

Mechanism Design with Limited Commitment*

Laura Doval[†] Vasiliki Skreta[‡]

First version: April 8, 2018. This version: May 1, 2019

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 payoffs 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 one: To the usual truth-telling 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

*Vasiliki Skreta acknowledges funding by the European Research Council (ERC) consolidator grant “Frontiers In Design.” Other acknowledgements to be added.

[†]California Institute of Technology, Pasadena, CA 91125. E-mail: ldoval@caltech.edu

[‡]University of Texas at Austin, University College London, and CEPR. E-mail: vskreta@gmail.com

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 done 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 only commit 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 \(2018\)](#) 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 payoffs of the mechanism-selection game. Second, we show how this insight can be used to transform the designer’s problem into a constrained optimization one: To the usual truth-telling 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 *communication device* and an *allocation rule* as

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

illustrated in Figure 1:²




Figure 1: Mechanisms: communication device, (M, β, S) , and allocation rule, α

Having observed her private information (her type, $v \in V$), the agent privately reports an *input* message, $m \in M$, into the mechanism; this then determines the distribution, $\beta(\cdot|m)$, from which an *output* message, $s \in S$, is drawn. In turn, the output message determines the distribution, $\alpha(\cdot|s)$, from which the allocation is drawn. The output message and the allocation are publicly observable: They constitute the contractible parts of the mechanism.

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

It is then clear why restricting attention to mechanisms that satisfy properties (i)-(iii) and truth-telling equilibria is with loss of generality under limited commitment: Upon observing the output message, the designer learns the agent’s type

²Myerson (1982) allows the designer to choose, as a function of the input message m , any joint distribution over $S \times A$. It is a consequence of Theorem 3.1 that the mechanisms in Figure 1 are without loss of generality; see Appendix II for a proof. Since the formulation of a mechanism in Figure 1 allows us to highlight the role of the communication device separately from that of the allocation, we opt for this formulation for pedagogical purposes.

report and hence her type. Then the agent may have an incentive to misreport if the designer cannot commit not to 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. The authors restrict attention to mechanisms in which the cardinality of the set of input and output messages is the same and β is “invertible.” 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 restriction on the class of mechanisms (i.e., (ii) and (iii) above) and show in a one-period model that focusing on mechanisms in 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 not clear whether taking input messages to be type reports is without loss when the designer and the agent interact repeatedly (see the discussion after [Theorem 3.1](#)).

The main contribution of this paper is to show that, under limited commitment, taking the set of output messages to be the set of posterior beliefs of the designer about the agent’s type, that is, $S = \Delta(V)$, is without loss of generality. [Theorem 3.1](#) shows that in a general mechanism-selection game between an uninformed designer and an informed agent introduced in [Section 2](#), any equilibrium payoff can be replicated by an equilibrium in which (a) the designer uses mechanisms in which input messages are type reports and output messages are beliefs, (b) the agent always participates in the mechanism, and (c) input and output messages have a *literal* meaning: The agent reports her type truthfully, and if the mechanism outputs a given posterior, this posterior coincides with the belief the designer holds about the agent’s type given the agent’s strategy and the mechanism. Given that any equilibrium payoff 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 3.1](#) implies that when the designer is subject to sequential rationality constraints, 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 in-

formation 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.³ 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.

Another contribution of our analysis is to show that to characterize equilibrium payoffs of the game between the designer and the agent, it suffices to consider a simpler game, denoted the *canonical game*. We record this result in Proposition 3.1. In the canonical game—studied in Section 3.2—the designer is restricted to offer mechanisms in which input messages are type reports and output messages are beliefs over the agent’s type. Theorem 3.1 (trivially) implies an equilibrium outcome of the canonical game can be achieved by strategy profiles in which the principal employs mechanisms that induce the agent to truthfully report his type and to always participate. However, the principal has *fewer deviations* in the canonical game and an equilibrium strategy may not be an equilibrium if the principal can deviate to *any* mechanism, as he can in the mechanism-selection game. One may then wonder whether analyzing the canonical game gives, unintentionally, some commitment power to the principal.

Proposition 3.1 shows this is not the case: Leveraging the construction used to establish Theorem 3.1, we show that, without loss of generality, the *best deviation* in the mechanism-selection game is equivalent to a deviation to a canonical mechanism that induces the agent to report truthfully and to participate with probability one. In a finite horizon setting, Proposition 3.1 justifies writing the designer’s problem as a sequence of maximization problems over canonical mechanisms subject to the agent’s participation and incentive compatibility constraints and the designer’s sequential rationality constraints.

³See Section 5.1 for a formal discussion of the approach and its potential issues.

Section 4 illustrates the methodology for the case of transferable utility and preferences that satisfy increasing differences in distributions. The resulting program allows us to highlight the connection between our problem and the literature on information design; after all, the designer can be thought of as a sender who designs the information structure for a receiver, who happens to be his future self. However, there are differences. In our setting, the first-period principal (the sender in [Kamenica and Gentzkow \(2011\)](#)) also takes an action for each posterior he induces. In addition, the first-period principal’s objective function depends on the prior as well as the posterior, whereas in [Kamenica and Gentzkow \(2011\)](#), it only depends on the posterior. Finally, the first-period principal cannot implement any Bayes’ plausible distribution over posteriors, but only those that satisfy the incentive compatibility and participation constraints of the agent.

An important difference between the mechanisms used by [Hart and Tirole \(1988\)](#), [Laffont and Tirole \(1988\)](#), [Freixas et al. \(1985\)](#), and [Bester and Strausz \(2001\)](#) and the ones considered here is that whereas in the former papers, the principal observes the agent’s choice out of a menu of *contracts*, here, the agent’s input into the communication device is not observed. Under the assumptions of Section 4, Proposition 5.2 in Section 5.2 characterizes the mechanisms (i.e., the communication device and allocation pairs) that can be implemented with the agent making a choice out of a menu. The result is useful for the following reasons. First, by checking whether the solution to the program studied in Section 4 satisfies the conditions in Proposition 5.2, we can understand whether the modeling of a mechanism as a menu of contracts in the aforementioned works is without loss. Second, when the solution to the program does satisfy the conditions, it allows the analyst to propose a “simple” implementation of the optimum.

The paper contributes to the literature on mechanism design with limited commitment, referenced throughout the introduction.⁴ 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](#)

⁴A designer’s lack of commitment can take various forms, not considered in this paper, but that have been studied in other papers. See, for instance, [McAdams and Schwarz \(2007\)](#), [Vartiainen \(2013\)](#), and [Akbarpour and Li \(2018\)](#), in which the designer cannot commit even to the obey the rules of the current mechanism.

(1981)) study price dynamics and establish (under some conditions) Coase’s conjecture whereby a monopolist essentially loses all profits if he lacks commitment. In an analogous vein, [Burguet and Sakovics \(1996\)](#), [McAfee and Vincent \(1997\)](#), [Caillaud and Mezzetti \(2004\)](#), and [Liu et al. \(2018\)](#) 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.

Mechanism-selection in a dynamic environment with limited commitment is considered in [Deb and Said \(2015\)](#). The authors study a model of sequential screening, in which new buyers arrive over time. Like in [Skreta \(2006\)](#) and [Skreta \(2015\)](#), [Deb and Said \(2015\)](#) consider general mechanisms but a finitely long interaction. Infinitely long contract-selection games are studied in [Strulovici \(2017\)](#) and [Gerardi and Maestri \(2018\)](#). The former studies renegotiation and finds that equilibrium allocations become efficient as the parties become arbitrarily patient. In [Gerardi and Maestri \(2018\)](#), however, the limit allocation is inefficient whenever firing the agent—what the authors refer to “firing allocation”—is not a solution when there is commitment.

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 \(2017\)](#) and [Best and Quigley \(2017\)](#), who study models of direct communication between an informed sender and an uninformed receiver.⁵ [Lipnowski and Ravid \(2017\)](#) show how the posterior approach of [Kamenica and Gentzkow \(2011\)](#) can be used to characterize equilibrium outcomes, and study their properties in the cheap talk model of [Crawford and Sobel \(1982\)](#) (the leading model of communication without commitment), when the sender’s preferences do not depend on the state of the world.⁶ Finally, given that the search for the best equilibrium often reduces to solving a constrained information design problem we relate to, among others, [Le Treust and Tomala \(2017\)](#), [Georgiadis and Szentes \(2018\)](#), and [Boleslavsky and Kim \(2018\)](#).

⁵[Salamanca \(2016\)](#) studies mediated communication in [Kamenica and Gentzkow \(2011\)](#).

⁶[Golosov and Iovino \(2016\)](#) study a social insurance model with a continuum of agents, where private information is not persistent across stages. They leverage the resulting repeated-game structure to solve for the best equilibrium.

The rest of the paper is organized as follows. Section 2 describes the model and notation. Section 2.1 analyzes a simple version of the model in Skreta (2006); it allows us to introduce the main ideas of the paper in a simple and well-known setting. Section 2.2 discusses the modeling assumptions. Section 3 introduces the main theorem and provides a sketch of the proof. Section 4 specializes the results to the two-period model of Bester and Strausz (2007) with transferable utility and single-crossing preferences. We compare the solution of the ‘relaxed’ problem to the information design model of Kamenica and Gentzkow (2011). Section 5.1 discusses using recommendations as output messages. Section 5.2 studies implementation when the principal observes the agent’s choice. Section 5.3 discusses an example with multiple agents. All proofs are relegated to the Appendix. The supplementary material (Sections I-IV) contains omitted proofs and extensions discussed throughout the main text.

