + \int_{\theta}^{\bar{\theta}} q(x)dx \right) dF^*(\theta) \right\},$$

where $V(\theta) = p(\theta) - \theta \cdot q(\theta)$ for $\theta \in [\underline{\theta}, \bar{\theta}]$.

Note F^* has a point-mass of $\lim_{\theta' \to \underline{\theta}^+} F^*(\theta')$ which we, for notational convenience, equate to $F^*(\underline{\theta})$ at $\theta = \underline{\theta}$. But it is otherwise absolutely continuous and has a corresponding density function. We denote the cumulative function without the point-mass at $\theta = \underline{\theta}$ by F^*_- with corresponding density $f^*_-(\theta) = \frac{d}{d\theta} F^*(\theta)$ such that $F^*(\theta) = F^*(\underline{\theta}) + F^*_-(\theta)$ for $\theta \in [\underline{\theta}, \bar{\theta}]$.

Then, we can rewrite the objective function, denoted OBJ, as follows. Note that

$$\begin{aligned} \text{OBJ} = \mathbb{E}_{\theta \sim F^*}[\bar{R}(\theta)] + V(\bar{\theta}) - F^*(\underline{\theta}) \left(R(q(\underline{\theta})) - \underline{\theta} \cdot q(\underline{\theta}) - \int_{\underline{\theta}}^{\bar{\theta}} q(x)dx \right) \\ - \int_{\underline{\theta}}^{\bar{\theta}} f^*_-(\theta) \left(R(q(\theta)) - \theta \cdot q(\theta) - \int_{\theta}^{\bar{\theta}} q(x)dx \right) d\theta \,. \end{aligned}$$

The last term is equivalently $-\int_{\underline{\theta}}^{\bar{\theta}} f^*_-(\theta) \left(R(q(\theta)) - \theta \cdot q(\theta) \right) d\theta + \int_{\underline{\theta}}^{\bar{\theta}} F^*_-(\theta) q(\theta) d\theta$ where we changed the order of integrals. Then,

$$\begin{aligned} \text{OBJ} = \mathbb{E}_{\theta \sim F^*}[\bar{R}(\theta)] + V(\bar{\theta}) - F^*(\underline{\theta}) \left(R(q(\underline{\theta})) - \underline{\theta} \cdot q(\underline{\theta}) \right) \\ - \int_{\underline{\theta}}^{\bar{\theta}} f^*_-(\theta) \left(R(q(\theta)) - \left(\theta + \frac{F^*(\theta)}{f^*_-(\theta)} \right) \cdot q(\theta) \right) d\theta \,. \end{aligned}$$

We now show that S^* minimizes OBJ pointwise. Since $q^*(\underline{\theta}) = \bar{q}(\theta)$, the third term is minimized. Note $\frac{\partial}{\partial \theta} \ln F^*(\theta) = \frac{f^*_-(\theta)}{F^*(\theta)} = \frac{1}{R'(q^*(\theta)) - \theta}$ for $\theta \in [\underline{\theta}, \kappa]$. Then, $R'(q^*(\theta)) = \theta + \frac{F^*(\theta)}{f^*_-(\theta)}$ for $\theta \in [\underline{\theta}, \kappa]$ which is exactly the support of f^*_- and where q^* is nonnegative. It follows that the integrand in the fourth term is minimized pointwise and, thus, the fourth term is minimized. Finally, note that $V(\underline{\theta}) = 0$ for S^* and that the second term is minimized. Overall, the objective is minimized and S^* is an optimal solution. This completes the proof that S^* and F^* form a saddle point.

To compute the minimax regret, we recall from the analysis for the first part of the saddle-point result that the integrand, equivalently, $\widehat{\text{Regret}}(S^*, \theta)$, is constant for $\theta \in [\underline{\theta}, \kappa]$ and decreases for $\theta > \kappa$. It follows that the minimax regret is equal to the integrand evaluated at $\theta = \underline{\theta}$ which is $\int_{\underline{\theta}}^{\bar{\theta}} q^*(x)dx$. The minimax regret is clearly nonnegative because q^* is nonnegative. We argue it is strictly greater than 0. For the sake of contradiction, assume that $\int_{\underline{\theta}}^{\bar{\theta}} q^*(x)dx = 0$. As q^* is non-increasing and nonnegative, it would follow that $q^*(\theta) = 0$ for $\theta > \underline{\theta}$. Together with $q^*(\underline{\theta}) = \bar{q}(\theta) > 0$, it would contradict the continuity of q^* which is implied by the well-definedness of S^*. □

D Additional Materials for Section 5.1

D.1 Proof of Proposition 5

We characterize an optimal mechanism using the relax and verify approach of Kakade et al. (2013). Using that shocks are multiplicatively separable, we can write $\theta_t = \tau\gamma_t$ with γ_t drawn i.i.d. from G. We consider a relaxed environment in which γ_t are public and observable by the principal—the agent's only private information is the parameter τ. By the revelation principle, we can restrict attention to direct mechanisms in which the agent reports the parameter τ in Round 0. We denote by $x_t(\tau, \gamma_{1:t})$ and $p_t(\tau, \gamma_{1:t})$ the allocation and payment, respectively, in Round t when the report is τ and the γ-component of the shocks are $\gamma_{1:t} = (\gamma_1, \ldots, \gamma_t)$. Under the ex-ante participation constraint, the optimal performance achievable when the agent's private distribution $F(\cdot;\tau)$ is known (i.e., the parameter τ is known) is $\mathrm{OPT}(F(\cdot;\tau), T) = T\mathbb{E}_{\theta\sim F(\cdot;\tau)}[\theta] = T\tau\mathbb{E}[\gamma]$ because the principal can simply charge an entry-fee equal to the agent's expected value and then allocate the items for free over the rounds. For the entry fee, the principal can, for example, require the evenly-split fixed constant payment $\tau\mathbb{E}[\gamma]$ over the rounds that the agent has to pay upon participating. Therefore, the minimax regret in the multi-round problem is lower bounded as $\mathrm{Regret}(T) \geq \mathrm{Regret}^{\mathrm{RELAX}}(T)$ where

$$\begin{aligned}
\mathrm{Regret}^{\mathrm{RELAX}}(T) = \inf_{(x,p)} \sup_{\tau\in[0,1]} & \left\{ T\tau\mathbb{E}[\gamma] - \mathbb{E}_{\gamma_{1:T}}\left[\sum_{t=1}^{T} p_t(\tau,\gamma_{1:t})\right]\right\} \\
\text{s.t. } V(\tau) = & \ \mathbb{E}_{\gamma_{1:T}}\left[\sum_{t=1}^{T} \tau\gamma_t x_t(\tau,\gamma_{1:t}) - p_t(\tau,\gamma_{1:t})\right] \geq 0\,, \quad \forall \tau\in[0,1]\,, \\
V(\tau) \geq & \ \mathbb{E}_{\gamma_{1:T}}\left[\sum_{t=1}^{T} \tau\gamma_t x_t(\tau',\gamma_{1:t}) - p_t(\tau',\gamma_{1:t})\right]\,, \quad \forall \tau,\tau'\in[0,1]\,, \\
0 \leq & \ x_t(\tau,\gamma_{1:t}) \leq 1\,, p_t(\tau,\gamma_{1:t}) \in \mathbb{R}\,.
\end{aligned}$$

The first constraint is an individual rationality constraint that guarantees that the ex-ante utility of the agent when his true parameter is τ, which is denoted by $V(\tau)$, is non-negative. The second constraint is an incentive compatibility constraint imposing that the agent is better off reporting his true parameter.

