European Economic Review 167 (2024) 104777

Contents lists available at ScienceDirect

European Economic Review

journal homepage: www.elsevier.com/locate/eer

Mechanism design with limited commitment: Markov environments[☆]

Laura Doval[a], Vasiliki Skreta[b,*]

[a] Columbia Business School and CEPR, United States of America
[b] University of Texas at Austin, University College London, and CEPR, United States of America

ARTICLE INFO

JEL classification:
D84
D86

Keywords:
Mechanism design
Limited commitment
Revelation principle
Direct blackwell mechanisms
Information design
Short-term mechanisms
Markov environments
Dynamic mechanism design

ABSTRACT

We provide a revelation principle for a class of single-agent dynamic mechanism design settings in which the agent's private information evolves stochastically over time and the designer can only commit to short-term mechanisms. We restrict attention to *Markov environments*, in which (i) the agent's type in period $t+1$ depends only on her period-t type and the period-t allocation, (ii) the designer's and the agent's payoffs are time-separable, and (iii) their period-t payoffs depend only on period-t type and the period-t allocation. We show all equilibrium payoffs can be attained with the designer using *flow direct Blackwell* mechanisms, which consist of a mapping from the agent's current type report to posterior beliefs about the current type, and a mapping from these beliefs to allocations. Furthermore, all equilibrium payoffs can be attained with strategies in which the agent participates and truthfully reports her type, and the beliefs that result from the mechanism correspond to the designer's equilibrium beliefs. This result greatly simplifies the search of optimal dynamic and sequentially rational mechanisms in dynamic mechanism design problems, which include dynamic Mirrlees and social insurance models.

1. Introduction

We provide a revelation principle for a class of single-agent dynamic mechanism design settings in which the agent's private information evolves stochastically over time and the designer can only commit to short-term mechanisms. Asymmetric information and misaligned incentives are pervasive in a wide range of repeated interactions in industrial organization, managerial economics, political economy, and public finance. The recent literatures on dynamic mechanism design and dynamic public finance highlight the importance for applications of allowing for privately informed agents who learn their private information over time:[1] buyers learn about their valuation of a good over time, sometimes as the result of consumption; managers' productivity changes as they learn on the job (Garrett and Pavan, 2012); individuals' productivity evolves along their life cycle, which drives income uncertainty (Farhi and Werning, 2013). Recent contributions to these literatures study the implications of evolving — rather than persistent — private information for the design of optimal policies.

☆ We thank the Editor, David Levine, and two anonymous referees for feedback that has greatly improved the paper. Vasiliki Skreta is grateful for generous financial support through the ERC consolidator grant 682417 "Frontiers in design." This research is supported by grants from the National Science Foundation (Doval: SES-2131706; Skreta: SES-1851729)

* Corresponding author.
E-mail addresses: laura.doval@columbia.edu (L. Doval), vskreta@gmail.com (V. Skreta).

[1] See, e.g., Pavan et al. (2014) and the references therein for dynamic mechanism design, and Golosov et al. (2006) and Stantcheva (2020) for a review of the new dynamic public finance literature.

https://doi.org/10.1016/j.euroecorev.2024.104777
Received 20 March 2023; Received in revised form 24 May 2024; Accepted 31 May 2024
Available online 1 June 2024
0014-2921/

While considering a more flexible information environment, the literatures on dynamic mechanism design and dynamic public finance retain the standard assumption that the designer fully commits to the sequence of mechanisms the agent is faced with. Instead, assuming the designer can only commit to today's mechanism, but not the sequence of mechanisms the agent faces later on, is natural. As firms learn about consumers' willingness to pay, they may renege on previous price commitments to engage in price discrimination; governments may wish to revise their tax policy as they learn the income distribution from current payments; as a firm learns its manager productivity, they may wish to revise their assigned tasks. It is well-known that deviations from the commitment assumption introduce difficulties for mechanism design, as the revelation principle fails to hold. This paper fills this gap: By identifying a suitable extension of the framework in Doval and Skreta (2022), we provide a revelation principle for dynamic mechanism design with evolving types and commitment to short-term mechanisms.

Formally, we study the following family of mechanism-selection games. A designer interacts with a privately informed agent over a possibly infinite time horizon. In every period $t \geq 1$, the agent privately learns her current period type $\theta_t \in \Theta$, after which the designer offers the agent a mechanism, which determines the allocation a_t for that period. As in Doval and Skreta (2022), a mechanism is defined as a tuple (M, φ, S), where M is a set of *input* messages, S is a set of *output* messages, and φ is a mapping associating to each input message $m \in M$, a distribution $\varphi(\cdot|m)$ over output messages, S, and allocations, A. Given her private information (her sequence of types through period t, $\theta^t = (\theta_1, \ldots, \theta_t) \in \Theta^t$) and faced with the mechanism (M, φ, S), the agent privately reports an input message, $m \in M$, to the mechanism, which then determines the distribution, $\varphi(\cdot|m)$, from which an output message, $s \in S$, and an allocation, $a \in A$, are drawn. The output message and the allocation are publicly observable. Because the designer can only commit to short-term mechanisms, we study the payoffs the designer and the agent can achieve under Perfect Bayesian equilibrium.

Our previous work, Doval and Skreta (2022), provides a revelation principle for mechanism-selection games in which the designer faces an agent with fully-persistent private information. In particular, we show that it is without loss of generality to restrict the designer to choosing mechanisms in which the designer asks the agent to report her private information and together with the allocation, the mechanism outputs a belief over the agent's type. Furthermore, it is without loss to restrict attention to equilibrium strategies in which the messages can be taken literally: the agent *truthfully* reports her private information and the output belief corresponds to the principal's equilibrium belief about the agent's type.

Extending the logic of Doval and Skreta (2022) to the case in which the agent's type evolves over time brings forth two challenges. The first relates to the set of inputs into the mechanism. When the agent's information evolves over time, the agent's period-t private information is given by her profile of types through period t, $\theta^t = (\theta_1, \ldots, \theta_t)$. The result in Doval and Skreta (2022) would then imply the designer needs to elicit the agent's multidimensional type θ^t, with the message space growing over time.[2] To deal with this challenge, we restrict attention to *Markov* environments, as in Pavan et al. (2014). These are environments in which (i) the agent's type in period $t+1$ depends on her period-t type θ_t and the period-t allocation a_t, and (ii) the principal and the agent's payoffs are time-separable, and their period-t payoffs depend only on (a_t, θ_t). The *Markov* assumption implies the agent's incentives to report only depend on her period-t type. As we discuss below, this property allows for a simpler version of the revelation principle, much more suitable for applications. The second challenge relates to the set of outputs of the mechanism. After all, when the agent's type evolves over time two candidates for the output messages exist: the principal's belief at the end of period t about the agent's type θ^t and the principal's belief at the beginning of period $t+1$ about the agent's type, $\theta^{t+1} = (\theta^t, \theta_{t+1})$.

Theorem 1 proves an analogue of the revelation principle in Doval and Skreta (2022) for Markov environments. Indeed, Theorem 1 identifies a set of mechanisms, and hence a mechanism-selection game, that is enough to replicate any equilibrium payoff of any mechanism-selection game in our family. In this game, which we denote the canonical game, the designer can only offer mechanisms in which input messages are *current* type reports and output messages are beliefs. Moreover, Theorem 1 shows that any equilibrium payoff of the canonical game can be replicated by a *canonical equilibrium* in which the agent always participates in the mechanisms offered in equilibrium by the designer, and input and output messages have a *literal* meaning: the agent truthfully reports her *current* type, and if the mechanism outputs a given posterior, this posterior coincides with the designer's equilibrium beliefs about the agent's *current* type. Furthermore, in a canonical equilibrium, the designer only offers the agent *flow direct Blackwell mechanisms*, in which conditional on the output message, the allocation is drawn independently of the agent's type report. Thus, Theorem 1 implies that to characterize the equilibrium payoffs that can be achieved in some equilibrium in some mechanism-selection game, it is without loss of generality to restrict attention to the analysis of the canonical equilibria of the canonical game. As our companion paper (Doval and Skreta, 2023) illustrates, Theorem 1 reduces the search of the designer optimal mechanism to the solution to a constrained optimization program.

Theorem 1 highlights two simplifications brought forth by the restriction to Markov environments. First, the designer only needs to elicit the agent's current type, θ_t, as opposed to the agent's type profile, θ^t. Second, note Theorem 1 identifies the principal's beliefs at the end of period t as the mechanism's canonical output messages. Moreover, the Markov assumption on the environment implies the mechanism only needs to keep track of the principal's belief about θ_t–as opposed to the belief over the type profile θ^t. These two simplifications are key for applications: Our companion paper (Doval and Skreta, 2023) showcases them in the context of an industrial organization application.

[2] Because the revelation principle in Doval and Skreta (2022) has the agent submit a type report in each period, the result resembles more that of Townsend (1988) than that of Myerson (1982, 1986). Indeed, the logic of Myerson (1982, 1986) implies that with fully persistent types, the agent only submits a type report once at the beginning. For that reason, Myerson's revelation principle only ensures truthtelling and obedience along truthful histories. Instead, re-reporting in Townsend (1988) restores truthtelling on and off the path of play.

Related literature. The paper contributes to the literatures in dynamic mechanism design, dynamic public finance, and mechanism design with limited commitment. After the pioneering work of Baron and Besanko (1984), Courty and Li (2000), and Battaglini (2005) further the dynamic mechanism design literature. Battaglini (2005) highlights that without fully persistent types optimal mechanisms are "less" time-inconsistent and extends the no distortion at the top and distortion at the bottom properties of static mechanism design with single-crossing preferences. Armed with Theorem 1, our companion work (Doval and Skreta, 2023) echoes the observation that optimal mechanisms under commitment can be implemented with limited commitment when types are not fully persistent. However, in contrast to Battaglini (2005) and Doval and Skreta (2023) also show the designer may find it optimal to distort the mechanism both at the top and at the bottom to soften ratcheting forces. Finally, Doval and Skreta (2023) shows how Theorem 1 can be used to obtain the analogue of the dynamic envelope condition in Pavan et al. (2014) for dynamic mechanism design with limited commitment.

Golosov et al. (2006) and Stantcheva (2020) review the dynamic public finance literature and the dynamic Mirrlees approach to taxation. Most papers in this literature assume time-separable payoffs and that the evolution of private information follows a first-order Markov process. Kapička (2013) uses the first-order approach to characterize efficient allocations in a dynamic economy where agents' types evolve over time. Assuming individual productivities evolve over time as is the case in the data, Farhi and Werning (2013) study optimal taxation over the life cycle, showing that optimal taxes are age-dependent. Stantcheva (2015) expands on the dynamic Mirrlees model by allowing for human capital accumulation, i.e., last period choices affect this period's productivity draw. Some contributions to this literature assume away the designer's ability to commit. In a two-period model with persistent types, Bisin and Rampini (2006) study how the market may discipline the fiscal authority when it cannot commit to long-term mechanisms. Assuming fully non-persistent types, Sleet and Yeltekin (2008), Farhi et al. (2012), and Golosov and Iovino (2021) study the design of optimal social insurance.[3] Theorem 1 contributes to this literature by providing a tool that can be used to characterize optimal policies without assuming the government can commit across periods.

Organization. The rest of the paper is organized as follows. Section 2 describes the model and notation. Section 3 introduces Theorem 1. Omitted statements and all proofs are in the appendix. Appendices A and C provide necessary definitions. Assuming the set of types is finite and mechanisms induce finite-support lotteries, Appendix B provides an easy-to-digest proof of Theorem 1. Appendix D provides the necessary formalisms to adapt the proof in Appendix B to the case of a continuum type space building on the results in Doval and Skreta (2022).

2. Model

To facilitate the comparison with Doval and Skreta (2022), we follow the model and the notation therein as much as possible in what follows:

Primitives. Two players, a principal (he) and an agent (she), interact over $T \leq \infty$ periods.[4] Each period, as a result of the interaction between the principal and the agent, an allocation $a \in A$ is determined. Let A^T denote the set $\times_{t=1}^T A$. We allow for the possibility that past allocations influence what the principal can offer the agent in the future. Thus, for each $t \geq 1$, a correspondence $\mathcal{A}_t : A^{t-1} \rightrightarrows A$ exists such that for every sequence of allocations up to period t, $a^{t-1} = (a_1, \ldots, a_{t-1})$, $\mathcal{A}_t(a^{t-1})$ describes the set of allocations the principal can offer in period t (with the convention that when $t = 1$, $a^0 = \{\emptyset\}$). Furthermore, we assume an allocation a^* exists such that a^* is always available. Below, allocation a^* plays the role of the agent's outside option.

In contrast to Doval and Skreta (2022), we consider a *Markov* environment, defined by two properties (c.f., Pavan et al., 2014). First, the agent's private information is described by a non-homogeneous Markov process: In each period $t \geq 1$, the agent's type θ_t is drawn from a set of types Θ according to a distribution $F_t(\cdot|\theta_{t-1}, a_{t-1})$, where (θ_{t-1}, a_{t-1}) denotes the agent's type and allocation in period $t-1$ (with the convention that when $t = 1$, $F_t(\cdot|\theta_{t-1}, a_{t-1}) \equiv F_1$). Second, the principal's and the agent's payoffs are time separable and their period-t flow payoffs only depends on the current allocation and the agent's period-t type. Formally, letting $(a^T, \theta^T) \in (A \times \Theta)^T$ denote the allocations and the agent's private information through period T, the principal and the agent's payoffs are given by

$$W(a^T, \theta^T) = \sum_{t=1}^{T} \delta^t w_t(a_t, \theta_t), \quad U(a^T, \theta^T) = \sum_{t=1}^{T} \delta^t u_t(a_t, \theta_t).$$

We impose some technical restrictions on our model.[5] The sets Θ and A are Polish, that is, completely metrizable, separable, topological spaces. They are endowed with their Borel σ-algebra. Throughout we assume that Θ is at most countable and discuss the case in which Θ is a continuum in Appendix D. We also assume Θ is compact. Endowing product sets with their product σ-algebra, we assume the principal and the agent's utility functions, W and U, are bounded measurable functions. Similarly, for each $t \geq 1$ and for each $a^{t-1} \in A^{t-1}$, the set $\mathcal{A}_t(a^{t-1})$ is a measurable set.

[3] Persistent private information is not the only source of time-inconsistency of optimal mechanisms. For instance, once individuals make capital investments, the government may prefer to tax capital because it does not distort contemporaneous incentives. Anticipating this, individuals' incentives to invest in capital in previous periods will be dampened.

[4] To simultaneously analyze the cases of finite and infinite horizon, we abuse notation as follows. When $T = \infty$, and notation of the form $t = 1, \ldots, T$, $\sum_{t=1}^T$, or $\times_{t=1}^T$, appears, we take this to mean $t = 1, \ldots$, $\sum_{t\in\mathbb{N}}$, or $\times_{t\in\mathbb{N}}$, respectively.

