Mechanism Design with Limited Commitment[*]

Laura Doval[†]　　Vasiliki Skreta[‡]

March 21, 2021

Abstract

We develop a tool akin to the revelation principle for mechanism design with limited commitment. We identify a *canonical* class of mechanisms rich enough to replicate the outcomes of any equilibrium in a mechanism-selection game between an uninformed designer and a privately informed agent. A cornerstone of our methodology is the idea that a mechanism should encode not only the rules that determine the allocation, but also the information the designer obtains from the interaction with the agent. Therefore, how much the designer learns, which is the key tension in design with limited commitment, becomes an explicit part of the design. We show how this insight can be used to transform the designer's problem into a constrained optimization problem: To the usual truthtelling and participation constraints, one must add the designer's sequential rationality constraint.

KEYWORDS: mechanism design, limited commitment, revelation principle, information design

JEL CLASSIFICATION: D84, D86

1 INTRODUCTION

The standard assumption in dynamic mechanism design is that the designer can commit to long-term contracts. This assumption is useful: It allows us to characterize the best possible payoff for the designer in the presence of adverse selection and/or moral hazard, and it is applicable in many settings. Often, however, this assumption is made for technical convenience. Indeed, when the designer can commit to long-term contracts, the mechanism-selection problem can be reduced to a constrained optimization problem thanks to the *revelation principle*.[1] However, as the literature starting with Laffont and Tirole (1987, 1988) shows, when the designer can commit only to short-term contracts, the tractability afforded by the revelation principle is lost. Indeed, mechanism design problems with limited commitment are difficult to analyze without imposing auxiliary assumptions either on the class of contracts the designer can choose from, as in Gerardi and Maestri (2020) and Strulovici (2017), or on the length of the horizon, as in Skreta (2006, 2015).

This paper provides a "revelation principle" for dynamic mechanism-selection games in which the designer can only commit to short-term contracts. We study a game between an uninformed designer and an informed agent with persistent private information. Although the designer can commit within each period to the terms of the interaction–the current mechanism–he cannot commit to the terms the agent faces later on, namely, the mechanisms that are chosen in the continuation game. First, we show there is a class of mechanisms that is sufficient to replicate all equilibrium outcomes of the mechanism-selection game. Second, we show how this insight can be used to transform the designer's problem into a constrained optimization problem: To the usual truthtelling and participation constraints, one must add the designer's sequential rationality constraint.

The starting point of our analysis is the class of mechanisms we allow the designer to select from. Following Myerson (1982) and Bester and Strausz (2007), we consider mechanisms defined by a *general communication device* as illustrated in Figure 1a:

---

[1]The "revelation principle" denotes a class of results in mechanism design; see Gibbard (1973), Myerson (1979), and Dasgupta et al. (1979).

$$\Theta \xrightarrow{\ r(\cdot|\theta)\ } M \xrightarrow{\ \beta(\cdot|m)\ } S \times A \qquad\qquad \Theta \xrightarrow{\ r(\cdot|\theta)\ } M \xrightarrow{\ \beta(\cdot|m)\ } S \xrightarrow{\ \alpha(\cdot|s)\ } A$$

Figure (a) General Mechanisms   Figure (b) Canonical Mechanisms:
$$M = \Theta, S = \Delta(\Theta)$$

Figure 1: Mechanisms

Having observed her private information (her type, $\theta \in \Theta$), the agent privately reports an *input* message, $m \in M$, to the mechanism; this then determines the distribution, $\beta(\cdot|m)$, 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*: They constitute the contractible parts of the mechanism.

When the designer has commitment, the revelation principle implies that, without loss of generality, we can restrict attention to mechanisms satisfying the following three properties: (i) $M = \Theta$, (ii) $|M| = |S|$, and (iii) β is such that by observing the output message, the designer learns the input message, in this case the agent's type report. Moreover, the revelation principle implies that we can restrict attention to equilibria in which the agent truthfully reports her type, which means that the designer not only learns the agent's type report upon observing the output message but also learns the agent's true type.

It is then clear why restricting attention to mechanisms that satisfy properties (i)-(iii) and truthtelling equilibria is with loss of generality under limited commitment: Upon observing the output message, the designer learns the agent's type report and hence her type. Then the agent may have an incentive to misreport if the designer cannot commit to not react to this information. This is precisely the intuition behind the main result in Bester and Strausz (2001), which is the first paper to provide a general analysis of optimal mechanism design with limited commitment. Instead of allowing the designer to choose *any* mechanism, the authors restrict attention to mechanisms such that the cardinality of the set of input and output messages is the same and β is such that by observing the output message, the designer learns the input message.[2] They show that to sustain payoffs in the Pareto frontier, mechanisms in which input messages are type reports are without loss of generality. However, focusing on truthtelling equilibria is with loss of generality. In a follow-up paper, Bester and Strausz (2007) lift the restrictions on the class of mechanisms (i.e., (ii) and (iii) above) and show in a one-period model that focusing on mechanisms in

---

[2]The class of mechanisms considered in Bester and Strausz (2001) encompasses the mechanisms considered by most papers in the literature on limited commitment starting from Laffont and Tirole (1988).

which input messages are type reports and truthtelling equilibria is without loss of generality. The authors, however, do not characterize the output messages. It is also unclear whether taking input messages to be type reports is without loss when the designer and the agent interact repeatedly (see the discussion after Theorem 1).

The main contribution of this paper is to show that, under limited commitment, it is without loss of generality to take the set of output messages to be the set of designer's posterior beliefs about the agent's type, that is, $S = \Delta(\Theta)$. Theorem 1 shows that the following two games between an uninformed designer and an informed agent implement the same set of equilibrium distributions over types and allocations. In the first, the *mechanism-selection game*, the principal can offer the agent mechanisms as in Figure 1a. In the second, the *canonical* game, the principal can offer the agent mechanisms in which input messages are type reports, output messages are beliefs, and, conditional on the output message, the allocation is drawn independently of the agent's type report (see Figure 1b). Moreover, Theorem 1 shows that any equilibrium of the canonical game can be replicated by an equilibrium in which (a) the agent always participates in the mechanism, and (b) input and output messages have a *literal* meaning: The agent truthfully reports her type, and if the mechanism outputs a given posterior, this posterior coincides with the belief that the designer holds about the agent's type given the agent's strategy and the mechanism. Given that any equilibrium distribution over types and allocations can be replicated by mechanisms in which input messages are type reports and output messages are beliefs about the agent's type, we call this class of mechanisms *canonical.*

Theorem 1 implies that in mechanism design with limited commitment, the mechanism serves a dual role within a period. On the one hand, it determines the allocation for that period. On the other hand, it determines the information about the agent that is carried forward in the interaction. An advantage of the language of posterior beliefs is that it avoids potential infinite-regress problems. Indeed, in a finite horizon problem, an alternative set of output messages could be a recommendation for an allocation today and a sequence of allocations from tomorrow on. In the final period, the revelation principle in Myerson (1982) pins down the implementable allocations. Therefore, the recommended allocations can be determined via backward induction. This idea cannot be carried to an infinite horizon setting: These sets of output messages would necessarily have to make reference to the continuation mechanisms, which are themselves defined by a set of output messages.

Theorem 1 affords the analyst two main simplifications. First, it follows from its

proof that it is without loss of generality to restrict attention to the analysis of the canonical game, since it implements the same set of distributions over types and allocations, and thus payoffs, as the mechanism-selection game. Second, it provides the researcher with a tractable way to analyze problems of mechanism design with limited commitment by making how much the principal learns about the agent an explicit part of the design. The three constraints that the mechanism must satisfy, the participation and truthtelling constraints for the agent, and the Bayes' plausibility constraint, provide us with a tractable representation both of the agent's behavior in a given period and of its impact on the mechanism offered in the next via the information that is generated about the agent's type in the given period. A major challenge in the received literature on limited commitment is how to keep track of how the agent's best response to the mechanism affects the information that the principal obtains from the interaction, which in turn affects the principal's incentives to offer the mechanism in the first place. Instead, our framework allows us to reduce the agent's best response to the principal's mechanism and its informational feedback to a familiar set of constraints that the mechanism must satisfy. This avoids having to consider complicated mixed strategies on the part of the agent (see, for instance, Laffont and Tirole (1988); Bester and Strausz (2001)) and transforms it instead into a program that combines elements of mechanism design and information design.

While Theorem 1 assumes that the agent's type is fully persistent, this is not necessary for its conclusion to hold. Theorem 2 extends Theorem 1 to a version of what Pavan et al. (2014) denote as *Markov environments*. These are settings where (i) the agent's private information follows a possibly nonhomogeneous Markov process, (ii) the principal and the agent's payoffs are time-separable, and their flow payoffs depend only on today's allocation and the agent's current type, and (iii) the transition probability may depend both on today's type and today's allocation. Theorem 2 shows that in Markov environments it is without loss of generality to restrict attention to the characterization of equilibrium payoffs of the canonical game and to strategy profiles where the agent participates and truthfully reports her *current* type to characterize the set of payoffs the designer can implement in the mechanism-selection game.

We illustrate how our results can be used to shed new light on seemingly well-understood problems with an example in Section 4.[3] Section 4 considers a seller, who owns one

---

[3]Theorem 1 also opens the door to the analysis of optimal mechanisms under limited commitment in infinite-horizon settings. We illustrate this in Doval and Skreta (2020a), where we solve an infinite-horizon binary-type version of the sale of a durable good.

unit of a durable good, and interacts over two periods with a buyer with persistent and private information, as in Skreta (2006). Whereas the seller in Skreta (2006) is allowed to offer any mechanism amongst those considered by Laffont and Tirole (1988); Bester and Strausz (2001), we allow the seller to choose any canonical mechanism. In contrast to the main result in Skreta (2006), we show in Proposition 1 that *rationing* can strictly dominate posted prices (Remark 2 discusses how canonical mechanisms differ from the mechanisms in Skreta (2006), which explains the difference in the results.) This allows us to connect the mechanism design literature on the sale of a durable good with the work in theoretical industrial organization on alternative strategies for a durable good monopolist, such as rationing (Denicolo and Garella (1999); Gilbert and Klemperer (2000); McAfee and Wiseman (2008)) and clearance sales (Nocke and Peitz (2007)). Indeed, Proposition 1 provides a micro-foundation for a mechanism first suggested by Denicolo and Garella (1999) and a new rationale for the use of clearance sales.

Our work brings forth a new application of information design by placing its tools at the service of characterizing optimal mechanisms under limited commitment. By highlighting the canonical role of beliefs as the signals employed by the mechanism, Theorem 1 and Theorem 2 underscore the importance of jointly determining the mechanism together with how information is used in the mechanism and transmitted across periods. In doing so, it marries information design, which studies the design of information structures in a given institution, with mechanism design, which generally studies institutional design within a given information structure.

Related Literature: The paper contributes to the literature on mechanism design with limited commitment with an informed agent with persistent private information, referenced throughout the introduction.[4] Following the seminal contribution of Bester and Strausz (2001), a body of work studies optimal mechanisms under limited commitment in settings with finitely many types and finite horizon (e.g., Bisin and Rampini (2006); Hiriart et al. (2011); Fiocco and Strausz (2015); Beccuti and Möller (2018)). Since the results in Bester and Strausz (2001) do not extend to settings with a continuum of types and/or infinite horizon, the characterization of optimal mechanisms under limited commitment in these settings has proven elusive. On the one hand, Skreta (2006); Deb and Said (2015); Skreta (2015) study mechanism-selection games with a continuum of types and finite horizon. All three papers lever-

---

[4]A designer's lack of commitment can take various forms that are not considered in this paper but have been studied in others. See, for instance, McAdams and Schwarz (2007), Vartiainen (2013), and Akbarpour and Li (2020), in which the designer cannot commit even to obeying the rules of the current mechanism.

age the assumption of finite horizon, which pins down the optimal mechanism in the final period, to characterize the implications of the principal's sequential rationality constraints for the set of incentive-feasible outcomes. On the other hand, the small set of papers that study infinite-horizon problems of design under limited commitment do so under assumptions on either the set of mechanisms the designer is allowed to offer (e.g., Acharya and Ortner (2017); Strulovici (2017); Gerardi and Maestri (2020)) or the solution concept (e.g., Acharya and Ortner (2017)). All these papers use the set of mechanisms in Laffont and Tirole (1988); Bester and Strausz (2001).

Due to the difficulties with the revelation principle, a large body of work in public finance, political economy and taxation considers optimal time-consistent policies in settings where private information is fully nonpersistent (see, for instance, Sleet and Yeltekin (2008); Farhi et al. (2012); Golosov and Iovino (2016)). Moreover, a large literature studies the effect of limited commitment within a specific class of "mechanisms": The papers in the durable-good monopolist literature (Bulow (1982); Gul et al. (1986); Stokey (1981)) study price dynamics and establish (under some conditions) Coase's conjecture whereby a monopolist essentially loses all profits if it lacks commitment. In an analogous vein, Burguet and Sakovics (1996), McAfee and Vincent (1997), Caillaud and Mezzetti (2004), and Liu et al. (2019) study equilibrium reserve-price dynamics without commitment in different setups. The common thread is, again, that the seller's inability to commit reduces monopoly profits.

By highlighting the role that the designer's beliefs about the agent play in mechanism design with limited commitment, our paper also relates to Lipnowski and Ravid (2020) and Best and Quigley (2017), who study models of direct communication between an informed sender and an uninformed receiver.

Organization: The rest of the paper is organized as follows. Section 2 describes the model and notation. Section 3 introduces the main theorem and provides a sketch of the proof. Section 4 analyzes a two-period version of the model in Skreta (2006) to illustrate how one can apply Theorem 1 in a setting with a continuum of types and shed new light on a classic problem. Section 5 presents Theorem 2, which extends Theorem 1 to *Markov environments*. All proofs are in Appendix B and the supplementary material, Doval and Skreta (2020c) (Appendices C-E).

2   MODEL

Primitives: Two players, a principal (he) and an agent (she), interact over $T \leq \infty$ periods. Before the game starts, the agent observes her type, $\theta \in \Theta$, which is dis-

tributed according to a full support distribution μ_0. Each period, as a result of the interaction between the principal and the agent, an allocation $a \in A$ is determined. Let A^{T+1} denote the set $\times_{t=0}^{T} A$. For the principal, assume that there exists a function $W : A^{T+1} \times \Theta \mapsto \mathbb{R}$ such that his payoff from allocation $a^{T+1} \in A^{T+1}$ when the agent's type is θ is given by $W(a^{T+1}, \theta)$. Similarly, for the agent, when her type is θ, her payoff from allocation $a^{T+1} \in A^{T+1}$ is given by $U(a^{T+1}, \theta)$.

For every $t \geq 1$ and every sequence of allocations $a^t = (a_0, a_1, \ldots, a_{t-1})$, the principal can only choose $a_t \in \mathcal{A}(a^t)$ in period t. That is, there is a correspondence $\mathcal{A} : \cup_{t=1}^{T} A^t \mapsto A$ such that for $t \in \{1, \ldots, T\}$ and $a^t \in A^t$, $\mathcal{A}(a^t)$ describes the set of allocations that the principal can offer in period t given the allocations he has offered through period $t - 1$. On the one hand, the correspondence $\mathcal{A}$ encodes that the set of feasible allocations may be time dependent, so that $\mathcal{A}$ depends on a^t only through the time index t. On the other hand, it allows for the case in which the past allocations restrict what the principal can offer the agent in the future, as in the application in Section 4. Assume that there exists an allocation $a^* \in A$ such that a^* is always available. Below, allocation a^* plays the role of the agent's outside option. Given the general structure of payoffs, it is without loss of generality to take it to be time-independent.

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. We also assume that Θ is compact. Endowing product sets with their product σ-algebra, we assume that the principal and the agent's utility functions, W and U, are bounded measurable functions. Similarly, the correspondence $\mathcal{A}$ is measurable.

Mechanisms: In each period, the principal offers the agent a mechanism, $\mathbf{M}_t = (M^{\mathbf{M}_t}, S^{\mathbf{M}_t}, \beta^{\mathbf{M}_t})$, where $M^{\mathbf{M}_t}$ and $S^{\mathbf{M}_t}$ are Polish, and $\beta^{\mathbf{M}_t}$ is a transition probability from $M^{\mathbf{M}_t}$ to $S^{\mathbf{M}_t} \times A$. We endow the principal with a collection $(M_i, S_i)_{i \in \mathcal{I}}$ of input and output message sets in which $|\Theta| \leq |M_i|$ and $|\Delta(\Theta)| \leq |S_i|$. Moreover, we assume that $(\Theta, \Delta(\Theta))$ is an element in that collection. Denote by $\mathcal{M}$ the set of all mechanisms with message sets $(M_i, S_i)_{i \in \mathcal{I}}$, i.e., $\{\beta : M_i \mapsto \Delta(S_j \times A) : i, j \in \mathcal{I}\}$.

---

[5] In what follows, we adopt the following notational conventions. First, all Polish spaces are endowed with their Borel σ-algebra. For a Polish space, X, $\mathcal{B}_X$ denotes its 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 (2013)). For any two measurable spaces X and Y, a mapping $\zeta : X \mapsto \Delta(Y)$ is a *transition probability from X to Y* if for any measurable $C \subseteq Y$, $\zeta(C|x) \equiv \zeta(x)(C)$ is a measurable real valued function of $x \in X$.

Three remarks are in order. First, the restriction that M_i has at least as many messages as types is without loss of generality. The principal can always replicate a mechanism with a smaller set of input messages by using a larger set of input messages.[6] Second, we restrict the principal to choosing input and output messages within the set $(M_i, S_i)_{i \in \mathcal{I}}$. This allows us to have a well-defined set of deviations for the principal, thereby avoiding set-theoretic issues related to self-referential sets. The analysis that follows shows that the choice of the collection plays no further role in the analysis. Finally, we note that all aspects of the environment, except the agent's type $\theta \in \Theta$, are common knowledge between the principal and the agent.

Timing: In each period t, the game proceeds as follows. The principal offers the agent a mechanism, $\mathbf{M}_t$, with the property that for all $m \in M^{\mathbf{M}_t}$, $\beta^{\mathbf{M}_t}(S^{\mathbf{M}_t} \times \mathcal{A}(a^t)|m) = 1$, where a^t describes the allocations implemented through period $t-1$. Observing the mechanism, the agent decides whether to participate in the mechanism ($p = 1$) or not ($p = 0$). If she does not participate in the mechanism, a^* is implemented and the game proceeds to $t+1$. 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 $\beta^{\mathbf{M}_t}(\cdot|m)$; the output message and the allocation are observed by both the principal and the agent.

The above defines an extensive-form game, which we dub the *mechanism-selection game* and denote by $G_{\mathcal{M}}$. Public histories in this game are

$$h^t = (\mathbf{M}_0, p_0, s_0, a_0, \dots, \mathbf{M}_{t-1}, p_{t-1}, s_{t-1}, a_{t-1}),$$

where $p_r \in \{0, 1\}$ denotes the agent's participation with the restriction that $p_r = 0 \Rightarrow s_r = \varnothing, a_r = a^*$.[7] Given a mechanism $\mathbf{M}_t$, let $z_\varnothing(\mathbf{M}_t)$, denote the tuple $\mathbf{M}_t, 0, \varnothing, a^*$ and let $z_{(s_t, a_t)}(\mathbf{M}_t)$, denote the tuple $\mathbf{M}_t, 1, s_t, a_t$. Note that any public history at the end of period t can be written as $(h^t, z_\varnothing(\mathbf{M}_t))$ or $(h^t, z_{(s_t, a_t)}(\mathbf{M}_t))$.