We now show that the relaxed problem can be further lower bounded by a single-round problem, i.e., $\mathrm{Regret}^{\mathrm{RELAX}}(T) \geq T\widehat{\mathrm{Regret}}^{\mathrm{RELAX}}$ where

$$\begin{aligned}
\widehat{\mathrm{Regret}}^{\mathrm{RELAX}} = \inf_{(\hat{x},\hat{p})} \sup_{\tau\in[0,1]} & \left\{\mathbb{E}[\gamma]\left(\tau - \hat{p}(\tau)\right)\right\} \\
\text{s.t. } \hat{V}(\tau) = & \ \tau\hat{x}(\tau) - \hat{p}(\tau) \geq 0\,, \quad \forall\tau\in[0,1]\,, \\
\hat{V}(\tau) \geq & \ \tau\hat{x}(\tau') - \hat{p}(\tau')\,, \quad \forall\tau,\tau'\in[0,1]\,, \\
0 \leq & \ \hat{x}(\tau) \leq 1\,, \hat{p}(\tau)\in\mathbb{R}\,.
\end{aligned}$$

We prove the above claim by showing that every feasible mechanism for the multi-round problem corresponding to $\mathrm{Regret}^{\mathrm{RELAX}}(T)$ induces a feasible mechanism for the single-round problem corresponding to $\widehat{\mathrm{Regret}}^{\mathrm{RELAX}}$ that achieves the same objective value (divided by T). For any multi-round mechanism (x, p), consider the single-round mechanism $(\hat{x}, \hat{p})$ given by

$$\hat{x}(\tau) = \frac{1}{T\mathbb{E}[\gamma]}\mathbb{E}_{\gamma_{1:T}}\left[\sum_{t=1}^{T}\gamma_t x_t(\tau, \gamma_{1:t})\right] \quad \text{and} \quad \hat{p}(\tau) = \frac{1}{T\mathbb{E}[\gamma]}\mathbb{E}_{\gamma_{1:T}}\left[\sum_{t=1}^{T}p_t(\tau, \gamma_{1:t})\right]$$

for report $\tau \in [0,1]$, which is obtained by averaging the multi-round mechanism over time and the public γ-components of the shocks. By construction, we have that $V(\tau) = T\mathbb{E}[\gamma](\tau\hat{x}(\tau) - \hat{p}(\tau))$ and the IR constraint holds for the single-round mechanism $(\hat{x}, \hat{p})$ because $V(\tau) \geq 0$ for all $\tau \in [0,1]$. The IC constraint similarly holds. Additionally, since $\gamma_t \geq 0$ and $x_t(\cdot,\cdot) \in [0,1]$, we have that $0 \leq \hat{x}(\tau) \leq 1$ for $\tau \in [0,1]$. Similarly, we have that $\hat{p}(\tau) \in \mathbb{R}$ for $\tau \in [0,1]$. For any $\tau \in [0,1]$, the inner regret objective value for the multi-round mechanism can be equivalently written as $T\mathbb{E}[\gamma]\,(\tau - \hat{p}(\tau))$ which equals the corresponding objective value for the constructed single-round mechanism (when divided by T), and the claim follows.

Note the single-round problem corresponding to $\widehat{\text{Regret}}^{\text{RELAX}}$, when divided by $\mathbb{E}[\gamma]$, is exactly the single-round problem (3) in the dynamic selling problem for revenue maximization in Section 4.1. Then, Proposition 1 used in the proof of Proposition 3 immediately implies that the regret of the single-round problem is $\widehat{\text{Regret}}^{\text{RELAX}} = \mathbb{E}[\gamma]/e$ and (x^*, p^*) as stated in Proposition 3 is an optimal single-round mechanism. The single-round mechanism (x^*, p^*) induces the following dynamic mechanism for the original, multi-round problem: screen the agent by the parameter τ and charge an entry-fee $Tp^*(\tau)$ in Round 0, and then allocate each item with probability $x^*(\tau)$ in the subsequent rounds. Equivalently, we can screen the agent by the parameter τ in Round 0, and then allocate each item with probability $x^*(\tau)$ and charge $p^*(\tau)$ in the subsequent rounds, which fits the general formulation given in Section 2. We now conclude by arguing that this multi-round mechanism is incentive compatible in the original, unrelaxed environment in which the γ components of the shocks are private. This is because the mechanism does not ask the agent to report the values γ_t and the agent can only influence the mechanism by misreporting his parameter τ, which is never optimal because the mechanism is incentive compatible with respect to parameter τ.

D.2 Proof of Proposition 6

We first reduce determining the optimal regret $\text{Regret}^{\mathcal{S}}(T)$ for direct static mechanisms to a single-round problem as follows. This reduction works for any arbitrary distribution G supported on $\mathbb{R}_+$. Note that

$$\begin{aligned}
\text{Regret}^{\mathcal{S}}(T) &= \inf_{S \in \mathcal{S}^{\times 1}} \sup_{\tau \in [0,1]} \text{Regret}(S^{\times T}, F(\cdot;\tau), T) \\
&= \inf_{S \in \mathcal{S}^{\times 1}} \sup_{\tau \in [0,1]} \left\{T\mathbb{E}_{\theta \sim F(\cdot;\tau)}[\theta] - \text{PrincipalUtility}(S^{\times T}, \sigma^{\text{TR}}, F(\cdot;\tau), T)\right\} \\
&= T \cdot \inf_{S \in \mathcal{S}^{\times 1}} \sup_{\tau \in [0,1]} \left\{\mathbb{E}_{\theta \sim F(\cdot;\tau)}[\theta] - \text{PrincipalUtility}(S, \sigma^{\text{TR}}, F(\cdot;\tau), 1)\right\} ,
\end{aligned}$$

where the second step follows because the known-distribution benchmark is $\text{OPT}(F, T) = T\mathbb{E}_{\theta \sim F}[\theta]$ for any distribution F and σ^{TR} denotes the recommended truthful reporting strategy for the agent, and the last step follows from Lemma 2 (in Appendix A.3.2). To see $\text{OPT}(F, T) = T\mathbb{E}_{\theta \sim F}[\theta]$ for any distribution F, note that we argued $\text{OPT}(F, T) \leq T\mathbb{E}_{\theta \sim F}[\theta]$ in Section 4.1, because the principal's revenue is at most the agent's surplus subject to the agent's participation. In fact, the principal can fully extract the agent's surplus with the knowledge of F via, say, what is commonly known as the bundling strategy. The principal can, for example, bundle all items and sell the bundle at the expected value (see, e.g., Bakos and Brynjolfsson 1999).[1]

We then use the outcome distribution representation described in Section 3.3 for a mechanism $S \in \mathcal{S}^{\times 1}$ and equivalently write the last optimization problem as follows via the same reasoning in the proof of Lemma 1:

$$T \cdot \inf_{S \in \Delta(\Omega)^{\Theta}} \left\{ \sup_{\tau \in [0,1]} \mathbb{E}_{\theta \sim F(\cdot;\tau)}[\theta] - \mathbb{E}_{\theta \sim F(\cdot;\tau), \omega \sim S_\theta}[u(\theta, \omega)] \ \text{ s.t. } \ \text{(IC)}, \text{(IR)} \right\},$$

where the IC/IR constraints are as formulated in (3). Using the interim rules (x, p) to describe single-round direct mechanisms as in Section 4.1, it follows that

$$\text{Regret}^{\mathcal{S}}(T) = T \cdot \inf_{(x,p)} \left\{ \sup_{\tau \in [0,1]} \mathbb{E}_{\theta \sim F(\cdot;\tau)} \left[\theta - p(\theta) \right] \ \text{ s.t. } \ \text{(IC)}, \text{(IR)} \right\}.$$

Now, we let G be the exponential distribution with mean 1 and solve for the resulting single-round problem above with $T = 1$. Let $\exp(\tau)$ denote the exponential distribution with mean τ for $\tau \in [0, 1]$. In particular, we show $\text{Regret}^{\mathcal{S}}(1) = 1 - \frac{1}{e}$. To see this, note that since $\tau = 1$ is a feasible parameter, we have:

$$\text{Regret}^{\mathcal{S}}(1) \geq \inf_{(x,p)} \left\{ \mathbb{E}_{\theta \sim \exp(1)} \left[\theta - p(\theta) \right] \ \text{ s.t. } \ \text{(IC)}, \text{(IR)} \right\} = 1 - 1/e\,,$$

because, from Myerson (1981), an optimal Bayesian mechanism that maximizes the revenue when the agent's value distribution is exponential with mean 1 is given by $x(\theta) = \mathbf{1}\{\theta \geq 1\}$ and $p(\theta) = \mathbf{1}\{\theta \geq 1\}$, i.e., a posted pricing mechanism with the price of 1. Furthermore, since this mechanism is feasible for the single-round problem, we obtain that

$$\text{Regret}^{\mathcal{S}}(1) \leq \sup_{\tau \in [0,1]} \mathbb{E}_{\theta \sim \exp(\tau)} \left[\theta - \mathbf{1}\{\theta \geq 1\} \right] = \sup_{\tau \in [0,1]} \left\{ \tau - e^{-\frac{1}{\tau}} \right\} = 1 - 1/e\,,$$

where the last equality follows because $\tau - e^{-\frac{1}{\tau}}$ is increasing in $\tau \in [0, 1]$ and the supremum is achieved at $\tau = 1$. It follows that $\text{Regret}^{\mathcal{S}}(1) = 1 - \frac{1}{e}$. The same argument shows that the deterministic posted pricing mechanism with the price of 1 is an optimal single-round direct mechanism in the single-round problem.

Going back to the multi-round problem with T rounds, it follows that $\text{Regret}^{\mathcal{S}}(T) = (1 - 1/e)T$ and an optimal direct static mechanism that achieves this minimax regret is one that repeats the deterministic posted pricing mechanism with the price of 1.