2 MODEL

Primitives There are two players: a principal (he) and an agent (she). They interact over $T \leq \infty$ periods. Before the game starts, the agent observes her type, $v \in V$. V is any finite set; however, the main insights extend to the case in which V is a Polish space (see Appendix IV). Each period, as a result of the interaction between the principal and the agent, an allocation $a \in A$ is determined. Assume A is a compact (possibly finite) space.

Given a sequence of allocations $a^t = (a_0, a_1, \dots, a_t)$, the principal can only choose $a_{t+1} \in \mathcal{A}(a^t)$. That is, there is a correspondence $\mathcal{A} : \bigcup_{n=0}^T A^n \mapsto A$ such that for $t \in \mathbb{N}$, $a^t \in A^t$, $\mathcal{A}(a^t)$ describes the set of allocations the principal can offer given the allocations he has offered in the past. Assume \mathcal{A} is compact-valued and there exists an allocation $a^* \in A$ such that a^* is always available.⁷

Payoffs are defined as follows. For the principal, assume there exists a function,

⁷We later use allocation a^* to model the agent’s participation decision within each period: If the agent prefers not to participate, allocation a^* is implemented automatically. For instance, in a trade model such as the one in Section 2.1, a^* corresponds to no trade and no transfers. The constraint correspondence \mathcal{A} also allows us to capture that the agent can walk away from the mechanism as in Gerardi and Maestri (2018): We could specify that the first time a^* is implemented, then this allocation is the only one available thereafter.

$W : A^T \times V \mapsto \mathbb{R}$ such that his payoff from allocation $a \in A^T$ when the agent's type is v is given by $W(a, v)$. Similarly for the agent, when her type is v , her payoff from allocation $a \in A^T$ is given by $U(a, v)$.

Mechanisms: In each period, the principal offers the agent a mechanism, $\mathbf{M}_t = \langle (M^{\mathbf{M}_t}, \beta^{\mathbf{M}_t}, S^{\mathbf{M}_t}), \alpha^{\mathbf{M}_t} \rangle$, which consists of a *communication device*, $(M^{\mathbf{M}_t}, \beta^{\mathbf{M}_t}, S^{\mathbf{M}_t})$, and an *allocation rule*, $\alpha^{\mathbf{M}_t}$, where

$$\begin{aligned}\beta^{\mathbf{M}_t} &: M^{\mathbf{M}_t} \mapsto \Delta^*(S) \\ \alpha^{\mathbf{M}_t} &: S^{\mathbf{M}_t} \mapsto \Delta^*(\mathcal{A}),\end{aligned}$$

and where $\Delta^*(C)$ denotes the set of distributions on C with *finite* support. We endow the principal with a collection $(M_i, S_i)_{i \in \mathcal{I}}$ of input and output message sets in which each M_i is finite, $|V| \leq |M_i|$, and $\Delta(M_i) \subseteq S_i$.⁸ Moreover, we assume $(V, \Delta(V))$ is an element in that collection. Denote by \mathcal{M} the set of all mechanisms with message sets $(M_i, S_i)_{i \in \mathcal{I}}$. A mechanism is *canonical* if $(V, \Delta(V))$ are its sets of input and output messages. Let \mathcal{M}^C denote the set of canonical mechanisms and let \mathbf{M}_t^C denote an element in that set.

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.⁹ Second, we restrict the principal to design $\beta^{\mathbf{M}_t}$ and $\alpha^{\mathbf{M}_t}$ to be distributions with finite support, thus allowing us to focus on the novel conceptual features of the environment, as opposed to dealing with measure-theoretic complications. To replicate any equilibrium of the game when the principal selects distributions with finite support using canonical mechanisms, we find the principal only needs to use distributions with finite support. This last observation, of course, would not be true if the set of types were not finite.^{10, 11} Finally, we re-

⁸Technically, we only need that S_i contains an image of $\Delta(M_i)$.

⁹To see this, suppose the principal would rather use a mechanism, \mathbf{M}'_t , with a message space $M^{\mathbf{M}'_t}$ with cardinality strictly less than $|V|$. Then he can choose a mechanism \mathbf{M}_t with $M^{\mathbf{M}_t} = V$, choose β 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.

¹⁰Appendix IV extends our result to the case in which V is a compact and metrizable space.

¹¹We conjecture, however, that the restriction to distributions with finite support is without

strict the principal to choose 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, 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.

Timing: In each period t ,

- The principal and the agent observe a draw from a correlating device $\omega \sim U[0, 1]$.
- The principal offers the agent a mechanism \mathbf{M}_t .
- The agent observes the mechanism and decides whether to participate ($p = 1$) or not ($p = 0$). If she does not participate, a^* is implemented and the game proceeds to $t + 1$.
- If she participates, she privately submits a report $m \in M^{\mathbf{M}_t}$.
- $s \in S^{\mathbf{M}_t}$ is drawn according to $\beta^{\mathbf{M}_t}(\cdot|m)$, which is publicly observed.
- $a \in A$ is drawn according to $\alpha^{\mathbf{M}_t}(\cdot|s)$, which is publicly observed.

This defines an extensive form game, which we dub the *mechanism-selection game*. If, instead, the principal can only choose mechanisms in \mathcal{M}^C , we denote it as the *canonical game*.

Public histories in this game are¹²

$$h^t = (\omega_0, \mathbf{M}_0, p_0, s_0, a_0, \dots, \omega_{t-1}, \mathbf{M}_{t-1}, p_{t-1}, s_{t-1}, a_{t-1}, \omega_t),$$

where $p_r \in \{0, 1\}$ denotes the agent's participation with the restriction that $p_r = 0 \Rightarrow s_r = \emptyset, a_r = a^*$. Public histories capture what the principal knows through period t . Let H^t denote the set of all period t public histories. A strategy for the principal is then given by $\Gamma_t : H^t \mapsto \Delta(\mathcal{M})$.

loss of generality when the set of types is finite.

¹²The restriction that the support of $\beta^{\mathbf{M}_t}(\cdot|m)$ is finite for $m \in M^{\mathbf{M}_t}$, together with the finiteness of $M^{\mathbf{M}_t}$, imply that there are output messages $s \in S^{\mathbf{M}_t}$ that can never arise. Thus, we can remove from the tree all the histories that are consistent with mechanism \mathbf{M}_t being offered and $s \in S^{\mathbf{M}_t}$ such that $\sum_{m \in M^{\mathbf{M}_t}} \beta^{\mathbf{M}_t}(s|m) = 0$, without affecting the equilibrium set. However, for tractability, we do not make this distinction in our notation.

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 element

$$h_A^t = (\omega_0, \mathbf{M}_0, m_0, p_0, s_0, a_0, \dots, \omega_{t-1}, \mathbf{M}_{t-1}, p_{t-1}, m_{t-1}, s_{t-1}, a_{t-1}, \omega_t).$$

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 $\{v\} \times H_A^t$ when her type is v . The agent's participation strategy is $\pi_v : H_A^t \times \mathcal{M}_t \mapsto [0, 1]$. Conditional on participating in the mechanism \mathbf{M}_t , her reporting strategy is a distribution $r_v(h_A^t, \mathbf{M}_t, 1) \in \Delta(M^{\mathbf{M}_t})$ for each of her types v and each $h_A^t \in H_A^t$.

A belief for the principal at the beginning of time t , history h^t , is a distribution $\mu(h^t) \in \Delta(V \times H_A^t(h^t))$, where $H_A^t(h^t)$ is the set of agent histories that are consistent with the public history h^t , which is observed by the principal. The principal is thus uncertain both about the agent's payoff-relevant type, v , and her payoff-irrelevant type, h_A^t .

Our focus is on studying the equilibria of the mechanism-selection and canonical games. By equilibrium, we mean Perfect Bayesian equilibrium (henceforth, PBE), defined as follows:

Definition 2.1. A *Perfect Bayesian Equilibrium* is a tuple $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$ such that for each h^t the following hold:

1. Given $\mu_t^*(h^t)$, $\Gamma_t^*(h^t)$ is sequentially rational given $(\pi_v^*, r_v^*)_{v \in V}$,
2. Given $\Gamma^*(h^t)$, $\pi_v^*(h_A^t, \cdot), r_v^*(h_A^t, \cdot, 1)$ are sequentially rational for all $h_A^t \in H_A^t(h^t)$,
3. $\mu^*(h^t)$ is derived via Bayes' rule whenever possible.

Implicit in the definition of PBE is the assumption that the principal does not update his beliefs about the agent following a deviation by the principal. That is, we assume beliefs are *pre-consistent* (see [Hendon et al. \(1996\)](#)).

Remark 2.1. [Belief updating depends only on the realized output message] Fix

a history h^t . Given $\mu \in \Delta(V \times H_A^t(h^t))$ and a mechanism $M^{\mathbf{M}_t}$, Bayesian updating depends on the agent's strategy and the communication device, but not on the allocation rule. To see this, suppose the agent participates with positive probability in the mechanism, the output message is $s \in S^{\mathbf{M}_t}$ and allocation a is observed; then the principal's belief about the agent being at history $(v, h_A^t, \mathbf{M}_t, 1, s, a)$ is given by

$$\frac{\mu(h^t)(v, h_A^t) \pi_v^*(h_A^t, \mathbf{M}_t) \sum_{m \in M^{\mathbf{M}_t}} r_v^*(h_A^t, \mathbf{M}_t, 1)(m) \beta^{\mathbf{M}_t}(s|m) \alpha^{\mathbf{M}_t}(a|s)}{\sum_{\tilde{v}, \tilde{h}_A^t} \mu(\tilde{v}, \tilde{h}_A^t) \pi_{\tilde{v}}^*(\tilde{h}_A^t, \mathbf{M}_t) \sum_{\tilde{m} \in M^{\mathbf{M}_t}} r_{\tilde{v}}^*(\tilde{h}_A^t, \mathbf{M}_t, 1)(\tilde{m}) \beta^{\mathbf{M}_t}(s|\tilde{m}) \alpha^{\mathbf{M}_t}(a|s)},$$

and all the terms concerning $\alpha^{\mathbf{M}_t}(a|s)$ drop out.