[5] In what follows, we adopt the following notational conventions. First, all Polish spaces are endowed with their Borel σ-algebra. Second, product spaces are endowed with their product σ-algebra. Third, for a Polish space, Y, we let $\Delta(Y)$ denote the set of all Borel probability measures over Y, endowed with the weak* topology. Thus, $\Delta(Y)$ is also a Polish space (Aliprantis and Border, 2006). For any two measurable spaces, X and Y, a transition probability from X to Y is a measurable function $\zeta : X \mapsto \Delta(Y)$. When integrating under the measure $\zeta(x)$, we use the notation $\zeta(\cdot|x)$.

Mechanisms. In each period, the allocation is determined by a mechanism $\mathbf{M}_t = (M^{\mathbf{M}_t}, S^{\mathbf{M}_t}, \varphi^{\mathbf{M}_t})$, where $M^{\mathbf{M}_t}$ and $S^{\mathbf{M}_t}$ are the mechanism's input and output messages and $\varphi^{\mathbf{M}_t}$ assigns to each $m \in M^{\mathbf{M}_t}$ a distribution over $S^{\mathbf{M}_t} \times A$. We endow the principal with a collection $\{(M_i, S_i)\}_{i\in\mathcal{I}}$ of input and output message sets, such that (i) M_i, S_i are Polish spaces, (ii) $|\Theta| \leq |M_i|$, M_i is at most countable, and (iii) $|\Delta(\Theta)| \leq |S_i|$. Moreover, we assume $(\Theta, \Delta(\Theta))$ is an element in that collection. Let $\mathcal{M}_{\mathcal{I}}$ denote the set of all mechanisms with message sets $(M_i, S_i)_{i\in\mathcal{I}}$ that is, $\mathcal{M}_{\mathcal{I}} = \cup_{i,j\in\mathcal{I}} \left\{\varphi : M_i \mapsto \Delta(S_j \times A) : \varphi \text{ is measurable }\right\}$.

Two remarks are in order. First, we restrict the principal to choosing mechanisms in $\mathcal{M}_{\mathcal{I}}$. This restriction allows us to have a well-defined strategy space for the principal, thereby avoiding set-theoretic issues related to self-referential sets. The analysis that follows shows the choice of the collection plays no further role in the analysis. Second, because each M_i is at most countable, the set of mechanisms $\mathcal{M}_{\mathcal{I}}$ is a Polish space. As we discuss in Appendix D, this property is key to being able to define a mechanism-selection game for a given collection $\mathcal{I}$ (see also Section 2).

Mechanism-selection game(s). Each collection $\mathcal{I}$ induces a mechanism-selection game, which we denote by $G_{\mathcal{I}}$, and is defined as follows. At the beginning of each period, both players observe the realization of a public randomization device, $\omega \sim U[0,1]$. The agent also privately observes her type θ_t. The principal then offers the agent a mechanism, $\mathbf{M}_t$, with the property that for all $m \in M^{\mathbf{M}_t}$, $\varphi^{\mathbf{M}_t}(S^{\mathbf{M}_t} \times \mathcal{A}_t(a^{t-1})|m) = 1$, where recall that a^{t-1} describes the allocations implemented through period $t-1$. Observing the mechanism, the agent decides whether to participate in the mechanism ($\pi = 1$) or not ($\pi = 0$). If she does not participate in the mechanism, a^* is implemented and the game proceeds to period $t+1$. Instead, if she chooses to participate, she sends a message $m \in M^{\mathbf{M}_t}$, which is unobserved by the principal. An output message and an allocation (s_t, a_t) are drawn according to $\varphi^{\mathbf{M}_t}(\cdot|m)$. The output message and the allocation are observed by both the principal and the agent, and the game proceeds to period $t+1$.

Histories. The game $G_{\mathcal{I}}$ has two types of histories: public and private. Public histories capture what the *principal* knows through period t: the past realizations of the public randomization device, his past choices of mechanisms, the agent's participation decisions, and the realized output messages and allocations. We let h^t denote a public history through period t and let H^t denote the set of all such histories. Instead, private histories capture what the *agent* knows through period t. First, the agent knows the public history of the game, her past types, and her input messages into the mechanism (henceforth, the agent history). Second, the agent also knows her current private information. We let h_A^t denote an agent's history through period t and let $H_A^t(h^t)$ denote the set of agent histories consistent with public history h^t. Thus, $H_A^t(h^t) \times \Theta$ denotes the set of private histories consistent with public history h^t.

Belief system and strategies. In Markov environments, it is important to keep track of two beliefs for the principal. The first is the belief he holds about the agent's type θ_t at the end of period t; the second is the belief he holds about the agent's type θ_{t+1} at the beginning of period $t+1$, after applying F_{t+1}. We denote the former by μ_{t+1} and the latter by ν_{t+1}. That is, $\mu_{t+1}(h_A^{t+1}|h^{t+1})$ is the probability the principal assigns to the agent being at information set h_A^{t+1} at the end of period t when h^{t+1} is the publicly available information, whereas

$$\nu_{t+1}(h_A^{t+1}, \theta_{t+1}|h^{t+1}) = \mu_{t+1}(h_A^{t+1}|h^{t+1})F_{t+1}(\theta_{t+1}|\theta_t, a_t),$$

is the probability the principal assigns to the agent being at information set $h_A^{t+1} = (h_A^t, \theta_t, \cdot)$ at the end of period t *and* her type being θ_{t+1} in period $t+1$, where θ_t is the agent's type in period t and a_t is the allocation in period t consistent with h_A^{t+1}.

A behavioral strategypage-strategies for the principal is a collection of measurable mappings $(\sigma_{Pt})_{t=1}^T$, where for each period t and each public history h^t, $\sigma_{Pt}(h^t)$ describes the principal's (possibly random) choice of mechanism at h^t.[6] Similarly, a behavioral strategy for the agent is a collection of measurable mappings $(\sigma_{At})_{t=1}^T \equiv (\pi_t, r_t)_{t=1}^T$, where for each period t, each private history (h_A^t, θ_t), and each mechanism, $\mathbf{M}_t$, $\pi_t(h_A^t, \theta_t, \mathbf{M}_t)$ describes the agent's participation decision, whereas $r_t(h_A^t, \theta_t, \mathbf{M}_t)$ describes the agent's choice of input messages in the mechanism, conditional on participation. The tuple $(\sigma_P, \sigma_A, \mu) \equiv (\sigma_{Pt}, \sigma_{At}, \mu_t)_{t=1}^T$ defines an *assessment.*

Equilibrium. For each collection $\mathcal{I}$, we study the equilibria of the mechanism-selection game $G_{\mathcal{I}}$. By equilibrium, we mean *Perfect Bayesian equilibrium* (henceforth, PBE), informally defined as follows. An assessment $(\sigma_P, \sigma_A, \mu)$ is a PBE if it is sequentially rational and the belief system satisfies Bayes' rule where possible. The formal statement is in Appendix A. For now, we note that if the principal's strategy space is finite, Θ is finite, and the mechanisms used by the principal have finite support, our definition of PBE coincides with that in Fudenberg and Tirole (1991).

Equilibrium outcomes and payoffs. The prior $\nu_1 \equiv F_1$ together with a strategy profile (σ_P, σ_A) determine a distribution over the terminal nodes H_A^{T+1}. We are interested instead in the distribution they induce over the payoff-relevant outcomes, $(\Theta \times A)^T$. We say $\eta \in \Delta\left((\Theta \times A)^T\right)$ is a PBE outcome if a PBE of the mechanism-selection game exists that induces η. Each PBE outcome η induces a payoff tuple, $\left(w, (u_\theta)_{\theta\in\Theta}\right)$, where $w = \mathbb{E}_\eta W$ and for each θ in Θ, $u_\theta = \mathbb{E}_{\eta|\theta} U$. We denote by $\mathcal{E}_{\mathcal{I}}^*$ the set of PBE payoffs of $G_{\mathcal{I}}$.

2.1. Canonical mechanisms and assessments

Theorem 1 singles out one mechanism-selection game and a class of assessments that allows us to replicate any equilibrium payoff of $G_{\mathcal{I}}$, for any collection $\mathcal{I}$ of input and output messages. Following Doval and Skreta (2022), we dub this extensive form the canonical game and the class of assessments, canonical assessments, which we formally define next.

[6] Because the set $\mathcal{M}_{\mathcal{I}}$ is Polish, the sets of public and private histories are the (at most) countable product of Polish spaces. Thus, the sets of public and private histories are Polish spaces.

Canonical game. The canonical game is a mechanism selection game in which the principal can only select mechanisms that use *current* type reports as input messages and distributions over the agent's current type as output messages. That is, in the canonical game, the principal can only offer mechanisms in which $M = \Theta$ and $S = \Delta(\Theta)$. We denote the set of equilibrium payoffs of the canonical game by $\mathcal{E}^*$.

Definition 1 (*Flow Direct Blackwell Mechanisms*)**.** A mechanism $(\Theta, \Delta(\Theta), \varphi)$ is a *flow direct Blackwell mechanism* (henceforth, f-DBM) if the mapping $\varphi : \Theta \mapsto \Delta(\Delta(\Theta) \times A)$ can be obtained as the *composition* of two mappings, $\beta : \Theta \mapsto \Delta(\Delta(\Theta))$ and $\alpha : \Delta(\Theta) \mapsto \Delta(A)$. Formally, for each θ and each pair of measurable subsets, $\widetilde{U} \subset \Delta(\Theta)$ and $\widetilde{A} \subset A$, $\varphi(\widetilde{U} \times \widetilde{A}|\theta) = \int_{\widetilde{U}} \alpha(\widetilde{A}|\mu)\beta(d\mu|\theta)$.

In a f-DBM, conditional on the output message, the allocation is drawn *independently* of the agent's type report. Interpreted as a Blackwell experiment, β encodes how much information the principal learns about the agent's type. Instead, the allocation rule α describes the mechanism's (possibly randomized) allocation, given the information that the principal learns about the agent's type. We let $\mathcal{M}_C$ denote the set of direct Blackwell mechanisms.

Remark 1 (*Comparison with Doval and Skreta (2022)*)**.** At first glance, the canonical game herein endows the principal with the same mechanisms as the canonical game in Doval and Skreta (2022). When the agent's type is fully-persistent, a mechanism that asks the agent to submit a type report is effectively allowing the agent to report all her payoff-relevant private information. Instead, when the agent's type is not fully-persistent, a mechanism that elicits a report in Θ does not necessarily allow the agent to submit all her payoff-relevant information. Indeed, the natural generalization of the canonical mechanisms in Doval and Skreta (2022) to the case of non-fully persistent types would have Θ^t as the set of input messages. For this reason, we refer to the analogue of DBMs in Doval and Skreta (2022) as *flow* DBMs to stress the mechanism only attempts to elicit the agent's *current* type. As we explain after Theorem 1, the restriction to Markov environments is responsible for the result that f-DBMs are without loss of generality.

Canonical assessments. A canonical assessment specifies behavior for the principal and the agent that is, in a sense, simple. First, the principal always chooses f-DBMs. Second, the agent best responds to the principal's equilibrium choice of mechanisms by participating. Third, input and output messages have *literal* meaning: Conditional on participating, the agent truthfully reports her type, and if the mechanism outputs $\mu \in \Delta(\Theta)$, μ coincides with the principal's updated beliefs about the agent's type. The notation μ signifies the mechanism uses the principal's beliefs at the end of period t as the output messages. Formally:

Definition 2 (*Canonical Assessments*)**.** An assessment $(\sigma_P, \sigma_A, \mu)$ of mechanism-selection game $G_{\mathcal{I}}$ is canonical if the following holds for all $t \geq 1$ and all public histories h^t:

1. The principal offers f-DBMs, that is, $\sigma_{Pt}(h^t)(\mathcal{M}_C) = 1$.
2. For all mechanisms $\mathbf{M}_t$ in the support of the principal's strategy at h^t,
   (a) For all types θ_t in the support of the principal's beliefs in period t, $v_t(h^t)$, $\pi_t(h^t_A, \theta_t, \mathbf{M}_t) = 1$,
   (b) For all types θ_t in Θ, $r_t(h^t_A, \theta_t, \mathbf{M}_t) = \delta_{\theta_t}$, and
   (c) The mechanism's output belief μ coincides with the principal's updated belief about the agent's type at the end of period t. Formally, for $h^{t+1} = (h^t, \mathbf{M}_t, \mu, \cdot)$, the marginal of $\mu_{t+1}(h^{t+1})$ on θ_t, $\mu_{t+1\Theta_t}(h^{t+1})$, coincides with μ.
3. The agent's strategy depends only on her private type in period t and the public history.[7]

We let $\mathcal{E}^C$ denote the set of equilibrium payoffs of the canonical game that are induced by canonical PBE assessments (henceforth, canonical PBE).

3. Revelation principle for Markov environments

Section 3 presents the paper's main result: To characterize the set of equilibrium outcomes that can arise in some mechanism-selection game, it is enough to characterize the canonical PBE outcomes of the canonical game. Formally,

Theorem 1 (*Revelation Principle for Markov Environments*)**.** *For any PBE assessment of any mechanism-selection game $G_{\mathcal{I}}$, a payoff-equivalent canonical PBE of the canonical game exists. That is,*

$$\bigcup_{\mathcal{I}} \mathcal{E}^*_{\mathcal{I}} = \mathcal{E}^* = \mathcal{E}^C.$$

Theorem 1 plays the same role in mechanism design with limited commitment as the revelation principle does in the commitment case. First, it identifies a well-defined set of mechanisms, $\mathcal{M}_C$, that, without loss of generality, the principal uses to implement any equilibrium outcome. Second, it simplifies the analysis of the behavior of the agent in the game induced by the mechanisms chosen by the principal: we can always restrict attention to assessments in which the agent participates and truthfully reports her type. As

[7] Whereas items 2a and 2b of Definition 2 imply the agent's strategy depends only on her *current* private type and the public history *on the path* of the equilibrium strategy starting at h^t, 3 implies this property also holds *off the path* of the equilibrium strategy starting at h^t, e.g., when the principal deviates and offers a mechanism not in the support of $\sigma_{P_t}(h^t)$.

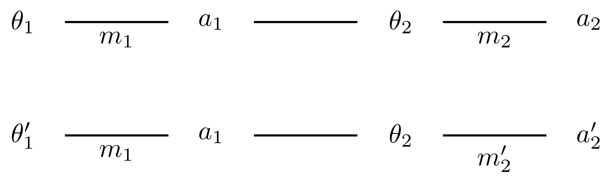

Fig. 1. Private-history independence in Markov environments.

we illustrate in Doval and Skreta (2023), this restriction allows us to reduce the agent's behavior to a set of constraints that the mechanism must satisfy, as in the case of commitment.