Public histories capture what the principal knows through period t. Let H^t denote the set of all period t public histories. A history for the agent consists of the *public* history of the game together with the agent's inputs into the mechanism (henceforth, the agent history) and her private information. Formally, an agent history is an ele-

---

[6]To see this, suppose that the principal would rather use a mechanism, $\mathbf{M}'_t$, with a message space $M^{\mathbf{M}'_t}$ with cardinality strictly less than $|\Theta|$. Then he can choose a mechanism $\mathbf{M}_t$ with $M^{\mathbf{M}_t} = \Theta$, choose $\beta^{\mathbf{M}_t}$ to coincide with $\beta^{\mathbf{M}'_t}$ on the first $|M^{\mathbf{M}'_t}|$ messages, and have $\beta^{\mathbf{M}_t}$ coincide with $\beta^{\mathbf{M}'_t}(\cdot|m'_1)$ for all remaining messages.

[7]While there is no output message when the agent does not participate in the mechanism, we denote this by $s = \varnothing$ to keep the length of all the histories the same.

ment

$$h_A^t = (\mathbf{M}_0, p_0, m_0, s_0, a_0, \ldots, \mathbf{M}_{t-1}, p_{t-1}, m_{t-1}, s_{t-1}, a_{t-1}),$$

with $p_r = 0 \Rightarrow m_r = \varnothing$. Given a public history h^t, let $H_A^t(h^t)$ denote the set of agent histories consistent with h^t. The agent also knows her type, and hence a history through period t is an element of $\{\theta\} \times H_A^t$ when her type is θ.

Strategies: Since the principal's action space, $\mathcal{M}$, is an uncountable set of functions, we model the principal's behavioral strategy, $(\sigma_{Pt})_{t=0}^T$, following Aumann (1964). That is, endow $[0,1]$ with the Borel σ-algebra and the Lebesgue measure, λ. Then, σ_{Pt} is defined as a jointly measurable function from $H^t \times [0,1]$ to $\mathcal{M}$.[8] We denote the collection $(\sigma_{Pt})_{t=0}^T$ by σ_P. The agent's participation strategy is a transition probability, π_t, from $\Theta \times H_A^t \times \mathcal{M}$ to $\{0,1\}$. Conditional on participating in the mechanism $\mathbf{M}_t$, her reporting strategy is a transition probability, r_t, from $\Theta \times H_A^t(h^t) \times \mathcal{M} \times \{1\}$ to $\cup_{i \in \mathcal{I}} M_i$ such that $r_t(\theta, h_A^t, \mathbf{M}_t) \in \Delta(M^{\mathbf{M}_t})$.[9] We denote the tuple (π_t, r_t) by σ_{At}, and the collection $(\sigma_{At})_{t=0}^T$ by σ_A.

The definitions above assume that one can define a measurable structure on $\mathcal{M}$ such that we can define the principal and the agent's strategies as measurable functions of the histories. As noted by Aumann (1961), this requires that $\mathcal{M}$ be a standard Borel space. As we explain in Appendix C, this is the case when Θ is finite or countable. Instead, when Θ is a continuum, the set $\mathcal{M}$ is not a standard Borel space. When Θ is a continuum, there are two approaches one may follow. The first approach is to allow the principal to choose from a subset $\mathcal{M}' \subset \mathcal{M}$ which is a standard Borel space, in which case it is correct to write the principal and the agent's strategies as conditioning on the past chosen mechanisms. The second approach relies on the idea that one can represent a very large set of mechanisms in terms of a standard Borel space, without ex-ante restricting the mechanisms the principal is allowed to choose from. This approach, which we implicitly use in our applications, requires defining the principal and the agent's strategies in a different way, so we relegate this discussion to Appendix C.[10]

---

[8]To keep notation simple, we do not add the restriction that if $\mathbf{M}_t$ is in the support of $\sigma_{Pt}(h^t)$, then $\beta^{\mathbf{M}_t}(S^{\mathbf{M}_t} \times \mathcal{A}(a^t)|m) = 1$, where a^t is the allocation up to period t according to h^t.

[9]While technically the agent's reporting strategy should be written $r_t(\theta, h_A^t, \mathbf{M}_t, 1)$ to account for the agent's decision to participate, we omit the 1 to simplify notation.

[10]The issue of choosing mechanisms at random also shows up in the competing principals literature, where it is typical to restrict attention to pure strategy equilibria of the mechanism-selection game.

A belief for the principal at the beginning of time t, history h^t, is a distribution $\mu_t(h^t) \in \Delta(\Theta \times H_A^t(h^t))$. The principal is thus uncertain both about the agent's payoff-relevant type, θ, and her payoff-irrelevant private history, h_A^t. The collection $(\mu_t)_{t=0}^T$ denotes the belief system.

An assessment is a tuple $(\sigma_{Pt}, \sigma_{At}, \mu_t)_{t=0}^T \equiv (\sigma_P, \sigma_A, \mu)$. Our focus is on studying the equilibria of the mechanism-selection game. By equilibrium, we mean Perfect Bayesian equilibrium (henceforth, PBE), defined as follows:

Definition 1. *A* Perfect Bayesian Equilibrium *is a tuple $(\sigma_P, \sigma_A, \mu)$ such that the following holds:*

1. *$(\sigma_P, \sigma_A, \mu)$ is sequentially rational (Definition A.1), and*

2. *The belief system satisfies Bayes' rule where possible (Definition A.2).*

The formal statement is in Appendix A. For now, we note that if the principal's strategy space were finite, Θ is finite, and the mechanisms used by the principal have finite support, then this coincides with the definition in Fudenberg and Tirole (1991).[11]

The prior μ_0 together with a strategy profile (σ_P, σ_A) determine a distribution over the terminal nodes $\Theta \times H_A^{T+1}$.[12] We are interested instead in the distribution they induce over the payoff-relevant outcomes, $\Theta \times A^{T+1}$. We say that $\gamma \in \Delta(\Theta \times A^{T+1})$ is a PBE outcome of the mechanism-selection game if there exists a PBE of the mechanism-selection game that induces γ. We denote by $\mathcal{O}_\mathcal{M}^*$ the set of PBE outcomes of $G_\mathcal{M}$. Our main result establishes the equivalence between the set of PBE outcomes of the mechanism-selection game and those of another game, which we dub the *canonical game* and introduce next.

The canonical game: The canonical game is essentially the same as the mechanism-selection game except for two features. First, in every period the principal can only choose *canonical* mechanisms, which we denote by $\mathcal{M}_C$. Canonical mechanisms differ from the mechanisms introduced above in two respects. First, the sets of input and output messages are given by $(\Theta, \Delta(\Theta))$. Second, if $(\Theta, \Delta(\Theta), \tilde{\beta}) \in \mathcal{M}_C$, then there exist two transition probabilities, β from Θ to $\Delta(\Theta)$ and α from $\Delta(\Theta)$ to A such

---

[11]The only difference between Bayes' rule where possible and consistency in sequential equilibrium is the following. 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.

[12]This distribution, which exists via the Ionescu Tulcea extension theorem (Tulcea (1949)), is formally defined in Appendix A.

that, for all $\theta \in \Theta$, and for all measurable subsets $\tilde{U} \times \tilde{A} \subseteq \Delta(\Theta) \times A$

$$\tilde{\beta}(\tilde{U} \times \tilde{A}|\theta) = \int_{\tilde{U}} \alpha(\tilde{A}|\mu)\beta(d\mu|\theta),$$

so that conditional on the output message $\mu \in \Delta(\Theta)$, the allocation is drawn *independently* of the type report. Third, at the beginning of each period both players observe the realization of a public randomization device $\omega \sim U[0,1]$. We denote the canonical game by G and its set of PBE outcomes by $\mathcal{O}^*$.

Remark 1 (An auxiliary game)**.** *The proofs of our results require translating strategy profiles from the mechanism-selection game to the canonical game, which is notationally involved. To facilitate the presentation of our results we rely on a third* auxiliary *game, $G_{\mathcal{M}}^A$, which is exactly like the mechanism-selection game, except that at the beginning of each period the principal and the agent observe a draw from a public randomization device $\omega \sim U[0,1]$. Note that one can trivially adapt any strategy profile σ of the mechanism-selection game to a strategy profile σ' of the auxiliary game simply by specifying that σ' follows σ for every realization of the public randomization device. Similarly, any strategy profile σ' of the auxiliary game in which the principal chooses a canonical mechanism at every history can trivially be adapted to a strategy profile $\tilde{\sigma}$ of the canonical game simply by specifying that $\tilde{\sigma}$ coincides with σ' only at histories where the principal has offered canonical mechanisms throughout.*

3   MAIN RESULT

Section 3 presents the main result of the paper. Theorem 1 shows that the mechanism-selection game and the canonical game have the same set of equilibrium outcomes. Moreover, any PBE assessment of the mechanism-selection game can be replicated by a PBE assessment of the canonical game in which (a) the agent always participates in the mechanism, and (b) input and output messages have a *literal* meaning: The agent truthfully reports her type, and if the mechanism outputs $\mu \in \Delta(\Theta)$ at the end of period t, then μ is indeed the belief the principal holds about the agent at the end of that period.

Theorem 1. *The mechanism-selection game and the canonical game have the same set of PBE outcomes, i.e., $\mathcal{O}_{\mathcal{M}}^* = \mathcal{O}^*$ for any collection of mechanisms, $\mathcal{M}$, with which we endow the principal.*

Moreover, for any PBE outcome of $G_{\mathcal{M}}$, there exists a PBE assessment $(\sigma_P, \sigma_A, \mu)$ of G that implements the same outcome and satisfies the following properties:

1. *The agent's strategy depends only on her private type and the public history.*

2. *For all public histories h^t, for all θ in the support of $\mu_t(h^t)$, the agent participates in the mechanism offered by the principal at that history and with probability one truthfully reports her type,*

3. *For all public histories h^t, if the mechanism offered by the principal at h^t outputs a posterior μ', the principal's updated equilibrium beliefs about the agent coincide with μ'. That is, for all measurable subsets $\tilde{\Theta}$, $\tilde{U}$, and $\tilde{A}$ of Θ, $\Delta(\Theta)$, and A,*

$$\int_\Theta \int_{\tilde{U}} \int_{\tilde{A}} \mu_{t+1}(\tilde{\Theta}|h^t, z_{(\mu', a_t)}) \alpha^{\mathbf{M}_t}(da_t|\mu') \beta^{\mathbf{M}_t}(d\mu'|\theta) \mu_t(d\theta|h^t) =$$
$$= \int_{\tilde{\Theta}} \int_{\tilde{U}} \alpha^{\mathbf{M}_t}(\tilde{A}|\mu') \beta^{\mathbf{M}_t}(d\mu'|\theta) \mu_t(d\theta|h^t) = \int_\Theta \int_{\tilde{U}} \int_{\tilde{A}} \mu'(\tilde{\Theta}) \alpha^{\mathbf{M}_t}(da_t|\mu') \beta^{\mathbf{M}_t}(d\mu'|\theta) \mu_t(d\theta|h^t).$$

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, $\mathcal{M}_C$, to which we can restrict the principal's choice set without loss of generality. 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 where the agent participates and truthfully reports her type. As is evidenced by the application in Section 4, this allows us to reduce the agent's behavior to a set of constraints that the mechanism must satisfy, exactly as in the case of commitment.

The proof that any PBE outcome of the mechanism-selection game can be achieved as a PBE outcome of the canonical game that satisfies the properties listed in Theorem 1 relies on four steps, which we review next.[13]

Input messages as type reports: To fix ideas, consider the typical proof for the standard revelation principle in static settings. The mechanism, $\mathbf{M}$, together with the agent's reporting strategy induces a transition probability from Θ to $S^{\mathbf{M}} \times A$. This allows us to conclude that we can replace the set of input messages with the set of type reports, as illustrated in Figure 2a.

In the dynamic setting, however, this argument would only allow us to conclude that we can rewrite the mechanism as a transition probability from $\Theta \times H_A^t(h^t)$ to $S^{\mathbf{M}_t} \times A$, as illustrated in Figure 2b. Indeed, to replicate the agent's reporting strategy, the

---

[13]In Doval and Skreta (2020b), we provide a proof of these four steps for the case in which Θ is finite and the principal can only offer mechanisms $\mathbf{M}$ such that, for all $m \in M^{\mathbf{M}}$, the support of $\beta^{\mathbf{M}}(\cdot|m)$ is finite. That proof mirrors the one in Appendix B, but is technically simpler and more accessible.
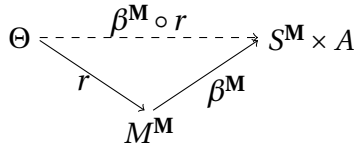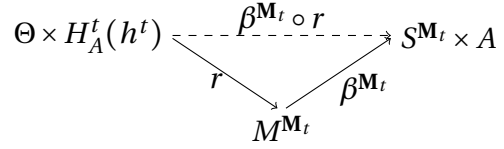
Figure (a) Static revelation principle

Figure (b) Mechanism selection game

Figure 2: Type reports as input messages

mechanism needs to obtain all the information on which the agent conditions her strategy, which potentially is (θ, h_A^t).

To circumvent this difficulty, we show that given a PBE in which the agent conditions her strategy on the payoff-irrelevant part of her private history at some public history h^t, there exists another outcome-equivalent PBE in which she does not and in which the principal obtains the same payoff after each continuation history consistent with h^t and the equilibrium strategy (see Proposition B.1). Thus, conditional on the public history h^t, the agent's reporting strategy and the mechanism induce a transition probability from Θ to $S^{\mathbf{M}_t} \times A$, so we can always take the set of input messages to be the set of type reports.[14] This result relies on two observations. First, because input messages are payoff irrelevant and unobserved by the principal, if the agent chooses different strategies at (θ, h_A^t) and $(\theta, \tilde{h}_A^t)$ with $h_A^t, \tilde{h}_A^t \in H_A^t(h^t)$, then she is indifferent between these two strategies. However, the principal may not be indifferent between these two strategies. The second step is to show that we can build an alternative strategy for the agent that conditions only on (θ, h^t) and yields the principal the same payoff.

This first step also gives us an important conceptual insight: The principal cannot peek into his past mechanisms. To do so, he would have to ask the agent to report her previous communication to him. Our result implies that this information cannot be elicited in any payoff-relevant way.

It follows from this first step that it is without loss of generality to focus on equilibrium assessments of the mechanism-selection game such that the principal offers mechanisms where the set of input messages are type reports, i.e., $M^{\mathbf{M}_t} = \Theta$, and the agent truthfully reports her type. In what follows, when we refer to a PBE assessment of $G_{\mathcal{M}}$, we mean one that satisfies these properties.

---

[14]This result is useful also in applications. It states that in our game the set of PBE payoffs coincides with the set of Public PBE payoffs (Athey and Bagwell (2008)). In games with time-separable payoffs, Public PBE payoffs have a recursive structure and are amenable to self-generation techniques, as in Abreu et al. (1990). For an example, see Doval and Skreta (2020a).

Output messages as beliefs: To understand the steps involved in showing that output messages can be taken without loss of generality to be the principal's beliefs about the agent's type, it is useful to consider the other uses the principal may have for the output messages beyond encoding information about the agent.

First, the principal could use $S^{\mathbf{M}_t}$ to encode randomizations on the allocation, e.g., two tuples, (s_t, a_t) and (s'_t, a'_t), may be associated with the same posterior belief. This is not an issue, however, since a canonical mechanism allows the principal to randomize on the allocation conditional on the posterior belief.

Second, the principal could use $S^{\mathbf{M}_t}$ to coordinate continuation play, e.g., two tuples $(s_t, a_t), (s'_t, a_t)$ may be associated with two different continuation equilibria, even if they induce the same posterior belief.[15] This is where canonical mechanisms and the property that output messages must coincide with the principal's updated beliefs may be more restrictive than other elements in $\mathcal{M}$: Beliefs are not a rich enough language to encode both the principal's updated beliefs and the suggested continuation play.[16] Nevertheless, the canonical game has a feature that the mechanism-selection game does not: the public randomization device. As we explain next, this allows us to subsume the second role of the output message.

The potential challenge in using the public randomization device to subsume the second role of the output message is that, by definition, the use of the public randomization device in the canonical game can only depend on publicly available information, while the output message in the mechanism-selection game is drawn as a function of the agent's type, since the agent is reporting truthfully. We leverage here that canonical mechanisms use *beliefs* as output messages. To see this, note that beliefs are a sufficient statistic for the information about the agent's type that is encoded in the output messages of the mechanism-selection game. Thus, conditional on the induced belief and the allocation, the selection of continuation play contains no further information about the agent's type. This is how we are able to decompose the mechanism in the mechanism-selection game into a mechanism in the canonical game that uses beliefs as output messages and a public randomization

---

[15]The same idea arises in the literature on Bayesian persuasion. Implicit in the result in Kamenica and Gentzkow (2011) that any experiment can be written as a distribution over posteriors is the assumption that the receiver breaks ties in favor of the sender. Unlike in Bayesian persuasion, it is not clear that the players may be indifferent between two continuation equilibria, so the public randomization device does not generally reduce to simple tie-breaking.

[16]Note that this is not a matter of cardinality, but a consequence of the restriction that suggested beliefs must coincide with equilibrium beliefs. Ultimately, by Kuratowski's theorem (see Parthasarathy (2005)), $\Delta(\Theta)$ is in bijection with $\Delta(\Theta) \times [0, 1]$ so there are enough messages to encode both the principal's updated beliefs and the suggested continuation play.

device.

The above argument also explains why in a canonical mechanism, conditional on the output message, the allocation can be drawn independently of the agent's type (report). Ultimately, conditional on the induced belief, the allocation contains no further information about the agent's type.

Formally, the proof of this result proceeds as follows (see Proposition B.3). Suppose that the principal offers $\mathbf{M}_t$ in period t. The principal's belief about the agent's type, together with the agent's reporting strategy and the mechanism $\beta^{\mathbf{M}_t}$ induces a joint distribution $\mathbb{P}$ over $\Theta \times S^{\mathbf{M}_t} \times A \times \Delta(\Theta)$.[17] Since conditional on the induced posterior, $(s_t, a_t) \in S^{\mathbf{M}_t} \times A$ carries no further information about the agent's type, this allows us to "split" the mechanism into a transition probability $\tilde{\beta}$ from Θ to $\Delta(\Theta)$, a transition probability α from $\Delta(\Theta)$ to A, and a transition probability ω from $\Delta(\Theta) \times A$ to $S^{\mathbf{M}_t}$. The transition probability α plays the first role of the output message and highlights the importance of allowing the principal to offer randomized allocations.[18] The transition probability ω corresponds to the public randomization device: By Kuratowski's theorem we can always embed $S^{\mathbf{M}_t}$ into $[0,1]$ (see Parthasarathy (2005)).[19]

Three conceptual insights arise from this result. First, when the mechanism is canonical, the principal can separate the design of the information that the mechanism encodes about the agent's type from the design of the allocation. Second, the allocation has to be measurable with respect to the information generated by the mechanism: The more the principal desires to tailor the allocation to the agent's type, the more he has to learn about the agent's type through the mechanism.[20] Third, it provides a microfoundation for the public randomization device in the canonical game: it represents the principal's attempt to coordinate play in the mechanism-selection game.