2.1 Example: two-period sale of a durable good.

To fix ideas, consider the following example. The principal is a seller who owns one unit of a durable good and assigns value 0 to it. The agent is a buyer whose valuation for the good is her private information. The buyer's valuation can take two values, $v \in \{v_L, v_H\}$, where $v_H - v_L > 0$. The seller's belief that $v = v_H$ is μ_1 . An allocation is a pair $(q, t) \in \{0, 1\} \times \mathbb{R}$, where q indicates whether a sale occurs ($q = 1$) or not ($q = 0$), and t is a transfer from the buyer to the seller. Utilities are quasilinear so that the buyer's utility is $u(q, t; v) = vq - t$ and the seller's is $w(q, t; v) = t$. Both players share a common discount factor $\delta \in (0, 1)$.

The timing is as follows: in each period $t \in \{1, 2\}$

- The seller chooses a mechanism.
- The buyer observes the mechanism and decides whether to participate.
 - If she does not participate, the good is not sold and no payments are made; if $t = 1$, we move to period 2.
 - If she participates, the mechanism determines the allocation.
- If the good is not sold and $t = 1$, move on to $t = 2$.

Because the horizon is finite, we can solve the game by backward induction. Then let $t = 2$ and denote by μ_2 the seller's posterior belief that $v = v_H$. In $t = 2$,

the seller has full commitment and the solution is routine. The seller posts a price equal to v_L when $\mu_2 < v_L/v_H \equiv \bar{\mu}$, a price equal to v_H when $\mu_2 > \bar{\mu}$, and at $\mu_2 = \bar{\mu}$, then the seller is indifferent between the two prices. Thus, the seller's revenue as a function of μ_2 is given by

$$R_2(\mu_2) = \begin{cases} v_L & \text{if } \mu_2 \leq \bar{\mu} \\ \mu_2 v_H & \text{otherwise} \end{cases} = \begin{cases} \mu_2 v_H + (1 - \mu_2) \hat{v}_L(\mu_2) & \text{if } \mu_2 \leq \bar{\mu} \\ \mu_2 v_H & \text{otherwise} \end{cases},$$

where $\hat{v}_L(\mu_2) = v_L - (\mu_2/(1 - \mu_2))(v_H - v_L)$ and the equality follows from noting that when the price is v_L , the seller leaves rents $v_H - v_L$ with probability μ_2 to the high type.

We now turn to period 1. Recall that μ_1 denotes the probability that the buyer's valuation is v_H . Consistent with the mechanism-selection game introduced in Section 2, we allow the seller to offer the buyer a mechanism that consists of a communication device (M_1, β_1, S_1) and an allocation rule $\alpha : S_1 \mapsto \Delta(\{0, 1\} \times \mathbb{R})$. The assumption of quasilinearity implies that, without loss of generality, the seller does not randomize on the transfers, so that $\alpha_1(q, t|s_1) = q(s_1) \times \mathbb{1}[t = t(s_1)]$.

Theorem 3.1 shows that, without loss of generality, input messages are type reports, $M_1 = V$, and output messages are the seller's beliefs about the buyer's valuation, $S_1 = \Delta(V)$. We now provide intuition for this in the context of the example.

1. To see that $M_1 = V$, note that β_1 together with the agent's reporting strategy induces another distribution on S_1 ,

$$\sum_{m_1 \in M_1} \beta_1(s_1|m_1) r_v(m_1) \equiv \beta^*(s_1|v).$$

If the seller offers $\langle (V, \beta^*, S_1), \alpha_1 \rangle$ to the buyer, then the buyer tells the truth (see also [Bester and Strausz \(2007\)](#)).

2. To see why S_1 can be taken to be $\Delta(V)$, note that upon the realization of s_1 , two things happen. First, the allocation $\alpha(s_1)$ is determined. Second, if the

allocation is no trade, s_1 is used to update the principal's beliefs as follows:

$$\mu_2(v = v_H | s_1) \left(\sum_{v \in V} \mu_1(v) \beta^*(s_1 | v) \right) = \mu_1(v_H) \beta^*(s_1 | v_H),$$

where we have already used that $M_1 = V$ and the buyer reports truthfully. Given the belief induced by s_1 , we know what happens in period 2; there is no use for s_1 beyond that. Thus, we can take $S_1 = \Delta(V)$. Thus, we write $\beta(\mu_2 | v), q(\mu_2), t(\mu_2)$ instead of $\beta(s_1 | v), q(s_1), t(s_1)$ thereafter.

With these observations, we can describe the seller's optimal mechanism in period 1 via the following program:

$$R_1(\mu_1) \equiv \max_{\beta, q, t} \sum_{\mu_2 \in \Delta(V)} \left(\sum_{v \in V} \mu_1(v) \beta(\mu_2 | v) \right) [t(\mu_2) + (1 - q(\mu_2)) \delta R_2(\mu_2)]$$

subject to for all $v \in \{v_L, v_H\}$:

$$\begin{aligned} \sum_{\mu_2 \in \Delta(V)} \beta(\mu_2 | v) (v q(\mu_2) - t(\mu_2) + \delta(1 - q(\mu_2)) u_B(\mu_2, v)) &\geq 0 \text{ (PC}_v\text{)} \\ \sum_{\mu_2 \in \Delta(V)} (\beta(\mu_2 | v) - \beta(\mu_2 | v')) (v q(\mu_2) - t(\mu_2) + \delta(1 - q(\mu_2)) u_B(\mu_2; v)) &\geq 0 \text{ (IC}_{v, v'}\text{)} \\ \mu_2(v_H) \left(\sum_{v \in V} \mu_1(v) \beta(\mu_2 | v) \right) &= \mu_1(v_H) \beta(\mu_2 | v_H) \text{ (BC}_{\mu_2}\text{)}. \end{aligned}$$

That is, the seller chooses β, q, t to maximize his profit subject to the agent's participation and incentive compatibility constraint and a Bayesian consistency constraint. The latter says that when the mechanism outputs μ_2 , then μ_2 is the belief that obtains via Bayesian updating. The buyer's participation and incentive compatibility constraints take into account her continuation values, denoted by $u_B(\mu_2, v)$: For low values of μ_2 , the high type is served at a low price in period 2.¹³

¹³Implicit in the buyer's participation constraint is that, if she does not participate in $t = 1$, the seller has belief $\mu_2 \geq \bar{\mu}$ and then sets a price of v_H in $t = 2$. Thus, both types of the agent earn a payoff of 0 in case they do not participate in $t = 1$. Given Theorem 3.1, this is without loss of optimality. Indeed, Theorem 3.1 shows that it is without loss of generality to have the agent participate with probability 1. Hence, not participating of the mechanism becomes an off-path event and beliefs are not pinned down by Bayes' rule in this case.

As usual, we can show that PC_{v_L} and IC_{v_H, v_L} bind, and these two constraints imply the others. Therefore, we can use them to replace the transfers in the seller's objective to obtain

$$R_1(\mu_1) \equiv \max_{\tau, q} \sum_{\mu_2 \in \Delta(V)} \tau(\mu_2) [q(\mu_2)(\mu_2 v_H + (1 - \mu_2)\hat{v}_L(\mu_1)) + \delta(1 - q(\mu_2))R_2(\mu_2; \mu_1)] \quad (1)$$

$$\text{s.t.} \quad \sum_{\mu_2 \in \Delta(V)} \tau(\mu_2)\mu_2 = \mu_1,$$

where $\tau(\mu_2) = \sum_{v \in V} \mu_1(v)\beta(\mu_2|v)$ is the probability that μ_2 is the induced posterior and

$$\delta R_2(\mu_2; \mu_1) = \begin{cases} \delta(\mu_2 v_H + (1 - \mu_2)\hat{v}_L(\mu_1)) & \text{if } \mu_2 < \bar{\mu} \\ \delta\mu_2 v_H & \text{if } \mu_2 > \bar{\mu} \end{cases} \quad (2)$$

is an adjusted version of the seller's period 2 revenue. We now explain equations (1) and (2) in detail. Equation (1) shows that the seller's period 1 problem can be solved by finding (i) a trade probability for each posterior and (ii) a distribution over posteriors that averages out to the prior. Given a posterior μ_2 , the trade probability, $q(\mu_2)$, is chosen to maximize a version of the virtual surplus, familiar from mechanism design with commitment. To see this, note that in equation (1), the probability of each type is evaluated using the posterior μ_2 , but the virtual value for the low type is computed using the prior μ_1 . This reflects that the seller in period 1 assigns probability μ_1 to $v = v_H$, and μ_1 is the rate at which he pays rents to the high type. Similarly, $\delta R_2(\mu_2; \mu_1)$ adjusts the revenues in period 2 by the rents the period 1 seller must leave to the buyer: In effect, should the period 1 seller induce $\mu_2 < \bar{\mu}$, the buyer obtains a rent of $\delta(v_H - v_L)$, which the seller in period 1 has to take into account.