The proof of Theorem 1 follows from similar steps to those in Doval and Skreta (2022), except for two features worth noting. First, Theorem 1 features a separation between the beliefs the output message represents (μ_{t+1}) and the beliefs the principal uses in the next period to select mechanisms (ν_{t+1}). This separation highlights that the principal in period t attempts to design his prior for period $t+1$ (which nature will then transform into ν_{t+1} through F_{t+1}.)[8]

Second, whereas the revelation principle in Doval and Skreta (2022) implies it is without loss of generality to restrict attention to strategies in which the agent does not condition her strategy on her past input messages, the history independence of the agent's strategy in Theorem 1 is stronger: The agent conditions her strategy on neither her past communication nor her past payoff-relevant types. This is where we more prominently employ the restriction to Markov environments. Assuming the environment is Markov and given a PBE in which the agent conditions her strategy either on her past input messages or her past types at some public history h^t, we show another payoff-equivalent PBE exists in which she does not (see Proposition B.1).

Proposition B.1 then affords an important simplification: In each period t, the principal only needs to elicit the agent's current payoff-relevant type θ_t, not the past realizations. We believe this simplification is important for applications. In more general environments, a similar result would obtain, but the principal may need to elicit the whole realization $(\theta_1, \ldots, \theta_t)$.

The stronger form of history independence also implies that contrary to the main theorem in Doval and Skreta (2022), the mechanism-selection and canonical games implement the same set of payoffs but not necessarily the same equilibrium outcomes.[9] To understand why we can replicate the equilibrium payoffs, but not the equilibrium outcomes, consider the following two-period illustration:[10]

Fig. 1 depicts two histories that may lead to the agent's type being θ_2 in $t = 2$. In the one depicted above, the agent's type is θ_1, she sends message m_1 and obtains allocation a_1 in period 1. In the one depicted below, the agent's type is θ_1', she sends message m_1 and obtains the allocation a_1 in period 1. In both cases, the agent is sending the same input message in $t = 1$. Since both histories are consistent with the same public history (in this case, $h^1 = (a_1)$), the agent faces the same mechanism in period 2. Finally, note that the agent's reporting strategy in period 2 depends on her type in period 1. If her period 1 type is θ_1, the agent sends message m_2 and thus obtains allocation a_2. Instead, if her period 1 type is θ_1', the agent sends message m_2' and obtains allocation a_2'.

When we make the agent's reporting strategy in $t = 2$ independent of both her past input messages and her past types, we are not able to preserve the outcome distribution in the original assessment. To see this, note that when we make the agent's strategy independent of the payoff-irrelevant part of her private history in the proof of Proposition B.1, we have the agent play a particular randomized strategy. In Fig. 1, this corresponds to the agent sending both m_2 and m_2' when her type is θ_2 in period 2, independently of whether her type is θ_1 or θ_1' in period 1. It follows that we are not able to replicate the distribution over $(\tilde{\theta}_1, \tilde{\theta}_2, \tilde{a}_1, \tilde{a}_2) \in \Theta^2 \times A^2$ implied by the original assessment. For instance, the sequence $(\theta_1, \theta_2, a_1, a_2')$ will now have positive probability.

However, since the agent's payoffs are time separable and only depend on her current type, we are able to replicate the expected payoffs the agent obtains in the original assessment. In Fig. 1, it must be the case that $u_2(a_2, \theta_2) = u_2(a_2', \theta_2)$; otherwise, it would not be optimal for the agent to send m_2 and m_2' with positive probability. This, in turn, implies that $u_1(a_1, \theta_1) + \delta u_2(a_2, \theta_2) = u_1(a_1, \theta_1) + \delta u_2(a_2', \theta_2)$ and similarly for θ_1'. Thus, when we have the agent mix over m_2 and m_2' in $t = 2$ conditional on θ_2, we do not affect the agent's payoffs, even if we do change the outcome distribution.

4. Concluding remarks

This paper provides a revelation principle for dynamic mechanism design in Markov environments in which a principal, who can only commit to short-term mechanisms, interacts with an agent whose private information evolves stochastically over time. Our tool opens up the study of optimal mechanisms under limited commitment with evolving types, which is relevant for the study of optimal taxation, social insurance, managerial economics, among other applications. Our model has certain features worth discussing.

First, throughout the paper we assume the set of types Θ is at most countable. In Appendix D, we provide the necessary formalisms to extend the revelation principle to the case of a continuum type space, a leading specification in mechanism design (see Theorem D.1). Our companion paper (Doval and Skreta, 2023) applies this extension to an industrial organization application.

[8] See Ely (2017) for another model where the same choice is made: the information designer designs μ_{t+1} and then nature transitions this "prior" using the Markov transition.

[9] The revelation principle-style arguments in Peters (2001), Hart et al. (2017) and Ben-Porath et al. (2019) are also in terms of payoff equivalence.

[10] For simplicity, Fig. 1 abstracts away from many details of our model. For instance, we omit the output message in the $t = 1$ mechanism.

Second, we assume the principal interacts with a single agent. Whereas the case of multiple agents is an extension worth pursuing, it is outside of this paper's scope. With multiple agents, the principal has discretion over what the agents observe about each others' output messages and allocations, which he nevertheless observes. Thus, the principal may become endogenously privately informed. Little is known about dynamic exogenously informed principal problems, even with commitment. We expect the limited commitment case in which information is endogenously obtained to bring forth additional challenges.

Third, we assume the principal can only commit to the mechanism for one period. By appropriately defining a period's length, this setting allows for the *changing principal* framework typically used to represent sequences of governments in the regulation and public finance literatures (c.f. Laffont and Tirole (1988)). Instead, our mechanism-selection game does not cover the case in which the principal loses his commitment power with some probability and this event is not observed by the agent. Such probabilistic weakening of commitment is employed in the information design settings of Fréchette et al. (2022) and Lipnowski et al. (2022). Understanding the implications of such weakenings of commitment for optimal mechanism design is left for future work.

Appendix A. Collected definitions and notation

Appendix A introduces the necessary notation to define the payoffs from an assessment and hence, the definition of Perfect Bayesian equilibrium. It also collects notation that is used in the proofs. Throughout Appendices A and B, we assume Θ is at most countable and the mechanisms used by the principal have finite support. The proof that Theorem 1 holds when either Θ is a continuum or the support of the principal's mechanism is not finite follows from the same steps as in Doval and Skreta (2022), so we omit it for simplicity. Appendix C defines the main objects needed to extend the proof presented herein to the case in which either Θ is a continuum or the support of the principal's mechanism is not finite.

Histories and strategies. Let $\mathcal{M}_{i,j}$ denote the set of transition probabilities from M_i to $S_j \times A$; since M_i is at most countable $\mathcal{M}_{i,j}$ is Polish. With this notation, $\mathcal{M}_{\mathcal{I}} = \cup_{i,j\in\mathcal{I}}\mathcal{M}_{i,j}$. To simplify notation, we do not explicitly include the agent's decision to participate in the mechanism in the histories of the game. Instead, we follow the convention that if the agent does not participate, the input message is $\emptyset$, the output message is $\emptyset$, and the allocation is a^*. Thus, when the principal offers a mechanism in $\mathcal{M}_{i,j}$, the possible private outcomes are $M_iS_jA_\emptyset \equiv \left(M_i \times S_j \times A\right) \cup \{(\emptyset,\emptyset,a^*)\}$, while the public outcomes are $S_jA_\emptyset \equiv \left(S_j\times A\right)\cup\{(\emptyset,a^*)\}$. We endow $M_iS_jA_\emptyset$ and $S_jA_\emptyset$ with the disjoint topology and we note that they are Polish sets under than topology.

With the above notation, an outcome at the end of period t is an element of $\mathsf{Z}_{A,t} \equiv \Theta\times\cup_{i,j\in\mathcal{I}}\left(\mathcal{M}_{i,j}\times M_iS_jA_\emptyset\right)\times\Omega$; the public component of the outcome in period t is an element of $\mathsf{Z} = \cup_{i,j\in\mathcal{I}}\left(\mathcal{M}_{i,j}\times S_jA_\emptyset\right)\times\Omega$. Since $\mathcal{I}$ is at most countable, Z_A and Z are Polish when endowed with the disjoint topology. For $t\geq 1$, public histories at the beginning of period t are $H^t = \Omega\times\mathsf{Z}^{t-1}$, while private histories are $H_A^t = \Omega\times\mathsf{Z}_A^{t-1}$, with the understanding that $\mathsf{Z}^0 = \mathsf{Z}_A^0 = \{\emptyset\}$ and $\emptyset$ denotes the empty history.

We write the agent private histories so that $h_A^t = (h_A^{t-1},\theta_{t-1},\mathbf{M}_{t-1},m_{t-1},s_{t-1},a_{t-1},\omega_t)$, with the convention that when $t=1$, $h_A^1 = \{\omega_1\}$ for some $\omega_1\in[0,1]$. Thus, an information set for the agent in period t is given by (h_A^t,θ_t), where h_A^t is the agent history through period t and θ_t the agent's realized type at the beginning of period t.

The principal's behavioral strategy is a collection $(\sigma_{Pt})_{t=1}^T$ where $\sigma_{Pt}: H^t\mapsto\Delta(\mathcal{M}_{\mathcal{I}})$ is a measurable function. The agent's behavioral strategy $(\sigma_{At})_{t=1}^T$ is a collection $\sigma_{At}\equiv(\pi_t,r_t)_{t=1}^T$ such that $\pi_t: H_A^t\times\Theta\mathcal{M}_{\mathcal{I}}\mapsto\Delta(\{0,1\})$ and $r_t: H_A^t\times\Theta\times\mathcal{M}_{\mathcal{I}}\mapsto\Delta(\cup_{i\in\mathcal{I}}M_i)$ are measurable and $r_t(h_A^t,\theta_t,\mathbf{M}_t)(M^{\mathbf{M}_t}) = 1$.

Shorthand notation. Given a mechanism $\mathbf{M}_t$, let $z_{(s_t,a_t)}(\mathbf{M}_t)$ denote the tuple $\mathbf{M}_t,s_t,a_t$, which summarizes the period-t outcomes from offering $\mathbf{M}_t$, where $(s_t,a_t)\in S^{\mathbf{M}_t}A_\emptyset$. Note that any public history at the beginning of period t can be written as $h^t = (h^{t-1},z_{(s_{t-1},a_{t-1})}(\mathbf{M}_{t-1}),\omega_t)$, with the convention that when $t=1$, $h^1=\{\omega_1\}$ for some $\omega_1\in[0,1]$.

Finally, given an assessment, (σ_P,σ_A,μ), it is useful to collapse the distribution on $M^{\mathbf{M}_t}S^{\mathbf{M}_t}A_\emptyset$, defined by

$$(1-\pi_t(h_A^t,\theta_t,\mathbf{M}_t))\mathbb{1}[(m_t,s_t,a_t)=(\emptyset,\emptyset,a^*)] + \pi_t(h_A^t,\theta_t,\mathbf{M}_t)r_t(h_A^t,\theta_t,\mathbf{M}_t)(m_t)\varphi^{\mathbf{M}_t}(s_t,a_t|m_t) \tag{A.1}$$

and we denote it by $\kappa_t^{\sigma_A}(m_t,s_t,a_t|h_A^t,\theta_t,\mathbf{M}_t)$.

Beliefs and payoffs. To define Perfect Bayesian equilibrium, we need to define the principal and the agent's payoff from a given assessment. To do so, fix an assessment, (σ_P,σ_A,μ). The prior F_1 and the strategy profile $\sigma=(\sigma_P,\sigma_A)$ induce a probability distribution over the terminal histories H_A^{T+1}, P^σ, via the Ionescu-Tulcea theorem (Pollard, 2002) (Appendix C formally defines this distribution). Moreover, fixing t and (h_A^t,θ_t), the measure $P^{\sigma|h_A^t,\theta_t}$ corresponds to the measure induced by drawing with probability 1 (h_A^t,θ_t) and then using the continuation strategy profile to determine the distribution over the continuation histories. Fix a public history h^t. Recall that the principal's prior belief at h^t $v_t(\cdot|h^t)\in\Delta(H_A^t(h^t)\times\Theta)$ is obtained from the belief assessment μ and the transition probability in period t, F_t, as follows. Let $h^t=(h^{t-1},z_{(s_{t-1},a_{t-1})}(\mathbf{M}_{t-1}),\omega_t)$ and let $h_A^t=(h_A^{t-1},\theta_{t-1},m_{t-1},z_{(s_{t-1},a_{t-1})}(\mathbf{M}_{t-1}),\omega_t)$. Then,

$$v_t(h_A^t,\theta_t|h^t) = \mu_t(h_A^t|h^t)F_t(\theta_t|\theta_{t-1},a_{t-1}). \tag{A.2}$$

The principal's continuation payoff at h^t is then given by[11]

$$W_t(\sigma,\mu|h^t) = \sum_{h_A^t\in H_A^t(h^t),\theta_t} v_t(h_A^t,\theta_t|h^t)\mathbb{E}_{\sigma_{Pt}(h^t)}\left[\mathbb{E}^{P^{\sigma|(h_A^t,\theta_t,\mathbf{M}_t)}}\left[\sum_{\tau=t}^T w_\tau(\cdot,\theta_t)\right]\right] \equiv \mathbb{E}_{\sigma_{Pt}(h^t)}\left[W_t(\sigma,\mu|h^t,\mathbf{M}_t)\right]. \tag{A.3}$$

[11] In a slight abuse of notation, we denote the principal's continuation payoff by W_t to signify that this is the expectation from period t onwards of the principal's payoff W defined in Section 2.

The principal's payoff from offering $\mathbf{M}_t$ at h^t, $W_t(\sigma,\mu|h^t,\mathbf{M}_t)$, depends on the belief system μ only through $F_t(\cdot)\mu_t(\cdot|h^t)\equiv \nu_t(\cdot|h^t)$.