[1] In the bundling mechanism, the principal lets the agent decide whether to continue or quit in Round 0 and then requires the agent to pay a one-time payment of $T\mathbb{E}_{\theta \sim F}[\theta]$ in Round 1 and allocates all items in Round 1 and future rounds. The agent would be indifferent between participating and not participating and when he decides to participate by continuing in Round 0, he would be bound by the bundling contract. The recommended strategy for the agent is to participate.

E Missing Proofs from Section 5.2

For any distribution F, recall that

$$\begin{aligned} \bar{u}(F) := \sup_{S\in\mathcal{S}^{\times 1}} & \int_\Theta \int_\Omega u(\theta,\omega)\mathrm{d}S_\theta(\omega)\mathrm{d}F(\theta) \\ \text{s.t.} & \int_\Theta \int_\Omega v(\theta,\omega)\mathrm{d}S_\theta(\omega)\mathrm{d}F(\theta) \geq 0\,. \end{aligned} \tag{E-3}$$

For $\mathbb{E}_{\theta\sim F}[\bar{u}(\delta_\theta)]$, we can equivalently write

$$\begin{aligned} \mathbb{E}_{\theta\sim F}[\bar{u}(\delta_\theta)] = \sup_{S\in\mathcal{S}^{\times 1}} & \int_\Theta \int_\Omega u(\theta,\omega)\mathrm{d}S_\theta(\omega)\mathrm{d}F(\theta) \\ \text{s.t.} & \int_\Omega v(\theta,\omega)\mathrm{d}S_\theta(\omega) \geq 0\,, \quad \forall\theta\in\Theta\,. \end{aligned} \tag{E-4}$$

We will use the following result in the proofs:

Proposition 3. *We have the following relations:*

1. *For any distribution $F\in\Delta(\Theta)$, $\mathrm{OPT}(F,T)\leq T\cdot\bar{u}(F)$.*
2. *For any $\theta\in\Theta$, $\mathrm{OPT}(\delta_\theta,T) = T\cdot\bar{u}(\delta_\theta)$.*
3. *For any distribution $F\in\Delta(\Theta)$, $\bar{u}(F)\geq \mathbb{E}_{\theta\sim F}[\bar{u}(\delta_\theta)]$.*

This is proved in Appendix E.3. The proofs of Propositions 7 and 8 from Section 5.2 are provided in Appendices E.1 and E.2, respectively.

E.1 Proof of Proposition 7

For any distribution $F\in\mathcal{F}$, we have

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

where the first step is by the first part of Proposition 3, the second by the linearity assumption on $\bar{u}(F)$, and the third by the second part of Proposition 3.

E.2 Proof of Proposition 8

In what follows, we show $\bar{u}(F)\leq\mathbb{E}_{\theta\sim F}[\bar{u}(\delta_\theta)]$ when the stated conditions hold. By Part 3 of Proposition 3, this suffices. Consequently, it would follow that Assumption 2 holds by Proposition 7.

Part 1): In words, the game is such that the payment is part of the outcome and enters linearly with coefficients of opposing signs into the utility functions of the principal and agent. Separating out the payment, the outcome space can be represented as $\Omega = \Omega^0\times\mathbb{R}$ where Ω^0 is the space of non-payment component of the outcomes and an outcome $\hat{\omega}$ is a pair $(\hat{\omega}^0, \hat{p})$ where $\hat{\omega}^0$ is the non-payment component and $\hat{p}$ is the payment. We use superscript 0 to denote the non-payment parts of the outcome and outcome space. Since the payment enters linearly into the utility functions of the principal and agent, we can represent $u(\theta, (\hat{\omega}^0, \hat{p})) = u^0(\theta, \hat{\omega}^0) + \alpha \cdot \hat{p}$ for some function $u^0 : \Theta \times \Omega^0 \to \mathbb{R}$ and scalar $\alpha \geq 0$ and, similarly, $v(\theta, (\hat{\omega}^0, \hat{p})) = v^0(\theta, \hat{\omega}^0) - \beta \cdot \hat{p}$ for some function $v^0 : \Theta \times \Omega^0 \to \mathbb{R}$ and scalar $\beta > 0$. Note we interpret a payment as a monetary transfer from the agent to the principal and this fixes the signs in front of α and β.

Fix an arbitrary distribution $F \in \mathcal{F}$. Let S be an arbitrary feasible solution for the optimization problem (E-3) defined for $\bar{u}(F)$. We define the payment offset $q_\theta = \frac{1}{\beta} \int_\Omega v(\theta, \omega) \mathrm{d}S_\theta(\omega)$ for all $\theta \in \Theta$. Now, consider a single-round direct mechanism S' where S'_θ is the outcome distribution S_θ modified with the fixed offset q_θ such that to realize an outcome $\hat{\omega} \sim S'_\theta$, we draw $(\hat{\omega}^0, \hat{p}) \sim S_\theta$ and set $\hat{\omega} = (\hat{\omega}^0, \hat{p} + q_\theta)$.

We show that S' is a feasible solution to the optimization problem (E-4) defining $\mathbb{E}_{\theta \sim F}[\bar{u}(\delta_\theta)]$ and that S' obtains the objective value in (E-4) that is at least that obtained by S in (E-3). As S was arbitrary, it would follow that $\bar{u}(F) \leq \mathbb{E}_{\theta \sim F}[\bar{u}(\delta_\theta)]$ and, as F was arbitrary, the proposition statement would follow.

For any $\theta \in \Theta$,

$$\begin{aligned}
\int_\Omega v(\theta, \omega) \mathrm{d}S'_\theta(\omega) &= \int_{\Omega^0 \times \mathbb{R}} v(\theta, (\omega^0, p + q_\theta)) \mathrm{d}S_\theta((\omega^0, p)) \\
&= \int_{\Omega^0 \times \mathbb{R}} \left(v^0(\theta, \omega^0) - \beta(p + q_\theta)\right) \mathrm{d}S_\theta((\omega^0, p)) \\
&= \int_{\Omega^0 \times \mathbb{R}} \left(v(\theta, (\omega^0, p)) - \beta q_\theta\right) \mathrm{d}S_\theta((\omega^0, p)) \\
&= \int_\Omega v(\theta, \omega) \mathrm{d}S_\theta(\omega) - \beta q_\theta \\
&= 0\,,
\end{aligned}$$

where the last step follows from how the payment offset is defined. Hence, S' is a feasible solution to (E-4).

Similarly, for any $\theta \in \Theta$,

$$\int_\Omega u(\theta, \omega) \mathrm{d}S'_\theta(\omega) = \int_{\Omega^0 \times \mathbb{R}} u(\theta, (\omega^0, p + q_\theta)) \mathrm{d}S_\theta((\omega^0, p)) = \int_\Omega u(\theta, \omega) \mathrm{d}S_\theta(\omega) + \alpha q_\theta\,.$$

Integrating the first and last expressions over Θ, we obtain

$$\begin{aligned}
\int_\Theta \int_\Omega u(\theta, \omega) \mathrm{d}S'_\theta(\omega) \mathrm{d}F(\theta) &= \int_\Theta \int_\Omega u(\theta, \omega) \mathrm{d}S_\theta(\omega) \mathrm{d}F(\theta) + \alpha \int_\Theta q_\theta \mathrm{d}F(\theta) \\
&= \int_\Theta \int_\Omega u(\theta, \omega) \mathrm{d}S_\theta(\omega) \mathrm{d}F(\theta) + \frac{\alpha}{\beta} \int_\Theta \int_\Omega v(\theta, \omega) \mathrm{d}S_\theta(\omega) \mathrm{d}F(\theta) \\
&\geq \int_\Theta \int_\Omega u(\theta, \omega) \mathrm{d}S_\theta(\omega) \mathrm{d}F(\theta)\,,
\end{aligned}$$

where the second-to-last step follows from the definition of the payment offset q_θ and the last step follows since S is a feasible solution to (E-3) and $\int_\Theta \int_\Omega v(\theta,\omega)\mathrm{d}S_\theta(\omega)\mathrm{d}F(\theta) \geq 0$. Therefore, S' obtains the objective value in (E-4) that is at least that obtained by S in (E-3). As S was arbitrary, $\bar{u}(F) \leq \mathbb{E}_{\theta\sim F}[\bar{u}(\delta_\theta)]$.

Part 2): By the assumption on the game, we have $v(\theta,\omega) \geq 0$ for all $\theta \in \Theta$ and $\omega \in \Omega$. Then, for any single-round direct mechanism S and shock $\theta \in \Theta$,

$$\int_\Omega v(\theta,\omega)\mathrm{d}S_\theta(\omega) \geq 0\,.$$

Clearly, for any $F \in \mathcal{F}$, any feasible solution to the optimization problem (E-3) is a feasible solution to the optimization problem (E-4) and obtains the same objective. The proposition follows.