Note that equation (2) does not specify what the seller's payoff is when $\mu_2 = \bar{\mu}$. This, in fact, depends on the prior μ_1 : When $\mu_2 = \bar{\mu}$, the period 2 seller is indifferent between prices v_H and v_L . The period 1 seller, however, is not indifferent; this fact is illustrated in Figures 2 and 3 below:




Figure 2: $\mu_1 \leq \bar{\mu}$




Figure 3: $\bar{\mu} < \mu_1$

If the seller's prior is such that he would sell to the low valuation buyer today ($\mu_1 \leq \bar{\mu}$), then he would rather have the period 2 seller also serve the low valuation type when indifferent in period 2, as illustrated in Figure 2. However, if the period 1 seller would prefer to exclude the buyer when her valuation is low, then he would prefer the low valuation buyer to be excluded in period 2 as well when $\mu_2 = \bar{\mu}$, as illustrated in Figure 3.

In what follows, we solve the seller's problem for the case in which $\mu_1 > \bar{\mu}$.¹⁴ Because the seller can choose $q(\mu_2)$ for each μ_2 , the best he can do is choose it to pointwise maximize the objective function in equation (1), as illustrated in Figures 4 and 5 below:

¹⁴The case in which $\mu_1 < \bar{\mu}$ is immediate: The seller can achieve the commitment solution by selling to both types of the buyer in period 1.




Figure 4: Value of setting $q(\mu_2) = 0$ (black) and $q(\mu_2) = 1$ (blue)




Figure 5: Pointwise maximum of the blue and black lines in Figure 4

This reduces the principal's problem to that of finding a distribution over posteriors to solve:

$$\max_{\tau} \sum_{\mu_2 \in \Delta(V)} \tau(\mu_2) \max\{\mu_2 v_H + (1 - \mu_2) \hat{v}_L(\mu_1), \delta R(\mu_2; \mu_1)\},$$

subject to the constraint that the distribution must average to the prior. Under some parameter values, the solution is as depicted in Figure 6 below:¹⁵

¹⁵When $\bar{\mu} < \mu_1$, there are two possible solutions depending on the parameter values. When $\mu_1 > \bar{\mu}$ is high enough, we obtain the solution depicted in Figure 6 and described in the main text. For lower values of $\mu_1 > \bar{\mu}$, we obtain the solution familiar to the literature on the ratchet effect (see Hart and Tirole (1988)). In this case, the seller sets a price of v_L in period 2, and a price of $v_H - \delta \Delta v$ in period 1; the buyer buys in period 1 when $v = v_H$ and in period 2 when $v = v_L$.




Figure 6: Optimal mechanism for period 1 seller: The black dashed line depicts the concavification of the function in Figure 5

The seller sets a price of v_H in periods 1 and 2. The buyer does not buy when her value is low, whereas the buyer randomizes between buying today and tomorrow when her value is high. The randomization is such that when the seller sees no sale at the end of period 1, he attaches probability $\bar{\mu}$ to $v = v_H$.

The example highlights both how the language of type reports and posterior beliefs is enough to replicate what the principal can obtain from any other mechanism and also how useful this language is to solve mechanism design problems with limited commitment. Indeed, it allows us to reduce the problem of finding the best equilibrium for the principal to a constrained optimization problem. In Section 4, we return to the setting of transferable utility and preferences that satisfy increasing differences and show that the connection between our problem and information design extends beyond the example.

However, the example does not allow us to highlight some features of the model, which we discuss in the next section. The reader eager to see the results can skip to Section 3; however, the discussion may be useful to follow the proof sketch of the main theorem.

2.2 Discussion: Randomized allocations and public correlating device

We now discuss two aspects of the model that do not seem to play a role in the example, but are important in what follows: The principal is allowed to offer a randomization over allocations, and the principal and the agent have access to a public correlating device.

Randomized allocations There are two reasons for allowing the principal to choose randomized allocations. First, randomized allocations are necessary for the set of input messages to be the set of type reports; this is inherited from the revelation principle with commitment (see [Strausz \(2003\)](#)). To see this, consider the situation illustrated in Figure 7 below. The mechanism is simple: If the agent reports m , then the output message is m and the allocation is a , whereas if she reports m' , the output is m' and the allocation is a' . Assume that when her type is v , the agent sends m and m' with probability p and $1 - p$, respectively; thus, she obtains a and a' with probability p and $1 - p$, respectively.




Figure 7: Agent of type v randomizes over m and m' generating a randomized allocation

If we restrict the principal to offer deterministic allocations, then he cannot replicate the agent's allocation just by asking for a truthful type report. However, if we allow the principal to offer a mechanism such that when the input is v and the output is v , the allocation is a randomization between a and a' , then he can replicate the allocation type v obtains just by soliciting a type report.

Second, randomized allocations are necessary for the set of output messages to be the set of distributions over the agent's type. To see this, note that two different output messages, s and s' , may be associated with two different allocations, a and a' , but with the same posterior belief, as illustrated in Figure 8 below:




Figure 8: Two output messages, s and s' , induce same posterior but different allocations

By allowing the principal to offer randomized allocations, we can collapse s and s' to one output message s'' associated to one posterior, μ .

Public correlating device The correlating device is important for output messages to be the principal's posterior beliefs about the agent's type. Note that two output messages, s and s' , may be associated with two different continuation equilibria, even if they induce the same allocation and posterior beliefs, as in Figure 9 below:




Figure 9: Two output messages, s and s' , induce same posterior and allocations, but different continuation equilibria

The correlating device allows us to collapse s and s' into one output message (and hence, one posterior belief) and coordinate continuation play with the correlating device, akin to what is done in repeated games. This feature arises, somewhat trivially, in Section 2.1. In the example, for each posterior belief different from $\bar{\mu}$, there is a unique continuation equilibrium in period 2, and hence there is no need to select amongst continuation equilibria. However, when $\mu_2 = \bar{\mu}$, there are two continuation equilibria; when we allowed the first-period principal to select between them, we implicitly made use of a (trivial) correlating device.

3 RESULTS

Section 3 presents the main results of the paper. Theorem 3.1 shows that any equilibrium payoff of the mechanism-selection game can be replicated by an equilibrium in which (a) the designer uses mechanisms in which input messages are type reports and output messages are beliefs, (b) the agent always participates in the mechanism, and (c) input and output messages have a *literal* meaning: The agent reports her type truthfully, and if the mechanism outputs $\mu \in \Delta(V)$ at the end of period t , then μ is indeed the belief the principal holds about the agent at the end of that period. Motivated by Theorem 3.1, Section 3.2 studies the PBE of the *canonical game*. It follows immediately from Theorem 3.1 that any equilibrium payoff of the mechanism-selection game is also an equilibrium payoff of the canonical game, after adapting the strategy profiles and systems of beliefs to the canonical game. Because the canonical game has a smaller set of deviations for the principal than the mechanism-selection game, one may conjecture that there are dynamic mechanisms consistent with equilibrium in the canonical game, which would not be consistent with equilibrium in the mechanism-selection game. Proposition 3.1 shows this conjecture is false.

3.1 Revelation Principle for Sequentially Optimal Mechanism Design

Theorem 3.1. Fix any PBE of the mechanism-selection game, $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$. Then there exists a payoff-equivalent PBE, $\langle \Gamma', (\pi'_v, r'_v)_{v \in V}, \mu' \rangle$, such that

1. At all histories, the principal offers canonical mechanisms, that is, $(\forall h^t)(\forall \mathbf{M}_t : \Gamma'(h^t)(\mathbf{M}_t) > 0), \mathbf{M}_t \in \mathcal{C}$.
2. At all histories where the principal assigns positive probability to the agent's type being v , the agent participates with probability 1 when her type is v , that is, $(\forall v, \forall h_A^t \in H_A^t)(\forall \mathbf{M}_t \in \text{supp } \Gamma'(h^t)) \pi'_v(h_A^t, \mathbf{M}_t) = 1$ whenever $\mu'(h^t)(v) > 0$.
3. At all histories, the agent reports her type truthfully, that is, $(\forall v, \forall h_A^t \in H_A^t)(\forall \mathbf{M}_t \in \text{supp } \Gamma'(h^t)), r'_v(h_A^t, \mathbf{M}_t, 1) = \delta_v$.

4. At all histories, recommended beliefs coincide with realized beliefs $t + 1$:

$$\mu'(h^t, \mathbf{M}_t, 1, \mu)(v) = \frac{\mu'(h^t)(v)\beta^{\mathbf{M}_t}(\mu|v)}{\sum_{v' \in V} \mu'(h^t)(v')\beta^{\mathbf{M}_t}(\mu|v')} = \mu(v).$$

The proof is in Appendix B. In what follows, we provide a sketch of the main steps in the proof.

The first main step shows that, without loss of generality, the agent's participation and reporting strategy conditions only on her type v and the public history. This step, which follows from Proposition A.1 in Appendix A, is key to showing that the set of canonical input messages is the set of type reports. If the agent conditioned her strategy on the payoff-irrelevant part of her private history, the principal would need to elicit h_A^t together with v in order to replicate the agent's behavior in the mechanism.

We now qualify what we mean by *without loss of generality*: 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 payoff-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. The proof of this consists of two parts. First, we observe that because the input messages are payoff irrelevant and unobserved by the principal, if the agent chooses different strategies at (v, h_A^t) and $(v, h_A^{t'})$, with $h_A^t, h_A^{t'} \in H_A^t(h^t)$, then she is indifferent between these two strategies. However, the principal may not be indifferent between these two strategies. The second part shows we can build an alternative strategy that does not condition on h_A^t beyond h^t and gives the principal the same continuation payoff.