We now show $W_t(\sigma,\mu|h^t,\mathbf{M}_t)$ has a recursive representation. To do so, we first show the principal's payoff from offering mechanism $\mathbf{M}_t$ at h^t conditional on the agent being at node (h^t_A,θ_t) can be written recursively:

$$W_t(\sigma|h^t_A,\theta_t,\mathbf{M}_t)= \sum_{(m_t,s_t,a_t)\in M^{\mathbf{M}_t}S^{\mathbf{M}_t}A_\emptyset}\kappa_t^{\sigma_A}(m_t,s_t,a_t|h^t_A,\theta_t,\mathbf{M}_t)\left(w_t(a_t,\theta_t)+\delta\mathbb{E}_\omega\left[\sum_{\theta_{t+1}}F_{t+1}(\theta_{t+1}|\theta_t,a_t)\mathbb{E}_{\sigma_{P_t}(h^{t+1})}W_{t+1}(\sigma|h^{t+1}_A,\theta_{t+1},\cdot)\right]\right). \tag{A.4}$$

Using Eq. (A.4), we can write the principal's payoff at $(h^t,\mathbf{M}_t)$ as follows

$$W_t(\sigma,\mu|h^t,\mathbf{M}_t)=\sum_{h^t_A,\theta_t}\nu_t(h^t_A,\theta_t|h^t)W_t(\sigma|h^t_A,\theta_t,\mathbf{M}_t).$$

Using Eq. (C.1), define

$$\Pr{}^{\nu,\sigma_A}_{t+1}(h^t,z_{(s_t,a_t)}(\mathbf{M}_t)|h^t)=\sum_{(h^t_A,\theta_t),m_t\in M^{\mathbf{M}_t}_\emptyset}\nu_t(h^t_A,\theta_t|h^t)\kappa_t^{\sigma_A}(m_t,s_t,a_t|h^t_A,\theta_t,\mathbf{M}_t). \tag{A.5}$$

With this notation at hand, we can express the principal's payoff $W_t(\sigma,\mu|h^t,\mathbf{M}_t)$ as follows

$$\begin{aligned}
&W_t(\sigma,\mu|h^t,\mathbf{M}_t)= \\
&\sum_{(s_t,a_t)}\Pr{}^{\nu,\sigma_A}_{t+1}(h^t,z_{(s_t,a_t)}(\mathbf{M}_t)|h^t)\mathbb{E}_{\mu_{t+1}(\cdot|h^t,z_{(s_t,a_t)}(\mathbf{M}_t))}\left[w_t(a_t,\cdot)+\delta\mathbb{E}_{\omega_{t+1}}\sum_{\theta_{t+1}}F_{t+1}(\theta_{t+1}|\cdot,a_t)\mathbb{E}_{\sigma_{Pt+1}}W_{t+1}(\sigma|h^{t+1}_A,\theta_{t+1},\cdot)\right] \\
&=\sum_{(s_t,a_t)\in S^{\mathbf{M}_t}A_\emptyset}\Pr{}^{\nu,\sigma_A}_{t+1}(h^t,z_{(s_t,a_t)}(\mathbf{M}_t)|h^t)\left[\begin{array}{l}\sum_{(h^t_A,\theta_t),m_t}\mu_{t+1}(h^t_A,\theta_t,m_t,z_{(s_t,a_t)}(\mathbf{M}_t)|h^t,z_{(s_t,a_t)}(\mathbf{M}_t))w_t(a_t,\theta_t)\\+\delta\mathbb{E}_{\omega_{t+1}}\mathbb{E}_{\sigma_{Pt+1}}\left[W_{t+1}(\sigma,\mu|h^t,z_{(s_t,a_t)}(\mathbf{M}_t),\omega_{t+1},\mathbf{M}_{t+1})\right]\end{array}\right],
\end{aligned} \tag{A.6}$$

which completes the recursion. In more compact notation, Eq. (A.4) implies that

$$W_t(\sigma,\mu|h^t,\mathbf{M}_t)= \sum_{h^{t+1}\in H^{t+1}(h^t,\mathbf{M}_t)}\Pr{}^{\nu,\sigma_A}_{t+1}(h^{t+1}|h^t)\left[\sum_{h^{t+1}_A\in H^{t+1}_A(h^{t+1})}\mu_{t+1}(h^{t+1}_A|h^{t+1})w_t(\cdot)+\delta\mathbb{E}_{\omega_{t+1},\sigma_P(h^{t+1})}W_{t+1}(\sigma,\mu|h^{t+1},\mathbf{M}_{t+1})\right].$$

Similarly, for the agent we have that in period t at node (h^t_A,θ_t) when the principal offers her mechanism $\mathbf{M}_t$, her payoff is given by:

$$\begin{aligned}
U_t(\sigma|h^t_A,\theta_t,\mathbf{M}_t)&=\sum_{(m_t,s_t,a_t)}\kappa_t^{\sigma_A}(m_t,s_t,a_t|h^t_A,\theta_t,\mathbf{M}_t)\left(u_t(a_t,\theta_t)+\delta\mathbb{E}^{P^{\sigma|h^t_A,\theta_t,\cdot}}\sum_{\tau\geq t}\delta^{\tau-1}u_\tau(\cdot)\right)\\
&=\sum_{(m_t,s_t,a_t)}\kappa_t^{\sigma_A}(m_t,s_t,a_t|h^t_A,\theta_t,\mathbf{M}_t)\left(u_t(a_t,\theta_t)+\delta\mathbb{E}_{\omega_{t+1}}\left[\sum_{\theta_{t+1}}F_{t+1}(\theta_{t+1}|\theta_t,a_t)\mathbb{E}_{\sigma_{Pt+1}}U_{t+1}(\sigma|h^{t+1}_A,\theta_{t+1},\mathbf{M}_t)\right]\right),
\end{aligned} \tag{A.7}$$

where in the above expression $h^{t+1}_A=(h^t_A,\theta_t,\mathbf{M}_t,m_t,s_t,a_t,\omega_{t+1})$. The second line in Eq. (A.7) highlights that the agent's payoff also has a recursive structure. We use this frequently in the proof of Proposition B.1.

Perfect Bayesian equilibrium. Having defined the principal and the agent's payoffs, we can formally define Perfect Bayesian equilibrium.

Definition A.1. An assessment (σ_P,σ_A,μ) is *sequentially rational* if for all $t\geq 1$ and public histories h^t, the following hold:

1. If $\mathbf{M}_t$ is in the support of $\sigma_{Pt}(h^t)$, then $W_t(\sigma,\mu|h^t,\mathbf{M}_t)\geq W_t(\sigma,\mu|h^t,\mathbf{M}'_t)$ for all $\mathbf{M}'_t\in\mathcal{M}_\mathcal{I}$,
2. For all $(h^t_A,\theta_t)\in H^t_A(h^t)\times\Theta$, and $\mathbf{M}_t$ in $\mathcal{M}_\mathcal{I}$, $U_t(\sigma|h^t_A,\theta_t,\mathbf{M}_t)\geq U_t(\sigma_P,\sigma'_A|h^t_A,\theta_t,\mathbf{M}_t)$ for all σ'_A.

Definition A.2. An assessment (σ_P,σ_A,μ) *satisfies Bayes' rule where possible* if for all public histories h^t and mechanisms $\mathbf{M}_t$ the following holds:

$$\begin{aligned}
&\mu_{t+1}(h^t_A,\theta_t,m_t,z_{(s_t,a_t)}(\mathbf{M}_t),\omega_{t+1}|h^t,z_{(s_t,a_t)}(\mathbf{M}_t),\omega_{t+1})\Pr{}^{\nu,\sigma_A}_{t+1}(h^t,z_{(s_t,a_t)}(\mathbf{M}_t)|h^t,\mathbf{M}_t)\\
&=\nu_t(h^t_A,\theta_t|h^t)\kappa_t^{\sigma_A}(m_t,s_t,a_t|h^t_A,\theta_t,\mathbf{M}_t),
\end{aligned} \tag{A.8}$$

where $\nu_t(h^t_A,\theta_t|h^t)$ is defined as in Eq. (A.2).

Definition A.3. An assessment (σ_P,σ_A,μ) is a *Perfect Bayesian equilibrium* if it is sequentially rational and satisfies Bayes' rule where possible.

The only difference between Bayes' rule where possible (Definition A.2) and consistency in sequential equilibrium is that under PBE, the principal can assign zero probability to a type and then, after the agent deviates, can assign positive probability to that same type.

Prunning. Given a mechanism $\mathbf{M}_t$, let $(S^{\mathbf{M}_t} \times A)_+ = \{(s_t, a_t) : (\exists m \in M^{\mathbf{M}_t})\varphi^{\mathbf{M}_t}(s_t, a_t|m) > 0\}$. The set $S^{\mathbf{M}_t} \times A \setminus (S^{\mathbf{M}_t} \times A)_+$ has zero probability *regardless* of the agent's strategy. Hence, if we remove from the tree those paths that are consistent with mechanism $\mathbf{M}_t$ and $(s, a) \notin (S^{\mathbf{M}_t} \times A)_+$, this does not change the set of equilibrium outcomes. Hereafter, these histories are removed from the tree.

Principal pure strategies: Lemma D.1 in Doval and Skreta (2022) shows it is without loss of generality to focus on PBE assessments of $G_{\mathcal{I}}$ in which the principal plays a pure strategy. In what follows, we thus focus on PBE assessments with this property.

Appendix B. Proof of Theorem 1

The proof of Theorem 1 follows from the proof of Propositions B.1–B.2 below.

Proposition B.1. *For every PBE assessment $(\sigma_P, \sigma_A, \mu)$ of $G_{\mathcal{I}}$, a payoff-equivalent PBE assessment $(\sigma'_P, \sigma'_A, \mu')$ exists such that for every $t \geq 1$ and every public history h^t, the agent's strategy only depends on her current type, θ_t, and the public history h^t.*

We relegate the proof of Proposition B.1 to Appendix B.1. In what follows, we focus on PBE of the mechanism-selection game that satisfy the properties of Proposition B.1 and abuse notation in the following two ways: First, we write the agent's strategy as a function of her private type and the public history alone, with the understanding that $\sigma_{At}(h^t_A, \theta_t) = \sigma_{At}(h^t, \theta_t)$ whenever $h^t_A \in H^t_A(h^t)$. Similarly, we write the belief system at history h^t as inducing distributions over Θ and not over $\Theta \times H^t_A(h^t)$.

Proposition B.2. *Fix $\mathcal{I}$ and let $(\sigma_P, \sigma_A, \mu)$ be a PBE assessment of $G_{\mathcal{I}}$ that satisfies Proposition B.1. Then, an outcome-equivalent canonical PBE assessment $(\sigma'_P, \sigma'_A, \mu')$ of $G_{\mathcal{I}}$ exists.*

Proof of Proposition B.2. Let $(\sigma_P, \sigma_A, \mu)$ be as in the statement of Proposition B.2. Let h^t be a public history and let $\mathbf{M}_t$ denote the mechanism that the principal offers at h^t under σ_{Pt}. Let Θ^+ denote the support of the principal's beliefs at h^t, $v_t(h^t)$.

For types in Θ^+, use Eq. (C.1) to define an auxiliary mapping $\varphi' : \Theta^+ \mapsto \Delta\left(S^{\mathbf{M}_t} A_\emptyset\right)$, as follows:

$$\varphi'(s_t, a_t|\theta_t) = \sum_{m \in M^{\mathbf{M}_t}_\emptyset} \kappa_t^{\sigma_A}(m_t, s_t, a_t|h^t, \theta_t, \mathbf{M}_t). \tag{B.1}$$

That is, φ' corresponds to the *direct* version of $\varphi^{\mathbf{M}_t}$ for $\theta_t \in \Theta^+$; we use it in what follows to construct an alternative mechanism for the principal, $\mathbf{M}'_t$, that uses message sets $(\Theta, \Delta(\Theta))$.

Omitting the dependence on (σ_A, v), recall that $\Pr_{t+1}(h^t, z_{(s_t,a_t)}(\mathbf{M}_t)|h^t, \mathbf{M}_t)$ denotes the probability of history $(h^t, z_{(s_t,a_t)}(\mathbf{M}_t))$ under the equilibrium strategy when the principal offers $\mathbf{M}_t$ at h^t (Eq. (A.5)). Eq. (A.5) implies we can write $\Pr_{t+1}$ using φ' as follows:

$$\Pr_{t+1}(h^t, z_{(s_t,a_t)}(\mathbf{M}_t)|h^t, \mathbf{M}_t) = \sum_{\theta_t \in \Theta^+} v_t(\theta_t|h^t)\varphi'(s_t, a_t|\theta_t).$$

In what follows, to simplify notation we omit the dependence of $\Pr_{t+1}(\cdot|h^t, \mathbf{M}_t)$ on $\mathbf{M}_t$.

The first step is to show the distribution over continuation histories $\Pr_{t+1}$ can be seen as inducing a distribution over posterior beliefs, allocations, and realizations of a public randomization device. To see this, for $\mu \in \Delta(\Theta)$, let $B(\mu)$ denote the set

$$B(\mu) = \left\{(s_t, a_t) \in S^{\mathbf{M}_t} A_\emptyset : \mu_{t+1}(\cdot|h^t, z_{(s_t,a_t)}(\mathbf{M}_t)) = \mu\right\},$$

and let $B_{a_t}(\mu)$ denote the projection of $B(\mu)$ onto $\{a_t\}$. In what follows, for any subset $B \subseteq S^{\mathbf{M}_t} A_\emptyset$, we abuse notation and write $\Pr_{t+1}(B|h^t)$ instead of $\sum_{(s_t,a_t)\in B} \Pr_{t+1}(h^t, z_{(s_t,a_t)}(\mathbf{M}_t)|h^t)$.

Let $\Delta(\Theta)^+$ denote the smallest subset of $\Delta(\Theta)$ such that $\Pr_{t+1}(B(\Delta(\Theta)^+)|h^t) = 1$. That is, $\Delta(\Theta)^+$ is the set of principal posterior beliefs that are pinned down via Bayes' rule. Using Eq. (B.2), define the principal's payoff conditional on $(h^t, z_{(s_t,a_t)}(\mathbf{M}_t))$ as follows:

$$W_t(\sigma, \mu_{t+1}(\cdot|h^{t+1})|h^t, z_{(s_t,a_t)}(\mathbf{M}_t)) =$$
$$\mathbb{E}_{\mu_{t+1}(\cdot|h^t, z_{(s_t,a_t)}(\mathbf{M}_t))}\left[w_t(a_t, \cdot) + \delta \mathbb{E}_{\omega_{t+1}} \sum_{\theta_{t+1}} F_{t+1}(\theta_{t+1}|\cdot, a_t)\mathbb{E}_{\sigma_{Pt+1}} W_{t+1}(\sigma|h^{t+1}_A, \theta_{t+1}, \cdot)\right].$$

Then, we can write the principal's payoff at history h^t when he offers $\mathbf{M}_t$ as follows:

$$\sum_{(s_t,a_t)} \Pr_{t+1}(h^t, z_{(s_t,a_t)}(\mathbf{M}_t)|h^t)W_t(\sigma, \mu_{t+1}(\cdot|h^t, z_{(s_t,a_t)}(\mathbf{M}_t))|h^t, z_{(s_t,a_t)}(\mathbf{M}_t)) \tag{B.2}$$
$$= \sum_{\mu \in \Delta(\Theta)^+} \Pr_{t+1}(B(\mu)|h^t) \sum_{a_t \in A} \underbrace{\frac{\Pr_{t+1}(B_{a_t}(\mu)|h^t)}{\Pr_{t+1}(B(\mu)|h^t)}}_{\text{allocation rule}} \sum_{s_t \in B_{a_t}(\mu)} \underbrace{\frac{\Pr_{t+1}(h^t, z_{(s_t,a_t)}(\mathbf{M}_t)|h^t)}{\Pr_{t+1}(B_{a_t}(\mu)|h^t)}}_{\text{public randomization}} W_t(\sigma, \mu|h^t, z_{(s_t,a_t)}(\mathbf{M}_t)).$$

The above equation shows two ways in which we can think of the distribution over continuation histories starting from h^t. The first is standard: we draw history $(h^t, z_{(s_t,a_t)}(\mathbf{M}_t))$ using the distribution induced by the equilibrium strategy, $\Pr_{t+1}(\cdot|h^t)$. The second is the one that delivers the direct Blackwell mechanisms: we first draw a belief μ using the distribution over continuation equilibrium beliefs induced by $\Pr_{t+1}(\cdot|h^t)$ and then we draw an allocation a_t, conditional on the continuation equilibrium belief coinciding with μ. The principal's posterior belief μ and the allocation a_t may still not be enough to pin down the continuation history, so we draw the output message s_t conditional on s_t being consistent with a_t and μ.