E.3 Proof of Proposition 3

Part 1): Fix an arbitrary distribution $F \in \Delta(\Theta)$. Assume the principal commits to an incentive compatible dynamic mechanism A and the agent plays the recommended strategy σ. Let $\{\omega_t\}_{t=1}^T$ be the resulting random sequence of realized outcomes. For each $\theta \in \Theta$, we define measure $\mu_\theta(Q) = \frac{1}{T}\sum_{t=1}^T \Pr(\omega_t \in Q \mid \theta_t = \theta)$ for any $Q \subseteq \Omega$ and let S_θ be the corresponding distribution over Ω such that $\omega \sim S_\theta$ means an outcome ω is realized with probability $\mu_\theta(\omega)$. Consider a single-round direct mechanism $S = \{S_\theta\}_{\theta\in\Theta}$ that given a report θ returns an outcome $\omega \sim S_\theta$. We note that

$$\begin{aligned}
\mathrm{PrincipalUtility}(A,\sigma,F,T) &= \mathbb{E}\left[\sum_{t=1}^T u(\theta_t,\omega_t)\right] \\
&= \sum_{t=1}^T \mathbb{E}_{\theta_t}\left[\mathbb{E}[u(\theta_t,\omega_t)|\theta_t]\right] \\
&= \sum_{t=1}^T \mathbb{E}_{\theta\sim F}\left[\mathbb{E}[u(\theta_t,\omega_t)|\theta_t=\theta]\right] \\
&= T\cdot \mathbb{E}_{\theta\sim F}\left[\frac{1}{T}\sum_{t=1}^T \mathbb{E}_{\omega_t|\theta_t=\theta}[u(\theta,\omega_t)]\right] \\
&= T\cdot \mathbb{E}_{\theta\sim F}\left[\mathbb{E}_{\omega\sim S_\theta}[u(\theta,\omega)]\right]\,,
\end{aligned}$$

where the second equality follows from the linearity of expectations and the tower rule, the third from that the idiosyncratic shocks are drawn independently and identically, and the last from the construction of S. Hence, we have

$$\mathrm{PrincipalUtility}(A,\sigma,F,T) = T\cdot\int_\Theta\int_\Omega u(\theta,\omega)\mathrm{d}S_\theta(\omega)\mathrm{d}F(\theta)\,.$$

Similarly, we have

$$\mathrm{AgentUtility}(A,\sigma,F,T) = T\cdot\int_\Theta\int_\Omega v(\theta,\omega)\mathrm{d}S_\theta(\omega)\mathrm{d}F(\theta)\,.$$

Since the recommended strategy σ is a utility-maximizing strategy for the agent and the agent can achieve the aggregate utility of 0 by not participating, it must be that $\mathrm{AgentUtility}(A,\sigma,F,T) \geq 0$. It follows that S is a feasible solution for $\bar{u}(F)$ and achieves the objective of $\frac{1}{T}\cdot\text{PrincipalUtility}(A, \sigma, F, T)$. As the dynamic mechanism $A \in \mathcal{A}$ was arbitrary, the first part follows.

Part 2): Since we have the first part, it suffices to show $\text{OPT}(\delta_\theta, T) \geq \bar{u}(\delta_\theta)$ for any $\theta \in \Theta$. Fix an arbitrary $\theta \in \Theta$ and let the agent's distribution be δ_θ. Note $\bar{u}(\delta_\theta)$ is equivalently

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

where G is an outcome distribution over Ω. For an arbitrary $\epsilon > 0$, let G_ϵ be an outcome distribution that satisfies the IR constraint in the above optimization problem and

$$\int_\Omega u(\theta, \omega) \mathrm{d}G_\epsilon(\omega) \geq \bar{u}(\delta_\theta) - \epsilon \,.$$

Consider the corresponding dynamic mechanism A_ϵ that repeatedly determines an outcome according to G_ϵ in each round. For the recommended strategy, we let the agent participate when his distribution is δ_θ and give the better of the choices of participating or not when his distribution is something else. When the agent's distribution is δ_θ, since G_ϵ satisfies the IR constraint, participating is a utility-maximizing strategy and the agent accepts the outcomes being drawn independently and identically from G_ϵ. The agent's only other option is to not participate which leads to the aggregate utility of 0. When the agent's distribution is not δ_θ, the recommended strategy is such that the agent still follows the strategy. That is, A_ϵ is incentive compatible.

By construction, the aggregate utility of the principal under A_ϵ is at least $T \cdot \bar{u}(\delta_\theta) - \epsilon \cdot T$. As ϵ was arbitrary, this implies $\text{OPT}(\delta_\theta, T) \geq T \cdot \bar{u}(\delta_\theta)$. Combined with the first part, $\text{OPT}(\delta_\theta, T) = T \cdot \bar{u}(\delta_\theta)$.

Part 3): Note we always have $\bar{u}(F) \geq \mathbb{E}_{\theta \sim F}[\bar{u}(\delta_\theta)]$ for all $F \in \Delta(\Theta)$ unconditionally. This is because a feasible solution in the optimization problem in (E-4) is a feasible solution in the optimization problem (E-3) and obtains the same objective value.

F Additional Materials for Section 5.4

We prove Theorem 2 in the next subsection. We prove Propositions 9 and 10 in Appendices F.2–F.4.

F.1 Proof of Theorem 2

We first derive a lower bound on the multi-round minimax regret Regret(T) via the same reasoning used in the proof of Lemma 2. Note that

$$\begin{aligned}
\text{Regret}(T) &= \inf_{A\in\mathcal{A}} \sup_{F\in\mathcal{F}} \text{Regret}(A, F, T) \\
&\geq \inf_{A\in\mathcal{A}} \sup_{\theta\in\Theta} \text{Regret}(A, \delta_\theta, T) \\
&= \inf_{A\in\mathcal{A}} \sup_{\theta\in\Theta} \{\text{OPT}(\delta_\theta, T) - \text{PrincipalUtility}(A, \sigma, \delta_\theta, T)\} \\
&= T \cdot \inf_{A\in\mathcal{A}} \sup_{\theta\in\Theta} \{\text{OPT}(\delta_\theta, 1) - \text{PrincipalUtility}(S(A), \sigma^{\text{TR}}, \delta_\theta, 1)\} \\
&= T \cdot \inf_{A\in\mathcal{A}} \sup_{\theta\in\Theta} \text{Regret}(S(A), \delta_\theta, 1) \\
&\geq T \cdot \inf_{S'\in\mathcal{S}^{\times 1}} \sup_{\theta\in\Theta} \text{Regret}(S', \delta_\theta, 1)\,,
\end{aligned} \tag{F-5}$$

where the second step follows since point-mass distributions are a subset of $\mathcal{F}$ in the inner maximization expression by Assumption 1; the third step is by the regret definition where σ is the corresponding recommended strategy under A; the fourth step is by Proposition 1 and Lemma 1 where $S(A)$ is the single-round direct IC/IR mechanism derived from A as described in the proof of Lemma 1; the second-to-last step is by the regret definition; and the last step follows because the single-round direct IC/IR mechanisms in $\mathcal{S}^{\times 1}$ are a superset of those mechanisms $S(A)$ derived from incentive compatible dynamic mechanisms, i.e., $\{S(A) \mid A \in \mathcal{A}\}$.

Now, define $\Delta = \sup_{F\in\mathcal{F}} \{\text{OPT}(F, T) - \mathbb{E}_{\theta\sim F}[\text{OPT}(\delta_\theta, T)]\}$. Let $\epsilon \geq 0$ be arbitrary and consider a single-round direct IC/IR mechanism $S \in \mathcal{S}^{\times 1}$ satisfying

$$\sup_{\theta\in\Theta} \text{Regret}(S, \delta_\theta, 1) \leq \inf_{S'\in\mathcal{S}^{\times 1}} \sup_{\theta\in\Theta} \text{Regret}(S', \delta_\theta, 1) + \frac{\epsilon}{T}\,. \tag{F-6}$$

Then,

$$\begin{aligned}
&\sup_{F\in\mathcal{F}} \text{Regret}(S^{\times T}, F, T) \\
&\quad= \sup_{F\in\mathcal{F}} \left\{\text{OPT}(F, T) - \text{PrincipalUtility}(S^{\times T}, \sigma^{\text{TR}}, F, T)\right\} \\
&\quad= \sup_{F\in\mathcal{F}} \{\text{OPT}(F, T) - T \cdot \text{PrincipalUtility}(S, \sigma^{\text{TR}}, F, 1)\} \\
&\quad\leq \sup_{F\in\mathcal{F}} \{\mathbb{E}_{\theta\sim F}[\text{OPT}(\delta_\theta, T)] - T \cdot \text{PrincipalUtility}(S, \sigma^{\text{TR}}, F, 1)\} + \Delta \\
&\quad= T \cdot \sup_{F\in\mathcal{F}} \{\mathbb{E}_{\theta\sim F}[\text{OPT}(\delta_\theta, 1)] - \text{PrincipalUtility}(S, \sigma^{\text{TR}}, F, 1)\} + \Delta \\
&\quad\leq T \cdot \sup_{\theta\in\Theta} \text{Regret}(S, \delta_\theta, 1) + \Delta\,,
\end{aligned} \tag{F-7}$$

where the first step is by the definition of Regret notion and σ^{TR} is the truthful reporting strategy recommended for direct mechanisms; the second step is by Lemma 2; the third step is because the supremum operator is sublinear; the second-to-last step is by Proposition 1; and the last step follows from Lemma 3.