This first step also gives us an important conceptual insight: The principal cannot peek into his past correlating devices. To do so, he would like to ask the agent to report to him what she did in the previous mechanisms. Proposition A.1 shows that this information cannot be elicited in any payoff-relevant way.

The second main step shows that, without loss of generality, there is a one-to-one mapping between the output messages used at history h^t and the posterior beliefs

of the principal in the PBE at history h^t (see Proposition A.3 in Appendix A). This step follows mainly from the observations we made in the discussion in Section 2.2. The principal may have two other uses for the output messages. On the one hand, because the allocation must be measurable with respect to the output messages, he may use them to offer a richer set of alternatives. On the other hand, he may use the output messages to coordinate continuation play. Proposition A.3 shows that randomized allocations and the access to the public correlating device can achieve these two goals, respectively.

These two steps deliver that, without loss of generality, input messages can be taken to be type reports and output messages can be taken to be the designer's beliefs about the agent's type. After all, knowing the agent's type is all that is needed to replicate her behavior within the mechanism, and hence the relevant beliefs for the principal are about the agent's payoff-relevant type.

Proposition A.2 in Appendix A shows that having the agent participate in the mechanism is without loss of generality (we discuss at the end of the section why Theorem 3.1 only requires this for types with positive probability). The logic is similar to the one 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.

With these preliminary steps at hand, the proof of Theorem 3.1 in Appendix B shows that any mechanism \mathbf{M}_t offered by the principal at history h^t can be replicated by a canonical mechanism $\mathbf{M}_t^C = \langle (V, \beta^{\mathbf{M}_t^C}, \Delta(V)), \alpha^{\mathbf{M}_t^C} \rangle$ as follows. The second step implies that there is an invertible mapping which maps each output message into the belief over types that it induces:

$$\sigma(\mathbf{M}_t)(s_t) = \sum_{h_A^t \in H_A^t(h^t), m_t \in M^{\mathbf{M}_t}} \mu^*(h^t, \mathbf{M}_t, 1, s_t)(\cdot, h_A^t, \mathbf{M}_t, 1, m_t, s_t).$$

Note that we obtain the belief over V by taking the marginal over all agent histories

consistent with the public history h^t . Using this, we can define a communication device $\beta^{\mathbf{M}_t^C} : V \mapsto \Delta^*(\Delta(V))$ and an allocation rule $\alpha^{\mathbf{M}_t^C} : \Delta(V) \mapsto \Delta^*(A)$ as follows:

$$\begin{aligned}\beta^{\mathbf{M}_t^C}(\mu|v) &= \sum_{m \in M^{\mathbf{M}_t}} \beta^{\mathbf{M}_t}(\sigma^{-1}(\mathbf{M}_t)(\mu)|m)r_v^*(h_A^t, \mathbf{M}_t, 1)(m) \\ \alpha^{\mathbf{M}_t^C}(\mu) &= \alpha^{\mathbf{M}_t}(\sigma^{-1}(\mathbf{M}_t)(\mu)).\end{aligned}$$

The proof then shows that when faced with this mechanism, the agent's best response is to participate and report truthfully and that when the principal observes an output of μ , his beliefs are indeed μ .

We have yet to discuss why Theorem 3.1 only requires that the agent participates with probability 1 is required for her types to which the principal assigns positive probability. Consider then a history h^t such that the principal's belief assigns probability 0 to the agent's type being v^* . Suppose the principal selects mechanism \mathbf{M}_t . Assume also the agent's strategy at v^* specifies sending an input message m^* , which assigns positive probability to an output s^* . Finally, assume that s^* has zero probability under all other $m \in M^{\mathbf{M}_t}$. PBE does not impose restrictions on the principal's belief when he observes s^* ; in particular, it could be that $\mu^*(h^t, \mathbf{M}_t, 1, s^*) = \mu^*(h^t, \mathbf{M}_t, 1, s')$, where s' is an output message with positive probability under the equilibrium strategy. This would, of course, break the one-to-one mapping between output messages and posterior beliefs.

To deal with the aforementioned issue, we show that, given a PBE, we can always modify the mechanisms chosen in equilibrium by the principal so that the agent does not have access to messages like m^* . Namely, he can make the distribution of the communication device for any such message the same as that of a message that is used on the path. The principal can always do this without affecting the incentives of those types that have positive probability; however, he may change the participation incentives of those types that have probability 0.

3.2 The Canonical Game

Theorem 3.1 shows that any equilibrium payoff of the game between the principal and the agent can be achieved with the principal selecting at each history a canonical mechanism such that the agent participates with probability one and reports her type truthfully. This observation motivates the analysis in this section where we study the equilibria of the canonical game.

An immediate corollary of Theorem 3.1 is the following:

Corollary 3.1. *Any PBE payoff of the mechanism-selection game can be achieved as a PBE payoff of the canonical game.*

Because it features a restricted set of choices for the principal, one may suspect that in the canonical game, the principal is able to implement more mechanisms than in the mechanism-selection game. However, this is not the case. Indeed, we show that given any equilibrium of the mechanism-selection game, without loss of generality, the *best* deviation for the principal after *any* history can be achieved by offering a canonical mechanism for that period and also in the continuation histories, whereas the agent participates with probability one and truthfully reports her type. This observation implies the canonical game contains all relevant deviations for the principal. It is not then possible to achieve payoffs in the canonical game that cannot be achieved in the mechanism-selection game. This is recorded in Proposition 3.1 below:

Proposition 3.1. *If $\langle \Gamma^{*^C}, (\pi_v^{*^C}, r_v^{*^C})_{v \in V}, \mu^{*^C} \rangle$ is a PBE of the canonical game, then there is an equilibrium of the mechanism-selection game $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$ that achieves the same payoff.*

Two important lessons follow from Proposition 3.1 and its proof. First, to characterize the equilibrium payoffs of the mechanism-selection game, it suffices to characterize the equilibrium payoffs of the canonical game.¹⁶ Second, the proof of

¹⁶It is not obvious that such a result should hold. To see this, we reason by analogy with the informed principal problem of Myerson (1983), where the principal is also a player, and focusing on deviations to direct and incentive compatible mechanisms is with loss of generality. Two of the equilibrium notions Myerson analyzes have analogues in our paper. In *expectational equilibria*, the principal can choose from *any* mechanism, as in the model in Section 2. *Undominated mechanisms* are direct incentive compatible mechanisms that weakly dominate any other direct

Proposition 3.1 highlights that it is enough for the principal to look among those canonical mechanisms that incentivize the agent to participate and truthfully report her type. This second observation is important. In finite-horizon settings, it justifies writing down the principal’s problem as a series of maximization problems subject to constraints: the participation and incentive compatibility constraints for the agent and the sequential rationality constraints for the principal. This provides a game-theoretic foundation for the programs studied by Bester and Strausz (2001) and Bester and Strausz (2007). In Doval and Skreta (2018b) we show how this approach also simplifies looking for the best equilibrium for the principal in an infinite-horizon problem.

4 TRANSFERABLE UTILITY AND INCREASING DIFFERENCES

Section 4 considers a simplified version of the game in Section 2. The purpose is to show how one can harness the results in Section 3 to solve for the principal’s optimal mechanism under limited commitment. In particular, our formulation of the canonical set of output messages as beliefs allows us to write the principal’s problem as a *constrained* information design problem. Using this formulation and an extension¹⁷ of the techniques in Le Treust and Tomala (2017), we characterize upper bounds on the set of posteriors used in an optimal mechanism. Along the way, we also highlight the differences between the problem considered here and the one introduced by Kamenica and Gentzkow (2011).

Consider the following simpler version of the game in Section 2:

- The agent observes her type $v_i \in \{v_1, \dots, v_N\}$. Let $\mu_i^0 = Pr(v = v_i)$.
- The principal offers the agent a mechanism $\mathbf{M} = \langle (M, \beta, S), \alpha \rangle$.
- The agent observes \mathbf{M} and decides whether to participate.
 - If she does not participate, a^* is implemented.
 - If she participates, she privately submits a report $m \in M$:

and incentive compatible mechanism, as in the canonical game in Section 3.2. Myerson shows that strong solutions, which is a strengthening of undominated mechanisms, are expectational equilibria, but the reverse does not necessarily hold.

¹⁷See Doval and Skreta (2018a) for further details.

- $s \in S$ is drawn according to $\beta(\cdot|m)$, which is publicly observed,
- $a \in A$ is drawn according to $\alpha(\cdot|s)$, which is publicly observed.
- The principal selects an action $y \in Y(a)$, where $Y(a) \subseteq Y$ is a compact (possibly finite) space.

The *non-contractible* action y captures in reduced form the principal’s limited commitment: In the example in Section 2.1, y corresponds to the choice of mechanism in the second period.¹⁸ The correspondence $Y(a)$ plays the role of the correspondence \mathcal{A} in Section 2: It captures how past allocations may affect the principal’s available choices in the continuation.

The above is the game that underlies the maximization problem analyzed by Bester and Strausz (2007). Theorem 3.1 and Proposition 3.1 provide a game-theoretic foundation for why the search for the principal’s best equilibrium can be cast in terms of such a program. To facilitate the comparison between the papers, we follow their notation as much as possible. In what follows, $w_i(a, y)$ denotes the principal’s utility when $v = v_i$; similarly, $u_i(a, y)$ is the agent’s utility when her type is v_i .