Conditional on the induced posterior belief μ, the second step shows (i) the probability that the allocation is a_t is independent of θ_t, and (ii) the probability that the output message is $s_t \in B_{a_t}(\mu)$ is independent of θ_t. To see this, note that for any belief $\mu \in \Delta(\Theta)^+$, for any $(s_t, a_t) \in B(\mu)$ and for any θ_t such that $\mu(\theta_t) > 0$, we have[12]

$$\mu(\theta_t) = \frac{v_t(\theta_t|h^t)\varphi'(s_t, a_t|\theta_t)}{\Pr_{t+1}(h^t, z_{(s_t,a_t)}(\mathbf{M}_t)|h^t)} = \frac{v_t(\theta_t|h^t)\sum_{s'_t \in B_{a_t}(\mu)} \varphi'(s'_t, a_t|\theta_t)}{\Pr_{t+1}(B_{a_t}(\mu)|h^t)} = \frac{v_t(\theta_t|h^t)\sum_{(s'_t,a'_t)\in B(\mu)} \varphi'(s'_t, a'_t|\theta_t)}{\Pr_{t+1}(B(\mu)|h^t)}. \tag{B.3}$$

That is, the principal updates to μ either when (i) he observes $(s_t, a_t) \in B(\mu)$, (ii) he learns that a_t is the realized allocation, that is, he learns that $s_t \in B_{a_t}(\mu)$, or (iii) he learns that the output message and the allocation belong to $B(\mu)$. Thus, for all θ_t in the support of μ, we have

$$\begin{aligned} \frac{\Pr_{t+1}(B_{a_t}(\mu)|h^t)}{\Pr_{t+1}(B(\mu)|h^t)} &= \frac{\sum_{s'_t\in B_{a_t}(\mu)} \varphi'(s'_t, a_t|\theta_t)}{\sum_{(s'_t,a'_t)\in B(\mu)} \varphi'(s'_t, a'_t|\theta_t)} \\ \frac{\Pr_{t+1}(h^t, z_{(s_t,a_t)}(\mathbf{M}_t)|h^t)}{\Pr_{t+1}(B_{a_t}(\mu)|h^t)} &= \frac{\varphi'(s_t, a_t|\theta_t)}{\sum_{s'_t\in B_{a_t}(\mu)} \varphi'(s'_t, a_t|\theta_t)}, \end{aligned} \tag{B.4}$$

where each of the equalities follows from applying Eq. (B.3). Eq. (B.4) shows (i) the probability that the allocation is a_t conditional on the induced belief being μ is independent of θ_t, and (ii) the probability that the output message is $s_t \in B_{a_t}(\mu)$ conditional on the allocation being a_t and the induced belief μ is independent of θ_t. It follows that for all $\mu \in \Delta(\Theta)^+$, $(s_t, a_t) \in B(\mu)$ and for all θ_t in the support of μ, we can *split* the auxiliary mapping as follows:

$$\begin{aligned} \varphi'(s_t, a_t|\theta_t) &= \left(\sum_{(s'_t,a'_t)\in B(\mu)} \varphi'(s'_t, a'_t|\theta_t)\right) \frac{\Pr_{t+1}(B_{a_t}(\mu)|h^t)}{\Pr_{t+1}(B(\mu)|h^t)} \frac{\Pr_{t+1}(h^t, z_{(s_t,a_t)}(\mathbf{M}_t)|h^t)}{\Pr_{t+1}(B_{a_t}(\mu)|h^t)} \\ &= \frac{\mu(\theta_t)}{v_t(\theta_t|h^t)} \Pr_{t+1}(B(\mu)|h^t) \frac{\Pr_{t+1}(B_{a_t}(\mu)|h^t)}{\Pr_{t+1}(B(\mu)|h^t)} \frac{\Pr_{t+1}(h^t, z_{(s_t,a_t)}(\mathbf{M}_t)|h^t)}{\Pr_{t+1}(B_{a_t}(\mu)|h^t)}, \end{aligned}$$

where the last equality follows from the last equality in Eq. (B.3).

Thus, the agent's payoff at history h^t, when the principal offers mechanism $\mathbf{M}_t$ and her type is $\theta_t \in \Theta^+$, can be written as follows:

$$\begin{aligned} &\sum_{(s_t,a_t)\in S^{\mathbf{M}_t} A_\emptyset} \varphi'(s_t, a_t|\theta_t) \left[u_t(a_t, \theta_t) + \delta \mathbb{E}^{P^{\sigma|h^t_A, \theta_t, z_{(s_t,a_t)}(\mathbf{M}_t)}} \sum_{\tau\geq t+1} u_\tau(\cdot)\right] = \\ &\sum_{\mu\in\Delta(\Theta)} \frac{\mu(\theta_t)}{v_t(\theta_t|h^t)} \Pr_{t+1}(B(\mu)|h^t) \sum_{a_t\in A} \frac{\Pr_{t+1}(B_{a_t}(\mu)|h^t)}{\Pr_{t+1}(B(\mu)|h^t)} \sum_{s_t\in B_{a_t}(\mu)} \frac{\Pr_{t+1}(h^t, z_{(s_t,a_t)}(\mathbf{M}_t)|h^t)}{\Pr_{t+1}(B_{a_t}(\mu)|h^t)} [u_t(a_t, \theta_t) \\ &+\delta \mathbb{E}^{P^{\sigma|h^t_A, \theta_t, z_{(s_t,a_t)}(\mathbf{M}_t)}} \sum_{\tau\geq t+1} u_\tau(\cdot)]. \end{aligned} \tag{B.5}$$

The difference between the principal and the agent's payoff in Eqs. (B.2) and (B.5) is that the agent cares only about the distribution over (μ, s_t, a_t) *conditional* on θ_t, whereas the principal's payoff is expressed in terms of the unconditional distribution. For this reason the agent's payoff features the term $\mu(\theta_t)/v_t(\theta_t|h^t)$.

We now define the direct Blackwell mechanism $\mathbf{M}^C_t = (\Theta, \Delta(\Theta), \varphi^{\mathbf{M}^C_t})$: First, for $\theta_t \in \Theta^+$,

$$\varphi^{\mathbf{M}^C_t}(\mu, a_t|\theta_t) = \underbrace{\frac{\mu(\theta_t)}{v_t(\theta_t|h^t)} \Pr_{t+1}(B(\mu)|h^t)}_{\beta^{\mathbf{M}^C_t}(\mu|\theta_t)} \underbrace{\frac{\Pr_{t+1}(B_{a_t}(\mu)|h^t)}{\Pr_{t+1}(B(\mu)|h^t)}}_{\alpha^{\mathbf{M}^C_t}(a_t|\mu)},$$

where the decomposition in terms of $\beta^{\mathbf{M}^C_t}, \alpha^{\mathbf{M}^C_t}$ is well-defined because of the independence properties highlighted after Eq. (B.4). Second, if $\theta_t \notin \Theta^+$, let $\theta^*_t(\theta_t)$ denote a maximizer of

$$\begin{aligned} &\sum_{(\mu,a_t)\in\Delta(\Theta)\times A} \varphi^{\mathbf{M}^C_t}(\mu, a_t|\tilde{\theta}_t) \sum_{s_t\in B_{a_t}(\mu)} \frac{\Pr_{t+1}(h^t, z_{(s_t,a_t)}(\mathbf{M}_t)|h^t)}{\Pr_{t+1}(B_{a_t}(\mu)|h^t)} \mathbb{E}^{P^{\sigma|(h^t, \theta_t, z_{(s_t,a_t)}(\mathbf{M}_t))}} [u_t(a_t, \theta_t) \\ &+\delta \mathbb{E}^{P^{\sigma|h^t_A, \theta_t, z_{(s_t,a_t)}(\mathbf{M}_t)}} \sum_{\tau\geq t+1} u_\tau(\cdot)], \end{aligned} \tag{B.6}$$

12 Note that if $\mu \in \Delta(\Theta)^+$ and θ_t is such that $\mu(\theta_t) > 0$, then $\theta_t \in \Theta^+$.

where $\tilde{\theta}_t \in \Theta^+$. Let $\varphi^{\mathbf{M}_t^C}(\cdot|\theta_t) = \varphi^{\mathbf{M}_t^C}(\cdot|\theta_t^*(\theta_t))$. Change the principal's strategy at h^t so that he offers $\mathbf{M}_t^C$ instead of $\mathbf{M}_t$. Change the agent's strategy so that, conditional on participating, the agent truthfully reports her type, $r_t'(h^t, \theta_t, \mathbf{M}_t^C) = \delta_{\theta_t}$.

For $\mu \in \Delta(\Theta)^+$ and allocation a_t, enumerate the output messages in $B_{a_t}(\mu)$ as $s_t^1, \ldots, s_t^K$. (We omit the dependence of K on μ and a_t to simplify notation.) Define the sequence $\{\omega_k\}_{k=0}^K$ such that $\omega_0 = 0, \omega_K = 1$ and for $k = 1, \ldots, K-1$,

$$\omega_k - \omega_{k-1} = \frac{\Pr_{t+1}(h^t, z_{(s_t^k, a_t)}(\mathbf{M}_t)|h^t)}{\Pr_{t+1}(B_{a_t}(\mu)|h^t)}.$$

Modify the continuation strategies as follows: for $k=1,\ldots,K$ and $\omega \in [\omega_{k-1}, \omega_k]$, let $\sigma|_{(h^t, z_{(\mu,a_t)}(\mathbf{M}_t^C), \omega)}$ coincide with $\sigma|_{(h^t, z_{(s_t^k, a_t)}(\mathbf{M}_t), \frac{\omega-\omega_{k-1}}{\omega_k-\omega_{k-1}})}$. Note these strategies imply the principal and the agent's payoffs remain the same as in the original equilibrium whenever $\theta_t \in \Theta^+$. Furthermore, modify the continuation strategies so that $\sigma|_{(h^t, z_{(\emptyset, a^*)}(\mathbf{M}_t^C))} = \sigma|_{(h^t, z_{(\emptyset, a^*)}(\mathbf{M}_t))}$.

For θ_t in Θ^+, set $\pi_t'(h^t, \theta_t, \mathbf{M}_t^C) = 1$. For types not in Θ^+, use Eq. (B.6) to compute $\pi_t'(h^t, \theta_t, \mathbf{M}_t^C)$ accordingly. Conditional on participating, the agent can guarantee at most the payoff from imitating the strategy followed by θ_t' for some $\theta_t' \in \Theta^+$. This strategy was already feasible in the original PBE, so the agent has no new deviations. It follows that the new assessment is a PBE of the auxiliary game. □

B.1. Proof of Proposition B.1

Proof of Proposition B.1. Fix a PBE assessment and let h^t denote a public history such that there exists a mechanism $\mathbf{M}_t$, a type θ_t, and two private histories $h_A^t, \overline{h}_A^t \in H_A^t(h^t)$ such that $\sigma_{At}(h_A^t, \theta_t, \mathbf{M}_t) \neq \sigma_{At}(\overline{h}_A^t, \theta_t, \mathbf{M}_t)$.

We make the following observations about the agent's payoff at node (h_A^t, θ_t), which are also true about $(\overline{h}_A^t, \theta_t)$. Fix $n \geq t$. Note that for any node (h_A^n, θ_n) that weakly succeeds (h_A^t, θ_t), there exists an equivalent node $(\overline{h}_A^n, \theta_n)$ that weakly succeeds $(\overline{h}_A^t, \theta_t)$.[13] Let σ_A' denote the strategy that coincides with σ_A everywhere except that $\sigma_A'(h_A^n, \theta_n) = \sigma_A(\overline{h}_A^n, \theta_n)$ for each (h_A^n, θ_n) that weakly succeeds (h_A^t, θ_t). Then, it must be the case that $U_t(\sigma|h_A^t, \theta_t, \mathbf{M}_t) \geq U_t(\sigma_P, \sigma_A'|h_A^t, \theta_t, \mathbf{M}_t)$. Otherwise, the agent would have a deviation at (h_A^t, θ_t) (to σ_A'). Note that $U_t(\sigma_P, \sigma_A'|h_A^t, \theta_t, \mathbf{M}_t) = U_t(\sigma|\overline{h}_A^t, \theta_t, \mathbf{M}_t)$. Swapping the roles of $h_A^t, \overline{h}_A^t$, we conclude that

$$U_t(\sigma|h_A^t, \theta_t, \mathbf{M}_t) = U_t(\sigma_P, \sigma_A'|h_A^t, \theta_t, \mathbf{M}_t), \tag{B.7}$$

and a similar indifference holds for the agent at information set $(\overline{h}_A^t, \theta_t)$.

Consider now Eq. (A.7) at $t-1$ and $(h_A^{t-1}, \theta_{t-1})$ such that (h_A^t, θ_t) weakly succeeds $(h_A^{t-1}, \theta_{t-1})$. Note that Eq. (B.7) implies that $U_{t-1}(\sigma|h_A^{t-1}, \theta_{t-1}, \mathbf{M}_{t-1}) = U_{t-1}(\sigma_P, \sigma_A'|h_A^{t-1}, \theta_{t-1}, \mathbf{M}_{t-1})$, where σ_A' is the strategy we constructed before. Working backwards from $n = t-1$ to $n = 1$ one can see that if at node (h_A^t, θ_t) the agent is indifferent between σ_A' and σ_A, so is the agent at any (h_A^n, θ_n) that precedes (h_A^t, θ_t).

Finally, take now Eq. (A.7). The above argument implies that if at a node $(h_A^{t+1}, \theta_{t+1})$ that succeeds (h_A^t, θ_t), the agent is indifferent between two continuation strategies, then so is the agent at (h_A^t, θ_t). Working forward from $n = t+1$ to any $n = \tau > t+1$, one concludes that if the agent is indifferent between two (continuation) strategies at (h_A^τ, θ_τ), then so is the agent at (h_A^t, θ_t).

Summing up, we have established that (a) for the same public history h^t, if the continuation strategy at (h_A^t, θ_t) differs from the continuation strategy at $(\overline{h}_A^t, \theta_t)$, then the agent is indifferent between both continuation strategies (and any mixture between them), (b) this indifference carries through to the nodes that precede both (h_A^t, θ_t) and $(\overline{h}_A^t, \theta_t)$, and (c) any indifference at nodes that succeed (h_A^t, θ_t) and $(\overline{h}_A^t, \theta_t)$ also holds at these nodes (and hence its predecessors by (b)).