---

[17]Formally, assume that the mechanism $\mathbf{M}_t$ is such that $M^{\mathbf{M}_t} = \Theta$ and the agent truthfully reports her type. Let $\mu \in \Delta(\Theta)$ denote the principal's *prior* belief at h^t. Then, the posterior beliefs μ' satisfy $\int_\Theta \int_{\tilde{S} \times \tilde{A}} \mu'(\tilde{\Theta}|s,a) \beta^{\mathbf{M}_t}(d(s,a)|\theta) \mu(d\theta) = \int_{\tilde{\Theta}} \beta^{\mathbf{M}_t}(\tilde{S} \times \tilde{A}|\theta) \mu(d\theta)$ for any measurable subsets $\tilde{\Theta}$, $\tilde{S}$, $\tilde{A}$ of Θ, $S^{\mathbf{M}_t}$, and A. Note that the posterior beliefs define a transition probability from $S^{\mathbf{M}_t} \times A$ to $\Delta(\Theta)$. Denote it by T. Then, the joint distribution $\mathbb{P}$ is defined by $\mathbb{P}(\tilde{\Theta} \times \tilde{S} \times \tilde{A} \times \tilde{U}) = \int_{\tilde{\Theta}} \int_{\tilde{S} \times \tilde{A}} \mathbb{1}[T(s,a) \in \tilde{U}] \beta^{\mathbf{M}_t}(d(s,a)|\theta) \mu(d\theta)$, where $\tilde{U}$ denotes a measurable subset of $\Delta(\Theta)$.

[18]Strausz (2003) also stresses the importance of allowing for randomized allocations for the standard revelation principle to hold.

[19]One can then apply the integral transform theorem to make the distribution $U[0,1]$.

[20]Contrast this with the case in which the principal has commitment, where we write a mechanism as a menu of options, one for each type of the agent. We do this even if the optimal mechanism offers the same allocation to a set of agent types. When the principal has commitment, it is irrelevant whether the allocation reveals more information beyond the set of types that receive that allocation, since additional information can always be ignored. Under limited commitment, however, this is not the case and the principal in general trades off tailoring the allocation to the agent's type and the information that is learned through this.

Bayes' rule and participation: Underlying the previous step is the assumption that the beliefs associated with the output messages are determined via Bayes' rule. In particular, the principal is never surprised by any output message he observes. To achieve this we show in Proposition B.2 that we can "eliminate" from the mechanism all input messages that are used only by types to whom the principal assigns 0 probability. This, of course, may change the participation decision for these types, which is why Theorem 1 only guarantees participation for those types in the support of the principal's beliefs.[21]

Finally, we show in Proposition B.4 that without loss of generality the agent participates in the mechanism whenever her type is in the support of the principal's beliefs. The logic is similar to that in the case of commitment: Whatever the agent obtains when she does not participate can be replicated by making her participate. However, there is a caveat: When the agent does not participate, her outcome is an allocation for today and a continuation mechanism for tomorrow. Therefore, we must guarantee that, when the agent participates, the principal still offers the same continuation as when she did not participate. We rely here on the mapping between output messages and posterior beliefs and the public randomization device. The first allows us to identify which output message one should associate the types that chose not to participate: the one that corresponds to the principal's updated belief conditional on nonparticipation.[22] The public randomization device allows us to replicate the distribution over continuations the agent faces in the PBE of the mechanism-selection game for those types that found it optimal to randomize between participating and not participating in the PBE of the mechanism-selection game.

The arguments so far describe why the mechanism-selection game implements no more PBE outcomes than the canonical game. While the canonical game has fewer deviations for the principal, it follows from our proof that the canonical game cannot sustain more PBE outcomes than the mechanism-selection game. The reason for this is that our construction applies to each history in the mechanism-selection game, not only those that are on the path of the equilibrium assessment under consideration. This essentially shows that whatever deviation the principal could entertain, he can also achieve it employing canonical mechanisms.[23] Thus, the smaller

---

[21]If an agent's type has zero probability at a specific public history, this means that she can only reach it through a deviation from σ_A. This change to the mechanism actually makes the deviation less attractive, and hence this "disincentivizes" the agent from deviating in the first place.

[22]Starting from an equilibrium in which the mechanism is rejected with positive probability, this belief is also determined via Bayes' rule

[23]This stands in contrast to the literature on the informed principal and on competing principals, where oftentimes revelation principle-style arguments apply on the path of play, but not off the path.

set of deviations in the canonical game fails to provide the principal with "more commitment."

4   EXAMPLE: SALE OF A DURABLE GOOD

We now apply Theorem 1 to study the sale of a durable good in two periods with a continuum of types. The advantage of this setting is that we know what the seller's optimal mechanism is within the set of mechanisms considered previously in the literature on limited commitment: Skreta (2006) shows that posted prices are optimal. The main result of this section, Proposition 1 shows that once we allow the seller to select from a richer set of mechanisms, posted prices may no longer be optimal. In particular, we show that *rationing* may dominate posted prices, allowing us to draw connections with the literature in industrial organization that studies alternative strategies to posted prices for durable good monopolists (e.g., Denicolo and Garella (1999); Gilbert and Klemperer (2000); Nocke and Peitz (2007); McAfee and Wiseman (2008)). Indeed, Proposition 1 provides a microfoundation for a mechanism first suggested by Denicolo and Garella (1999) and a new rationale for the use of clearance sales as in Nocke and Peitz (2007).

Formally, we consider the following special case of the model in Section 2. A seller (the principal) and a buyer (the agent) interact over two periods. The seller owns one unit of a durable good and assigns value 0 to it. The buyer has private information: Before her interaction with the seller starts, she observes her valuation $\theta \in \Theta = [\underline{\theta}, \overline{\theta}]$. Let F_1 denote the seller's prior belief over Θ.[24] We assume that F_1 has full support and is such that the *virtual values,* $\varphi(\theta, F_1) \equiv \theta - (1 - F_1(\theta))/f_1(\theta)$ are nondecreasing. An allocation is a pair $(q, x) \in \{0, 1\} \times \mathbb{R} \equiv A$, where q indicates whether the good is sold ($q = 1$) or not ($q = 0$), and x is a payment from the buyer to the seller. If the good is sold in period 1, the game ends.[25] Moreover, if the buyer rejects the mechanism, the good is not sold and no payments are made, that is, $a^* = (0, 0)$. Payoffs are as follows. If in period $t \in \{1, 2\}$, the allocation is (q, x), the flow payoffs are $u(q, x, \theta) = \theta q - x$ and $w(q, x, \theta) = x$. The buyer and the seller share a common discount factor $\delta \in (0, 1)$ and maximize the expected discounted sum of payoffs.

We proceed as follows. First, we show that we can characterize the seller-optimal PBE as the solution to a constrained optimization problem that only involves the seller (Equation 1). This is already in stark contrast to the existing work in mecha-

---

For an eloquent discussion of this, see Peters (2001).

[24]We use the standard cdf notation, F_1, instead of μ_1 to denote the principal's prior belief since unlike in Theorem 1, we are now assuming that Θ is a subset of $\mathbb{R}$.

[25]Formally, $\mathcal{A}((1, x)) = \{0\} \times \mathbb{R}$ and $\mathcal{A}((0, x)) = \{0, 1\} \times \mathbb{R}$.

nism design with limited commitment, which needs to keep track of how the buyer's best response to the seller's mechanism determines the information that the seller obtains from the interaction, which in turn affects the seller's incentives to offer the mechanism in the first place. Second, Proposition 1 characterizes the solution to that program under the restriction that the mechanism induces at most one posterior, F_{2D}, such that the seller does not sell the good when the posterior is F_{2D} (*delay*). It only remains to verify that it is indeed optimal for the seller to choose such a distribution over posteriors. While we do not pursue this here, we conjecture based on our previous work, Doval and Skreta (2020a), that this is indeed the optimal information structure for the seller.

To arrive at the program that characterizes the seller's maximum revenue, we appeal to Theorem 1. First, in what follows, we restrict attention to the canonical game. Second, it is without loss of generality to consider assessments where the buyer's strategy does not depend on the payoff-irrelevant part of the private history. In particular, in period 2, the seller's optimal mechanism only needs to elicit the buyer's payoff relevant type, θ. Let F_2 denote the seller's belief in period 2. Since the seller has commitment in period 2, Proposition 2 in Skreta (2006) implies that the optimal mechanism in period 2 is a posted price *regardless* of the properties of F_2. We denote by $\hat{\theta}_2(F_2)$ a solution to his maximization problem.[26]

Third, it is without loss of generality to consider assessments where (i) the buyer's best response to the seller's optimal choice of mechanism in period 1 is to participate and truthfully report her type with probability 1, and (ii) when the output message is F_2, the seller updates his belief to F_2. Moreover, the assumption of quasilinearity implies that, without loss of generality, the seller does not randomize on the transfers: Below $x(F_2)$ denotes the expected payment conditional on F_2 and $q(F_2) \in [0,1]$ denotes the probability with which the good is sold. Thus, we can write the seller's problem in period 1 as follows:

$$\max_{(q,x,\beta)} \int_{\Theta} \int_{\Delta(\Theta)} \left(x(F_2) + (1 - q(F_2))\delta\hat{\theta}_2(F_2)\mathbb{1}[\theta \geq \hat{\theta}_2(F_2)] \right) \beta(dF_2|\theta)F_1(d\theta) \tag{1}$$

$$\text{s.t.} \begin{cases} (\forall\theta \in \Theta) \quad U(\theta) = \int_{\Delta(\Theta)} (\theta q(F_2) - x(F_2) + (1-q(F_2))\delta u^*(\theta, F_2))\beta(dF_2|\theta) \geq 0 \\ (\forall\theta, \tilde{\theta} \in \Theta)U(\theta) \geq \int_{\Delta(\Theta)} (\theta q(F_2) - x(F_2) + (1-q(F_2))\delta u^*(\theta, F_2))\beta(dF_2|\tilde{\theta}) \\ (\forall\tilde{\Theta} \times \tilde{U} \in \mathcal{B}_{\Theta \otimes \Delta(\Theta)}) \int_{\Theta} \int_{\tilde{U}} F_2(\tilde{\Theta})\beta(dF_2|\theta)F_1(d\theta) = \int_{\tilde{\Theta}} \beta(\tilde{U}|\theta)F_1(d\theta) \end{cases} .$$

That the seller's belief about the buyer's type updates to F_2 when the output message is F_2 appears twice in the above expression: first, in the third constraint, which is

---

[26]It may be that the seller is indifferent among several prices. We determine the tie-breaking rule as a solution to the problem in period 1.

the Bayes plausibility constraint and, second, in the objective function, where the seller's payoff in period 2 when the agent's type is θ and his belief is F_2 corresponds to whether θ buys the good at a price of $\hat{\theta}_2(F_2)$.

The two remaining constraints are the buyer's participation and incentive compatibility constraints. The buyer's payoff in the mechanism, $U(\theta)$, is determined as follows. For each F_2 in the support of $\beta(\cdot|\theta)$, she receives the good with probability $q(F_2)$ and makes a payment of $x(F_2)$; with the remaining probability, there is no trade, and she obtains a continuation payoff, $u^*(\theta, F_2)$, which describes her optimal decision of whether to buy the good at $\hat{\theta}_2(F_2)$. The participation constraint states that the buyer has to earn a payoff of at least 0 by participating. Indeed, since non-participation is a 0 probability event, we can specify that upon rejection of the mechanism the seller believes that the buyer's valuation is $\overline{\theta}$, so that in period 2 the seller chooses a price of $\overline{\theta}$ when the buyer chooses not to participate. The incentive compatibility constraint states that when her type is θ the buyer cannot obtain a higher payoff by reporting that her type is $\tilde{\theta} \neq \theta$. When the buyer reports $\tilde{\theta}$, she obtains a different distribution over output messages $\beta(\cdot|\tilde{\theta})$; however, in period 2, she still chooses optimally whether to buy the good, which explains the term $u^*(\theta, F_2)$.

The three constraints in Equation 1 provide us with a tractable representation of both the buyer's behavior and its impact on the mechanism offered in period 2 via the information that is generated about the buyer's type in period 1. This allows us to characterize the seller-optimal PBE by focusing only on the period 1's seller choice of mechanism, knowing that as long as the mechanism satisfies the constraints we are able to find a buyer's strategy in the game to fully specify the PBE assessment that implements the seller's maximum revenue.[27]

As we show in Appendix D, the incentive constraints deliver the envelope representation of the buyer's payoffs, so we can replace the transfers out of the seller's payoff and reduce Equation 1 to the following program. The seller in period 1 chooses a distribution over posteriors, $P_{\Delta(\Theta)}$, and, for each posterior he induces, a probability of trade $q(F_2)$, to solve

$$\max_{P_{\Delta(\Theta)}, q} \int_{\Delta(\Theta)} \left[q(F_2) \int_{\underline{\theta}}^{\overline{\theta}} \varphi(\theta, F_1) F_2(d\theta) + (1 - q(F_2))\delta \int_{\hat{\theta}_2(F_2)}^{\overline{\theta}} \varphi(\theta, F_1) F_2(d\theta) \right] P_{\Delta(\Theta)}(dF_2),$$

$$(2)$$

---

[27]The same idea applies in our application to the infinite horizon analysis of the sale of the durable good in Doval and Skreta (2020a).

subject to (i) $P_\Delta(\Theta)$ must be Bayes' plausible given F_1 and (ii) a monotonicity condition, which states that, in expectation, higher types must trade with higher probability (see Equation D.4 in Appendix D). Equation 2 describes the seller's payoff in terms of the distribution over posteriors induced by the mechanism. If at posterior F_2 the seller sells the good ($q(F_2) = 1$), he obtains the expected virtual surplus, where the expectation is calculated using F_2, but the virtual values are calculated using F_1. This reflects that the probability with which the seller pays rents to a buyer of type θ is measured by the probability $F_1(\theta)$ that buyer types below θ receive the good. Instead, if at F_2 the seller does not sell the good ($q(F_2) = 0$) he obtains the (discounted) expected virtual surplus of selling the good at price $\hat{\theta}_2(F_2)$. While the posted price in period 2 is optimal with respect to the *posterior* virtual values $\varphi(\theta, F_2)$, it may not be for the *prior* virtual values. This reflects the conflict between the period 1 and period 2 sellers: if they hold different beliefs about the buyer's type, they pay rents with different probabilities, and therefore may prefer different mechanisms.

Instead of fully solving the problem in Equation 1, we restrict attention to mechanisms where the seller induces at most one posterior at which $q(F_2) = 0$ and denote this posterior by F_{2D}. While we do not show that this is optimal, the analysis of this simpler class of mechanisms already expands on what is known about optimal mechanisms for this particular setting. In a slight abuse of notation, let $\beta_D(\theta)$ denote the probability that a buyer of type θ is delayed, i.e., $\beta_D(\theta) = \beta(F_{2D}|\theta)$. Then, the seller's problem reduces to[28]

$$\max_{\beta_D} \int_{\underline{\theta}}^{\overline{\theta}} \varphi(\theta, F_1)(1 - \beta_D(\theta)) F_1(d\theta) + \delta \int_{\hat{\theta}_2(F_{2D})}^{\overline{\theta}} \varphi(\theta, F_1) \beta_D(\theta) F_1(d\theta), \quad (\text{OPT}_{\beta_D})$$

subject to the constraint that β_D is nonincreasing.[29] Proposition 1 shows that the solutions to OPT_{β_D} are indexed by three parameters $\tilde{\theta}_1, \tilde{\theta}_2, \gamma$, with $\gamma \le 1$ so that

$$\beta_D(\theta) = \begin{cases} 0 & \text{if } \theta \ge \tilde{\theta}_2 \\ \gamma & \text{if } \theta \in (\tilde{\theta}_1, \tilde{\theta}_2) \\ 1 & \text{otherwise} \end{cases}. \tag{3}$$

Note that when $\gamma > 0$, the seller observes whether the good is sold in period 1, but not whether the buyer's type is above or below $\tilde{\theta}_1$. Equation 3 encompasses three selling

---

[28]Because the objective function in Equation 2 is linear in F_2 when $q(F_2) = 1$, it is without loss to group all the terms where $q(F_2) = 1$ into one, even if, as we show below, there is more than one posterior at which the seller sets $q(F_2) = 1$.

[29]Under the restriction that there is at most one belief at which the seller sets $q(F_2) = 0$, β_D nonincreasing is equivalent to the monotonicity condition (see Equation D.4).

strategies for durable good monopolists considered in the literature. First, when $\gamma = 0$ or $\tilde{\theta}_1 = \tilde{\theta}_2$, this corresponds to a posted price mechanism, with types above $\tilde{\theta}_1$ obtaining the good in period 1. Under a posted price mechanism we necessarily obtain a decreasing sequence of prices: Conditional on the good not being sold in period 1, the seller learns that the buyer's type is below $\tilde{\theta}_1$, so the period 2 demand, $1 - F_{2D}(\theta)$, is lower than the period 1 demand, $1 - F_1(\theta)$. Second, when $\gamma > 0$ and $\tilde{\theta}_2 = \overline{\theta}$, this corresponds to what Denicolo and Garella (1999) denote as *proportional rationing*. Under proportional rationing, buyer types above $\tilde{\theta}_1$ obtain the good with probability less than 1 in period 1. This allows the seller to induce a stronger demand in period 2 and hence avoid decreasing the period 2 price, at the cost of selling the good less often to high-valuation buyers in period 1. For discount factors close to 1, this cost is small compared to the benefit of inducing higher prices in period 2. Finally, when $\gamma > 0$ and $\tilde{\theta}_2 < \overline{\theta}$, the mechanisms in Equation 3 correspond to what Nocke and Peitz (2007) denote as *clearance sales*: the seller sets two prices in period 1, one at which the good is sold with probability 1 and another at which it is rationed. Under clearance sales, the seller cannot induce higher prices in period 2 than in period 1, since $\tilde{\theta}_2 < \overline{\theta}$ implies that the period 2 demand is lower than the period 1 demand. The proof of Proposition 1 suggests, however, why such a mechanism may be optimal: It allows the period 1 seller to satisfy the period 2 sequential rationality constraint while simultaneously maximizing the probability with which high buyer types are served in period 1.

Proposition 1. *Any solution to* OPT_{β_D} *is as in Equation 3. Moreover,*

1. *If* $\hat{\theta}_2(F_{2D}) < \tilde{\theta}_1$, *it is optimal to set* $\gamma = 0$.

2. *Otherwise, the seller chooses* $\tilde{\theta}_1, \tilde{\theta}_2, \gamma$ *subject to the constraint that* $\tilde{\theta}_1 \leq \hat{\theta}_2(F_{2D})$.

Proposition 1 shows that we can reduce the search for the solution to OPT_{β_D} to the comparison between (i) the revenue-maximizing posted price mechanism, and (ii) the revenue-maximizing "rationing" mechanism such that the period 2 price excludes at least as many buyer types as the period 1 mechanism.

Proposition 1 clarifies and provides a foundation for the analysis in Denicolo and Garella (1999). The authors consider a two-period model of limited commitment where in each period the seller can choose both a price at which to sell the good, and a probability, $\gamma(\theta)$, with which a buyer of type θ, who wants to buy the good at the posted price, receives the good. Implicit in their analysis is that the seller only observes whether the good is sold but not whether the buyer is willing to buy the

good at the posted price. They focus on the case of proportional rationing and provide conditions under which it dominates posted prices. While thought-provoking, their analysis neither shows how the seller can actually implement this alternative mechanism, nor that this is the optimal mechanism when the seller only observes whether the good is sold.[30] Our analysis allows us to interpret the seller's strategy set in Denicolo and Garella (1999) as a special case of a mechanism that induces at most one posterior at which the seller sets $q(F_2) = 0$. While Proposition 1 shows that proportional rationing is among the candidate mechanisms, it highlights that clearance sales may also be optimal. Nocke and Peitz (2007) show that clearance sales may dominate posted prices in a model where, while the seller has commitment, he faces uncertainty about demand. As we explain below, clearance sales arise as a candidate mechanism under limited commitment as an instrument for the period 1 seller to deal with the period 2 sequential rationality constraint.