In standard mechanism design fashion, we focus here on the case of transferable utility and *increasing differences*, leaving the full analysis to Section I in the Supplementary Material.¹⁹ First, we assume $A = Q \times \mathbb{R}$, where $q \in Q$ denotes the physical part of the allocation and $t \in \mathbb{R}$ denotes a monetary transfer from the agent to the principal. Hereafter, we take $Y(q, t) = Y(q)$, and in a slight abuse of notation, we denote $u_i(a, y) = u_i(q, y) - t$, $w_i(a, y) = w_i(q, y) + t$.²⁰ Second, because mechanisms in our setting determine lotteries over outcomes, the appropriate notion of increasing differences is the one in Kartik et al. (2017):

¹⁸The beauty of the simple model is that the non-contractible part of the allocation y may stand for other forms of contractual incompleteness, such as renegotiation. From this point of view, we believe the techniques presented herein could be used to understand optimal contracting in other environments of interest, where distortions arise, for instance, from the need to renegotiate contracts or hold-up problems.

¹⁹We comment at the end of this section how the results we obtain translate to the case of non-transferable utility

²⁰Under transferable utility, if the action y represents the choice of a continuation mechanism, then the assumption that $Y(\cdot)$ does not depend on t is innocuous. Section I in the supplementary material does not restrict how the correspondence $Y(\cdot)$ depends on a .

Definition 4.1 (Kartik et al. (2017)). The family $\{u_i\}_{i=1}^N$ satisfies *monotonic expectational differences* if for any two distributions $P, Q \in \Delta(A \times Y)$ $\int u_i(\cdot) dP - \int u_i(\cdot) dQ$ is monotonic in i .

Kartik et al. (2017) show that u satisfies monotonic expectational differences if, and only if, it takes the form $u_i(a, y) = g_1(a, y)f_1(i) + g_2(a, y) + c(i)$, where g_1, g_2 are finitely integrable and f_1 is monotonic. In what follows, we assume that $f_1(i)$ is increasing, so that types with higher labels correspond to “higher” types.

The analysis in Section 3 implies the solutions to the following program characterize the PBE of the aforementioned game:²¹

$$\begin{aligned} \max_{\beta, \alpha, y} \sum_{i,h} \mu_i^0 \beta_{i,h} \left[\sum_{(q,t)} \alpha_h(q,t) (w_i(q, y_h(q)) + t) \right] & \quad (\mathcal{P}) \\ \text{s.t.} \left\{ \begin{array}{l} \sum_h \beta_{i,h} \sum_{(q,t)} \alpha_h(q,t) [u_i(q, y_h(q)) - t] \geq 0, i \geq 1 \\ \sum_h (\beta_{i,h} - \beta_{k,h}) \sum_{(q,t)} \alpha_h(q,t) [u_i(q, y_h(q)) - t] \geq 0, i, k \in \{1, 2, \dots, N\}, i \neq k \\ y_h(q) \in y^*(\mu_h, q) \equiv \arg \max_{y \in Y(q)} \sum_{i=1}^N \mu_{h,i} w_i(q, y) \\ \mu_{h,i} \sum_{j=1}^N \mu_j^0 \beta_{j,h} = \mu_i^0 \beta_{i,h}, \end{array} \right. \end{aligned}$$

where $\beta_{i,h} = \beta(\mu_h | v_i)$, $\alpha_h = \alpha(\cdot | \mu_h)$ and $\mathcal{H} = \{1, \dots, H\}$ indexes the posteriors. That is, the principal selects the best canonical mechanism from among the ones that (i) induce participation with probability 1, (ii) induce truth-telling with probability 1, and (iii) satisfy the principal’s sequential rationality constraints. Implicit in this program is that the number of posteriors induced by the principal is also a

²¹In the event that the agent does not participate in the mechanism, allocation a^* gets implemented. Moreover, the principal chooses $y^* \in Y(a^*)$ to maximize his expected utility, where expectations are taken with respect to his beliefs after observing the agent does not participate of the mechanism, which is an off-path event. Implicit in the agent’s participation constraint in program (\mathcal{P}) is that the above choices can be made so that the agent obtains a payoff of 0 *regardless* of her type. This allows us to focus on the issues of limited commitment, without having to worry of the issue of type dependent participation constraints.

Moreover, we assume that $g_1(a^*, y^*) = \min_{(a,y): y \in \mathcal{Y}(a)} g_1(a, y)$. This allows us to conclude that whenever the lowest type, v_1 participates of the mechanism, then all types participate of the mechanism. We thank an anonymous referee who made us realize we omitted stating this assumption in a previous version of the paper, which can be accessed [here](#).

variable of choice.

The rest of the section proceeds as follows. First, Proposition 4.1 shows how, under our assumptions, we can simplify the number of constraints in program (\mathcal{P}) . Second, we show how to cast the simpler program as a constrained information design one. Finally, we use this connection to characterize an upper bound on the number of posteriors in the optimal mechanism.

Transferable utility implies that focusing on mechanisms that do not randomize on transfers is without loss of generality. Hereafter, we replace t with its expectation, denoted by t_h . Like increasing differences in mechanism design with commitment, monotonic expectational differences implies the solutions to (\mathcal{P}) coincide with the solutions to a simpler program, which imposes only a subset of the incentive compatibility constraints. Finally, both assumptions imply that the participation constraint of the lowest type binds. The above remarks are recorded in Proposition 4.1.

Proposition 4.1. *If $\{u_i\}_{i=1}^N$ satisfies monotonic expectational differences, then to characterize the solution to (\mathcal{P}) , it suffices to guarantee the following hold:*

1. *The agent's participation constraint binds when her type is v_1 .*
2. *Adjacent incentive constraints are satisfied.*

See Appendix D for a proof. In mechanism design with commitment, we could simplify (\mathcal{P}) further by showing the *downward-looking*²² incentive constraints always bind at the optimum. This then justifies the study of the so-called *relaxed program*:

$$\begin{aligned} \max_{\beta, \alpha, y} \sum_{i,h} \mu_i^0 \beta_{i,h} \left[\sum_q \alpha_h(q) w_i(q, y_h(q)) + t_h \right] & \quad (\mathcal{R}) \\ \text{s.t.} \begin{cases} \sum_h \beta_{1,h} \left[\sum \alpha_h(q) u_1(q, y_h(q)) - t_h \right] = 0, \\ \sum_h (\beta_{i,h} - \beta_{i-1,h}) \left[\sum \alpha_h(q) u_i(q, y_h(q)) - t_h \right] = 0, i \in \{2, \dots, N\} \\ y_h(q) \in y^*(\mu_h, q) \equiv \arg \max_{y \in Y(q)} \sum \mu_{h,i} w_i(q, y) \\ \left[\sum_{j=1}^N \mu_j^0 \beta_{j,h} \right] \mu_{h,i} = \mu_i^0 \beta_{i,h} \end{cases} \end{aligned}$$

²²That is, the constraints that say v_i does not report her type is v_{i-1} .

obtained by dropping the *monotonicity constraints*:²³

$$\sum_h (\beta_{i,h} - \beta_{i-1,h}) \sum_q \alpha_h(q) (u_i(q, y_h(q)) - u_{i-1}(q, y_h(q))) \geq 0, i \in \{2, \dots, N\}. \quad (\text{M})$$

In mechanism design with commitment, it suffices to check that the solution to the relaxed program satisfies the monotonicity constraints, (M), to show it is the solution to (P) (see the discussion in footnote 23).

However, in mechanism design with limited commitment, the solution to the relaxed program is not necessarily a solution to (P) even if it satisfies the monotonicity constraints, when the type space is finite and there are three or more types. This is illustrated in Example 2 in Appendix D. Whereas in the relaxed program the binding downward-looking incentive constraints together with v_1 's participation constraint impose N restrictions on the transfers (t_h), the solution to the relaxed program might use less than N posteriors. Therefore, finding transfers t_h that satisfy all constraints may not be possible.²⁴ Alternatively, not all downward-looking constraints may bind in the optimal mechanism.

Fortunately, the above is not an issue when there are two types or a continuum of types. In both cases, it is possible to show that downward looking constraints bind (see Appendices D and IV). Because most of the literature focuses on one of these cases, and because the relaxed program provides a useful benchmark, the rest of this section studies its properties.

We can use the binding constraints to substitute the transfers out of the principal's program and obtain the following:

²³The constraints in equation (M) are obtained from combining the restriction that v_i does not want to report v_{i-1} and v_{i-1} does not want to report v_i . Under Definition 4.1, the binding downward-looking incentive constraints together with the monotonicity constraints imply the local constraints in Proposition 4.1.

²⁴This is never an issue in mechanism design with commitment: Without loss of generality, we can always have one transfer for each type.

$$\begin{aligned} \max_{\tau, \alpha, y} \sum_h \tau(\mu_h) \sum_{i=1}^N \mu_{h,i} \sum_q \alpha_h(q) [w_i(q, y_h(q)) + u_{i,h} - \frac{1 - \sum_{n \leq i} \mu_n^0}{\mu_i^0} (u_{i+1,h} - u_{i,h})] \\ \text{s.t. } \sum \tau(\mu_h) \mu_h = \mu^0, \end{aligned}$$

where $\tau(\mu_h) = \sum_i \mu_i^0 \beta_{i,h}$ and $u_{i,h} = u_i(q, y_h(q))$.