Consider the following strategy which, by the above arguments, is payoff equivalent to σ_A. For $\tau \geq t$, let h^τ denote a strategy on the path of h^t and let $\mathbf{M}_\tau$ denote a mechanism chosen by the principal at h^τ. For all $h_A^\tau \in H_A^\tau(h^\tau)$ and all $\theta_\tau \in \Theta_\tau$, the following is also an optimal strategy at $(h_A^\tau, \theta_\tau, \mathbf{M}_\tau)$

$$\pi_\tau'(h_A^\tau, \theta_\tau, \mathbf{M}_\tau) = \sum_{\overline{h}_A^\tau \in H_A^\tau(h^\tau)} \frac{\nu_\tau(\overline{h}_A^\tau, \theta_\tau|h^\tau)}{\sum_{\tilde{h}_A^\tau \in H_A^\tau(h^\tau)} \nu_\tau(\tilde{h}_A^\tau, \theta_\tau|h^\tau)} \pi_\tau(\overline{h}_A^\tau, \theta_\tau, \mathbf{M}_\tau),$$

whenever $\sum_{\tilde{h}_A^\tau \in H_A^\tau(h^\tau)} \nu_\tau(\tilde{h}_A^\tau, \theta_\tau|h^\tau) > 0$ and

$$r_\tau'(h_A^\tau, \theta_\tau, \mathbf{M}_\tau) = \sum_{\overline{h}_A^\tau \in H_A^\tau(h^\tau)} \frac{\nu_\tau(\overline{h}_A^\tau, \theta_\tau, \mathbf{M}_\tau)\pi_\tau(\overline{h}_A^\tau, \theta_\tau, \mathbf{M}_\tau)}{\sum_{\tilde{h}_A^\tau \in H_A^\tau(h^\tau)} \nu_\tau(\tilde{h}_A^\tau, \theta_\tau|h^\tau)\pi_\tau(\tilde{h}_A^\tau, \theta_\tau, \mathbf{M}_\tau)} r_\tau(\overline{h}_A^\tau, \theta_\tau, \mathbf{M}_\tau),$$

whenever $\sum_{\tilde{h}_A^\tau \in H_A^\tau(h^\tau)} \nu_\tau(\tilde{h}_A^\tau, \theta_\tau|h^\tau)\pi_\tau(\tilde{h}_A^\tau, \theta_\tau, \mathbf{M}_\tau) > 0$. Let $\tau \geq t$ and fix h^τ on the path of h^t. Consider now types θ_τ such that $\sum_{h_A^\tau \in H_A^\tau(h^\tau)} \nu_\tau(h_A^\tau, \theta_\tau|h^\tau) = 0$. Then, for any $h_A^\tau \in H_A^\tau(h^\tau)$, they participate with probability

$$\pi_\tau'(h_A^\tau, \theta_\tau, \mathbf{M}_\tau) = \sum_{\overline{h}_A^t \in H_A^\tau(h^\tau)} \frac{\pi_\tau(\overline{h}_A^\tau, \theta_\tau, \mathbf{M}_\tau)}{|H_A^\tau(h^\tau)|}.$$

[13] Formally, if $n \geq t+1$ and (h_A^n, θ_n) succeeds (h_A^t, θ_t), then $h_A^n = (h_A^t, \theta_t, z_\cdot(\mathbf{M}_t), \ldots, z_\cdot(\mathbf{M}_n))$. Then, $\overline{h}_A^n = (\overline{h}_A^t, \theta_t, z_\cdot(\mathbf{M}_t), \ldots, z_\cdot(\mathbf{M}_n))$. For $n = t$, $(h_A^n, \theta_n) = (h_A^t, \theta_t)$, so that $(\overline{h}_A^n, \theta_n) = (\overline{h}_A^t, \theta_t)$.

If public history $(h^\tau, \mathbf{M}_\tau, 1)$ in on the path of h^t, then modify the agent's reporting strategy so that she reports according to

$$r'_\tau(h^\tau_A, \theta_\tau, \mathbf{M}_\tau) = \sum_{\overline{h}^\tau_A \in H^\tau_A(h^\tau)} \frac{\pi_\tau(\overline{h}^\tau_A, \theta_\tau, \mathbf{M}_\tau)}{\sum_{\tilde{h}^\tau_A \in H^\tau_A(h^\tau)} \pi_\tau(\tilde{h}^\tau_A, \theta_\tau, \mathbf{M}_\tau)} r_\tau(\overline{h}^\tau_A, \theta_\tau, \mathbf{M}_\tau),$$

whenever $\sum_{\tilde{h}^\tau_A \in H^\tau_A(h^\tau)} \pi_\tau(\tilde{h}^\tau_A, \theta_\tau, \mathbf{M}_\tau) > 0$ and with probability

$$r'_\tau(h^\tau_A, \theta_\tau, \mathbf{M}_\tau) = \sum_{\overline{h}^t_A \in H^\tau_A(h^\tau)} \frac{r_\tau(\overline{h}^\tau_A, \theta_\tau, \mathbf{M}_\tau)}{|H^\tau_A(h^\tau)|},$$

otherwise. Fix $\tau \geq t$ and a public history h^τ on the path of the strategy (σ_P, σ'_A) starting from h^t. Under the new strategy, Bayes' rule implies that beliefs at the end of period τ, when the public history is $(h^\tau, z_{(s_\tau,a_\tau)}(\mathbf{M}_\tau))$ that the agent's private history is $(h^\tau_A, \theta_\tau, m, z_{(s_\tau,a_\tau)}(\mathbf{M}_\tau))$ are given by

$$\mu'_{\tau+1}(h^\tau_A, \theta_\tau, m, z_{(s_\tau,a_\tau)}(\mathbf{M}_\tau))|(h^\tau, z_{(s_\tau,a_\tau)}(\mathbf{M}_\tau)) = \frac{v'_\tau(h^\tau_A, \theta_\tau|h^\tau)\kappa^{\sigma'_A}_\tau(m_\tau, s_\tau, a_\tau|h^\tau_A, \theta_\tau, \mathbf{M}_\tau)}{\sum_{(\overline{h}^\tau_A, \overline{\theta}_\tau, \overline{m}_\tau)} v'_\tau(\overline{h}^\tau_A, \overline{\theta}_\tau|h^\tau)\rho^{\sigma'_A}(\overline{m}_\tau, s_\tau, a_\tau|\overline{h}^\tau_A, \overline{\theta}_\tau, \mathbf{M}_\tau)}.$$

We show recursively that for each $\tau \geq t$ and each θ_τ, the following holds

$$\sum_{\overline{h}^\tau_A, \overline{m}_\tau} \mu'_{\tau+1}(\overline{h}^\tau_A, \theta_\tau, \overline{m}_\tau, z_{(s_\tau,a_\tau)}(\mathbf{M}_\tau)|h^\tau, z_{(s_\tau,a_\tau)}(\mathbf{M}_\tau)) = \sum_{\overline{h}^\tau_A, \overline{m}_\tau} \mu_{\tau+1}(\overline{h}^\tau_A, \theta_\tau, \overline{m}_\tau, z_{(s_\tau,a_\tau)}(\mathbf{M}_\tau)|h^\tau, z_{(s_\tau,a_\tau)}(\mathbf{M}_\tau)). \tag{B.8}$$

A corollary of this is that

$$\sum_{\overline{h}^\tau_A, \overline{m}_\tau} v'_{\tau+1}(\overline{h}^\tau_A, \theta_\tau, \overline{m}_\tau, z_{(s_\tau,a_\tau)}(\mathbf{M}_\tau)|h^\tau, z_{(s_\tau,a_\tau)}(\mathbf{M}_\tau)) = \sum_{\overline{h}^\tau_A, \overline{m}_\tau} v_{\tau+1}(\overline{h}^\tau_A, \theta_\tau, \overline{m}_\tau, z_{(s_\tau,a_\tau)}(\mathbf{M}_\tau)|h^\tau, z_{(s_\tau,a_\tau)}(\mathbf{M}_\tau)). \tag{B.9}$$

To see that Eq. (B.8) implies Eq. (B.9), notice that

$$v'_{\tau+1}(h^{\tau+1}_A, \theta_{\tau+1}|h^{\tau+1}) = \mu'_{\tau+1}(h^{\tau+1}_A|h^{\tau+1})F_{\tau+1}(\theta_{\tau+1}|\theta_\tau(h^{\tau+1}_A), a(h^{\tau+1})),$$

where $\theta_\tau(h^{\tau+1}_A)$ is the agent's private type at the end of period τ, when the history is $h^{\tau+1}_A$, while $a(h^{\tau+1})$ is the allocation at the end of period τ, when the public history is $h^{\tau+1}$. Then, Eq. (B.8) implies that

$$\begin{aligned}
\sum_{h^\tau_A} v'_{\tau+1}(h^{\tau+1}_A, \theta_{\tau+1}|h^{\tau+1}) &= \sum_{h^\tau_A} \mu'_{\tau+1}(h^{\tau+1}_A|h^{\tau+1})F_{\tau+1}(\theta_{\tau+1}|\theta_\tau(h^{\tau+1}_A), a^*) \\
&= \sum_{h^\tau_A} \mu'_{\tau+1}(h^\tau_A, \theta_\tau(h^{\tau+1}_A), z_{(s_t,a_t)}(\mathbf{M}_t)|h^{\tau+1})F_{\tau+1}(\theta_{\tau+1}|\theta_\tau(h^{\tau+1}_A), a(h^{\tau+1})) \\
&= \sum_{h^\tau_A} \mu_{\tau+1}(h^\tau_A, \theta_\tau(h^{\tau+1}_A), z_{(s_t,a_t)}(\mathbf{M}_t)|h^{\tau+1})F_{\tau+1}(\theta_{\tau+1}|\theta_\tau(h^{\tau+1}_A), a(h^{\tau+1})) \\
&= \sum_{h^\tau_A} v_{\tau+1}(h^\tau_A, \theta_\tau(h^{\tau+1}_A), z_{(s_t,a_t)}(\mathbf{M}_t)|h^{\tau+1}).
\end{aligned}$$

To show Eq. (B.8) holds, we proceed by induction. Suppose that we have established that Eq. (B.8) holds for $t \leq n < \tau$. We now show it holds for $n = \tau$. To see this, consider $h^{\tau+1}_A = (h^\tau_A, \theta_\tau, m, z_{(s_\tau,a_\tau)}(\mathbf{M}_\tau))$ and note that the denominator on the RHS in the first line of Eq. (B.8) is simply $\Pr^{v',\sigma'_A}_{\tau+1}(h^\tau, z_{(s_\tau,a_\tau)}(\mathbf{M}_\tau)|h^\tau)$. Moreover,

$$\begin{aligned}
\Pr^{v',\sigma'_A}_{\tau+1}(h^\tau, z_{(s_\tau,a_\tau)}(\mathbf{M}_\tau)|h^\tau) &= \sum_{(\overline{h}^\tau_A, \overline{\theta}_\tau, \overline{m}_\tau)} v'_\tau(\overline{h}^\tau_A, \overline{\theta}_\tau|h^\tau)\rho^{\sigma'_A}(\overline{m}_\tau, s_\tau, a_\tau|\overline{h}^\tau_A, \overline{\theta}_\tau, \mathbf{M}_\tau) \\
&= \sum_{(\overline{\theta}_\tau, \overline{m}_\tau)} \kappa^{\sigma'_A}_\tau(\overline{m}_\tau, s_\tau, a_\tau|\overline{h}^\tau_A, \overline{\theta}_\tau, \mathbf{M}_\tau) \sum_{\overline{h}^\tau_A} v'_\tau(\overline{h}^\tau_A, \overline{\theta}_\tau|h^\tau) \\
&= \sum_{(\overline{\theta}_\tau, \overline{m}_\tau)} \kappa^{\sigma'_A}_\tau(\overline{m}_\tau, s_\tau, a_\tau|\overline{h}^\tau_A, \overline{\theta}_\tau, \mathbf{M}_\tau) \sum_{\overline{h}^\tau_A} v_\tau(\overline{h}^\tau_A, \overline{\theta}_\tau|h^\tau) = \Pr^{v,\sigma_A}_{\tau+1}(h^\tau, z_{(s_\tau,a_\tau)}(\mathbf{M}_\tau)|h^\tau).
\end{aligned} \tag{B.10}$$

The last equality follows from the definition of σ'_A. The key equality is the third one which employs the inductive hypothesis.[14] To see that it holds, consider the case in which $h^\tau = (h^{\tau-1}, z_{s_{\tau-1},a_{\tau-1}}(\mathbf{M}_{\tau-1}))$. Then

$$\sum_{\overline{h}^\tau_A \in H^\tau_A(h^\tau)} v'_\tau(\overline{h}^\tau_A, \overline{\theta}_\tau|h^\tau) = \sum_{(\overline{h}^{\tau-1}_A, \overline{\theta}_{\tau-1}, \overline{m}_{\tau-1})} v'_\tau(\overline{h}^{\tau-1}_A, \overline{\theta}_{\tau-1}, \overline{m}_{\tau-1}, z_{(s_{\tau-1},a_{\tau-1})}(\mathbf{M}_{\tau-1}), \overline{\theta}_\tau|h^\tau)$$

[14] Note that it trivially holds for $\tau = t$.

$$= \sum_{\overline{\theta}_{\tau-1}} \sum_{\overline{h}_A^{\tau-1}, \overline{m}_{\tau-1}} \nu'_\tau(\overline{h}_A^{\tau-1}, \overline{\theta}_{\tau-1}, \overline{m}_{\tau-1}, z_{(s_{\tau-1},a_{\tau-1})}(\mathbf{M}_{\tau-1}), \overline{\theta}_\tau | h^\tau)$$
$$= \sum_{\overline{\theta}_{\tau-1}} \sum_{\overline{h}_A^{\tau-1}, \overline{m}_{\tau-1}} \mu'_\tau(\overline{h}_A^{\tau-1}, \overline{\theta}_{\tau-1}, \overline{m}_{\tau-1}, z_{(s_{\tau-1},a_{\tau-1})}(\mathbf{M}_{\tau-1}) | h^\tau) F_\tau(\overline{\theta}_\tau | \overline{\theta}_{\tau-1}, a_{\tau-1})$$
$$= \sum_{\overline{\theta}_{\tau-1}} \sum_{\overline{h}_A^{\tau-1}, \overline{m}_{\tau-1}} \mu_\tau(\overline{h}_A^{\tau-1}, \overline{\theta}_{\tau-1}, \overline{m}_{\tau-1}, z_{(s_{\tau-1},a_{\tau-1})}(\mathbf{M}_{\tau-1}) | h^\tau) F_\tau(\overline{\theta}_\tau | \overline{\theta}_{\tau-1}, a_{\tau-1})$$
$$= \sum_{\overline{\theta}_{\tau-1}} \sum_{\overline{h}_A^{\tau-1}, \overline{m}_{\tau-1}} \nu_\tau(\overline{h}_A^{\tau-1}, \overline{\theta}_{\tau-1}, \overline{m}_{\tau-1}, z_{(s_{\tau-1},a_{\tau-1})}(\mathbf{M}_{\tau-1}), \theta_\tau | h^\tau) = \sum_{\overline{h}_A^\tau} \nu_\tau(\overline{h}_A^\tau, \overline{\theta}_\tau | h^\tau),$$

where the first three equalities are definitional, the fourth uses the inductive hypothesis that Eq. (B.8) holds at $\tau - 1$ and the rest are definitional.