The proof of Proposition 1 proceeds in two steps. To explain them, we introduce one final piece of notation. For any incentive-compatible β_D, let $\hat{\theta}$ denote the highest buyer type such that $\beta_D(\theta) = 1$. The first step shows that *any* incentive-compatible β_D that induces a period 2 price below $\hat{\theta}$ is dominated by a posted price mechanism.[31] That is, if rationing is not conducive to more exclusion in period 2 than in period 1, the seller prefers to post a price in period 1 to reap the benefits of trading with the high-valuation buyers. Thus, for a non-posted price mechanism to be optimal it must be that $\hat{\theta} \leq \hat{\theta}_2(F_{2D})$.

The second step is to show that for any incentive-compatible policy β_D such that $\hat{\theta}_2(F_{2D}) \geq \hat{\theta}$ there exists an alternative incentive-compatible β'_D of the form described in Equation 3 which induces the same price in period 2, but leads to a higher revenue in period 1. This is where the assumption that virtual values, $\varphi(\theta, F_1)$, are monotone matters the most. Whenever β_D is not as in Equation 3, we show that the seller can change the probability with which the buyer's different types are delayed to in-

---

[30]Indeed, in footnote 3 Denicolo and Garella (1999) write the following:

> "One issue we do not analyze is the determination of the optimal rationing function $\gamma(v)$. To implement such a function, the monopolist would have to know the customer's type-which presumably would allow him to engage in perfect static price discrimination. Alternatively, a self-selection constraint should be imposed if rationing is not proportional."

[31]Denicolo and Garella (1999) make a similar observation in Proposition 1 of their paper. However, as we explain in Appendix D, they do not account for the incentive costs of implementing β_D when they argue that a posted price dominates. Indeed, a posted price mechanism leaves more rents to the buyer than other nonincreasing policies β_D. Our proof shows that despite this, a posted price dominates when $\hat{\theta}_2(F_{2D}) < \hat{\theta}$.

crease the probability with which he sells to high types and reduce the probability with which he sells to low types in period 1, while also implementing the same period 2 price. Thus, the candidate mechanisms in Equation 3 satisfy the period 2 sequential rationality constraint at the lowest cost for the period 1 seller by allowing the period 1 seller to maximize the probability with which he trades with high buyer types.

Figure 3 illustrates the seller's revenue from the best posted price ($\gamma = 0$) and the optimal mechanism as a function of the discount factor in the case in which F_1 is $U[0,1]$. As anticipated above, proportional rationing and clearance sales can only be optimal when the discount factor is close to 1. Under this parametrization, proportional rationing is never optimal. Note that when $\delta = 0$ and $\delta = 1$, posted prices achieve the commitment payoff, which is 0.25 when F_1 is $U[0,1]$.
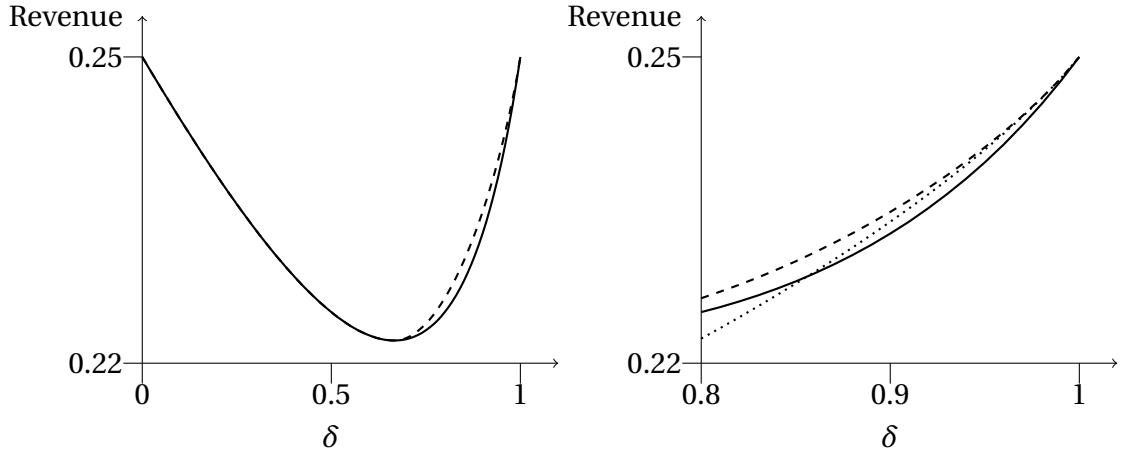

Figure 3: Seller's optimal revenue when F_1 is $U[0,1]$: posted prices (solid), optimal mechanism (dashed), proportional rationing (right panel, dotted). The right panel reproduces the left for $\delta \in (0.8,1)$

We conclude Section 4 by discussing the connection with previous analysis of this problem in the literature, particularly its difference with the main result in Skreta (2006). A reader interested in the case in which the agent's information evolves over time and the application to nonlinear pricing can proceed to Section 5 without loss of continuity.

Remark 2. *[Comparison with Skreta (2006)] To understand the difference between the result in Proposition 1 and that in Skreta (2006), it is instructive to compare the incentive constraints in Equation 1 to those implied by mechanisms where the seller observes the buyer's choice of input message as in, for instance, Hart and Tirole (1988);*

24

Laffont and Tirole (1988); Bester and Strauss (2001); Skreta (2006). While not expressed in the language of type reports or beliefs, the incentive constraints in Skreta (2006) require that for each F_2 in the support of $\beta(\cdot|\theta)$, the buyer prefers the tuple $(q(F_2), x(F_2), u^*(\theta, F_2))$ to any other tuple $(q(F_2'), x(F_2'), u^*(\theta, F_2'))$ in the mechanism. In particular, the buyer must be indifferent between any two tuples that she chooses with positive probability. Contrast this with the incentive constraints in Equation 1, where the buyer does not necessarily have to be indifferent between the tuples $(q(F_2), x(F_2), u^*(\theta, F_2))$ in the support of $\beta(\cdot|\theta)$, although in expectation, the lottery she faces over such tuples under truthtelling must be better than the one she faces by lying. Indeed, when posted prices fail to be optimal, the seller exploits the weaker incentive constraints in Equation 1: Buyer types in $(\tilde{\theta}_1, \tilde{\theta}_2)$ are not indifferent between receiving the good in period 1 at the rationing price and receiving the good in period 2 at price $\hat{\theta}_2(F_{2D})$.

Note, however, that the property that the seller attains a higher payoff when he chooses canonical mechanisms than that he obtains in the two-period version of the model in Skreta (2006) is an artifact of the two-period model. Indeed, for any belief that the seller may have in period 2, his payoff is the same in both models conditional on the good not being sold in period 1. Proceeding via backward induction, the seller in our model chooses from a larger set of mechanisms in period 1, while facing the exact same continuation values. For longer horizons, however, the comparison of the seller's payoffs in the two models is not obvious since the larger set of canonical mechanisms also implies that the seller has a larger set of deviations in our model than in the model in Skreta (2006).

The previous discussion highlights that under limited commitment the principal may benefit from employing mechanisms where the output message (and hence, the allocation) does *not* reveal the input message that the agent submitted into the mechanism. This is in contrast with the standard revelation principle for the case of commitment when the principal faces a privately informed agent (adverse selection): As we explained in the introduction, it follows from the result in Myerson (1982) that it is without loss of generality in that case to consider mechanisms where the principal learns the input message from observing the realization of the output message. Instead, Myerson (1982) shows that adding "noise" to the communication may be essential when the principal also faces an agent whose actions are not contractible (moral hazard). Indeed, it may be beneficial to pool in the same output message different types of the privately informed agent to incentivize the agent whose action is not contractible to follow the recommendation. Mechanism design with

limited commitment is closer to the hybrid model of adverse selection and moral hazard in Myerson (1982) than it is to the model of pure adverse selection. Indeed, note that in a given period the principal faces, in a sense, *two* agents whose incentives he needs to manage: the privately informed agent (adverse selection) and his future self, whose choice of mechanism is not contractible (moral hazard). That is, today's principal needs to elicit the agent's information while simultaneously ensuring that his future behavior is sequentially rational. In the same way that output messages are key in the presence of moral hazard in Myerson (1982), they feature prominently in our framework.

Nevertheless, employing mechanisms where the output message does not reveal the input message the agent submitted into the mechanism might come at a cost: As we have argued, the allocation has to be measurable with respect to the information that the mechanism generates about the agent. Thus, the principal trades off the short-term gains of tailoring the allocation to the agent's type against the long-term costs in terms of his future sequential rationality constraints of releasing too much information about the agent. Contrast this with a model where, while the principal is allowed to choose a mechanism in each period, he does not observe the outcome of the mechanism, but only observes whether the relationship with the agent is still ongoing. In this case, within each period the principal could perfectly tailor the allocation to the agent's type without having to learn this information. Thus, the principal can potentially implement the commitment solution, unless the event that the relationship with the agent is still ongoing reveals information about the agent's type. An example of such a model is that in Correia-da Silva (2020), who analyzes a version of the model in this section under the assumption that the seller *only* observes whether the good is unsold in period 2 and shows that in that case the mechanisms in Proposition 1 are optimal. The same way that our two-period model provides an upper-bound on the seller's payoff in the model of Skreta (2006), the model of Correia-da Silva (2020) provides an upper bound on the seller's payoff in our two-period model.

5 MARKOV ENVIRONMENTS

The case in which the agent's private information is fully persistent is the cornerstone of the literature on mechanism design with limited commitment for good reason. There is, in a sense, a fixed amount of information to be learned about the agent and the principal needs to trade off the short-term gains and the future losses from the use of this information. In contrast, when the agent's type is less than fully persistent, nature renews the principal's uncertainty about the agent's type. Thus, as observed

in Battaglini (2005), dynamic mechanisms are more often time-consistent when the agent's type evolves over time.

Nevertheless, the case in which the agent's information evolves over time is relevant for applications as evidenced in, amongst other contributions, the recent public finance applications of Farhi and Werning (2013); Kapička (2013); Stantcheva (2015) and the burgeoning literature in dynamic mechanism design (see Pavan et al. (2014) for references). Thus, we show in this section that Theorem 1 extends to the case in which the agent's private information evolves over time. We do so for a version of what Pavan et al. (2014) denote as Markov environments, which we define next.[32]

The environment is *Markov* if the following holds. First, the agent's private information is described by a nonhomogenous Markov process with states in Θ and transitions $F_t : \Theta \times A \mapsto \Delta(\Theta)$ so that $F_t(\tilde{\Theta}|\theta_{t-1}, a_{t-1})$ describes the probability that the agent's type in period t is in $\tilde{\Theta}$ when her type in $t-1$ is θ_{t-1} and the allocation is a_{t-1}.[33] Second, the principal and the agent's payoffs are time separable and their period-t flow payoff only depends on the current allocation and the agent's period t type. That is, if $(a_t, \theta_t)_{0 \leq t \leq T}$ describes the allocations and agent's private information through period T, then

$$W((a_t, \theta_t)_{0 \leq t \leq T}) = \sum_{t=0}^{T} \delta^t w_t(a_t, \theta_t), \quad U((a_t, \theta_t)_{0 \leq t \leq T}) = \sum_{t=0}^{T} \delta^t u_t(a_t, \theta_t),$$

denote the payoffs to the principal and the agent, respectively. Everything else is as in the model in Section 2. However, to keep track of the agent's time-varying private information, we index the agent's private history differently. Namely, the agent's private history through period t, h_A^t corresponds to a sequence $(\theta_0, \mathbf{M}_0, p_0, m_0, s_0, a_0, \ldots, \theta_{t-1}, \mathbf{M}_{t-1}, p_{t-1}, m_{t-1}, s_{t-1}, a_{t-1})$. Since the agent knows her type at the beginning of period t, we index her information sets by (h_A^t, θ_t). To simplify notation, we write the agent's behavioral strategy $\sigma_{At}(h_A^t, \theta_t, \mathbf{M}_t) = (\pi_t(h_A^t, \theta_t, \mathbf{M}_t), r_t(h_A^t, \theta_t, \mathbf{M}_t)) \in [0,1] \times \Delta(M^{\mathbf{M}_t})$.

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 at the end of period t; the second is the belief he holds at the beginning of period $t+1$, after applying F_{t+1}. We denote the former by μ_{t+1}, anticipating that, as in Theorem 1, this is the belief that will be used as an output message.[34] We denote the latter by ν_{t+1}: This is the belief

---

[32] We comment at the end of this section on how our results extend outside Markov environments.

[33] Note that assuming that the set of states is time-invariant is without loss of generality.

[34] See Ely (2017) for another model where the same choice is made.

that is used to determine which types have positive probability in period $t+1$. That is, $\mu_{t+1}(h_A^t, \theta_t, m, z | h^t, z)$ is the probability that the principal assigns to the agent being at information set (h_A^t, θ_t) and sending message m at the end of period t when z is the outcome of the interaction in period t, while $v_{t+1}(h_A^t, \theta_t, m, z, \theta_{t+1} | h^t, z) = \mu_{t+1}(h_A^t, \theta_t, m, z | h^t, z) F_{t+1}(\theta_{t+1} | \theta_t, a_t)$ is the probability that the principal assigns to the agent being at information set (h_A^t, θ_t) in period t *and* her type being θ_{t+1} in period $t+1$ when the outcome of the interaction in period t is z, where a_t is the allocation in period t consistent with z.

To introduce Theorem 2, we need one last piece of notation. Let $\mathcal{E}_{\mathcal{M}}^*$ and $\mathcal{E}^*$ denote the set of equilibrium payoffs of the mechanism-selection game and the canonical game, respectively.

Theorem 2. *The mechanism-selection game and the canonical game have the same set of equilibrium payoffs, i.e., $\mathcal{E}_{\mathcal{M}}^* = \mathcal{E}^*$ for any collection $\mathcal{M}$ with which we endow the principal.*

Moreover, for any equilibrium payoff in $\mathcal{E}_{\mathcal{M}}^*$, there exists a PBE assessment $(\sigma_P, \sigma_A, \mu)$ of G that achieves the same payoff and satisfies the following properties:

1. *The agent's strategy depends on her current payoff-relevant type and the public history. That is, for all periods t, all public histories h^t, all $h_A^t, \overline{h}_A^t \in H_A^t(h^t)$, and all types $\theta_t \in \Theta$, $\sigma_{At}(h_A^t, \theta_t, \mathbf{M}_t) = \sigma_{At}(\overline{h}_A^t, \theta_t, \mathbf{M}_t)$*

2. *For all public histories h^t, if θ_t is in the support of $v_t(\cdot | h^t)$, then the agent participates in the mechanism offered by the principal at that history and with probability one truthfully reports her current type,*

3. *For all public histories h^t, if the mechanism offered by the principal at h^t outputs a posterior μ', the principal's updated equilibrium beliefs about the agent coincide with μ'. That is, for all measurable subsets $\tilde{\Theta}, \tilde{U}, \tilde{A}$ of $\Theta, \Delta(\Theta)$, and A,*

$$\int_\Theta \int_{\tilde{U}} \int_{\tilde{A}} \mu_{t+1}(\tilde{\Theta} | h^t, z_{(\mu', a_t)}) \alpha^{\mathbf{M}_t}(da_t | \mu') \beta^{\mathbf{M}_t}(d\mu' | \theta) v_t(d\theta | h^t) =$$
$$= \int_{\tilde{\Theta}} \int_{\tilde{U}} \alpha^{\mathbf{M}_t}(\tilde{A} | \mu') \beta^{\mathbf{M}_t}(d\mu' | \theta) v_t(d\theta | h^t) = \int_\Theta \int_{\tilde{U}} \int_{\tilde{A}} \mu'(\tilde{\Theta}) \alpha^{\mathbf{M}_t}(da_t | \mu') \beta^{\mathbf{M}_t}(d\mu' | \theta) v_t(d\theta | h^t).$$

Two remarks are in order. First, the history independence in Theorem 2 is, in a sense, stronger than that of Theorem 1. The agent does not condition her strategy on either her past communication or her past payoff-relevant types. This is where we more

prominently employ the restriction to Markov environments. This 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 $(\theta_1,\ldots,\theta_t)$. Note that the stronger form of history independence also implies that contrary to Theorem 1 the mechanism-selection and canonical games implement the same set of payoffs but not necessarily the same equilibrium outcomes.[35] Second, the separation between the beliefs that the output message represents (μ_{t+1}) and the beliefs that the principal uses in the next period to select mechanisms (ν_{t+1}) highlights that the principal in period t attempts to design his prior for period $t+1$ (which nature will then push forward using F_{t+1}.)

The proof of Theorem 2 is in Appendix E. Except for the proof of the stronger form of history independence, the remaining steps closely follow the proof of Theorem 1, and thus, we omit them. Moreover, since the extension to the case in which the set Θ is uncountable or the support of $\beta^{\mathbf{M}_t}$ is uncountable should be immediate from the proof of Theorem 1, we only present the proof under the assumption that Θ is finite and the principal can only choose mechanisms, $\mathbf{M}$, such that for all $m \in M^{\mathbf{M}}$, the support of $\beta^{\mathbf{M}}(\cdot|m)$ is finite.

REFERENCES

ABREU, D., D. PEARCE, AND E. STACCHETTI (1990): "Toward a theory of discounted repeated games with imperfect monitoring," *Econometrica*, 58, 1041–1063.

ACHARYA, A. AND J. ORTNER (2017): "Progressive learning," *Econometrica*, 85, 1965–1990.

AKBARPOUR, M. AND S. LI (2020): "Credible auctions: A trilemma," *Econometrica*, 88, 425–467.

ALIPRANTIS, C. AND K. BORDER (2013): *Infinite Dimensional Analysis: A Hitchhiker's Guide*, Springer-Verlag Berlin and Heidelberg GmbH & Company KG.

ATHEY, S. AND K. BAGWELL (2008): "Collusion with persistent cost shocks," *Econometrica*, 76, 493–540.

AUMANN, R. J. (1961): "Borel structures for function spaces," *Illinois Journal of Mathematics*, 5, 614–630.

---

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

——— (1964): "Mixed and behavior strategies in infinite extensive games," Tech. rep., Princeton University.

BATTAGLINI, M. (2005): "Long-term contracting with Markovian consumers," *American Economic Review*, 95, 637–658.

BECCUTI, J. AND M. MÖLLER (2018): "Dynamic adverse selection with a patient seller," *Journal of Economic Theory*, 173, 95–117.

BEN-PORATH, E., E. DEKEL, AND B. L. LIPMAN (2019): "Mechanisms with evidence: Commitment and robustness," *Econometrica*, 87, 529–566.

BERTSEKAS, D. P. AND S. E. SHREVE (1978): "Stochastic optimal control, volume 139 of Mathematics in Science and Engineering," .

BEST, J. AND D. QUIGLEY (2017): "Persuasion for the long run," *Working Paper*.

BESTER, H. AND R. STRAUSZ (2001): "Contracting with imperfect commitment and the revelation principle: the single agent case," *Econometrica*, 69, 1077–1098.

——— (2007): "Contracting with imperfect commitment and noisy communication," *Journal of Economic Theory*, 136, 236–259.

BISIN, A. AND A. A. RAMPINI (2006): "Markets as beneficial constraints on the government," *Journal of public Economics*, 90, 601–629.

BULOW, J. I. (1982): "Durable-goods monopolists," *Journal of Political Economy*, 90, 314–332.

BURGUET, R. AND J. SAKOVICS (1996): "Reserve prices without commitment," *Games and Economic Behavior*, 15, 149–164.

CAILLAUD, B. AND C. MEZZETTI (2004): "Equilibrium reserve prices in sequential ascending auctions," *Journal of Economic Theory*, 117, 78–95.

CORREIA-DA SILVA, J. (2020): "Optimal priority pricing by a durable goods monopolist," Tech. rep.

CRAUEL, H. (2002): *Random probability measures on Polish spaces*, vol. 11, CRC press.