Define:

$$\begin{aligned} \hat{u}_i(q, y; \mu^0) &\equiv u_i(q, y) - \frac{1 - \sum_{n \leq i} \mu_n^0}{\mu_i^0} (u_{i+1}(q, y) - u_i(q, y)), \\ w_i(\alpha, y_\mu(\alpha)) + \hat{u}_i(\alpha, y_\mu(\alpha); \mu^0) &\equiv \sum_q \alpha(q) (w_i(q, y_\mu(q)) + \hat{u}_i(q, y_\mu(q); \mu^0)) \\ \tilde{w}(\alpha, \mu; \mu^0) &\equiv \mathbb{E}_\mu [w_i(\alpha, y_\mu(\alpha)) + \hat{u}_i(\alpha, y_\mu(\alpha); \mu^0)], \end{aligned}$$

where \hat{u}_i is type i 's virtual utility from (q, y) and \tilde{w} is the expectation according to μ of the virtual surplus at $\{\alpha, y_\mu(\cdot)\}$, for some selection $y_\mu(q) \in y^*(\mu, q)$ (see Remark 4.1). Moreover, we drop the index h because thinking about these objects as functions of beliefs μ in what follows is useful.

Program (\mathcal{R}) is then equivalent to

$$\begin{aligned} \max_{\tau, \alpha, y} \mathbb{E}_\tau \tilde{w}(\alpha, \mu; \mu^0) \\ \text{s.t. } \mathbb{E}_\tau \mu = \mu^0 \end{aligned} \quad (3)$$

That is, the solution to the relaxed problem is obtained by maximizing a version of the virtual surplus, represented by \tilde{w} , and then choosing a distribution over posteriors that averages out to the prior. Equation (3) generalizes the program obtained in Section 2.1. The following remark is in order:

Remark 4.1 (Tie-breaking in favor of the principal). So far, we have remained silent about how $y_\mu(q)$ is chosen, beyond the restriction that $y_\mu(q) \in y^*(\mu, q)$. We can use the function $\tilde{w}(q, \mu; \mu^0)$ to determine how to break the possible ties in y^* and make the principal's objective upper-hemicontinuous. In fact, if $y, y' \in$

$y^*(\mu, q)$, then in the relaxed problem, y is selected as long as

$$\mathbb{E}_\mu[w_i(q, y) + \hat{u}_i(q, y; \mu^0)] \geq \mathbb{E}_\mu[w_i(q, y') + \hat{u}_i(q, y'; \mu^0)].$$

In other words, ties are broken in favor of the virtual surplus.

We now illustrate how to solve the program in (3). Towards this, fix the selection y as in Remark 4.1. Because the program is separable in the allocation, α , across posteriors, the solution can be obtained in two steps. First, for each posterior μ , we maximize $\tilde{w}(\cdot, \mu; \mu^0)$ with respect to α . Denote the value of this problem $\hat{w}(\mu; \mu^0)$. Second, we choose τ to maximize $\hat{w}(\mu; \mu^0)$ subject to the constraint that the posteriors must average out to the prior, μ^0 . This separability between the choice of the allocation rule, α , and the communication device, β , is afforded by ignoring the monotonicity constraints in (M). The latter may impose additional restrictions on how the allocation varies across different posteriors.

This discussion implies the solution to (3) can be obtained by solving:

$$\begin{aligned} \max_{\tau} \mathbb{E}_{\tau} \underbrace{\max_{\alpha} \tilde{w}(\alpha, \mu; \mu^0)}_{\hat{w}(\mu; \mu^0)} & \quad (4) \\ \text{s.t. } \mathbb{E}_{\tau} \mu = \mu^0 & \end{aligned}$$

An advantage of the above formulation is that a straightforward application of Carathéodory's theorem (see Rockafellar (1970)) implies that in (4), the solution never uses more than N posteriors.²⁵

Proposition 4.2. *The solution to (R) uses at most N posteriors.*

Then, if the solution to the relaxed program satisfies the monotonicity constraints *and* it is possible to find transfers (t_h) that satisfy the downward looking binding incentive constraints, we have found a solution to the principal's problem, (P).

In many instances, however, the solution to (R) will fail to satisfy the mono-

²⁵Bester and Strausz (2007) derive this result using methods in semi-infinite linear programming.

tonicity constraints, (M). As we show next, adding as many posteriors as binding monotonicity constraints at the optimum may be necessary:

Proposition 4.3. *Consider the program obtained by adding the monotonicity constraints (M) to the relaxed program (R). The solution to the new program uses at most $N + K$ posteriors, where K is the number of binding constraints at the optimum.*

The proof is in Appendix D and follows from extending the techniques in [Le Treust and Tomala \(2017\)](#) to our setting, where we have multiple inequality constraints and equality constraints.

Finally, we note the connection between our problem and a constrained information design problem holds beyond the case of transferable utility, as illustrated in Section I in the supplementary material. In particular, we show the assumption of monotonic expectational differences also reduces the problem to the analysis of the local incentive constraints. Moreover, we can again bound the number of posteriors by $3N - 1$.

Whereas the above formulation harnesses the connection between our problem and the one studied in information design, we close the section by highlighting two conceptual differences with this literature. The reader eager to see the results in the next section can skip it without loss of continuity.

First, the function $\hat{w}(\mu; \mu^0)$ in equation (4) stands for the sender’s objective function, $\hat{v}(\mu)$, in [Kamenica and Gentzkow \(2011\)](#). Recall that in [Kamenica and Gentzkow \(2011\)](#), $\hat{v}(\mu)$ is the sender’s expected utility of the receiver’s optimal action when the posterior is μ , where expectations are taken with respect to μ . Two differences are worth pointing out. First, in our setting, the first-period principal (the sender in [Kamenica and Gentzkow \(2011\)](#)) also takes an action for each posterior he induces, because he chooses the allocation α .

Second, the principal’s objective function in equation (3) depends both on the posterior, μ , and the prior, μ_0 , whereas in [Kamenica and Gentzkow \(2011\)](#), it only depends on the posterior. We already saw an instance of this in the example studied in Section 2.1. In fact, we saw that the virtual values are calculated using

the prior distribution, because this distribution is the one the principal uses to calculate the probability with which he leaves rents to the different types of the agent. That the principal's payoff depends both on the prior *and* the posterior may come as a surprise because when a distribution F can be written as a convex combination of distributions F_s , so that $F = \sum_{s=1}^S \lambda_s F_s$, then

$$\int \frac{(1-F)}{f} dF = \sum_{s=1}^S \lambda_s \int \frac{1-F_s}{f_s} dF_s. \quad ^{26}$$

That is, the posterior *information handicap* averages out to the prior information handicap. Thus, we may have expected that the information handicap in $\tilde{w}(\alpha, \mu; \mu_0)$ could be written solely as a function of μ . Only when the allocation is the same for all induced posteriors, and hence no type of the agent obtains rents, we can think of the principal's objective as only depending on the posterior.

5 DISCUSSION

5.1 Recommendations as output messages

As discussed in the introduction, in the finite-horizon case, there is another candidate for a canonical language: In period t , each output message could be associated to an allocation for period t and a recommended allocation for the subsequent periods. We use the formulation in Section 4 to discuss this formally.

Section 4 illustrated how the relaxed program can be formulated as an information design problem, where the principal in period 1 designs both the allocation for the agent and the information structure for the principal in period 2 (see equation (3)).

We borrow the terminology in [Kamenica and Gentzkow \(2011\)](#) and say a mechanism is *straightforward* if $S \subseteq \cup_{q \in Q} \{q\} \times Y(q)$ and after message $s = (q_s, y_s)$, the principal chooses q_s in period 1 to maximize $\tilde{w}(\alpha, \mu_s; \mu^0)$ and y_s in period 2 to maximize $\sum \mu_{i,s} w_i(q_s, y)$ in period 2, where μ_s are the beliefs generated by output message s .

²⁶Laura Doval thanks René Leal Vizcaino for a discussion that led to this observation.

Proposition 5.1. *The following are equivalent:*

1. *There exists a mechanism $\langle (V, \beta, S), \alpha \rangle$ and a mapping $y : S \times Q \mapsto \cup_{q \in Q} Y(q)$ that solves (\mathcal{R}) .*
2. *There exists a mechanism $\langle (V, \beta, \Delta(V)), \alpha \rangle$ and a mapping $y : \Delta(V) \times Q \mapsto \cup_{q \in Q} Y(q)$ that solves (\mathcal{R}) .*
3. *There exists a straightforward mechanism that solves (\mathcal{R}) .*

The proof is in Appendix E. Item 3 highlights that in the relaxed program (\mathcal{R}) , the set of output messages can also be taken to be recommendations for both incarnations of the principal, as in sender-receiver models of information design. The proposition uses both the separability between the allocation, α , and the information policy, τ , discussed in Section 4 and that there is a final period in which the principal takes the non-committed action, y .²⁷ The separability guarantees the first-period principal chooses to implement allocations q that are optimal for each posterior he induces for period 2. However, in the case of infinite horizon, the language of recommendations is *self-referential*: The principal would need to recommend the continuation mechanisms, which themselves involve a set of output messages. Thus, an advantage of the approach we advocate is that we can always resort to beliefs regardless of the game at hand.