Consider now the principal's payoff in Eq. (A.6). Working recursively using the identities established in Eqs. (B.8)–(B.10), it follows that the principal's payoff from σ_P remains the same under the new assessment. Furthermore, Eq. (B.9) implies that σ_P is still sequentially rational. Thus, $(\sigma_P, \sigma'_A, \mu')$ constitute a payoff equivalent PBE assessment. □

Appendix C. Induced distributions

Appendix C formally defines the distributions over terminal nodes induced by the principal's and the agent's strategy. The definitions presented herein apply when Θ is a continuum or the principal's mechanism does not have finite support. The defined distributions together with the steps in Appendix B allow us to extend the proof in Appendix B to the case of a continuum of types and without any restrictions in the support of the principal's mechanism, along the lines of that in Doval and Skreta (2022).

Given a strategy profile $\sigma = (\sigma_P, \sigma_A)$, the transition probabilities F_{t+1}, and a node (θ, h_A^t), we define transition probabilities from H_A^t to Θ, from $H_A^t \times \Theta$ to $\mathcal{M}_\mathcal{I}$, from $H_A^t \times \Theta \times \mathcal{M}_\mathcal{I}$ to $\cup_{i,j\in\mathcal{I}} M_i S_j A_\emptyset$ and from $H_A^t \times \Theta \times \mathcal{M}_\mathcal{I} \times \cup_{i,j\in\mathcal{I}} M_i S_j A_\emptyset$ to Ω as follows:

$$\begin{aligned}
\kappa_t^F(\widetilde{\Theta}|h_A^t) &= F_t(\widetilde{\Theta}|\theta_{t-1}(h_A^t), a_{t-1}(h_A^t)) \\
\kappa_t^{\sigma_P}(\cup_{i,j\in\mathcal{I}}\widetilde{\mathcal{M}}_{i,j}|h_A^t, \theta_t) &= \sum_{i,j\in\mathcal{I}} \sigma_{Pt}(\widetilde{\mathcal{M}}_{i,j}|h^t) \\
\kappa_t^{\sigma_A}(\widetilde{M_i S_j A_\emptyset}|h_A^t, \theta_t, \mathbf{M}_t) &= \left(1 - \pi_t(h_A^t, \theta_t, \mathbf{M}_t)\right) \mathbb{1}[(\emptyset, \emptyset, a^*) \in \widetilde{M_i S_j A_\emptyset}] \\
&\quad + \pi_t(h_A^t, \theta_t, \mathbf{M}_t) \int_{\widetilde{M_i S_j A_\emptyset}} r_t(h_A^t, \theta_t, \mathbf{M}_t) \otimes \varphi^{\mathbf{M}_t}(d(m_t, s_t, a_t)) \\
\kappa_{t+1}^\omega(\widetilde{\Omega}|h_A^t, \theta_t, \mathbf{M}_t, m_t, s_t, a_t) &= \int_{\widetilde{\Omega}} l(d\omega_{t+1})
\end{aligned} \tag{C.1}$$

where (i) $(\theta_{t-1}(h_A^t), a_{t-1}(h_A^t))$ denote the period-$t-1$ type and allocation consistent with h_A^t, (ii) h^t denotes the projection of (θ, h_A^t) onto $\Omega \times Z^{t-1}$, and (iii) the notation presumes that $\mathbf{M}_t \in \mathcal{M}_{i,j}$. Note that $\kappa_t^\sigma \equiv \kappa_t^F \otimes \kappa_t^{\sigma_P} \otimes \kappa_t^{\sigma_A} \otimes \kappa_{t+1}^\omega$ defines a transition probability from H_A^t to Z_A.[15]

The Ionescu-Tulcea extension theorem (Pollard, 2002) guarantees the existence of a sequence of probability measures $P_t^\sigma = \kappa_1^\omega \otimes \bigotimes_{\tau=1}^{t-1} \kappa_\tau^\sigma$ defined on the product sets $\left(H_A^t\right)_{t=1}^{T+1}$ and a probability measure P^σ on $(H_A^{T+1}, \mathcal{B}_{H_A^{T+1}})$ such that for each $t \geq 1$, the marginal of P^σ on H_A^t is P_t^σ.

Note that P^σ determines a distribution over H_A^{T+1}. The principal and the agent's payoffs, however, are defined over $(\Theta \times A)^T$. Definition C.1 formally defines the distribution on $(\Theta \times A)^T$ induced by P^σ.

Definition C.1. Fix an assessment $(\sigma_P, \sigma_A, \mu)$. The distribution $\eta^\sigma \in \Delta\left((\Theta \times A)^T\right)$ induced by the assessment is defined as follows:

$$\eta^\sigma(\widetilde{\Theta^T} \times \widetilde{A^T}) = \int_{H_A^{T+1}} \mathbb{1}[\mathrm{proj}_{(\Theta\times A)^T}(\theta, h_A^{T+1}) \in \widetilde{\Theta^T} \times \widetilde{A^T}] P^\sigma(dh_A^{T+1}),$$

for any measurable subsets $\widetilde{\Theta^T}$ of Θ^T and $\widetilde{A^T}$ of A^T.

Thus, the principal's payoff under assessment $(\sigma_P, \sigma_A, \mu)$, $W(\sigma, \mu)$, is given by

$$\int_{H_A^{T+1}} W(\mathrm{proj}_{(\Theta\times A)^T}(h_A^{T+1})) P^\sigma(dh_A^{T+1}) = \int_{(\Theta\times A)^T} W(a^T, \theta^T) \eta^\sigma(d(\theta^T, a^T)), \tag{C.2}$$

[15] Given two Polish spaces, X and Y, a transition probability from X to Y is a measurable map $\kappa : X \mapsto \Delta(Y)$. If κ is a transition probability from X to Y, then we denote by $\kappa(\cdot|x)$ the measure on Y induced by κ evaluated at x. If κ is a transition probability from X to Y and κ' is a transition probability from Y to Z, then their composition $\kappa \otimes \kappa'$ is the transition probability from X to $Y \times Z$ such that

$$\left(\kappa \otimes \kappa'\right)(\widetilde{Y} \times \widetilde{Z}|x) = \int_{\widetilde{Y}} \kappa'(\widetilde{Z}|y)\kappa(dy|x).$$

while the agent's payoff when her type is θ, $U(\sigma,\mu,\theta_1)$, is given by

$$\int_{H_A^{T+1}} U(\text{proj}_{(\Theta\times A)^T} h_A^{T+1}) P^{\sigma|\theta_1}(dh_A^{T+1}) = \int_{(\Theta\times A)^T} U(a^T,\theta'^T)\eta^\sigma_{\theta_1}(d(\theta'^T,a^T)), \tag{C.3}$$

where (i) $P^{\sigma|\theta_1}$ is the induced probability over H_A^{T+1} conditional on the agent's period-1 type being θ_1, and (ii) $\eta^\sigma_{\theta_1}$ is the proj_Θ-disintegration of η^σ.

Appendix D. Continuum type spaces

Appendix D extends Theorem 1 to the case in which Θ is a continuum. As we explain in Doval and Skreta (2022), when Θ is a continuum, the issues pointed out in Aumann (1961, 1963, 1964) prevent us from relying on the game-theoretic formulation of the mechanism-selection game to describe the set of outcomes, and hence, payoffs that can be achieved under limited commitment. In Doval and Skreta (2022), we overcome these challenges by providing a recursive representation of the set of *PBE-feasible* outcomes, that is, those outcome distributions that can be achieved under limited commitment. We extend the solution concept in Doval and Skreta (2022) to the Markovian environments and provide a recursive formulation of the set of PBE-feasible *payoffs*.

(Continuation) payoffs and beliefs. Because the definition is recursive, it is useful to have explicit notation for (i) the principal and the agent's (continuation) payoffs starting at any period t, and (ii) belief transitions from period t to period $t+1$. To that end, we define two functions, $U_t, W_t : (A\times\Theta)^{T-t} \mapsto \mathbb{R}$, that describe the principal and agent's payoffs from the beginning of period t. Formally,

$$U_t\left((a_\tau,\theta_\tau)_{\tau\geq t}\right) = \sum_{\tau=t}^{T}\delta^{\tau-t}u_\tau(a_\tau,\theta_\tau), W_t\left((a_\tau,\theta_\tau)_{\tau\geq t}\right) = \sum_{\tau=t}^{T}\delta^{\tau-t}w_\tau(a_\tau,\theta_\tau). \tag{D.1}$$

Given a distribution over (continuation) outcomes, $\eta\in\Delta\left((\Theta\times A)^{T-t}\right)$, we can use the functions (W_t,U_t) to define the principal and the agent's expected payoffs under η, $\mathbf{V}_t(\eta) = (\mathbf{W}_t(\eta),\mathbf{U}_t(\eta,\cdot))$. Formally,

$$\mathbf{W}_t(\eta) = \int_{(\Theta\times A)^{T-t}} W_t d\eta, \quad \mathbf{U}_t(\eta,\theta_t) = \int_{(\Theta\times A)^{T-t}} U_t d\eta_{\theta_t}, \tag{D.2}$$

where $\{\eta_{\theta_t} : \theta_t\in\Theta\}$ is the proj_Θ-disintegration of η (Kallenberg, 2017).

Finally, given a distribution $\mu_{t+1}\in\Delta(\Theta)$ and a period-t allocation, a_t, define the linear map $f_{t+1,a_t}(\mu_t)\in\Delta(\Theta)$ as follows: for all measurable subsets $\widetilde{\Theta}$ of Θ,

$$f_{t+1,a_t}(\mu_t)(\widetilde{\Theta}) = \int_\Theta F_{t+1}(\widetilde{\Theta}|\theta_t,a_t)\mu_{t+1}(d\theta_t). \tag{D.3}$$

That is, $f_{t+1,a_t} : \Delta(\Theta)\mapsto\Delta(\Theta)$ describes the principal's beliefs about the agent's type in period $t+1$, when his beliefs at the end of period t are μ_{t+1} and the allocation is a_t.

PBE-feasible payoffs. We now formally define the set of PBE-feasible payoffs (Definition D.5). By analogy with the mechanism-selection game, we keep the notation $\mathcal{E}^*_{\mathcal{I}}$ to denote PBE-feasible payoffs. Contrary to the mechanism-selection game, $\mathcal{E}^*_{\mathcal{I},t}$ is now a collection of correspondences describing the set of PBE-feasible payoffs for each period $t\geq 1$, each principal's belief v_t and each sequence of allocations up to period t, a^{t-1}, $\mathcal{E}^*_{\mathcal{I},t}(v_t,a^{t-1})$. As mentioned above, the set of PBE-feasible payoffs is recursive: what is PBE-feasible today depends on what is PBE-feasible from tomorrow onwards. Thus, we fix a period $t\geq 1$, and a pair (μ_t,a^{t-1}) throughout.

Definition D.5 consists of three components, which we introduce first: (i) the sequence of mechanisms offered by the principal (Definition D.1), (ii) optimal behavior by the agent within those mechanisms, and (iii) the payoffs the principal anticipates upon a deviation (Definition D.3). For simplicity, we assume the principal has one set of input and output messages, M and S, and we use the shorthand notation $SA_\emptyset$ to denote the set $(S\times A)\cup\{(\emptyset,a^*)\}$.

Dynamic mechanisms. We describe the analogue of the principal's strategy via a *dynamic mechanism*, defined as follows:

Definition D.1 (*Dynamic Mechanisms*). For $t\geq 1$ and $a^{t-1}\in A^{t-1}$, a dynamic mechanism given a^{t-1}, $(\varphi_\tau)_{\tau\geq t}$, is a sequence of measurable mappings[16] $\varphi_\tau : (SA_\emptyset\times\Omega)^{\tau-t}\times M\mapsto\Delta(S\times A)$, such that for all $\tau\geq t$ and all $(s^{\tau-t},a^{\tau-t},\omega^{\tau-t})$,

1. $\varphi_\tau(s^{\tau-t},a^{\tau-t},\omega^{\tau-t},\cdot) : M\mapsto\Delta(S\times A)$ is a measurable function, and
2. for all $m\in M$, $\varphi_\tau(s^{\tau-t},a^{\tau-t},\omega^{\tau-t})\left(\mathcal{A}_t(a^{t-1},a^{\tau-t})|m\right) = 1$.

When $t=1$, a dynamic mechanism describes the mechanism the agent faces in period 1, φ_1, the mechanism the agent faces in period 2 as a function of the agent's participation decision (i.e., whether $(s_1,a_1)\neq(\emptyset,a^*)$), and the realization of the public randomization device, $\varphi_2(s_1,a_1,\omega_2)$, and so on. Consider now $t>1$ and suppose the allocation so far is a^{t-1}. Then, we require that only allocations that are feasible given a^{t-1} are in the support of the mechanism φ_t.

As we explain next, a dynamic mechanism defines an extensive-form game for the agent:

[16] Below, Ω denotes the set of possible realizations of the public randomization device.

Agent-extensive form. Given (ν_t, a^{t-1}), a dynamic mechanism $(\varphi_\tau)_{\tau\geq t}$ defines an extensive-form game for the agent, $\Gamma(\nu_t,a^{t-1},(\varphi_\tau)_{\tau\geq t})$, as follows. First, nature draws the agent's type according to ν_t. Having observed her type, suppose that in stage $\tau-t$, the public history is $h_t^\tau=(s^{\tau-t},a^{\tau-t},\omega^{\tau-t})$. Then, faced with $\varphi_\tau(h_t^\tau)$, the agent decides whether to participate and, conditional on participating, her reporting strategy. If the agent rejects φ_τ, the "output message" is $\emptyset$ and the allocation is a^*. Instead, if she accepts $\varphi_\tau(h_t^\tau)$, she chooses an input message $m\in M$ that determines the distribution from which the output message and the allocation are drawn, $\varphi_\tau(h_t^\tau)(\cdot|m)$. In both cases, we proceed to stage $\tau+1-t$, where nature draws the agent's type using the appropriate transition $F_{\tau+1}(\cdot|\cdot)$.

In the agent-extensive form $\Gamma(\nu_t,a^{t-1},(\varphi_\tau)_{\tau\geq t})$, there are two types of histories. The public history h_t^τ encodes the agent's participation in the mechanism, the realized output messages and allocations, and the realizations of the public randomization device. The private histories encode everything the agent knows: her payoff-relevant type θ_t, the public history h_t^τ, and her past types and input messages. In a slight abuse of notation, we denote by $H_{At}^\tau(h_t^\tau)\times\Theta$ the agent's private histories consistent with h_t^τ.

Finally, we note that the agent evaluates the payoffs of a strategy in $\Gamma(\nu_t,a^{t-1},(\varphi_\tau)_{\tau\geq t})$ using U_t. We are now ready to define optimal play by the agent in $\Gamma(\nu_t,a^{t-1},(\varphi_\tau)_{\tau\geq t})$:

Agent-PBE. Together with the agent strategy σ_A, we can also define a system of beliefs $\mu\equiv(\mu_{\tau+1})_{\tau\geq t}^T$, which describes for each period $\tau\geq t$ and for each public history $h_t^{\tau+1}$, the principal's beliefs over the private histories, $\mu_{\tau+1}(h_t^{\tau+1})\in\Delta\left(H_{At}^{\tau+1}(h_t^{\tau+1})\right)$.