Algorithm 1. *Mechanism $A^*(F,T)$*

1. Round 0: The report space is $\mathcal{M} = \{\mathsf{CONTINUE}, \mathsf{QUIT}\}$. If the report is QUIT, the outcome is $\emptyset$ for all rounds. If the report is CONTINUE, we continue as follows.
2. Rounds 1–T: The report space is $\mathcal{M} = \{\theta^1, \theta^2\}$. If the report is θ^1, the outcome is ω^1 with probability q and $\emptyset$ with probability $1-q$. If the report is θ^2, the outcome is ω^2 with probability 1. Note $q = \begin{cases} 1 & \text{, if } f_1 < \frac{1}{2} \\ \frac{1-f_1}{f_1} & \text{, if } f_1 \geq \frac{1}{2} \end{cases}$.

By the above bounds, (F-5) and (F-7), and the property (F-6) of the single-round direct IC/IR mechanism S, it follows that

$$\begin{aligned} \sup_{F\in\mathcal{F}} \mathrm{Regret}(S^{\times T}, F, T) &\leq T \cdot \sup_{\theta\in\Theta} \mathrm{Regret}(S, \delta_\theta, 1) + \Delta \\ &\leq T \cdot \inf_{S'\in\mathcal{S}^{\times 1}} \sup_{\theta\in\Theta} \mathrm{Regret}(S', \delta_\theta, 1) + \Delta + \epsilon \\ &\leq \mathrm{Regret}(T) + \Delta + \epsilon\,. \end{aligned}$$

F.2 Proof of Proposition 9

We first show that the optimal performance achievable is $\mathrm{OPT}(F,T) = T \cdot \bar{u}(F)$ and then characterize the single-round full information benchmark $\bar{u}(F)$. Recall $F = (f_1, f_2)$ is the agent's private distribution over Θ where the shock is θ^i with probability f_i for $i = 1,2$ with $f_1 + f_2 = 1$.

Part 1 (Optimal performance). By Proposition 3, note $\mathrm{OPT}(F,T) \leq T \cdot \bar{u}(F)$ for any game. For the game in Table 1b, we show $\mathrm{OPT}(F,T) \geq T \cdot \bar{u}(F)$ and it would follow that $\mathrm{OPT}(F,T) = T \cdot \bar{u}(F)$.

Consider $A^*(F,T)$ in Algorithm 1. We show reporting CONTINUE in Round 0 and then reporting truthfully in Rounds 1–T is optimal (i.e., utility-maximizing) for the agent. Given the agent participates in a round (i.e., Rounds 1–T), truthful reporting is optimal on the per-round basis. If the shock is θ^1, reporting θ^1 yields $-q$ and reporting θ^2 yields $-\infty$. If the shock is θ^2, reporting θ^1 yields $-\infty$ (or 0 if $q = 0$) and reporting θ^2 yields 1. Hence, truthful reporting is optimal in each round. If the agent reports CONTINUE in Round 0 and participates in all the remaining rounds, the overall utility is $T \cdot (-q \cdot f_1 + f_2)$. The overall utility is $T \cdot (-f_1 + f_2)$ if $f_1 < \frac{1}{2}$ and 0 if $f_1 \geq \frac{1}{2}$, which is at least 0 for any distribution F. If the agent reports QUIT in Round 0 and does not participate in the remaining rounds, the utility is 0. Hence, reporting CONTINUE followed by truthful reporting is optimal over the entire horizon.

We use σ^{TR} to denote the agent's utility-maximizing strategy of reporting CONTINUE in Round 0 and then reporting truthfully in Rounds 1–T. For the mechanism $A^*(F,T)$, we let σ^{TR} be the recommended strategy for the agent. As we argued above, $A^*(F,T)$ with the recommended strategy σ^{TR} is incentive compatible. Given that the agent plays σ^{TR}, the principal's utility in each round (i.e., Rounds 1–T) is $q \cdot f_1 = \bar{u}(F)$ and

$$\mathrm{PrincipalUtility}(A^*(F,T), \sigma^{\mathrm{TR}}, F, T) = T \cdot \bar{u}(F)\,.$$

Then, we have

$$
\begin{aligned}
\mathrm{OPT}(F,T) &= \sup_{A\in\mathcal{A}} \mathrm{PrincipalUtility}(A,\sigma,F,T) \\
&\geq \mathrm{PrincipalUtility}(A^*(F,T),\sigma^{\mathrm{TR}},F,T) \\
&= T\cdot \bar{u}(F)\,,
\end{aligned}
$$

where σ is the corresponding recommended strategy for the incentive compatible mechanism A in the first step.

Part 2 (Single-round full information benchmark). Note for distribution F, the single-round full information benchmark is

$$
\begin{aligned}
\bar{u}(F) = \max_{x\in[0,1]} \quad & f_1\cdot x \\
\text{s.t.} \quad & -f_1\cdot x + f_2\cdot 1 \geq 0\,.
\end{aligned}
$$

To see this, let S be a single-round direct mechanism in the optimization problem defining $\bar{u}(F)$ (in Section 5.2) such that each S_{θ^i} is a distribution over Ω with probabilities given by $S_{\theta^i}(\emptyset)$, $S_{\theta^i}(\omega^1)$, and $S_{\theta^i}(\omega^2)$, and let $x = S_{\theta^1}(\omega^1)$ ($y = S_{\theta^2}(\omega^1)$) be the probability that outcome ω^1 is selected when the shock is θ^1 (θ^2). On the one hand, the mechanism maximizes $\mathbb{E}_{\theta\sim F,\omega\sim S_\theta}[u(\theta,\omega)] = f_1\cdot x + f_2\cdot y$ if S_{θ^1} and S_{θ^2} place as much probability mass as possible on outcome ω^1. We assume $y = 0$ without loss; if $f_2 > 0$, only $y = 0$ is feasible because the agent utility will be $-\infty$ and the ex-ante IR constraint will be violated if $y > 0$, and if $f_2 = 0$, y can be any value and we can set $y = 0$ without affecting the principal utility. On the other hand, setting x too large might violate the ex-ante IR constraint because outcome ω^1 gives the utility of -1 to the agent when the shock is θ^1. Setting $S_{\theta^2}(\omega^2) = 1$ allows for higher values of x at no cost to the principal.

Now, we determine $\bar{u}(F)$. Since $f_1 + f_2 = 1$, we see that the ex-ante IR constraint only binds when $f_1 \geq \frac{1}{2}$. Therefore, the optimal solution is $x = 1$ when $f_1 < \frac{1}{2}$ and $x = \frac{1-f_1}{f_1} \leq 1$ when $f_1 \geq \frac{1}{2}$, or, more succinctly, $x = \min\left\{1, \frac{1-f_1}{f_1}\right\}$. This implies

$$
\bar{u}(F) = \min\{f_1, 1-f_1\} = \begin{cases} f_1\,, & \text{if } f_1 < \frac{1}{2}\,, \\ 1-f_1\,, & \text{if } f_1 \geq \frac{1}{2}\,. \end{cases}
$$

F.3 First Part of Proposition 10

We prove the first part in this section and the second part in Appendix F.4.