DASGUPTA, P., P. HAMMOND, AND E. MASKIN (1979): "The implementation of social choice rules: Some general results on incentive compatibility," *The Review of Economic Studies*, 46, 185–216.

DEB, R. AND M. SAID (2015): "Dynamic screening with limited commitment," *Journal of Economic Theory*, 159, 891–928.

DENICOLO, V. AND P. G. GARELLA (1999): "Rationing in a durable goods monopoly," *The RAND Journal of Economics*, 44–55.

DOVAL, L. AND V. SKRETA (2020a): "Optimal mechanism for the sale of a durable good," *arXiv preprint arXiv:1904.07456*.

——— (2020b): "Simple proof of Theorem 1 in Doval and Skreta (2020)," Click here.

——— (2020c): "Supplement to "Mechanism Design with Limited Commitment"," Click here.

ELY, J. C. (2017): "Beeps," *American Economic Review*, 107, 31–53.

FARHI, E., C. SLEET, I. WERNING, AND S. YELTEKIN (2012): "Non-linear capital taxation without commitment," *Review of Economic Studies*, 79, 1469–1493.

FARHI, E. AND I. WERNING (2013): "Insurance and taxation over the life cycle," *Review of Economic Studies*, 80, 596–635.

FIOCCO, R. AND R. STRAUSZ (2015): "Consumer standards as a strategic device to mitigate ratchet effects in dynamic regulation," *Journal of Economics & Management Strategy*, 24, 550–569.

FUDENBERG, D. AND J. TIROLE (1991): "Perfect Bayesian equilibrium and sequential equilibrium," *journal of Economic Theory*, 53, 236–260.

GERARDI, D. AND L. MAESTRI (2020): "Dynamic contracting with limited commitment and the ratchet effect," *Theoretical Economics*, 15, 583–623.

GIBBARD, A. (1973): "Manipulation of voting schemes: a general result," *Econometrica*, 587–601.

GILBERT, R. J. AND P. KLEMPERER (2000): "An equilibrium theory of rationing," *The RAND Journal of Economics*, 1–21.

GOLOSOV, M. AND L. IOVINO (2016): "Social Insurance, Information Revelation, and Lack of Commitment," *Working Paper*.

GUL, F., H. SONNENSCHEIN, AND R. WILSON (1986): "Foundations of dynamic monopoly and the coase conjecture," *Journal of Economic Theory*, 39, 155 – 190.

HART, O. D. AND J. TIROLE (1988): "Contract renegotiation and Coasian dynamics," *The Review of Economic Studies*, 55, 509–540.

HART, S., I. KREMER, AND M. PERRY (2017): "Evidence games: Truth and commitment," *American Economic Review*, 107, 690–713.

HIRIART, Y., D. MARTIMORT, AND J. POUYET (2011): "Weak enforcement of environmental policies: a tale of limited commitment and limited fines," *Annals of Economics and Statistics/Annales d'Économie et de Statistique*, 25–42.

KALLENBERG, O. (2017): *Random measures, theory and applications*, Springer.

KAMENICA, E. AND M. GENTZKOW (2011): "Bayesian persuasion," *American Economic Review*, 101, 2590–2615.

KAPIČKA, M. (2013): "Efficient allocations in dynamic private information economies with persistent shocks: A first-order approach," *Review of Economic Studies*, 80, 1027–1054.

LAFFONT, J.-J. AND J. TIROLE (1987): "Comparative statics of the optimal dynamic incentive contract," *European Economic Review*, 31, 901–926.

——— (1988): "The dynamics of incentive contracts," *Econometrica*, 1153–1175.

LEHMANN, E. ET AL. (1988): "Comparing Location Experiments," *The Annals of Statistics*, 16, 521–533.

LIPNOWSKI, E. AND D. RAVID (2020): "Cheap talk with transparent motives," *Econometrica*, 88.

LIU, Q., K. MIERENDORFF, X. SHI, AND W. ZHONG (2019): "Auctions with limited commitment," *American Economic Review*, 109, 876–910.

MCADAMS, D. AND M. SCHWARZ (2007): "Credible sales mechanisms and intermediaries," *American Economic Review*, 97, 260–276.

MCAFEE, R. P. AND D. VINCENT (1997): "Sequentially optimal auctions," *Games and Economic Behavior*, 18, 246–276.

MCAFEE, R. P. AND T. WISEMAN (2008): "Capacity choice counters the Coase conjecture," *The Review of Economic Studies*, 75, 317–331.

MYERSON, R. B. (1979): "Incentive compatibility and the bargaining problem," *Econometrica*, 61–73.

——— (1982): "Optimal coordination mechanisms in generalized principal–agent problems," *Journal of Mathematical Economics*, 10, 67–81.

NOCKE, V. AND M. PEITZ (2007): "A theory of clearance sales," *The Economic Journal*, 117, 964–990.

PARTHASARATHY, K. R. (2005): *Probability measures on metric spaces*, vol. 352, American Mathematical Soc.

PAVAN, A., I. SEGAL, AND J. TOIKKA (2014): "Dynamic mechanism design: A myersonian approach," *Econometrica*, 82, 601–653.

PETERS, M. (2001): "Common agency and the revelation principle," *Econometrica*, 69, 1349–1372.

SERFOZO, R. (1982): "Convergence of Lebesgue integrals with varying measures," *Sankhyā: The Indian Journal of Statistics, Series A*, 380–402.

SKRETA, V. (2006): "Sequentially optimal mechanisms," *The Review of Economic Studies*, 73, 1085–1111.

——— (2015): "Optimal auction design under non-commitment," *Journal of Economic Theory*, 159, 854–890.

SLEET, C. AND Ş. YELTEKIN (2008): "Politically credible social insurance," *Journal of monetary Economics*, 55, 129–151.

STANTCHEVA, S. (2015): "Learning and (or) doing: Human capital investments and optimal taxation," Tech. rep., National Bureau of Economic Research.

STOKEY, N. L. (1981): "Rational expectations and durable goods pricing," *The Bell Journal of Economics*, 112–128.

STRAUSZ, R. (2003): "Deterministic mechanisms and the revelation principle," *Economics Letters*, 79, 333–337.

STRULOVICI, B. (2017): "Contract negotiation and the Coase conjecture: A strategic foundation for renegotiation-proof contracts," *Econometrica*, 85, 585–616.

TULCEA, C. I. (1949): "Mesures dans les espaces produits," *Atti Acad. Naz. Lincei Rend. Cl Sci. Fis. Mat. Nat*, 8.

VARTIAINEN, H. (2013): "Auction design without commitment," *Journal of the European Economic Association*, 11, 316–342.

A COLLECTED DEFINITIONS AND NOTATION

Histories and strategies: We review in Appendix C how to endow the set $\mathcal{M}$ with a measurable structure so that it is a standard Borel space (Aumann (1961)). In what follows, we take this as given.

Denote by M_t the set $\cup_{i \in \mathcal{I}_t} M_{ti} \cup \{\varnothing\}$ and similarly let S_t denote the set $\cup_{i \in \mathcal{I}_t} S_{ti} \cup \{\varnothing\}$. Let $\mathsf{Z}_t = \mathcal{M} \times \{0,1\} \times \mathsf{S}_t \times A$ and $\mathsf{Z}_{A,t} = \mathcal{M} \times \{0,1\} \times \mathsf{M}_t \times \mathsf{S}_t \times A$. By construction, Z_t and $\mathsf{Z}_{A,t}$ are standard Borel spaces.

Public histories through period t are $H^t = \times_{\tau=0}^{t-1} \mathsf{Z}_\tau$ and private histories are $\Theta \times H_A^t = \Theta \times \times_{\tau=0}^{t-1} \mathsf{Z}_{A,\tau}$. Note that the information sets of the principal can be described by a measurable function $\zeta_{Pt} : \Theta \times H_A^t \mapsto H^t$ where ζ_{Pt} is the projection of $\Theta \times H_A^t$ onto $\times_{\tau=0}^{t-1} \mathsf{Z}_\tau$. Similarly, the information sets of the agent can be described by a measurable function $\zeta_{At} : \Theta \times H_A^t \mapsto \Theta \times H_A^t$ where ζ_{At} is simply the identity.

Induced distributions and payoffs: Fix a mechanism $\mathbf{M}_t$. The strategy profile together with the mechanism define a transition probability from $\Theta \times H_A^t \times \mathcal{M}$ to $\mathsf{M}_t \times \mathsf{S}_t \times A$ as follows:

$$\rho^{\sigma_A}(\tilde{M} \times \tilde{S} \times \tilde{A} | \theta, h_A^t, \mathbf{M}_t) = \int_{\tilde{M}} \beta^{\mathbf{M}_t}(\tilde{S} \times \tilde{A} | m) r_t(\theta, h_A^t, \mathbf{M}_t)(dm),$$

for all measurable subsets $\tilde{M}, \tilde{S}, \tilde{A}$ of $\mathsf{M}_t, \mathsf{S}_t, A$.

Given $\sigma = (\sigma_P, \sigma_A)$ and (θ, h_A^t) we can define transition probabilities from $\Theta \times H_A^t$ to $\mathsf{Z}_{A,t}$ as follows:

$$\kappa_t^\sigma(\tilde{\mathcal{M}} \times \{0\} \times \{\varnothing\} \times \{\varnothing\} \times \{a^*\} | \theta, h_A^t) = \int_{\{x : \sigma_{Pt}(\zeta_{Pt}(\theta, h_A^t), x) \in \tilde{\mathcal{M}}\}} (1 - \pi_t(\theta, h_A^t, \sigma_P(\zeta_{Pt}(\theta, h_A^t), x))) \lambda(dx)$$

for any measurable subset $\tilde{\mathcal{M}}$ of $\mathcal{M}$, and

$$\kappa_t^\sigma(\tilde{\mathcal{M}} \times \{1\} \times \tilde{M} \times \tilde{S} \times \tilde{A} | \theta, h_A^t)$$
$$= \int_{\{x : \sigma_{Pt}(\zeta_{Pt}(\theta, h_A^t), x) \in \tilde{\mathcal{M}}\}} \rho^{\sigma_A}(\tilde{M} \times \tilde{S} \times \tilde{A} | \theta, h_A^t, \sigma_{Pt}(\zeta_{Pt}(\theta, h_A^t), x)) \pi_t(\theta, h_A^t, \sigma_{Pt}(\zeta_{Pt}(\theta, h_A^t), x)) \lambda(dx)$$

on the measurable rectangles $\tilde{\mathcal{M}} \times \tilde{M} \times \tilde{S} \times \tilde{A} \in \mathcal{B}_{\mathcal{M} \otimes \mathsf{M}_t \otimes \mathsf{S}_t \otimes A}$.

Let μ_0 denote the initial distribution on Θ. The Ionescu-Tulcea extension theorem (Tulcea (1949)) guarantees the existence of a sequence of probability measures $P_t^\sigma = \mu_0 \otimes \otimes_{\tau=0}^{t-1} \kappa_\tau^\sigma$ defined on the product spaces $(\Theta \times H_A^t)_{t=0}^T$ and a probability measure P^σ on $(\Theta \times H_A^{T+1}, \mathcal{B}_\Theta \otimes \otimes_{\tau=0}^T \mathcal{B}_{\mathsf{Z}_{A,\tau}})$ so that for any measurable $\tilde{\Theta} \times \tilde{H}_A^t \subseteq \Theta \times H_A^t$, $P_t^\sigma(\tilde{\Theta} \times \tilde{H}_A^t) = P^\sigma(\tilde{\Theta} \times \tilde{H}_A^t \times \prod_{\tau=t+1}^T \mathsf{Z}_{A,\tau})$.

Then, the principal's payoff, $W(P^\sigma)$, is given by

$$\int_{\Theta \times H_A^{T+1}} W(\text{proj}_{\Theta \times A^{T+1}}(\theta, h_A^{T+1})) P^\sigma(d(\theta, h_A^{T+1})) = \int_{\Theta \times A^{T+1}} W(a^{T+1}, \theta)(P^\sigma \circ \text{proj}_{\Theta \times A^{T+1}}^{-1})(d(\theta, a^{T+1})),$$

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

$$\int_{\Theta \times H_A^{T+1}} U(\text{proj}_{\Theta \times A^{T+1}}(\overline{\theta}, h_A^{T+1})) P^{\sigma|\theta}(d(\overline{\theta}, h_A^{T+1})) = \int_{\Theta \times H_A^{T+1}} U(a^{T+1}, \theta)(P^{\sigma|\theta} \circ \text{proj}_{\Theta \times A^{T+1}}^{-1})(d(\overline{\theta}, a^{T+1})),$$

where $P^{\sigma|\theta}$ is the induced probability over $\Theta \times H_A^{T+1}$ determined by drawing $\overline{\theta} = \theta$ with probability one and drawing the terminal history using P^σ.

Fix a measurable strategy σ. Fix a period t, a belief $p_t \in \Delta(\Theta \times H_A^t(h^t))$ and mechanism $\mathbf{M}_t$. Define the one-step ahead prediction equations on the measurable rectangles as follows:

$$f_t(p_t, \mathbf{M}_t)(\tilde{\Theta} \times \tilde{H}_A^t \times z_\varnothing(\mathbf{M}_t)) = \int_{\tilde{\Theta} \times \tilde{H}_A^t} (1 - \pi_t(\theta, h_A^t, \mathbf{M}_t)) p_t(d(\theta, h_A^t)), \tag{A.1}$$

$$f_t(p_t, \mathbf{M}_t)(\tilde{\Theta} \times \tilde{H}_A^t \times \mathbf{M}_t \times \{1\} \times \tilde{M} \times \tilde{S} \times \tilde{A}) = \int_{\tilde{\Theta} \times \tilde{H}_A^t} \rho^{\sigma_A}(\tilde{M} \times \tilde{S} \times \tilde{A} | \theta, h_A^t, \mathbf{M}_t) \pi_t(\theta, h_A^t, \mathbf{M}_t) p_t(d(\theta, h_A^t)).$$

The mapping $f_t : \Delta(\Theta \times H_A^t) \times \mathcal{M} \mapsto \Delta(\Theta \times H_A^t \times Z_{A,t}) \equiv \Delta(\Theta \times H_A^{t+1})$ is Borel measurable.

Let $p_{t+1} \in \Delta(\Theta \times H_A^{t+1})$, with marginal $p_{t+1H^{t+1}}$ over H^{t+1}, we can construct

$$\int_{\tilde{H}^{t+1}} q(\tilde{\Theta} \times \tilde{H}_A^{t+1}|h^{t+1}, p) p_{t+1H^{t+1}}(dh^{t+1}) = \int_{\tilde{\Theta} \times \tilde{H}_A^{t+1}} \mathbb{1}[(\theta, h_A^{t+1}) \in H'] p_{t+1}(d(\theta, h_A^{t+1})),$$

on the measurable rectangles $\tilde{\Theta} \times \tilde{H}_A^{t+1} \subset \Theta \times H_A^{t+1}$, $\tilde{H}^{t+1} \subset H^{t+1}$. Note that as a function of (h^{t+1}, p), q is measurable.

Fix a public history h^t, a mechanism $\mathbf{M}_t$, and a belief $\mu_t(h^t) \in \Delta(\Theta \times H_A^t(h^t))$. Define transition probabilities from $H^\tau \times \mathcal{M}$ to $\Theta \times H_A^{\tau+1}$ for $\tau \geq t$, recursively as follows:

$$\nu_t(h^t, \mathbf{M}_t) = f_t(\mu_t(h^t), \mathbf{M}_t),$$
$$\nu_\tau(h^\tau, \mathbf{M}_\tau) = f_\tau(q(h^\tau, \nu_{\tau-1}(h^{\tau-1}, \cdot)), \mathbf{M}_\tau).$$

Fix a belief $p \in \Delta(\Theta \times H_A^{t+1})$. This, together with the kernels $(\kappa_\tau^\sigma)_{\tau \geq t+1}$, induces a sequence of distributions $P_\tau^{\sigma|p} = p \otimes \bigotimes_{n=t+1}^\tau \kappa_n^\sigma$ on the product spaces $(\Theta \times H_A^\tau)_{\tau \geq t+1}$ and a measure $P^{\sigma|p}$ on $(\Theta \times H_A^{T+1}, \mathcal{B}_{\Theta \times H_A^T})$, by the Ionescu-Tulcea Theorem. Proposition 7.26, Lemma 7.28, and Corollary 7.29.1 in Bertsekas and Shreve (1978) imply that the mappings $p \mapsto P_\tau^{\sigma|p}$ and $p \mapsto P^{\sigma|p}$ are Borel measurable. The following lemma is a consequence of Lemma 10.4 in Bertsekas and Shreve (1978):

Lemma A.1. *For every Borel measurable strategy profile σ and Borel measurable subset $\tilde{\Theta} \times \tilde{H}_A^\tau \subseteq \Theta \times H_A^\tau$*

$$P_\tau^{\sigma|\mu_t(h^t)}(\tilde{\Theta} \times \tilde{H}_A^\tau|h^\tau, \mathbf{M}_\tau) = \nu_\tau(h^\tau, \mathbf{M}_\tau)(\tilde{\Theta} \times \tilde{H}_A^\tau) \quad P^{\sigma|\mu_t(h^t)} - \text{almost everywhere.}$$

Note that $\nu_t(h^t, \mathbf{M}_t) = f_t(\mu_t(h^t), \mathbf{M}_t)$ defines a transition probability from $H^t \times \mathcal{M}$ to $\Theta \times H_A^{t+1}$. Then, $P^{\sigma|\cdot} \circ \nu_t$ defines a transition probability from $H^t \times \mathcal{M}$ to $\Delta(\Theta \times H_A^{T+1})$. It follows from this that we can define a bounded measurable function from $H^t \times \mathcal{M}$ to $\mathbb{R}$ as follows:

$$W(\sigma, \nu|h^t, \mathbf{M}_t) = \int_{\Theta \times H_A^{t+1}} \mathbb{E}^{P^{\sigma|(\theta, h_A^{t+1})}} [W(a^t, \cdot, \theta)] \nu_t(d(\theta, h_A^{t+1})|h^t, \mathbf{M}_t).$$

Lemma A.1 implies that $W(\sigma, \nu|h^t, \mathbf{M}_t)$ represents the principal's payoff conditional on information set h^t and having chosen $\mathbf{M}_t$ when his beliefs are given by $\mu_t(h^t)$. $U(\sigma|\theta, h_A^t, \mathbf{M}_t)$ can be defined analogously.