We say that (σ_A,μ) is an agent-PBE of the agent-extensive form $\Gamma(\nu_t,a^{t-1},(\varphi_\tau)_{\tau\geq t})$ if the agent's strategy is sequentially rational (under payoffs $U_t(\cdot)$) and the belief system satisfies Bayes' rule where possible. Although the belief system is not needed to test whether the agent's strategy is optimal in the extensive-form game, it is needed to test the optimality of the principal's choice of mechanism.

Proposition B.1 applies verbatim, allowing us to conclude that for every agent-PBE (σ_A,μ) of $\Gamma(\nu_t,a^{t-1},(\varphi_\tau)_{\tau\geq t})$, a payoff-equivalent (σ_A',μ') exists, in which the agent's strategy only conditions on her type and the public history. This property is responsible for the recursive nature of the set of PBE-feasible payoffs here and also in the mechanism-selection game. Hereafter, when we say agent-PBE, we mean one that satisfies the above property.

(Continuation) payoffs. An agent-PBE (σ_A,μ) of $\Gamma(\nu_t,a^{t-1},(\varphi_\tau)_{\tau\geq t})$ defines a distribution over $(\Theta\times A)^{T-t}$, $\eta^{(\varphi_\tau)_{\tau\geq t},\sigma_A}$, that satisfies that the marginal on Θ is ν_t.

Definition D.2 (*Induced Payoffs*)**.** The payoff vector $\left(w_t,(u_{\theta_t})_{\theta_t\in\Theta}\right)$ is implemented by $((\varphi_\tau)_{\tau\geq t},(\sigma_A,\mu))$ at ν_t if the following hold:

1. $\left(w_t,(u_{\theta_t})_{\theta_t\in\Theta}\right)$ are the payoffs the principal and the agent obtain under $\eta^{(\varphi_\tau)_{\tau\geq t},\sigma_A}$, that is, $\left(w_t,(u_{\theta_t})_{\theta_t\in\Theta}\right)=\mathbf{V}_t(\eta^{(\varphi_\tau)_{\tau\geq t},\sigma_A})$.
2. (σ_A,μ) is an agent-PBE of $\Gamma(\nu_t,a^{t-1},(\varphi_\tau)_{\tau\geq t})$.

Furthermore, at any $\tau\geq t$ and at any history $h_t^{\tau+1}$, the belief assessment together with the dynamic mechanisms and the agent's strategy, defines a *continuation* outcome, $\eta^{(\varphi_\tau)_{\tau\geq t},\sigma_A|h_t^{\tau+1}}\in\Delta\left((\Theta\times A)^{T-(\tau+1)}\right)$, whose marginal on Θ coincides with $\nu_{\tau+1}(h_t^{\tau+1})\equiv f_{\tau+1,a_\tau(h_t^{\tau+1})}(\mu_{\tau+1}(h_t^{\tau+1}))$. We can similarly define a vector of continuation payoffs $\mathbf{V}_{\tau+1}(\eta^{(\varphi_\tau)_{\tau\geq t},\sigma_A|h_t^{\tau+1}})$.

Principal's sequential rationality. Fix a dynamic mechanism given a^{t-1}, $(\varphi_\tau)_{\tau\geq t}$, and an agent-PBE (σ_A,μ) of $\Gamma(\nu_t,a^{t-1},(\varphi_\tau)_{\tau\geq t})$. Suppose that the principal considers offering mechanism φ_t' instead of φ_t. In order to determine whether the principal wishes to deviate to φ_t', we need to determine the payoffs that can follow φ_t'. We denote this set by $\mathsf{D}_{\mathcal{E}_\mathcal{I}^*}(\nu_t,a^{t-1},\varphi_t')$ and is defined as follows:

Definition D.3 (*Deviant Payoffs*)**.** The set $\mathsf{D}_{\mathcal{E}_\mathcal{I}^*}(\nu_t,a^{t-1},\varphi_t')\subset\mathbb{R}$ consists of the payoffs w_t' such that the following holds:

1. A dynamic mechanism $(\varphi_\tau')_{\tau\geq t}$, where $(\varphi_\tau')_{\tau\geq t}=(\varphi_t',(\varphi_\tau')_{\tau\geq t+1})$, and an agent-PBE (σ_A',μ') of $\Gamma(\nu_t,(\varphi_\tau')_{\tau\geq t},a^{t-1})$ exist that implement $\left(w_t',(u_{\theta_t}')_{\theta_t\in\Theta}\right)$ at ν_t,
2. For all $h_t^{t+1}\equiv(s',a',\omega')\in SA_\emptyset\times\Omega$, $\mathbf{V}_{t+1}(\eta^{(\varphi_\tau')_{\tau\geq t},\sigma_A'|h_t^{t+1}})\in\mathcal{E}_{\mathcal{I},t+1}^*\left(f_{t+1,a'}(\mu_{t+1,\Theta}(h_t^{t+1})),a^{t-1},a'\right)$.

In words, a principal's payoff w' is in $\mathsf{D}_{\mathcal{E}_{\mathcal{I},\cdot}^*}(\cdot,\varphi_t')$ if it satisfies two properties. First, $w'=\mathbf{W}_t(\eta^{(\varphi_\tau')_{\tau\geq t},\sigma_A'})$ for a dynamic mechanism $(\varphi_\tau')_{\tau\geq t}$ such that φ_t' is the period t-mechanism and (σ_A',μ') is an agent-PBE given $(\varphi_\tau')_{\tau\geq t}$. Second, *continuation* payoffs are PBE-feasible, which means that the *punishment* for deviating to φ_t' is credible. The reason that we are able to require that continuation payoffs are PBE-feasible is that whenever the agent does not condition her strategy on the payoff-irrelevant part of the private history, the following holds: If (σ_A',μ') is an agent-PBE of $\Gamma(\nu_t,a^{t-1},(\varphi_\tau')_{\tau\geq t})$, then for all $h_t^{t+1}=(s',a',\omega')$, $(\sigma_A',\mu')|_{h_t^{t+1}}$ is an agent-PBE of $\Gamma(f_{a'}(\mu_{t+1,\Theta}'(h_t^{t+1})),(a^{t-1},a'),(\varphi_\tau'(h_t^{t+1},\cdot))_{\tau\geq t+1})$.

While we can use the set $\mathsf{D}_{\mathcal{E}_\mathcal{I}^*}(\nu_t,a^{t-1},\cdot)$ to test whether the principal has a deviation from $(\varphi_\tau)_{\tau\geq t}$ at the root of $\Gamma(\nu_t,a^{t-1},(\varphi_\tau)_{\tau\geq t})$, Definition D.4 also describes how we test for sequential rationality at later points in the agent-extensive form:

Definition D.4 (*Sequential Rationality*)**.** Fix $t\geq 1$, (ν_t,a^{t-1}), a dynamic mechanism $(\varphi_\tau)_{\tau\geq t}$ given a^{t-1}, and an agent-PBE (σ_A,μ) of $\Gamma(\nu_t,a^{t-1},(\varphi_\tau)_{\tau\geq t})$. $(\varphi_\tau)_{\tau\geq t}$ is *sequentially rational* given (σ_A,μ) if the following hold:

1. For all $\varphi_t':M\mapsto\Delta(S\times A)$, a payoff $w'\in\mathsf{D}_{\mathcal{E}_\mathcal{I}^*}(\cdot,\varphi_t')$ exists such that $\mathbf{W}_t(\eta^{(\varphi_\tau)_{\tau\geq t},\sigma_A})\geq w'$,
2. For all $h_t^{t+1}=(s',a',\omega')\in(SA_\emptyset\times\Omega)$, $\mathbf{V}_{t+1}(\eta^{(\varphi_\tau)_{\tau\geq t},\sigma_A|h_t^{t+1}})\in\mathcal{E}_{\mathcal{I},t+1}^*(f_{t+1,a'}(\mu_{t+1}(h_t^{t+1})),a^{t-1},a')$.

The first part of Definition D.4 states that the principal has no deviations in period t. The second part says that the principal has no deviations in periods $\tau \geq t+1$: the continuation payoffs induced by $(\varphi_\tau)_{\tau\geq t}$ and (σ_A, μ) are PBE - feasible continuation payoffs.

We are now ready to define the set of PBE feasible payoffs at (v_t, a^{t-1}), $\mathcal{E}^*_{\mathcal{I},t}(v_t, a^{t-1})$:

Definition D.5 (*PBE-feasible Payoffs*)**.** Fix $t \geq 1$, $(v_t, a^{t-1}) \in \Delta(\Theta) \times A^{t-1}$. The payoff vector $\left(w, (u_\theta)_{\theta\in\Theta}\right)$ is PBE-feasible at (v_t, a^{t-1}) if a dynamic mechanism $(\varphi_\tau)_{\tau\geq t}$ given a^{t-1} and an agent-PBE (σ_A, μ) of $\Gamma(v_t, a^{t-1}, (\varphi_\tau)_{\tau\geq t})$ exist such that

1. $\left(w, (u_\theta)_{\theta\in\Theta}\right) = \mathbf{V}_t(\eta^{(\varphi_\tau)_{\tau\geq t}, \sigma_A})$
2. $(\varphi_\tau)_{\tau\geq t}$ is sequentially rational given (σ_A, μ).

$\mathcal{E}^*_{\mathcal{I},t}(v_t, a^{t-1})$ denotes the set of PBE-feasible payoffs at (v_t, a^{t-1}).

By varying $\mathcal{I}$, we can define the set of PBE-feasible outcomes when the principal can offer mechanisms whose input and output messages are $(\Theta, \Delta(\Theta))$. Like in Theorem 1, our interest is in the canonical-PBE-feasible payoffs, that is, those payoffs that are induced by canonical dynamic mechanisms $\left(\varphi^C_\tau\right)_{\tau\geq t}$ (Definition 1) and canonical-agent PBE of the extensive-form game $\Gamma(v_t, a^{t-1}, \left(\varphi^C_\tau\right)_{\tau\geq t})$. In a slight abuse of notation, let $\mathcal{E}^C_t(\cdot)$ denote the correspondence of canonical-PBE-feasible outcomes when $\mathcal{I} = \{(\Theta, \Delta(\Theta))\}$.

Theorem D.1. *For all $t \geq 1$ and pairs $(v_t, a^{t-1}) \in \Delta(\Theta) \times A^{t-1}$, $\mathcal{E}^*_{\mathcal{I},t}(v_t, a^{t-1}) = \mathcal{E}^C_t(v_t, a^{t-1})$.*

We omit the proof of Theorem D.1, as it follows from adapting that of Doval and Skreta (2022) following the arguments in Appendix B.

References

Aliprantis, C., Border, K., 2006. Infinite Dimensional Analysis: A Hitchhiker's Guide. Springer-Verlag Berlin and Heidelberg GmbH & Company KG.
Aumann, R.J., 1961. Borel structures for function spaces. Illinois J. Math. 5 (4), 614–630.
Aumann, R.J., 1963. On choosing a function at random. In: Symposium on Ergodic Theory. Academic Press New Orleans, pp. 1–20.
Aumann, R.J., 1964. Mixed and Behavior Strategies in Infinite Extensive Games. Tech. Rep., Princeton University.
Baron, D.P., Besanko, D., 1984. Regulation and information in a continuing relationship. Inf. Econom. Policy 1 (3), 267–302.
Battaglini, M., 2005. Long-term contracting with Markovian consumers. Amer. Econ. Rev. 95 (3), 637–658.
Ben-Porath, E., Dekel, E., Lipman, B.L., 2019. Mechanisms with evidence: Commitment and robustness. Econometrica 87 (2), 529–566.
Bisin, A., Rampini, A.A., 2006. Markets as beneficial constraints on the government. J. Public Econom. 90 (4–5), 601–629.
Courty, P., Li, H., 2000. Sequential screening. Rev. Econ. Stud. 67 (4), 697–717.
Doval, L., Skreta, V., 2022. Mechanism design with limited commitment. Econometrica 90 (4), 1463–1500.
Doval, L., Skreta, V., 2023. Purchase history and product personalization. RAND (forthcoming).
Ely, J.C., 2017. Beeps. Amer. Econ. Rev. 107 (1), 31–53.
Farhi, E., Sleet, C., Werning, I., Yeltekin, S., 2012. Non-linear capital taxation without commitment. Rev. Econ. Stud. 79 (4), 1469–1493.
Farhi, E., Werning, I., 2013. Insurance and taxation over the life cycle. Rev. Econ. Stud. 80 (2), 596–635.
Fréchette, G.R., Lizzeri, A., Perego, J., 2022. Rules and commitment in communication: An experimental analysis. Econometrica 90 (5), 2283–2318.
Fudenberg, D., Tirole, J., 1991. Perfect Bayesian equilibrium and sequential equilibrium. J. Econom. Theory 53 (2), 236–260.
Garrett, D.F., Pavan, A., 2012. Managerial turnover in a changing world. J. Polit. Econ. 120 (5), 879–925.
Golosov, M., Iovino, L., 2021. Social insurance, information revelation, and lack of commitment. J. Polit. Econ. 129 (9), 2629–2665.
Golosov, M., Tsyvinski, A., Werning, I., Diamond, P., Judd, K.L., 2006. New dynamic public finance: A user's guide [with comments and discussion]. NBER Macroeconom. Annu. 21, 317–387.
Hart, S., Kremer, I., Perry, M., 2017. Evidence games: Truth and commitment. Amer. Econ. Rev. 107 (3), 690–713.
Kallenberg, O., 2017. Random Measures, Theory and Applications. Springer.
Kapička, M., 2013. Efficient allocations in dynamic private information economies with persistent shocks: A first-order approach. Rev. Econ. Stud. 80 (3), 1027–1054.
Laffont, J.-J., Tirole, J., 1988. The dynamics of incentive contracts. Econometrica 1153–1175.
Lipnowski, E., Ravid, D., Shishkin, D., 2022. Persuasion via weak institutions. J. Polit. Econ. 130 (10), 2705–2730.
Myerson, R.B., 1982. Optimal coordination mechanisms in generalized principal–agent problems. J. Math. Econom. 10 (1), 67–81.
Myerson, R.B., 1986. Multistage games with communication. Econometrica 323–358.
Pavan, A., Segal, I., Toikka, J., 2014. Dynamic mechanism design: A myersonian approach. Econometrica 82 (2), 601–653.
Peters, M., 2001. Common agency and the revelation principle. Econometrica 69 (5), 1349–1372.
Pollard, D., 2002. A User's Guide to Measure Theoretic Probability, 8, Cambridge University Press.
Sleet, C., Yeltekin, Ş., 2008. Politically credible social insurance. J. Monet. Econom. 55 (1), 129–151.
Stantcheva, S., 2015. Learning and (or) Doing: Human Capital Investments and Optimal Taxation. Tech. Rep., National Bureau of Economic Research.
Stantcheva, S., 2020. Dynamic taxation. Annu. Rev. Econ. 12, 801–831.
Townsend, R.M., 1988. Information constrained insurance: the revelation principle extended. J. Monetary Econ. 21 (2–3), 411–450.