By Lemma 2, for any arbitrary single-round direct IC/IR mechanism $S \in \mathcal{S}^{\times 1}$ and distribution $F \in \mathcal{F}$,

$$
\begin{aligned}
\mathrm{Regret}(S^{\times T},F,T) &= \mathrm{OPT}(F,T) - \mathrm{PrincipalUtility}(S^{\times T},F,T) \\
&= \mathrm{OPT}(F,T) - T\cdot \mathrm{PrincipalUtility}(S,F,1)\,.
\end{aligned}
$$

Since $\mathrm{OPT}(F,T) = T\cdot\bar{u}(F)$ by Proposition 9,

$$
\mathrm{Regret}(S^{\times T},F,T) = T\cdot\left(\bar{u}(F) - \mathrm{PrincipalUtility}(S,F,1)\right)\,.
$$

Algorithm 2. *Dynamic mechanism A parametrized in terms of T_1, T_2, δ, and q_0.*

1. Round 0: The report space is $\mathcal{M} = \{\mathsf{CONTINUE}, \mathsf{QUIT}\}$. If the report is QUIT, the outcome is $\emptyset$ for all rounds. If the report is $\mathsf{CONTINUE}$, we continue.
2. Phase 1 (T_1 rounds): The report space is $\mathcal{M} = \{\theta^1, \theta^2\}$. If the report is θ^1, the outcome is ω^1 with probability q_0 and $\emptyset$ with probability $1 - q_0$. If the report is θ^2, the outcome is ω^2 with probability 1.
3. Round $T_1 + 1$: The report space is $\mathcal{M} = \{\mathsf{CONTINUE}, \mathsf{QUIT}\}$. If the report is QUIT, the outcome is $\emptyset$ for the current and all remaining rounds. If the report is $\mathsf{CONTINUE}$, the outcome is $\emptyset$ for the current round and we continue.
4. Phase 2 (T_2 rounds):
   - From Phase 1, compute the fraction $\hat{f}_1$ of reports of θ^1. Let $\tilde{f}_1 = \hat{f}_1 + \delta$ and $\tilde{q} = \begin{cases} 1 & \text{, if } \tilde{f}_1 < \frac{1}{2} \\ \frac{1-\tilde{f}_1}{\tilde{f}_1} & \text{, if } \tilde{f}_1 \geq \frac{1}{2} \end{cases}$.
   - The report space is $\mathcal{M} = \{\theta^1, \theta^2\}$. If the report is θ^1, the outcome is ω^1 with probability $\tilde{q}$ and $\emptyset$ with probability $1 - \tilde{q}$. If the report is θ^2, the outcome is ω^2 with probability 1.

It suffices to show that

$$\inf_{S \in \mathcal{S}^{\times 1}} \sup_{F \in \mathcal{F}} \{\bar{u}(F) - \text{PrincipalUtility}(S, F, 1)\} = \frac{1}{2}\,. \tag{F-8}$$

Fix an arbitrary single-round direct IC/IR mechanism $S \in \mathcal{S}^{\times 1}$. Using the outcome distribution representation, let each S_{θ^i} be a distribution over Ω with probabilities $\alpha_0 = S_{\theta^1}(\emptyset)$, $\alpha_1 = S_{\theta^1}(\omega^1)$, and $\alpha_2 = S_{\theta^1}(\omega^2)$ and probabilities $\beta_0 = S_{\theta^2}(\emptyset)$, $\beta_1 = S_{\theta^2}(\omega^1)$, and $\beta_2 = S_{\theta^2}(\omega^2)$. Since S satisfies the (IR) constraint for shock θ^1, we have $0 \cdot \alpha_0 - 1 \cdot \alpha_1 - \infty \cdot \alpha_2 \geq 0$ and it follows that $\alpha_1 = \alpha_2 = 0$. Similarly, the (IR) constraint for shock θ^2 implies we have $0 \cdot \beta_0 - \infty \cdot \beta_1 + 1 \cdot \beta_2 \geq 0$ and it follows that $\beta_1 = 0$. Note $\text{PrincipalUtility}(S, F, 1) = \mathbb{E}_{\theta \sim F, \omega \sim S_\theta}[u(\theta, \omega)] = f_1 \alpha_1 + f_2 \beta_1$ for any $F = (f_1, f_2)$. From the above observations, $\text{PrincipalUtility}(S, F, 1) = 0$ for all F. Then, we have

$$\sup_{F \in \mathcal{F}} \{\bar{u}(F) - \text{PrincipalUtility}(S, F, 1)\} = \sup_{F \in \mathcal{F}} \bar{u}(F) = \sup_{F \in \mathcal{F}} \min\{f_1, 1 - f_1\} = \frac{1}{2}\,,$$

where the second step follows by Proposition 9.

As $S \in \mathcal{S}^{\times 1}$ was arbitrary and $\mathcal{S}^{\times 1}$ is not empty (for example, we can take $\alpha_0 = 1$ and $\beta_1 = 1$),

$$\inf_{S \in \mathcal{S}^{\times 1}} \sup_{F \in \mathcal{F}} \{\bar{u}(F) - \text{PrincipalUtility}(S, F, 1)\} = \frac{1}{2}\,,$$

which is (F-8).

F.4 Second Part of Proposition 10

Consider A in Algorithm 2 which is parametrized in terms of T_1, T_2, δ and q_0. We choose $T_1 = T^{2/3}$, $T_2 = T - T_1 - 1$, $\delta = \sqrt{\frac{\ln T_1}{4T_1}}$ and $q_0 = \frac{1}{\sqrt{T_1}}$. For ease of presentation, we mostly use T_1, T_2, δ and q_0 as parameters and use their values when necessary. We prove the result in three steps. First, we show that truthful reporting is a utility-maximizing strategy for the agent when his distribution

$F = (f_1, f_2)$ satisfies $f_1 \in [0, \frac{1}{1+q_0}]$. Second, we lower bound the principal's utility when the agent reports truthfully. Finally, we analyze the regret of the dynamic mechanism A.

Step 1. Assume the agent's distribution $F = (f_1, f_2)$ satisfies $f_1 \in [0, \frac{1}{1+q_0}]$. Let σ^{TR} denote the strategy that reports CONTINUE in Round 0, reports truthfully during Phase 1, reports CONTINUE if $-\tilde{q} \cdot f_1 + 1 \cdot f_2 \geq 0$ in Round $T_1 + 1$, and then reports truthfully during Phase 2 if the game continues to Phase 2. We note that the agent utility is at least 0 under σ^{TR}. In Phase 1, truthful reporting leads to the utility of $T_1 \cdot (-q_0 \cdot f_1 + 1 \cdot f_2) \geq 0$, since $f_1 \leq \frac{1}{1+q_0}$ implies $-q_0 \cdot f_1 + 1 \cdot (1 - f_1) = 1 - (1 + q_0) \cdot f_1 \geq 1 - (1 + q_0) \cdot \frac{1}{1+q_0} = 0$. If the game continues to Phase 2, it must be that $-\tilde{q} \cdot f_1 + 1 \cdot f_2 \geq 0$ and σ^{TR} leads to the utility of $T_2 \cdot (-\tilde{q} \cdot f_1 + 1 \cdot f_2) \geq 0$ in Phase 2. If the game does not continue, then it leads to the utility of 0 in Phase 2. In expectation, σ^{TR} leads to the utility of at least 0 in Phase 2. Hence, the overall utility is at least 0.

In fact, σ^{TR} is a utility-maximizing strategy for the agent. To see this, we first note truthful reporting is optimal on the per-round basis in each round in Phase 1 and in Phase 2 (for any value of $\tilde{q}$). That is, given the agent participates in a round, truthful reporting is a utility-maximizing strategy for the agent in that round. In each round in Phase 1, if the shock is θ^1, reporting θ^1 yields $-q_0$ and reporting θ^2 yields $-\infty$. If the shock is θ^2, reporting θ^2 yields 1 and reporting θ^1 yields $-\infty$. In each round in Phase 2, for any value of $\tilde{q}$, truthful reporting is optimal for the agent. If the shock is θ^1, reporting θ^1 yields $-\tilde{q}$ and reporting θ^2 yields $-\infty$. If the shock is θ^2, reporting θ^2 yields 1 and reporting θ^1 yields $-\infty$ (or 0 if $\tilde{q} = 0$).

Then, we note the only way for the agent to influence the principal's mechanism A is through reports in Phase 1 which determine the probability $\tilde{q}$ in Phase 2. Intuitively, the agent may consider some non-truthful reporting strategy in Phase 1 and continue to Phase 2 with $\tilde{q}$ determined favorably to benefit himself. Non-truthful reporting can only lead to lower per-round utilities during Phase 1 and each misreport of the shock costs $-\infty$. The agent needs to continue to Phase 2 in order to gain from such non-truthful reporting, but the cost overwhelms the potential gain from Phase 2. It follows that it is optimal for the agent to truthfully report during Phase 1 and given this observation, reporting CONTINUE in Round $T_1 + 1$ if $-\tilde{q} \cdot f_1 + 1 \cdot f_2 \geq 0$ can only benefit the agent because truthfully reporting is optimal on the per-round basis in Phase 2 and the utility from Phase 2 is $T_2 \cdot (-\tilde{q} \cdot f_1 + 1 \cdot f_2)$ given the game continues to Phase 2. Note the agent cannot influence the principal's mechanism during Phase 2. Hence, it is not possible to realize a greater utility overall than that achieved under σ^{TR}.