With this we can formally define Perfect Bayesian equilibrium:

Definition A.1. *An assessment $(\sigma_P, \sigma_A, \mu)$ is sequentially rational if for all t and public histories h^t,*

1. *If $\mathbf{M}_t \in supp\, \sigma_{Pt}(h^t)$, $W(\sigma, v | h^t, \mathbf{M}_t) \geq W(\sigma, v | h^t, \mathbf{M}'_t)$ for all $\mathbf{M}'_t \in \mathcal{M}$,*

2. *For all $\mathbf{M}_t \in \mathcal{M}$, $U(\sigma | \theta, h_A^t, \mathbf{M}_t) \geq U(\sigma_P, \sigma'_A | \theta, h_A^t, \mathbf{M}_t)$ for all $\theta \in \Theta, h_A^t \in H_A^t(h^t), \sigma'_A$.*

Definition A.2. *A system of beliefs satisfies Bayes' rule where possible if for all t, all public histories h^t, all measurable subsets $\tilde{\Theta}, \tilde{H}_A^t$ of $\Theta, H_A^t(h^t)$,*

$$\mu_{t+1}(\tilde{\Theta} \times \tilde{H}_A^t \times z_\varnothing(\mathbf{M}_t) | h^t, z_\varnothing(\mathbf{M}_t)) \int_{\Theta \times H_A^t} (1 - \pi_t(\theta, h_A^t, \mathbf{M}_t)) \mu_t(d(\theta, h_A^t) | h^t)$$

$$= \int_{\tilde{\Theta} \times \tilde{H}_A^t} (1 - \pi_t(\theta, h_A^t, \mathbf{M}_t)) \mu_t(d(\theta, h_A^t) | h^t),$$

and for all measurable subsets $\tilde{\Theta}, \tilde{H}_A^t, \tilde{M}, \tilde{S}, \tilde{A}$ of $\Theta, H_A^t(h^t), M^{\mathbf{M}_t}, S^{\mathbf{M}_t}, A$,

$$\int_{\Theta \times H_A^t} \int_{\tilde{S} \times \tilde{A}} \mu_{t+1}(\tilde{\Theta} \times \tilde{H}_A^t \times \tilde{M} \times z_{(s_t, a_t)}(\mathbf{M}_t) | h^t, z_{(s_t, a_t)}(\mathbf{M}_t)) \rho_{S^{\mathbf{M}_t} \times A}^{\sigma_A}(d(s_t, a_t) | \theta, h_A^t) \pi_t(\theta, h_A^t, \mathbf{M}_t) \mu_t(d(\theta, h_A^t) | h^t)$$

$$= \int_{\tilde{\Theta} \times \tilde{H}_A^t} \left(\int_{\tilde{M}} \beta^{\mathbf{M}_t}(\tilde{S} \times \tilde{A} | m) r_t(\theta, h_A^t, \mathbf{M}_t)(dm) \right) \pi_t(\theta, h_A^t, \mathbf{M}_t) \mu_t(d(\theta, h_A^t) | h^t).$$

Definition A.3. *An assessment $(\sigma_P, \sigma_A, \mu)$ is a* Perfect Bayesian equilibrium *if it is sequentially rational and satisfies Bayes' rule where possible.*

If the system of beliefs satisfies Bayes' rule where possible, then letting $v_t(h^t, \mathbf{M}_t) = f_t(\mu_t(\cdot | h^t), \mathbf{M}_t)$ and $v_{t H^{t+1}}$ denote the marginal of v_t on H^{t+1}, we have that

$$\int_{\tilde{H}^{t+1}} \mu_{t+1}(\tilde{\Theta} \times \tilde{H}_A^{t+1} | h^{t+1}) v_{t H^{t+1}}(dh^{t+1}) = v_t(h^t, \mathbf{M}_t)(\tilde{\Theta} \times \tilde{H}_A^{t+1}),$$

on the measurable rectangles $\tilde{\Theta} \times \tilde{H}_A^{t+1} \in \mathcal{B}_{\Theta \otimes H_A^{t+1}}$ and $\tilde{H}^{t+1} \in \mathcal{B}_{H^{t+1}}$. Working forward through the one-step ahead prediction equations we have that $q(h^\tau, v_{\tau-1}(h^{\tau-1}, \cdot)) = \mu_\tau(h^\tau)$ for those histories in the support of $P^{\sigma | \mu_t(h^t)}$. Note the following:

$$W(\sigma, v | h^t, \mathbf{M}_t) = \int_{\Theta \times H_A^{t+1}} \mathbb{E}^{P^\sigma | (\theta, h_A^{t+1})} \left[W(a^t, \cdot, \theta) \right] v_t(d(\theta, h_A^{t+1}) | h^t, \mathbf{M}_t) \tag{A.2}$$

$$= \int_{\Theta \times H_A^t} \left(\begin{array}{l} (1 - \pi_t(\theta, h_A^t, \mathbf{M}_t)) \mathbb{E}^{P^{\sigma | (\theta, h_A^t, z_\varnothing(\mathbf{M}_t))}} \left[W(a^t, a^*, \cdot, \theta) \right] + \pi_t(\theta, h_A^t, \mathbf{M}_t) \\ \int_{M^{\mathbf{M}_t} \times S^{\mathbf{M}_t} \times A} \mathbb{E}^{P^{\sigma | (\theta, h_A^t, m, z_{(s_t, a_t)}(\mathbf{M}_t))}} \left[W(a^t, a_t, \cdot, \theta) \right] \rho^{\sigma_A}(d(m, s_t, a_t) | \theta, h_A^t, \mathbf{M}_t)) \end{array} \right) \mu_t(d(\theta, h_A^t) | h^t)$$

$$= \int_{\{h^{t+1} \in H^{t+1} : h^t \prec h^{t+1}\}} \int_{\Theta \times H_A^{t+1}} \mathbb{E}^{P^\sigma | (\theta, h_A^{t+1})} \left[W(a^t, \cdot, \theta) \right] \mu_{t+1}(d(\theta, h_A^{t+1}) | v_t, h^{t+1}) v_{t H^{t+1}}(dh^{t+1} | h^t, \mathbf{M}_t),$$

where $h^t \prec h^{t+1}$ denotes that public history h^t precedes h^{t+1}.

Disintegration: The proofs in this appendix frequently make use of the notion of a

disintegration. For any two measurable spaces, X and Y, and a Borel measure ν on $X \times Y$, ν_X and ν_Y denote the marginals of ν on Y and X, respectively. Given a product space $X \times Y$, proj_Y denotes the projection of $X \times Y$ onto Y.

Given two Polish spaces, X and Y, and a joint measure ν on $X \times Y$, the (ν_X, proj_Y)-disintegration of ν is the collection of measures on $\mathcal{B}_{X \times Y}$, $\{\eta_x : x \in X\}$, where

(i) η_x is concentrated on $X = x$, i.e. $\eta_x(\{X \neq x\}) = 0$ ν_X– almost everywhere,

and for each non negative measurable function f on $\mathcal{B}_{X \times Y}$:

(ii) $x \mapsto \int_Y f(x,y)\eta_x(dy)$ is measurable,

(iii) $\int_{X \times Y} f(x,y)\nu(d(x,y)) = \int_X \int_Y f(x,y)\eta_x(dy)\nu_X(dx)$.

Proposition 3.6 in Crauel (2002) ensures that $\{\eta_x : x \in X\}$ exists and is unique ν_X-almost everywhere.

In what follows, we denote the principal's beliefs conditional on h^t and the agent participating in mechanism $\mathbf{M}_t$ by $\mu_t^+(\cdot|h^t) \in \Delta(\Theta \times H_A^t(h^t))$. That is, for any measurable subset $\tilde{\Theta} \times \tilde{H}_A^t \subset \Theta \times H_A^t(h^t)$,

$$\mu_t^+(\tilde{\Theta} \times \tilde{H}_A^t|h^t, \mathbf{M}_t) = \int_{\tilde{\Theta} \times \tilde{H}_A^t} \pi_t(\theta, h_A^t, \mathbf{M}_t)\mu_t(d(\theta, h_A^t)|h^t),$$

Note that when $\mu_t^+(\Theta \times H_A^t(h^t)|h^t, \mathbf{M}_t) \neq 0$, we can actually take it to be a probability measure by normalizing it appropriately.

B PROOF OF THEOREM 1

The proof of Theorem 1 follows from the proof of Propositions B.1-B.4 below.

Proposition B.1. *Fix a PBE assessment* $(\sigma_P, \sigma_A, \mu)$ *of* $G_{\mathcal{M}}$ *and a public history* h^t, *and a mechanism* $\mathbf{M}_t \in \mathcal{M}$. *Then, there exists a continuation strategy* σ_A' *such that:*

1. *For all public histories* h^τ *that succeed* h^t, $\sigma_{A\tau}'(\theta, h_A^\tau, \mathbf{M}_\tau) = \sigma_{A\tau}'(\theta, \overline{h}_A^\tau, \mathbf{M}_\tau)$ *for all* $h_A^\tau, \overline{h}_A^\tau \in H_A^\tau(h^\tau), \mathbf{M}_\tau \in \mathcal{M}$, *for all* $\theta \in \Theta$,

2. *For all public histories* h^τ *that succeed* h^t, $U(\mathbf{M}_\tau, \sigma_P, \sigma_A'|\theta, h_A^\tau) = U(\mathbf{M}_\tau, \sigma_P, \sigma_A|\theta, h_A^\tau)$ *for all* $h_A^\tau \in H_A^\tau(h^\tau), \mathbf{M}_\tau \in \mathcal{M}, \theta \in \Theta$,

3. *There exists a belief system* $(\mu')_{t=0}^T$ *such that* $(\sigma_P, \sigma_A', \mu')$ *is also a PBE, and*

4. *For all public histories* h^τ *on the equilibrium path of* (σ_P, σ_A) *starting at* h^t, $W(\mathbf{M}_\tau, \sigma_P, \sigma_A', \mu'|h^\tau) = W(\mathbf{M}_\tau, \sigma_P, \sigma_A, \mu|h^\tau)$.

Proof of Proposition B.1. Fix a public history h^t, a mechanism $\mathbf{M}_t$ and suppose that the set of agent types for which $\sigma_{At}(\theta, h_A^t, \mathbf{M}_t)$ is not $\mathcal{B}_{\Theta \otimes H^t}$-measurable has positive measure under $\mu_t(\cdot)$. Note that the continuation strategy of (θ, h_A^t) is feasible for $(\theta, \overline{h}_A^t)$ and vice versa. Thus, conditional on participating of $\mathbf{M}_t$, the agent at (θ, h_A^t) is not only indifferent between all the messages in the support of $r_t(\theta, h_A^t, \mathbf{M}_t)(\cdot)$, but is also indifferent between all messages in the support of $r_t(\theta, \overline{h}_A^t, \mathbf{M}_t)(\cdot)$. Therefore, the agent at (θ, h_A^t) is indifferent between $r_t(\theta, h_A^t, \mathbf{M}_t)$ and any randomization between $r_t(\theta, h_A^t, \mathbf{M}_t)$ and $r_t(\theta, \overline{h}_A^t, \mathbf{M}_t)$. This indifference also holds taking into account the decision to participate. Moreover, this is true for any continuation public history that is reached from h^t for the same reasons. That is, for any $\tau \geq t$ and h^τ that succeeds $(h^t, \mathbf{M}_t)$ and for any $h_A^\tau, \overline{h}_A^\tau$ that succeed h_A^t and $\overline{h}_A^t$, respectively, the agent of type θ is indifferent between her continuation strategy at $(h_A^\tau, \mathbf{M}_\tau)$ and that at $(\overline{h}_A^\tau, \mathbf{M}_\tau)$. We now construct a new strategy for the agent which, by the above arguments, is payoff equivalent to σ_A.

The principal's belief together with the agent's participation strategy induce a measurable map from $H^t \times \mathcal{M}$ to $\Delta(\Theta \times \{0,1\})$ given by

$$\mathbb{P}_\pi(\tilde{\Theta} \times \{1\}|h^t, \mathbf{M}_t) = \int_{\tilde{\Theta} \times H_A^t} \pi_t(\theta, h_A^t, \mathbf{M}_t)\mu_t(d(\theta, h_A^t)|h^t),$$

with marginal $\mu_{t\Theta}(\cdot|h^t)$ over the set of types. The disintegration theorem implies that we can write:

$$\mathbb{P}_\pi(\tilde{\Theta} \times \{1\}|h^t, \mathbf{M}_t) = \int_{\tilde{\Theta}} \pi_t'(\theta, h^t, \mathbf{M}_t)(\{1\})\mu_{t\Theta}(d\theta|h^t),$$

where, in a slight abuse of notation, $\{\pi_t'(\theta, \cdot)(\{1\}) : \theta \in \Theta\}$ is the $(\mu_{t\Theta}, \text{proj}_{\{0,1\}})$-disintegration of $\mathbb{P}_\pi(\cdot|h^t, \mathbf{M}_t)$.

Similarly, if $\int_{\Theta \times H_A^t(h^t)} \pi_t(\theta, h_A^t, \mathbf{M}_t)\mu_t(d(\theta, h_A^t)|h^t) > 0$, the principal's belief together with the agent's reporting strategy induce a measurable map from $H^t \times \mathcal{M}$ to $\Delta(\Theta \times M^{\mathbf{M}_t})$ given by

$$\mathbb{P}_r(\tilde{\Theta} \times \tilde{M}|h^t, \mathbf{M}_t) = \int_{\tilde{\Theta} \times H_A^t} r_t(\theta, h_A^t, \mathbf{M}_t)(\tilde{M})\mu_t^+(d(\theta, h_A^t)|h^t, \mathbf{M}_t),$$

with marginal $\mu_{t\Theta}^+(\cdot|h^t, \mathbf{M}_t)$ over the set of types. The disintegration theorem implies that we can write

$$\mathbb{P}_r(\tilde{\Theta} \times \tilde{M}|h^t, \mathbf{M}_t) = \int_{\tilde{\Theta}} r_t'(\theta, h^t, \mathbf{M}_t)(\tilde{M})\mu_{t\Theta}^+(d\theta|h^t, \mathbf{M}_t),$$

where $\{r'_t(\theta, h^t, \mathbf{M}_t) : \theta \in \Theta\}$ is the $(\mu^+_{t\Theta}, \mathrm{proj}_{M^{\mathbf{M}_t}})$-disintegration of $\mathbb{P}_r(\cdot|h^t, \mathbf{M}_t)$. Since π'_t, r'_t are $(\Theta \times H^t \times \mathcal{M})$- measurable, they are a fortiori $(\Theta \times H^t_A \times \mathcal{M})$-measurable.

We can similarly redefine π'_τ, r'_τ for $\tau \geq t+1$. Indeed, we have

$$\mathbb{P}_\pi(\tilde{\Theta} \times \{1\}|h^\tau, \mathbf{M}_\tau) = \int_{\tilde{\Theta} \times H^\tau_A} \pi_\tau(\theta, h^\tau_A, \mathbf{M}_\tau)\mu_\tau(d(\theta, h^\tau_A)|h^\tau) = \int_{\tilde{\Theta}} \pi'_\tau(\theta, h^\tau, \mathbf{M}_\tau)(\{1\})\mu_{\tau\Theta}(d\theta|h^\tau),$$

and, whenever $\int_{\Theta \times H^\tau_A} \pi_\tau(\theta, h^\tau_A, \mathbf{M}_\tau)\mu_\tau(d(\theta, h^\tau_A)|h^\tau) > 0$, define:[36]

$$\mathbb{P}_r(\tilde{\Theta} \times \tilde{M}|h^\tau, \mathbf{M}_\tau) = \int_{\tilde{\Theta} \times H^\tau_A} r_\tau(\theta, h^\tau_A, \mathbf{M}_\tau)(\tilde{M})\mu^+_\tau(d(\theta, h^\tau_A)|h^\tau, \mathbf{M}_\tau) = \int_{\tilde{\Theta}} r'_\tau(\theta, h^\tau, \mathbf{M}_\tau)(\tilde{M})\mu^+_{\tau\Theta}(d\theta|h^\tau, \mathbf{M}_\tau).$$

Note that the disintegration theorem automatically implies that the above equations determine well-defined strategies for the agent.

Given the new strategies, the one-step ahead equations (see Equation A.1) become:

$$f'_\tau(\mu'_\tau, \mathbf{M}_\tau)(\tilde{\Theta} \times \tilde{H}^\tau_A \times z_\varnothing(\mathbf{M}_\tau)) = \int_{\tilde{\Theta} \times \tilde{H}^\tau_A} (1 - \pi'_\tau(\theta, h^\tau_A, \mathbf{M}_\tau))\mu'_\tau(d(\theta, h^\tau_A)|h^\tau),$$

$$f'_\tau(\mu'_\tau, \mathbf{M}_\tau)(\tilde{\Theta} \times H'_A \times \mathbf{M}_\tau \times \{1\} \times \tilde{M} \times \tilde{S} \times \tilde{A}) = \int_{\tilde{\Theta} \times \tilde{H}^\tau_A} \rho^{\sigma'_A}(\tilde{M} \times \tilde{S} \times \tilde{A}|\theta, h^\tau_A, \mathbf{M}_\tau)\pi'_\tau(\theta, h^\tau_A, \mathbf{M}_\tau)\mu'_\tau(d(\theta, h^\tau_A)|h^\tau).$$

We now use these equations inductively to show that under the new strategies the distribution over $\Theta \times H^{T+1}$ induced by $P^{\sigma|\cdot} \circ \mu_t(\cdot|h^t)$ is preserved.[37] Suppose that we have shown that for $t \leq \tau' \leq \tau$, $\mu'_{\tau'}(\cdot \times H^{\tau'}_A(h^{\tau'})|h^{\tau'})$ coincides with $\mu_{\tau'}(\cdot \times H^{\tau'}_A(h^{\tau'})|h^{\tau'})$ for $P^{\sigma|\cdot} \circ \mu_t(\cdot|h^t)$ almost all $h^{\tau'}$. We now show that this holds for $\tau+1$. To see this, first note that for all measurable subset $\tilde{\Theta}$ of Θ, we have

$$\mu'^+_\tau(\tilde{\Theta} \times H^\tau_A(h^\tau)|h^\tau, \mathbf{M}_t) = \int_{\tilde{\Theta} \times H^\tau_A(h^\tau)} \pi'_\tau(\theta, h^\tau_A, \mathbf{M}_\tau)\mu'_\tau(d(\theta, h^\tau_A)|h^\tau) = \int_{\tilde{\Theta}} \pi'_\tau(\theta, h^\tau_A, \mathbf{M}_\tau)\mu'_{\tau\Theta}(d\theta|h^\tau)$$

$$= \int_{\tilde{\Theta}} \pi'_\tau(\theta, h^\tau_A, \mathbf{M}_\tau)\mu_{\tau\Theta}(d\theta|h^\tau) = \int_{\tilde{\Theta} \times H^\tau_A(h^\tau)} \pi_\tau(\theta, h^\tau_A, \mathbf{M}_\tau)\mu_\tau(d(\theta, h^\tau_A)|h^\tau) = \mu^+_\tau(\tilde{\Theta} \times H^\tau_A(h^\tau)|h^\tau, \mathbf{M}_\tau),$$

$$(\text{B.1})$$

where the second equality follows from Fubini's theorem and π'_t being $\mathcal{B}_{H^t}$-measurable and the third from the inductive hypothesis. This automatically implies that

$$\nu'_\tau(\mu_\tau(\cdot|h^\tau), \mathbf{M}_\tau)(\tilde{\Theta} \times H^\tau_A(h^\tau) \times z_\varnothing(\mathbf{M}_\tau)) = \nu_\tau(\mu_\tau(\cdot|h^\tau), \mathbf{M}_\tau)(\tilde{\Theta} \times H^\tau_A(h^\tau) \times z_\varnothing(\mathbf{M}_\tau)),$$

---

[36] The definition of (π_τ, r_τ) for θ not in the support of the principal's beliefs is irrelevant.

[37] While we change the agent's strategy at all histories which succeed $(h^t, \mathbf{M}_t)$, Bayes' rule where possible ties the beliefs at h^t and the beliefs at h^τ only at those histories h^τ that are on the path of the (agent's) strategy. This is why when we check that the principal's beliefs over Θ have not changed we do so along the path of the strategy profile starting at h^t. This is enough to check that we have not changed the principal's (continuation) payoffs.

and hence the updated belief $\mu_{\tau+1}(\cdot \times H_A^\tau(h^\tau)|h^\tau, \mathbf{M}_\tau, z_\varnothing(\mathbf{M}_\tau))$ remains unchanged if the event that the mechanism was rejected has positive probability in the original strategy profile. Moreover, for any measurable subset $\tilde{\Theta}$ of Θ, and for any measurable subset $\tilde{M}$ of $M^{\mathbf{M}_\tau}$ we have

$$\int_{\tilde{\Theta} \times H_A^\tau(h^\tau)} r_\tau'(\theta, h_A^\tau, \mathbf{M}_\tau)(\tilde{M}) \mu_\tau'^+(d(\theta, h_A^\tau)|h^\tau, \mathbf{M}_\tau) = \int_{\tilde{\Theta}} r_\tau'(\theta, h_A^\tau, \mathbf{M}_\tau)(\tilde{M}) \mu_{\tau\Theta}'^+(d\theta|h^\tau, \mathbf{M}_\tau)$$

$$= \int_{\tilde{\Theta}} r_\tau'(\theta, h_A^\tau, \mathbf{M}_\tau)(\tilde{M}) \mu_{\tau\Theta}^+(d\theta|h^\tau, \mathbf{M}_\tau) = \int_{\tilde{\Theta} \times H_A^\tau(h^\tau)} r_\tau(\theta, h_A^\tau, \mathbf{M}_\tau)(\tilde{M}) \mu_\tau^+(d(\theta, h_A^\tau)|h^\tau, \mathbf{M}_\tau),$$

where the first equality uses Fubini's theorem, the second equality uses Equation B.1, and the third equality uses the definition of $r_t'(\cdot)$. This implies that

$$f_\tau'(\mu_\tau'(\cdot|h^\tau), \mathbf{M}_\tau)(\tilde{\Theta} \times H_A^\tau(h^\tau) \times \mathbf{M}_\tau \times \{1\} \times M^{\mathbf{M}_\tau} \times \tilde{S} \times \tilde{A}) = f_\tau(\mu_\tau(\cdot|h^\tau), \mathbf{M}_\tau)(\tilde{\Theta} \times H_A^\tau(h^\tau) \times \mathbf{M}_\tau \times \{1\} \times M^{\mathbf{M}_\tau} \times \tilde{S} \times \tilde{A})$$

on the measurable rectangles $\tilde{\Theta} \times M^{\mathbf{M}_\tau} \times \tilde{S} \times \tilde{A} \in \mathcal{B}_{\Theta \otimes \mathsf{M}_\tau \times \mathsf{S}_\tau \times A}$. Thus, the (marginal) updated beliefs $\{\mu_{\tau+1}'(\cdot \times H_A^\tau(h^\tau) \times \mathbf{M}_\tau \times M^{\mathbf{M}_\tau} \times \cdot|h^\tau, \mathbf{M}_t, z) : z \in \mathsf{Z}_\tau\}$ coincide with $\{\mu_{t+1}(\cdot \times H_A^t(h^t) \times \mathbf{M}_t \times M^{\mathbf{M}_t} \times \cdot|h^t, \mathbf{M}_t, z) : z \in \mathsf{Z}_\tau\}$, $P^{\sigma|\mu_t(h^t)}$- almost surely. It follows that the principal's payoff remains the same (see Equation A.2). $\qquad\square$

Remark B.1. *Note that the PBE assessment one obtains from Proposition B.1 satisfies that $P^{\sigma'}(proj_{\Theta \times A^{T+1}} \Theta \times H_A^{T+1}) = P^\sigma(proj_{\Theta \times A^{T+1}} \Theta \times H_A^{T+1})$. It follows that the set of PBE outcomes of the mechanism-selection game is the same as the set of PBE outcomes of the mechanism-selection game when the agent's strategy only depends on her payoff relevant type and the public history.*

The outcome-equivalent PBE assessment one obtains from Proposition B.1 satisfies the following property. On the equilibrium path, the principal's beliefs over the agent's payoff-relevant type, $\theta \in \Theta$, do not depend on her payoff-irrelevant history, h_A^t. However, at a public history h^t reached after a deviation by the agent, the requirements of PBE do not rule out that the principal's *updated* beliefs depend nontrivially on both θ and h_A^t. It follows from Proposition B.1 that without loss of generality, we can assume that when the principal observes a deviation by the agent, his updated beliefs do not depend on h_A^t. The proof is available upon request.

Given a mechanism $\mathbf{M}_t$, let

$$(S^{\mathbf{M}_t} \times A)_+ = \bigcup_{m \in M^{\mathbf{M}_t}} \mathrm{supp}\, \beta^{\mathbf{M}_t}(\cdot|m). \tag{B.2}$$

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.

Fix a PBE assessment of $G_{\mathcal{M}}$ and a mechanism $\mathbf{M}_t$ and define a measure on $M^{\mathbf{M}_t} \times S^{\mathbf{M}_t} \times A$ as follows:

$$\mathbb{R}_{h^t, \mathbf{M}_t}(\tilde{M} \times \tilde{S} \times \tilde{A}) = \int_{\Theta \times H_A^t(h^t)} \int_{\tilde{M}} \beta(\tilde{S} \times \tilde{A}|m) r_t(\theta, h_A^t, \mathbf{M}_t)(dm) \pi_t(\theta, h_A^t, \mathbf{M}_t) \mu_t(d(\theta, h_A^t)|h^t).$$

(B.3)

Proposition B.2. *Fix a PBE assessment $(\sigma_P, \sigma_A, \mu)$ of $G_{\mathcal{M}}$ that satisfies Proposition B.1. Then, there exists an outcome-equivalent PBE assessment $(\sigma_P', \sigma_A', \mu')$ such that for all public histories h^t, for all mechanisms $\mathbf{M}_t$ on the path of the equilibrium strategy at h^t such that the agent participates with positive probability, $M^{\mathbf{M}_t} = \Theta$, $(S^{\mathbf{M}_t} \times A)_+$ is the support of $\mathbb{R}_{h^t, \mathbf{M}_t}(\Theta \times \cdot)$. Moreover, the agent truthfully reports her type.*

Proof of Proposition B.2. Fix a PBE assessment $(\sigma_P, \sigma_A, \mu)$ of $G_{\mathcal{M}}$ that satisfies Proposition B.1, a history h^t, and a mechanism $\mathbf{M}_t$ on the support of the principal's strategy at h^t, such that the agent participates with positive probability. Let Θ_+ denote the support of $\mu_t^+(\cdot|h^t, \mathbf{M}_t)$. By definition, Θ_+ is closed.

Define a new mechanism $\overline{\mathbf{M}}_t$ as follows. Let $(M^{\overline{\mathbf{M}}_t}, S^{\overline{\mathbf{M}}_t}) = (\Theta, S^{\mathbf{M}_t})$ (Recall that Θ is a feasible set of input messages.) Define the transition probability from Θ_+ to $S^{\mathbf{M}_t} \times A$ by specifying its values on the measurable rectangles, $\tilde{S} \times \tilde{A} \in \mathcal{B}_{S^{\mathbf{M}_t} \otimes A}$,

$$\beta^{\overline{\mathbf{M}}_t}(\tilde{S} \times \tilde{A}|\theta) = \int_{M^{\mathbf{M}_t}} \beta^{\mathbf{M}_t}(\tilde{S} \times \tilde{A}|m) r_t(\theta, h_A^t, \mathbf{M}_t)(dm).$$

By composition of measurable functions and since r_t does not depend on h_A^t, this defines a measurable mapping from Θ_+ to $\Delta(S^{\mathbf{M}_t} \times A)$. Note that this modification does not alter the principal's beliefs about the agent's type conditional on observing $(s_t, a_t) \in (S^{\mathbf{M}_t} \times A)_+$.

For $\theta \notin \Theta_+$, let

$$Q^*(\theta) = \arg \max_{q \in \Delta(\Theta_+)} \int_{\Theta_+} \left[\int_{S^{\mathbf{M}_t} \times A} \mathbb{E}^{P^{\sigma|(\theta, h_A^t, \cdot, z_{(s_t, a_t)}(\mathbf{M}_t))}} [U(a^t, a_t, \cdot, \theta)] \beta^{\overline{\mathbf{M}}_t}(d(s_t, a_t)|\tilde{\theta}) \right] q(d\tilde{\theta}),$$

where the payoff on the right-hand side of the above expression corresponds to the payoff from reporting (possibly at random) a type in Θ_+, and then conditional on (s_t, a_t), play proceeding as in the original strategy profile. The objective is contin-

uous in q[38] and $\Delta(\Theta_+)$ is compact since Θ_+ is compact (Theorem 15.11 in Aliprantis and Border (2013)). Then, the maximization is well-defined. Theorem 18.19 in Aliprantis and Border (2013) implies that there exists a measurable selector $q^*(\theta) \in Q^*(\theta)$. Use this to define $\beta^{\overline{\mathbf{M}}_t}(\cdot|\theta)$ for $\theta \notin \Theta_+$ on the measurable rectangles $\tilde{S} \times \tilde{A}$ of $S^{\mathbf{M}_t} \times A$ as follows:

$$\beta^{\overline{\mathbf{M}}_t}(\tilde{S} \times \tilde{A}|\theta) = \int_{\Theta_+} \beta^{\overline{\mathbf{M}}_t}(\tilde{S} \times \tilde{A}|\tilde{\theta}) q^*(\theta)(d\tilde{\theta}).$$

Note that this defines $\beta^{\overline{\mathbf{M}}_t}$ as a transition probability from Θ to $S^{\mathbf{M}_t} \times A$.[39]

Modify the principal's strategy so that he offers $\overline{\mathbf{M}}_t$ instead of $\mathbf{M}_t$. Modify the continuation strategies so that for any public history that succeeds $(h^t, \overline{\mathbf{M}}_t)$ play follows what would have transpired if instead $\mathbf{M}_t$ had been played. Set $r'_t(\theta, h^t_A, \overline{\mathbf{M}}_t) = \delta_\theta$. For $\theta \in \Theta_+$, it is a best response to set $\pi'_t(\theta, h^t_A, \overline{\mathbf{M}}_t) = \pi_t(\theta, h^t_A, \mathbf{M}_t)$. Types not in Θ_+ may not find it optimal to participate in the mechanism; recompute their participation strategies accordingly. Finally, use Equation A.1 equations to modify the belief system. It is immediate to check that the new assessment is also a PBE.

To finalize the proof, we need to show that the distribution $\mathbb{R}_{h^t, \overline{\mathbf{M}}_t} \in \Delta(M^{\overline{\mathbf{M}}_t} \times S^{\overline{\mathbf{M}}_t} \times A) \in \Delta(\Theta \times S^{\mathbf{M}_t} \times A)$ satisfies that the support of $\mathbb{R}_{h^t, \overline{\mathbf{M}}_t}(\Theta \times \cdot)$ is $(S^{\mathbf{M}_t} \times A)_+$. Clearly, $\operatorname{supp} \mathbb{R}_{h^t, \overline{\mathbf{M}}_t}(\Theta_+ \times \cdot) \subseteq (S^{\mathbf{M}_t} \times A)_+$. Suppose the inclusion is strict and let $(s_0, a_0) \in (S^{\mathbf{M}_t} \times A)_+ \smallsetminus \operatorname{supp} \mathbb{R}_{h^t, \overline{\mathbf{M}}_t}(\Theta_+ \times \cdot)$. Then, there exists an open neighborhood of (s_0, a_0), N_0 such that $\mathbb{R}_{h^t, \overline{\mathbf{M}}_t}(\Theta \times N_0) = 0$. We claim that $\operatorname{supp} \mathbb{R}_{h^t, \overline{\mathbf{M}}_t}(\cdot \times S^{\mathbf{M}_t} \times A) \neq \Theta_+$. Towards a contradiction, assume that $\operatorname{supp} \mathbb{R}_{h^t, \overline{\mathbf{M}}_t}(\cdot \times S^{\mathbf{M}_t} \times A) = \Theta_+$. Then, for all $\theta \in \Theta_+$ there exists an open neighborhood $\theta \in V_\theta$ such that $\mathbb{R}_{h^t, \overline{\mathbf{M}}_t}(V_\theta \times S^{\mathbf{M}_t} \times A) > 0$. Then, $\int_{\Theta_+} \beta^{\overline{\mathbf{M}}_t}(N_0|\theta) \mu_t^+(d\theta|h^t, \overline{\mathbf{M}}_t) = 0$ implies that $\beta^{\overline{\mathbf{M}}_t}(N_0|\theta) = 0$ for all $\theta \in \Theta_+$. Thus, $(s_0, a_0) \notin (S^{\mathbf{M}_t} \times A)_+$, a contradiction. □

Two corollaries follow from Proposition B.2. First, from now on, we can focus on PBE assessments $(\sigma_P, \sigma_A, \mu)$ that satisfy Proposition B.1 and where the principal offers mechanisms with input messages equal to the set of types and the agent truthfully reports her type conditional on participating. Second, if the agent participates in the mechanism offered by the principal, the principal is *never* surprised by the tuples (s_t, a_t) that come out of the mechanism. This, instead, means that if the agent par-

---

[38]The term in brackets is bounded above by the payoff the agent of type θ obtains in equilibrium. The results in Serfozo (1982) imply then continuity of the objective in q.

[39]To see this, fix a measurable subset C of $\Delta(S^{\mathbf{M}_t} \times A)$ and let B denote a measurable subset of $[0,1]$. Then, the set $\{\theta \in \Theta : \beta^{\overline{\mathbf{M}}_t}(C|\theta) \in B\} = \{\theta \in \Theta_+ : \beta^{\overline{\mathbf{M}}_t}(C|\theta) \in B\} \cup \{\theta \in \Theta \smallsetminus \Theta_+ : \beta^{\overline{\mathbf{M}}_t}(C|\theta) \in B\}$. Each set is in $\mathcal{B}_\Theta$ by construction and therefore their union is in $\mathcal{B}_\Theta$.

ticipates in the mechanism with positive probability, then beliefs, $\mu_{t+1}(\cdot, z_{(s_t,a_t)}(\mathbf{M}_t)|h^t, z_{(s_t,a_t)}(\mathbf{M}_t))$, are determined via Bayes' rule where possible.

The following two propositions require lifting a PBE assessment, $(\sigma_P, \sigma_A, \mu)$, from $G_{\mathcal{M}}$ to one in the auxiliary game. Each proposition finds an outcome-equivalent PBE assessment, $(\sigma'_P, \sigma'_A, \mu')$, of the auxiliary game with certain properties.

Proposition B.3. *Fix a PBE assessment $(\sigma_P, \sigma_A, \mu)$ of $G_{\mathcal{M}}$ that satisfies the properties listed in Propositions B.1-B.2. Then, there exists an outcome-equivalent PBE assessment of the* auxiliary game, *$(\sigma'_P, \sigma'_A, \mu')$ such that for all h^t, and $\mathbf{M}_t$ in the support of $\sigma_{Pt}(h^t)$ such that the agent participates with positive probability, the following holds. First, $\mathbf{M}_t$ is a canonical mechanism. Second, the agent truthfully reports her type. Third, the principal's updated beliefs coincide with the output message.*

Proof of Proposition B.3. Let $(\sigma_P, \sigma_A, \mu)$ be as in the statement of Proposition B.3. In a slight abuse of notation, lift $(\sigma_P, \sigma_A, \mu)$ so that it is a PBE assessment in the auxiliary game. Let h^t be a public history and let $\mathbf{M}_t$ denote a mechanism on the path of the equilibrium strategy starting at h^t. Let Θ_+ denote the support of $\mu_t^+(\cdot|h^t, \mathbf{M}_t)$.

The kernel $\nu_t(h^t, \mathbf{M}_t)$ defines a joint probability on $\Theta \times S^{\mathbf{M}_t} \times A$ via

$$\nu_t(h^t, \mathbf{M}_t)(\tilde{\Theta} \times \tilde{S} \times \tilde{A}) = \int_{\tilde{\Theta}} \beta^{\mathbf{M}_t}(\tilde{S} \times \tilde{A}|\theta) \mu_{t\Theta}^+(d\theta|h^t, \mathbf{M}_t),$$

while the updated beliefs satisfy

$$\nu_t(h^t, \mathbf{M}_t)(\tilde{\Theta} \times \tilde{S} \times \tilde{A}) = \int_{\tilde{S} \times \tilde{A}} \mu_{t+1}(\tilde{\Theta} \times H_A^{t+1}(h^t, z_{(s_t,a_t)}(\mathbf{M}_t))|h^t, z_{(s_t,a_t)}(\mathbf{M}_t)) \nu_{tS^{\mathbf{M}_t} \times A}(h^t, \mathbf{M}_t)(d(s_t, a_t)).$$

Recall that we can write the principal's payoff as follows (Equation A.2):

$$\int_{\Theta \times S^{\mathbf{M}_t} \times A} \mathbb{E}^{P^{\sigma|(\theta, h_A^t, z_{(s_t,a_t)}(\mathbf{M}_t))}} \left[W(a^t, a_t, \cdot, \theta) \right] \beta^{\mathbf{M}_t}(d(s_t, a_t)|\theta) \mu_{t\Theta}^+(d\theta|h^t, \mathbf{M}_t)$$
$$= \int_{S^{\mathbf{M}_t} \times A} \int_{\Theta} \mathbb{E}^{P^{\sigma|(\theta, h_A^t, z_{(s_t,a_t)}(\mathbf{M}_t))}} \left[W(a^t, a_t, \cdot, \theta) \right] \mu_{t+1\Theta}(d\theta|h^t, z_{(s_t,a_t)}(\mathbf{M}_t)) \nu_{tS^{\mathbf{M}_t} \times A}(d(s_t, a_t)|h^t, \mathbf{M}_t).$$

By Kuratowski's theorem, there exists a bijection $\omega : S^{\mathbf{M}_t} \mapsto [0,1]$ (see Parthasarathy (2005)). Define the measurable function $\mathcal{W} : A \times \Delta(\Theta) \times [0,1] \mapsto \mathbb{R}$ as follows:

$$\mathcal{W}(\mu, \omega(s_t), a_t) = \int_{\Theta} \mathbb{E}^{P^{\sigma|(\theta, h_A^t, z_{(s_t,a_t)}(\mathbf{M}_t))}} \left[W(a^t, a_t, \cdot, \theta) \right] \mu(d\theta).$$

We allow $\mathcal{W}$ to explicitly depend on s_t since continuation payoffs may depend on s_t beyond its impact on beliefs. Define the measurable map $T : S^{\mathbf{M}_t} \times A \mapsto \Delta(\Theta) \times [0,1] \times$

A, so that $T(s_t, a_t) = (\mu_{t+1}(\cdot|h^t, z_{(s_t,a_t)}(\mathbf{M}_t)), \omega(s_t), a_t)$. Define a measure over $\Theta \times \Delta(\Theta) \times [0,1] \times A$ by specifying it on the measurable rectangles:

$$\mathbb{P}(\tilde{\Theta} \times \tilde{U} \times \tilde{\Omega} \times \tilde{A}) = \nu_t(h^t, \mathbf{M}_t)(\tilde{\Theta} \times T^{-1}(\tilde{U} \times \tilde{\Omega} \times \tilde{A})), \qquad (B.4)$$

where we denote by $\tilde{\Omega}$ an element of $\mathcal{B}_{[0,1]}$ anticipating that this part of the output message will become the public randomization device. Note that we can write:

$$\int_{S^{\mathbf{M}_t} \times A} \mathcal{W}(\mu(h^t, z_{(s_t,a_t)}(\mathbf{M}_t)), \omega(s_t), a_t) \nu_{tS^{\mathbf{M}_t} \times A}(d(s_t, a_t)|h^t, \mathbf{M}_t) \qquad (B.5)$$

$$= \int_{S^{\mathbf{M}_t} \times A} \mathcal{W}(T(s_t, a_t)) \nu_{tS^{\mathbf{M}_t} \times A}(d(s_t, a_t)|h^t, \mathbf{M}_t) = \int_{T(S^{\mathbf{M}_t} \times A)} \mathcal{W}(\mu, \omega, a_t) \nu_{tS^{\mathbf{M}_t} \times A} \circ T^{-1}(d(\mu, \omega, a_t))$$

$$= \int_{\Delta(\Theta) \times [0,1] \times A} \mathcal{W}(\mu, \omega, a_t) \mathbb{P}_{\Delta(\Theta) \times [0,1] \times A}(d(\mu, \omega, a_t)).$$

Let $\{\eta_{(\mu,\omega,a)} : (\mu, \omega, a) \in \Delta(\Theta) \times [0,1] \times A\}$ denote the $(\mathbb{P}_{\Delta(\Theta) \times [0,1] \times A}, \text{proj}_\Theta)$ disintegration of $\mathbb{P}$. We have that on the measurable rectangles,

$$\int_{\tilde{U} \times \tilde{\Omega} \times \tilde{A}} \eta_{(\mu,\omega,a)}(\tilde{\Theta}) \mathbb{P}_{\Delta(\Theta) \times [0,1] \times A}(d(\mu, \omega, a)) = \mathbb{P}(\tilde{\Theta} \times \tilde{U} \times \tilde{\Omega} \times \tilde{A}) = \nu_t(h^t, \mathbf{M}_t)(\tilde{\Theta} \times T^{-1}(\tilde{U} \times \tilde{\Omega} \times \tilde{A}))$$

$$= \int_{T^{-1}(\tilde{U} \times \tilde{\Omega} \times \tilde{A})} T_{\Delta(\Theta)}(s_t, a_t)(\tilde{\Theta}) \nu_{tS^{\mathbf{M}_t} \times A}(d(s_t, a_t)|h^t, \mathbf{M}_t) = \int_{\tilde{U} \times \tilde{\Omega} \times \tilde{A}} \mu(\tilde{\Theta}) \mathbb{P}_{\Delta(\Theta) \times [0,1] \times A}(d(\mu, \omega, a)),$$

where $T_{\Delta(\Theta)}$ is the first coordinate of T. It follows from Proposition 3.6 in Crauel (2002) that $\eta_{(\mu,\omega,a)} = \mu$ $\mathbb{P}_{\Delta(\Theta) \times [0,1] \times A}$-almost everywhere. Intuitively this is just saying that "when the output message is (μ, ω)", the principal updates his beliefs to μ.