Step 2. Assume the agent's distribution $F = (f_1, f_2)$ satisfies $f_1 \in [0, \frac{1}{1+q_0}]$. Let U_1^{TR} and U_2^{TR} be the principal utility from Phases 1 and 2, respectively, when the principal's mechanism is A and the agent's utility-maximizing strategy is σ^{TR}, such that

$$\mathrm{PrincipalUtility}(A, \sigma^{\mathrm{TR}}, F, T) = U_1^{\mathrm{TR}} + U_2^{\mathrm{TR}} .$$

Note

$$U_1^{\mathrm{TR}} = T_1 \cdot q_0 \cdot f_1 \geq 0 . \tag{F-9}$$

We now consider U_2^{TR}. Let $\mathcal{E}$ be the event that the game continues to Phase 2 under σ^{TR}, or equivalently, $-\tilde{q} \cdot f_1 + 1 \cdot f_2 \geq 0$ and $\mathbf{1}_{\mathcal{E}}$ be the indicator that equals to 1 if the event occurs, and 0 otherwise; so, $\mathbf{1}_{\mathcal{E}} = \mathbf{1}\{-\tilde{q} \cdot f_1 + 1 \cdot f_2 \geq 0\}$. We have $U_2^{\mathrm{TR}} = T_2 \cdot \mathbb{E}[\tilde{q} \cdot f_1 \cdot \mathbf{1}_{\mathcal{E}}]$. Note that $\frac{1}{1+q_0} \geq \frac{1}{2}$ because $q_0 \leq 1$. We consider the following cases depending on whether f_1 is above or below $\frac{1}{2}$.

If $f_1 \le 1/2$, then the event $\mathcal{E}$ always occurs because $f_2 = 1 - f_1 \ge f_1$ and $\tilde{q} \in [0, 1]$. Therefore,

$$U_2^{\mathrm{TR}} = f_1 \cdot T_2 \cdot \mathbb{E}\left[\min\left\{1, \frac{1 - \tilde{f}_1}{\tilde{f}_1}\right\}\right], \tag{F-10}$$

where $\tilde{f}_1$ is always strictly positive because $\tilde{f}_1 = \hat{f}_1 + \delta \ge \delta > 0$.

If $f_1 > 1/2$, whenever the event $\mathcal{E}$ occurs we have $\tilde{f}_1 \ge 1/2$ and, consequently, $\tilde{q} = \frac{1-\tilde{f}_1}{\tilde{f}_1}$. To see this, note that if $\tilde{f}_1 < 1/2$, then $\tilde{q} = 1$, which implies that $-\tilde{q} \cdot f_1 + 1 \cdot f_2 = -f_1 + f_2 < 0$ and event $\mathcal{E}$ does not occur. Then, $-\tilde{q} \cdot f_1 + 1 \cdot f_2 \ge 0$ is equivalent to $\frac{1-f_1}{f_1} \ge \tilde{q} = \frac{1-\tilde{f}_1}{\tilde{f}_1}$. Since the transformation $x \mapsto \frac{1-x}{x}$ is decreasing, the event $\mathcal{E}$ can be equivalently written as $\mathcal{E} = \{\tilde{f}_1 \ge f_1\}$. This implies that

$$\begin{aligned} U_2^{\mathrm{TR}} &= f_1 \cdot T_2 \cdot \mathbb{E}\left[\mathbf{1}\{\tilde{f}_1 \ge f_1\} \cdot \frac{1 - \tilde{f}_1}{\tilde{f}_1}\right] \\ &= T_2 \cdot \left((1 - f_1) \cdot \Pr(\tilde{f}_1 \ge f_1) - \mathbb{E}\left[\max\left\{1 - \frac{f_1}{\tilde{f}_1}, 0\right\}\right]\right), \end{aligned} \tag{F-11}$$

where we used that $\frac{f_1(1-\tilde{f}_1)}{\tilde{f}_1} = (1 - f_1) - \left(1 - \frac{f_1}{\tilde{f}_1}\right)$ and $\mathbf{1}\{\tilde{f}_1 \ge f_1\} \cdot \left(1 - \frac{f_1}{\tilde{f}_1}\right) = \max\left\{1 - \frac{f_1}{\tilde{f}_1}, 0\right\}$.

The following result will be used in bounding the principal's utility. The proof is provided in Appendix F.5.

Lemma 5. *The following hold:*

1. *If $f_1 \le 1/2$, then* $\mathbb{E}\left[\min\left\{1, \frac{1-\tilde{f}_1}{\tilde{f}_1}\right\}\right] \ge 1 - 4\delta - \frac{2}{\sqrt{T_1}}$.
2. *If $f_1 > 1/2$, then* $\mathbb{E}\left[\max\left\{1 - \frac{f_1}{\tilde{f}_1}, 0\right\}\right] \le 2\delta + \frac{1}{\sqrt{T_1}}$.
3. $\Pr(\tilde{f}_1 \ge f_1) \ge 1 - e^{-2\delta^2 T_1}$.

Step 3. We next bound the regret of the dynamic mechanism A. For the recommended strategy σ under A, we let σ be the truthful reporting strategy σ^{TR} defined above for $F = (f_1, f_2)$ with $f_1 \in [0, \frac{1}{1+q_0}]$ and σ be any arbitrary utility-maximizing strategy for the agent corresponding to F for F with $f_1 \in (\frac{1}{1+q_0}, 1]$. Note A with the recommended strategy σ is incentive compatible by the argument in Step 1 and the construction of choosing utility-maximizing strategies. Note

$$\begin{aligned} \mathrm{Regret}(A, F, T) &= \mathrm{OPT}(F, T) - \mathrm{PrincipalUtility}(A, \sigma, F, T) \\ &= T \cdot \bar{u}(F) - \mathrm{PrincipalUtility}(A, \sigma, F, T), \end{aligned}$$

by Proposition 9. We upper bound the regret in three separate cases depending on the agent's distribution $F = (f_1, f_2)$. Note $\frac{1}{1+q_0} \ge \frac{1}{2}$ for $T_1 \ge 1$.

If $f_1 \in [0, \frac{1}{2}]$, then $\bar{u}(F) = f_1$. From (F-9) and (F-10), we have

$$\begin{aligned}\text{PrincipalUtility}(A, \sigma, F, T) &= \text{PrincipalUtility}(A, \sigma^{\text{TR}}, F, T) \\ &\geq f_1 \cdot T_2 \cdot \mathbb{E}\left[\min\left\{1, \frac{1-\tilde{f}_1}{\tilde{f}_1}\right\}\right].\end{aligned}$$

Then,

$$\begin{aligned}\frac{1}{T}\text{Regret}(A, F, T) &\leq f_1\left(1 - \frac{T_2}{T} \cdot \mathbb{E}\left[\min\left(1, \frac{1-\tilde{f}_1}{\tilde{f}_1}\right)\right]\right) \\ &\leq f_1\left(1 - \left(1 - \frac{T_1+1}{T}\right) \cdot \left(1 - 4\delta - \frac{2}{\sqrt{T_1}}\right)\right) \\ &= f_1\left(1 - \left(1 - \frac{T_1+1}{T} - 4\delta - \frac{2}{\sqrt{T_1}} + \frac{T_1+1}{T} \cdot \left(4\delta + \frac{2}{\sqrt{T_1}}\right)\right)\right) \\ &\leq f_1\left(\frac{T_1+1}{T} + 4\delta + \frac{2}{\sqrt{T_1}}\right) \\ &\leq 2\delta + \frac{T_1}{T} + \frac{1}{\sqrt{T_1}}, \qquad \text{(F-12)}\end{aligned}$$

where the second inequality follows from $T_2 = T - T_1 - 1$ and Part 1 of Lemma 5; the second-to-last inequality follows from dropping the negative term in the resulting expression in the parentheses; and the last inequality follows from $f_1 \leq \frac{1}{2}$ and $T_1 \geq 1$ which implies $T_1 + 1 \leq 2T_1$.