Now, let $\{\eta_\mu : \mu \in \Delta(\Theta)\}$ denote the $(\mathbb{P}_{\Delta(\Theta)}, \text{proj}_{\Theta \times [0,1] \times A})$-disintegration of $\mathbb{P}$. For any measurable subset $\tilde{U}$ of $\Delta(\Theta)$, we have that on the measurable rectangles $\tilde{\Theta} \times \tilde{\Omega} \times \tilde{A}$ of $\Theta \times [0,1] \times A$,

$$\int_{\tilde{U}} \eta_\mu(\tilde{\Theta} \times \tilde{A} \times \tilde{\Omega}) d\mathbb{P}_{\Delta(\Theta)} = \mathbb{P}(\tilde{\Theta} \times \tilde{U} \times \tilde{A} \times \tilde{\Omega}) = \int_{\tilde{U} \times \tilde{A} \times \tilde{\Omega}} \mu(\tilde{\Theta}) d\mathbb{P}_{\Delta(\Theta) \times A \times \Omega} = \qquad (B.6)$$

$$\int_{\tilde{U} \times A \times \Omega} \mathbb{1}[(a, \omega) \in \tilde{A} \times \tilde{\Omega}] \mu(\tilde{\Theta}) d\mathbb{P}_{\Delta(\Theta) \times A \times \Omega} = \int_{\tilde{U}} \int_{A \times \Omega} \mu(\tilde{\Theta}) \mathbb{1}[(a, \omega) \in \tilde{A} \times \tilde{\Omega}] \overline{\eta}_\mu(d(a, \omega)) d\mathbb{P}_{\Delta(\Theta)}$$

$$= \int_{\tilde{U}} \mu(\tilde{\Theta}) \int_{A \times \Omega} \mathbb{1}[(a, \omega) \in \tilde{A} \times \tilde{\Omega}] \overline{\eta}_\mu(d(a, \omega)) d\mathbb{P}_{\Delta(\Theta)} = \int_{\tilde{U}} \mu(\tilde{\Theta}) \overline{\eta}_\mu(\tilde{A} \times \tilde{\Omega}) d\mathbb{P}_{\Delta(\Theta)},$$

where the first and second equalities follow from the disintegration property, the third equality is a rewriting of the integral, the fourth uses the $(\mathbb{P}_{\Delta(\Theta)}, \pi_{A \times \Omega})$-disintegration of $\mathbb{P}_{\Delta(\Theta) \times A \times \Omega}$, $\{\overline{\eta}_\mu : \mu \in \Delta(\Theta)\}$, and the fifth equality uses that "conditional on μ", $\mu(\tilde{\Theta})$ is constant. It follows from this that $\Theta \perp\!\!\!\perp (A, \Omega) | \Delta(\Theta)$.

Now, let $\{\beta_\theta : \theta \in \Theta\}$ denote the $(\mathbb{P}_\Theta, \text{proj}_{\Delta(\Theta) \times \Omega \times A})$-disintegration of $\mathbb{P}$. Theorem 1.25 in Kallenberg (2017) implies that there exist two transition probabilities $p : \Theta \mapsto \Delta(\Theta)$,

$q: \Theta \times \Delta(\Theta) \mapsto [0,1] \times A$ such that $\beta = p \otimes q$ and hence,

$$\int_{\tilde{\Theta}} \beta_\theta(\tilde{U} \times \tilde{\Omega} \times \tilde{A}) \mathbb{P}_\Theta(d\theta) = \int_{\tilde{\Theta}} \left(\int_{\tilde{U}} q_{(\theta,\mu)}(\tilde{\Omega} \times \tilde{A}) p_\theta(d\mu) \right) \mathbb{P}_\Theta(d\theta).$$

Equation B.6 and Theorem 1.27 in Kallenberg (2017) imply that $q_{(\theta,\mu)} = \overline{\eta}_\mu$ $\mathbb{P}_{\Delta(\Theta)}$-almost everywhere, so that we can write[40]

$$\int_{\tilde{\Theta}} \beta_\theta(\tilde{U} \times \tilde{\Omega} \times \tilde{A}) \mathbb{P}_\Theta(d\theta) = \int_{\tilde{\Theta}} \left(\int_{\tilde{U}} \overline{\eta}_\mu(\tilde{\Omega} \times \tilde{A}) p_\theta(d\mu) \right) \mathbb{P}_\Theta(d\theta).$$

Theorem 1.25 in Kallenberg (2017) implies that we can write $\overline{\eta}_\mu$ as the composition of two transition probabilities, α from $\Delta(\Theta)$ to A and γ from $\Delta(\Theta) \times A$ to $[0,1]$, so that

$$\int_{\tilde{\Theta}} \beta_\theta(\tilde{U} \times \tilde{\Omega} \times \tilde{A}) \mathbb{P}_\Theta(d\theta) = \int_{\tilde{\Theta}} \left(\int_{\tilde{U}} \left(\int_{\tilde{A}} \gamma_{(\mu,a)}(\tilde{\Omega}) \alpha_\mu(da) \right) p_\theta(d\mu) \right) \mathbb{P}_\Theta(d\theta).$$

Finally, note that

$$\int_{\tilde{\Theta}} \beta_\theta(\tilde{U} \times \tilde{\Omega} \times \tilde{A}) \mathbb{P}_\Theta(d\theta) = \nu_t(h^t, \mathbf{M}_t)(\tilde{\Theta} \times T^{-1}(\tilde{U} \times \tilde{\Omega} \times \tilde{A})) = \int_{\tilde{\Theta}} \beta^{\mathbf{M}_t}(T^{-1}(\tilde{U} \times \tilde{\Omega} \times \tilde{A})|\theta) \mathbb{P}_\Theta(d\theta).$$

$$(\text{B.7})$$

Let $\mathbf{M}'_t$ be such that $(M^{\mathbf{M}'_t}, S^{\mathbf{M}'_t}) = (\Theta, \Delta(\Theta))$ and for $\theta \in \Theta_+$, define $\beta^{\mathbf{M}'_t}(\cdot|\theta)$ as the composition of the transition probabilities p from Θ to $\Delta(\Theta)$ and α from $\Delta(\Theta)$ to A.

Continuation strategies are modified so that when the outcome of the mechanism is (μ, a_t), we draw $\omega \in [0,1]$ according to $\gamma_{(\mu, a_t)}$ and we play the continuation corresponding to $(h^t, z_{(s_t, a_t)}(\mathbf{M}_t))$ where $T(s_t, a_t) = (\mu, a_t, \omega)$.[41] That is,

$$(\sigma'_P, \sigma'_A)\big|_{(h^t, z_{(\mu, a_t)}(\mathbf{M}'_t), \omega, \cdot)} = (\sigma_P, \sigma_A)\big|_{(h^t, z_{T^{-1}(\mu, a_t, \omega)}(\mathbf{M}_t), \cdot)}.$$

With the continuation strategies at hand, for types not in Θ_+, use the same argument as in the proof of Proposition B.2 to extend $\beta^{\mathbf{M}'_t}$ to all of Θ. This completes the spec-

---

[40]To facilitate checking the application of Theorem 1.27 in Kallenberg (2017) to our setting, we now use his notation. For any sets Y, X, and Z, and joint measure ν on $Y \times X \times Z$, let $\eta_{YX|Z}$ denote the $(\nu_Z, \text{proj}_{Y \times X})-$ disintegration of ν. Then Equation B.6 shows that $\eta_{\Theta A[0,1]|\Delta(\Theta)} = \eta_{\Theta|\Delta(\Theta)} \otimes \eta_{A[0,1]|\Delta(\Theta)}$, $\mathbb{P}_{\Delta(\Theta)}$-almost everywhere. By Theorem 1.25 in Kallenberg (2017), $\eta_{\Theta A[0,1]|\Delta(\Theta)} = \eta_{\Theta|\Delta(\Theta)} \otimes \eta_{A[0,1]|\Theta|\Delta(\Theta)}$, which means that $\eta_{A[0,1]|\Theta|\Delta(\Theta)} = \eta_{A[0,1]|\Delta(\Theta)}$ $\mathbb{P}_{\Delta(\Theta)}$-almost everywhere. Theorem 1.27 in Kallenberg (2017) shows that $\eta_{A[0,1]|\Delta(\Theta)|\Theta} = \eta_{A[0,1]|\Theta|\Delta(\Theta)}$ $\mathbb{P}_{\Theta\Delta(\Theta)}-$ almost everywhere. Together with the observation that $\eta_{\Delta(\Theta)A[0,1]|\Theta} = \eta_{\Delta(\Theta)|\Theta} \eta_{A[0,1]|\Delta(\Theta)|\Theta}$, completes the claim.

[41]At the risk of introducing more notation, one could use the probability integral transform and make the distribution on $[0,1]$ be the uniform distribution. Now, the probability integral transform requires that the distribution of ω be continuous. This can always be guaranteed by applying the result in Lehmann et al. (1988), which shows that for any (real-valued) random variable X one can always construct an information-equivalent random variable X^* the distribution of which is continuous.

ification of $\mathbf{M}'_t$. Equation B.7 implies that the agent receives the same payoff when her type is in Θ_+ by truthfully reporting. If $\theta \notin \Theta_+$, the agent's payoff from participating in the mechanism may be lower, so that we recompute the participation strategy accordingly. Equation B.5 implies that the principal receives the same payoff under $\mathbf{M}_t$ and under $\mathbf{M}'_t$. It follows that the new profile is also a PBE. □

Proposition B.4. *Fix a PBE assessment* $(\sigma_P, \sigma_A, \mu)$ *in* $G_{\mathcal{M}}$ *that satisfies* B.1-B.2. *Then, there is an outcome-equivalent PBE assessment* $(\sigma'_P, \sigma'_A, \mu')$ *of the* auxiliary game *that satisfies* Proposition B.3 *and such that the following holds. For every* $t \geq 0$, *for every public history* h^t *and mechanism* $\mathbf{M}_t$ *in the support of* $\sigma_{Pt}(h^t)$, $\pi_t(\theta, h^t_A, \mathbf{M}_t) = 1$ *for all types in the support of* $\mu_t(\cdot \times H^t_A | h^t)$.

Proof of Proposition B.4. Let $(\sigma_P, \sigma_A, \mu)$ denote a PBE assessment of $G_{\mathcal{M}}$ that satisfies Propositions B.1- B.2. In a slight abuse of notation, let $(\sigma_P, \sigma_A, \mu)$ denote the outcome-equivalent PBE of the auxiliary game as in Proposition B.3. Fix a public history h^t and a mechanism $\mathbf{M}_t$ on the path of σ_{Pt} at h^t such that

$$\int_{\Theta \times H^t_A(h^t)} \pi_t(\theta, h^t_A, \mathbf{M}_t) \mu_t(d(\theta, h^t_A) | h^t) < 1.$$

This implies that the belief $\mu_{t+1}(\cdot | h^t, z_\varnothing(\mathbf{M}_t)) \equiv \mu^\varnothing$ is defined via Bayes' rule using the equilibrium strategy profile. Let Θ_+ denote the support of $\mu_t(\cdot \times H^t_A | h^t)$.

If $\mathbf{M}_t$ is rejected with probability one, then modify the principal's strategy so that instead of offering $\mathbf{M}_t$, he offers $\mathbf{M}'_t$ such that for all $\theta \in \Theta$ $\beta^{\mathbf{M}'_t}(\cdot | \theta) = \delta_{(\mu^\varnothing, a^*)}$, where δ. denotes the Dirac measure. Modify the continuation strategies so that

$$(\sigma'_P, \sigma'_A)|_{(h^t, z_{(\mu^\varnothing, a^*)}(\mathbf{M}'_t))} = (\sigma_P, \sigma_A)|_{(h^t, z_\varnothing(\mathbf{M}_t))}.$$

Modify the agent's strategy so that $\pi'_t(\theta, h^t_A, \mathbf{M}'_t) = 1$, $r'_t(\theta, h^t_A, \mathbf{M}'_t) = \delta_\theta$ whenever θ is in Θ_+; otherwise, leave the agent's strategy unchanged.

Suppose now that the mechanism is accepted with positive probability, so that both $\mu^\varnothing$ and $\mu^+_t(\cdot | h^t, \mathbf{M}_t)$ are determined via Bayes' rule from the equilibrium strategy profile. By Proposition B.3, it is without loss of generality to assume that $(M^{\mathbf{M}_t}, S^{\mathbf{M}_t}) = (\Theta, \Delta(\Theta))$. Define a new mechanism $\mathbf{M}'_t$ as follows. Let $(M^{\mathbf{M}'_t}, S^{\mathbf{M}'_t}) = (M^{\mathbf{M}_t}, S^{\mathbf{M}_t})$ and define the transition probability $\beta^{\mathbf{M}'_t}$ from Θ_+ to $\Delta(\Theta) \times A$ on the measurable rectangles as follows

$$\beta^{\mathbf{M}'_t}(\tilde{U} \times \tilde{A} | \theta) = \pi_t(\theta, h^t_A, \mathbf{M}_t) \beta^{\mathbf{M}_t}(\tilde{S} \times \tilde{A} | \theta) + (1 - \pi_t(\theta, h^t_A, \mathbf{M}_t)) \mathbb{1}[\mu^\varnothing \in \tilde{U}, a^* \in \tilde{A}].$$

Define a joint probability over $\Theta \times \{0,1\} \times \Delta(\Theta) \times A$ as follows:

$$\mathbb{P}(\tilde{\Theta} \times \{0\} \times \tilde{U} \times \tilde{A}) = \mathbb{1}[\mu^\varnothing \in \tilde{U}, a^* \in \tilde{A}] \int_{\tilde{\Theta}} (1 - \pi_t(\theta, h_A^t, \mathbf{M}_t)) \mu_t(d(\theta, h_A^t)|h^t),$$

$$\mathbb{P}(\tilde{\Theta} \times \{1\} \times \tilde{U} \times \tilde{A}) = \int_{\tilde{\Theta}} \beta^{\mathbf{M}_t}(\tilde{U} \times \tilde{A}|\theta) \pi_t(\theta, h_A^t, \mathbf{M}_t) \mu_t(d(\theta, h_A^t)|h^t).$$

The disintegration theorem implies that we can write for $n \in \{0,1\}$

$$\mathbb{P}(\tilde{\Theta} \times \{n\} \times \tilde{U} \times \tilde{A}) = \int_{\tilde{U} \times \tilde{A}} \eta_{(\mu, a_t)}(\tilde{\Theta} \times \{n\}) \mathbb{P}_{\Delta(\Theta) \times A}(d(\mu, a_t)).$$

Let $q = \eta_{(\mu^\varnothing, a^*)}(\Theta \times \{0\})$. Modify the continuation strategy so that for $\omega \in [0, q)$[42]

$$(\sigma_P', \sigma_A')|_{(h^t, z_{\mu^\varnothing, a^*}(\mathbf{M}_t'), \omega)} = (\sigma_P, \sigma_A)|_{(h^t, z_\varnothing(\mathbf{M}_t), \frac{\omega}{q})},$$

while for $\omega \in (q, 1]$,

$$(\sigma_P', \sigma_A')|_{(h^t, z_{(\mu^\varnothing, a^*)}(\mathbf{M}_t'), \omega)} = (\sigma_P, \sigma_A)|_{(h^t, z_{(\mu^\varnothing, a^*)}(\mathbf{M}_t), \frac{\omega - q}{1 - q})}.$$

Modify the agent's strategy so that for types in Θ_+, $r_t'(\theta, h_A^t, \mathbf{M}_t') = \delta_\theta$ and $\pi_t'(\theta, h_A^t, \mathbf{M}_t') = 1$. For types not in Θ_+, extend $\beta^{\mathbf{M}_t'}$ to all of Θ as we did in the proof of Proposition B.2. Set $r_t'(\theta, h_A^t, \mathbf{M}_t') = \delta_\theta$. Their payoff from participating may be lower, so recompute their participation strategy accordingly. It is straightforward to check that the new assessment is also a Perfect Bayesian equilibrium. $\square$

Propositions B.1-B.4 imply that for any equilibrium assessment, $(\sigma_P, \sigma_A, \mu)$, of the mechanism-selection game $G_\mathcal{M}$, there exists an equilibrium assessment, $(\sigma_P', \sigma_A', \mu')$, of the *auxiliary game* that satisfies that for all periods t and public histories h^t,(i) the principal offers canonical mechanisms, (ii) the agent's strategy satisfies the properties listed in Theorem 1, and (iii) the beliefs employed by the mechanism coincide with the principal's equilibrium beliefs. Moreover, $(\sigma_P', \sigma_A', \mu')$ implements the same distribution over outcomes $\Theta \times A^{T+1}$ as $(\sigma_P, \sigma_A, \mu)$ does.

Remark 1 implies that we can lift $(\sigma_P', \sigma_A', \mu')$ to an outcome-equivalent assessment of the canonical game, $(\tilde{\sigma}_P, \tilde{\sigma}_A, \tilde{\mu})$. Furthermore, the construction highlights that whatever the principal can achieve starting at any public history h^t with mechanisms in $\mathcal{M}$, he can also alternatively achieve with canonical mechanisms. Thus, it follows that for any PBE outcome $\gamma \in \Delta(\Theta \times A^{T+1})$ of the canonical game, there is an

---

[42]The notation assumes that the public randomization device is drawn $U[0,1]$. This is without loss because of the probability integral transform (see the proof of Proposition B.3).

outcome-equivalent assessment $(\sigma_P, \sigma_A, \mu)$ of the mechanism-selection game. We do not include the proof of this construction since it readily follows from the above. We do note that in the mechanism-selection game the public randomization device in the canonical game must be subsumed in the mechanism. That this is feasible follows from Kuratowski's theorem since $S^{\mathbf{M}_t}$ and $\Delta(\Theta) \times [0, 1]$ have the same cardinality.