If $f_1 \in (\frac{1}{2}, \frac{1}{1+q_0}]$, then $\bar{u}(F) = 1 - f_1$. From (F-9) and (F-11), we have

$$\begin{aligned}\text{PrincipalUtility}(A, \sigma, F, T) &= \text{PrincipalUtility}(A, \sigma^{\text{TR}}, F, T) \\ &\geq T_2 \cdot \left((1-f_1) \cdot \Pr(\tilde{f}_1 \geq f_1) - \mathbb{E}\left[\max\left\{1 - \frac{f_1}{\tilde{f}_1}, 0\right\}\right]\right).\end{aligned}$$

Consequently, we obtain

$$\begin{aligned}\frac{1}{T}\text{Regret}(A, F, T) &\leq (1-f_1)\left(1 - \frac{T_2}{T} \cdot \Pr(\tilde{f}_1 \geq f_1)\right) + \frac{T_2}{T}\mathbb{E}\left[\max\left(1 - \frac{f_1}{\tilde{f}_1}, 0\right)\right] \\ &\leq (1-f_1)\left(1 - \left(1 - \frac{T_1+1}{T}\right) \cdot \left(1 - e^{-2\delta^2 T_1}\right)\right) + \frac{T_2}{T} \cdot \left(2\delta + \frac{1}{\sqrt{T_1}}\right) \\ &= (1-f_1)\left(1 - \left(1 - \frac{T_1+1}{T} - e^{-2\delta^2 T_1} + \frac{T_1+1}{T} \cdot e^{-2\delta^2 T_1}\right)\right) + \frac{T_2}{T} \cdot \left(2\delta + \frac{1}{\sqrt{T_1}}\right) \\ &\leq (1-f_1)\left(\frac{T_1+1}{T} + e^{-2\delta^2 T_1}\right) + \frac{T_2}{T} \cdot \left(2\delta + \frac{1}{\sqrt{T_1}}\right) \\ &\leq \frac{1}{2}e^{-2\delta^2 T_1} + 2\delta + \frac{T_1}{T} + \frac{1}{\sqrt{T_1}}, \qquad \text{(F-13)}\end{aligned}$$

where the second inequality follows from $T_2 = T - T_1 - 1$ and Parts 2 and 3 of Lemma 5; the second-to-last inequality follows from dropping the product term $\frac{T_1+1}{T} \cdot e^{-2\delta^2 T_1}$; and the last inequality follows because $1 - f_1 \leq \frac{1}{2}$, $T_2 \leq T$, and $T_1 \geq 1$ which implies $T_1 + 1 \leq 2T_1$.

If $f_1 \in (\frac{1}{1+q_0}, 1]$, then $\bar{u}(F) = 1 - f_1$. Note that the principal's utility is always at least 0 regardless of the agent's utility-maximizing strategy, i.e., $\mathrm{PrincipalUtility}(A,\sigma,F,T) \geq 0$. Using the last observation, we obtain

$$\frac{1}{T}\mathrm{Regret}(A,F,T) \leq 1 - f_1 \leq 1 - \frac{1}{1+q_0} = \frac{q_0}{1+q_0} \leq q_0\,, \tag{F-14}$$

where the last inequality follows because $q_0 \geq 0$.

Combining the upper bounds on the regret in above three cases, (F-12)–(F-14), and using that $q_0 = \frac{1}{\sqrt{T_1}}$, we obtain

$$\begin{aligned}
\frac{1}{T}\sup_{F\in\mathcal{F}} \mathrm{Regret}(A,F,T) &\leq 2\delta + \frac{1}{2}e^{-2\delta^2 T_1} + \frac{T_1}{T} + \frac{1}{\sqrt{T_1}} \\
&\leq \frac{(\ln T)^{1/2}}{T^{1/3}} + \frac{1}{2T^{1/3}} + \frac{1}{T^{1/3}} + \frac{1}{T^{1/3}} \\
&= \frac{(\ln T)^{1/2}}{T^{1/3}} + \frac{5}{2T^{1/3}}\,.
\end{aligned}$$

where the second inequality follows from our choices for δ and T_1.

F.5 Missing Proofs from Appendix F.4

Proof of Lemma 5. We prove each part at a time. For Part 1, note that the function $x \mapsto \frac{1}{x}$ is convex and a first-order expansion around $\frac{1}{2}$ yields the lower bound $\frac{1}{\tilde{f}_1} \geq 2 - 4\left(\tilde{f}_1 - \frac{1}{2}\right)$. Therefore,

$$\min\left\{1, \frac{1-\tilde{f}_1}{\tilde{f}_1}\right\} \geq \min\left\{1, 1-4\left(\tilde{f}_1 - \frac{1}{2}\right)\right\} = 1 - 4\max\left\{\tilde{f}_1 - \frac{1}{2}, 0\right\}$$

Because $\tilde{f}_1 = \hat{f}_1 + \delta$, $\delta \geq 0$, and $f_1 \leq \frac{1}{2}$, we have that

$$\max\left\{\tilde{f}_1 - \frac{1}{2}, 0\right\} \leq \delta + \max\left\{\hat{f}_1 - f_1, 0\right\} \leq \delta + |\hat{f}_1 - f_1|\,,$$

where the last inequality follows because $\max\{x,0\} \leq |x|$ for all $x \in \mathbb{R}$. Note that $T_1 \cdot \hat{f}_1$ is binomially distributed with T_1 trials and the success probability of f_1. Jensen's inequality and that $\mathbb{E}[\hat{f}_1] = f_1$ imply that

$$\mathbb{E}\left[|\hat{f}_1 - f_1|\right] \leq \sqrt{\mathrm{Var}(\hat{f}_1)} = \sqrt{\frac{f_1(1-f_1)}{T_1}} \leq \frac{1}{2\sqrt{T_1}}\,, \tag{F-15}$$

where the equality follows from the variance formula for a binomially distributed random variable and the last inequality follows because $f_1(1-f_1) \leq \frac{1}{4}$ for $f_1 \in [0,1]$. Putting everything together,

$$\begin{aligned}
\mathbb{E}\left[\min\left\{1, \frac{1-\tilde{f}_1}{\tilde{f}_1}\right\}\right] &\geq 1 - 4\mathbb{E}\left[\max\left\{\tilde{f}_1 - \frac{1}{2}, 0\right\}\right] \\
&\geq 1 - 4\delta - 4\mathbb{E}\left[|\hat{f}_1 - f_1|\right] \\
&\geq 1 - 4\delta - \frac{2}{\sqrt{T_1}}\,.
\end{aligned}$$

For Part 2, we use, again, that the function $x \mapsto \frac{1}{x}$ is convex to obtain that a first-order expansion around f_1 yields the lower bound $\frac{1}{\tilde{f}_1} \geq \frac{1}{f_1} - \frac{1}{f_1^2}(\tilde{f}_1 - f_1)$. Therefore,

$$\begin{aligned}
\max\left\{1 - \frac{f_1}{\tilde{f}_1}, 0\right\} &\leq \frac{1}{f_1}\max\left\{\tilde{f}_1 - f_1, 0\right\} \\
&\leq 2\delta + 2\max\left\{\hat{f}_1 - f_1, 0\right\} \\
&\leq 2\delta + 2|\hat{f}_1 - f_1|\,,
\end{aligned}$$

where the first inequality follows from the above lower bound; the second inequality follows from $f_1 > \frac{1}{2}$, $\tilde{f}_1 = \hat{f}_1 + \delta$, and $\delta \geq 0$; and the last is because $\max\{x, 0\} \leq |x|$ for all $x \in \mathbb{R}$. Taking expectations and using (F-15), we obtain

$$\mathbb{E}\left[\max\left\{1 - \frac{f_1}{\tilde{f}_1}, 0\right\}\right] \leq 2\delta + 2\mathbb{E}\left[|\hat{f}_1 - f_1|\right] \leq 2\delta + \frac{1}{\sqrt{T_1}}\,.$$

For Part 3, we use that $\tilde{f}_1 = \hat{f}_1 + \delta$ to obtain

$$\Pr(\tilde{f}_1 \geq f_1) = \Pr(\hat{f}_1 \geq f_1 - \delta) = 1 - \Pr(\hat{f}_1 < f_1 - \delta) \geq 1 - e^{-2\delta^2 T_1}\,,$$

where the last inequality follows from Hoeffding's inequality because $T_1 \cdot \hat{f}_1$ is binomially distributed with T_1 trials and success probability f_1. □

References

Yannis Bakos and Erik Brynjolfsson. Bundling information goods: Pricing, profits, and efficiency. *Management Science*, 45(12):1613–1630, 1999.

Dirk Bergemann and Karl H. Schlag. Pricing without priors. *Journal of the European Economic Association*, 6(2-3):560–569, 2008.

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

J.K. Hale. *Ordinary Differential Equations*. Pure and applied mathematics : a series of texts and monographs. Wiley-Interscience, 1969.

S. Kakade, I. Lobel, and H. Nazerzadeh. Optimal dynamic mechanism design and the virtual pivot mechanism. *Operations Research*, 61(3):837–854, 2013.

Vijay Krishna. *Auction Theory*. Elsevier, 2 edition, 2009.

Roger B. Myerson. Optimal auction design. *Math. Oper. Res.*, 6(1):58–73, February 1981. ISSN 0364-765X.