5.2 Implementation via contracts

Section 5.2 characterizes within the environment of Section 4 the tuples (β, q, y) that can be implemented using the contracts studied previously in the literature. An important difference between the mechanisms used by Hart and Tirole (1988), Laffont and Tirole (1988), Freixas et al. (1985), and Bester and Strausz (2001) and the ones considered here is that whereas in the former papers, the principal observes the agent's choice out of a menu, in the latter, the agent's input into the communication device is not observed. A consequence of this is that in the former setting, the agent has to be indifferent between all the elements of the menu

²⁷If the final period corresponded to a design problem, such as the sale of a durable good example, one could resort to the revelation principle in Myerson (1982) to reduce the principal's actions in the final period to the induced allocations.

that she selects with positive probability. By contrast, in our setting, the agent's incentive compatibility constraint has to hold in expectation: Although she may not be indifferent between all the allocations that arise with positive probability after she communicates with the mechanism, on average, they must be better than what she would obtain by reporting any other type.

Fix a *canonical* communication device $\beta : V \mapsto \Delta^*(\Delta(V))$ and a tuple $(q, y) : \Delta(V) \mapsto Q \times Y$, where we denote by $y(\mu) = y(q(\mu), \mu)$.²⁸ We want to find $t' : \Delta(V) \mapsto \mathbb{R}$ such that for all $v_i \in V$, for all $\mu : \beta(\mu|v_i) > 0$, and for all $\mu' : \sum_{v'} \beta(\mu'|v') > 0$,

$$u_i(q(\mu), y(\mu)) - t'(\mu) \geq u_i(q(\mu'), y(\mu')) - t'(\mu'). \quad (\text{DIC-P})$$

Note equation (DIC-P) corresponds to the definition of equivalence in [Mookherjee and Reichelstein \(1992\)](#).²⁹ Indeed, the problem we intend to solve is similar in spirit to the one analyzed in the literature that studies the equivalence between Bayesian and dominant-strategy implementation ([Manelli and Vincent \(2010\)](#) and [Gershkov et al. \(2013\)](#)). However, there are some differences that, although subtle, turn out to have important implications. First, in that literature, this problem only makes sense when there are multiple agents, whereas in our case, the source of randomness the agent faces (the randomization by the communication device) is endogenously chosen by the principal. Second, allocation and transfers in that setting depend on the agent's type, whereas here they depend on the belief induced when the principal observes the output message. As we show next, this implies joint restrictions on the communication device and the allocation rule.

As in Section 4, we assume the agent's preferences satisfy monotonic expectational differences. Thus, label the types so that $v_1 < \dots < v_N$. Note that if $u_i(q, y)$ satisfies Definition 4.1, then $u_i(q, y)$ has increasing differences. In effect, $u_i(q', y') - u_i(q, y) = f(i)(g_1(q', y') - g_1(q, y)) + g_2(q', y') - g_2(q, y)$, which is strictly

²⁸To keep notation simple, we ignore the possibility that q and y may be randomized allocations. It is immediate that this restriction is not necessary for the results.

²⁹[Mookherjee and Reichelstein \(1992\)](#) also require that t' raises the same revenue as t (see equation (DIC-T) below). We focus for now on the possibility of guaranteeing (DIC-P) holds and discuss the difficulties associated with guaranteeing the same revenue is collected at the end of Section 5.2.

increasing in i as long as $g_1(q', y') - g_1(q, y) \neq 0$. In what follows, we make the following assumption:

Assumption 1. For all $\mu, \mu' \in \Delta(V)$ such that $(\sum_{v \in V} \beta(\mu|v)) \times (\sum_{v \in V} \beta(\mu'|v)) > 0$, we assume $g_1(q(\mu), y(\mu)) - g_1(q(\mu'), y(\mu')) \neq 0$.

Given two beliefs μ and μ' , let

$$D_i(\mu, \mu') \equiv u_i(q(\mu), y(\mu)) - u_i(q(\mu'), y(\mu'))$$

denote the difference in payoffs from $(q(\mu), y(\mu))$ and $(q(\mu'), y(\mu'))$, when the agent type is v_i . The content of Assumption 1 is that $D_i(\mu, \mu')$ is strictly increasing in i .

We have the following:

Proposition 5.2. *Suppose the agent's Bernoulli utility function satisfies Definition 4.1 and (β, q, y) satisfies Assumption 1. A necessary and sufficient condition for (β, q, y) to satisfy (DIC-P) is that (β, q, y) satisfies*

1. For all $i \in \{1, \dots, N\}$ and $j > i$, for all μ and μ' ,

$$D_j(\mu', \mu) \geq D_i(\mu', \mu), \tag{DIC-M}$$

whenever $\mu'(v_j)\mu(v_i) > 0$.

2. β induces a monotone information structure: We can label the beliefs induced by (β, q, y) , $\{\mu_1, \dots, \mu_M\}$ so that

(a) If $i < j$, then $\max \text{supp } \mu_i \leq \min \text{supp } \mu_j$,

(b) For any i , there are at most three beliefs $\{\mu_i, \mu_{i+1}, \mu_{i+2}\}$ for which v_i has positive probability. Moreover, if v_i has positive probability in all three, then $\mu_{i+1}(v_i) = 1$.

The proof is in Appendix F. The first condition is the equivalent to the standard monotonicity condition for dominant-strategy incentive compatibility: For any two beliefs μ and μ' , the utility differential of the allocations (q, y) induced at these beliefs is higher for higher types. The second is new to our setting. Recall that, under Assumption 1, $D_i(\mu, \mu')$ is strictly increasing in i , which places restrictions

on the support of the beliefs induced by the principal's mechanism. In particular, to satisfy (DIC-P), when the agent's type is v_i , she must be indifferent between the allocations (q, y) that correspond to posteriors that assign positive probability to v_i . Monotonic expectational differences implies that when v_i is indifferent between $(q(\mu), y(\mu))$ and $(q(\mu'), y(\mu'))$ and $D_i(\mu, \mu') \geq 0$, then all types higher than v_i (weakly) prefer $(q(\mu), y(\mu))$ to $(q(\mu'), y(\mu'))$ (and the opposite holds for lower types). Assumption 1 then guarantees that higher types cannot be in the support of μ' (and the opposite holds for lower types).

Note that we cannot dispense with the assumption that β induces a monotone information structure. In the example in Appendix B in [Bester and Strausz \(2007\)](#), the agent can be of one of two types and her utility satisfies Definition 4.1; however, the optimal mechanism induces three posteriors under which both types have positive probability and having both types be indifferent between the three allocations induced by the mechanism is not possible.

Besides allowing us to connect the results in this paper with the previous literature on mechanism design with limited commitment, the result in Proposition 5.2 is also of practical value. Section 4 highlights that the characterization of S as the set of beliefs over the agent's type allows us to harness the tools of mechanism and information design to solve for the principal's optimal mechanism. Proposition 5.2 allows us then to check when the solution to the program in Section 4 is also a solution to the problem in which the principal observes the agent's choice out of a menu.

Proposition 5.2 is silent about whether the transfer scheme t' collects the same revenue as the original mechanism did, that is, whether t' also satisfies that

$$\sum_{v \in V} \beta(\mu|v)t'(\mu) = \sum_{v \in V} \beta(\mu|v)t(\mu), \quad (\text{DIC-T})$$

for all μ such that $\sum_{v \in V} \beta(\mu|v) > 0$.

Contrary to [Mookherjee and Reichelstein \(1992\)](#), we do not find that once (DIC-P) is satisfied, then (DIC-T) is satisfied. In particular, with two types, only when the solution features two beliefs, $\underline{\mu}(v_2) < \mu^0 < \bar{\mu}(v_2) = 1$, can one

satisfy both (DIC-P) and (DIC-T). We leave for future research the full analysis of the conditions under which both (DIC-P) and (DIC-T) hold.³⁰

5.3 Multiple agents: Example in *Bester and Strausz (2000)*

Bester and Strausz (2000) show that, with multiple agents, the result for the single-agent case in *Bester and Strausz (2001)* no longer holds. That is, if $M = S$ and β is deterministic, then there are equilibria with mechanisms in which $M \neq V$, whose payoffs cannot be replicated with canonical mechanisms. To keep the discussion self-contained, we replicate here their example and then explain why it does not invalidate the possibility of extending our techniques to the multi-agent case.

Example 1 (*Bester and Strausz (2000)*). There are two agents, labeled 1 and 2. Only agent 1 has private information; let $v \in \{\underline{v}, \bar{v}\}$ denote her type. The prior that the type is \underline{v} is denoted by $\mu \in [0, 1]$. The set of allocations $A = [0, 2]$. Payoffs are defined as follows:

$$\begin{aligned} W(a, \underline{v}) &= -a^2, W(a, \bar{v}) = -(2 - a)^2 \\ U_1(a, \underline{v}) &= -(0.5 - a)^2, U_1(a, \bar{v}) = -(1.5 - a)^2 \\ U_2(a) &= -10(1 - a)^2. \end{aligned}$$

That is, the principal's payoff depends on agent 1's type, whereas agent 2's payoff does not.

The timing is as follows. The principal selects a communication device for the agents, who then submit messages. Upon seeing the messages, the principal chooses $a \in A$.

The principal's payoffs are such that if, after seeing m , his posterior is $\mu(m)$, then he chooses allocation

$$a^*(m) = 2(1 - \mu(m)).$$

³⁰Note that *Mookherjee and Reichelstein (1992)*, *Manelli and Vincent (2010)*, *Gershkov et al. (2013)* use the representation of the agent's utility function obtained via the envelope theorem to pin down transfers. As discussed in Section 4, we cannot guarantee downward-looking constraints bind at the optimum when there are three or more types.

Bester and Strausz (2000) construct an equilibrium with three messages $\{m_a, m_b, m_c\}$ that cannot be replicated with messages $\{\underline{v}, \bar{v}\}$. Let $M = \{m_a, m_b, m_c\}$. Then there is a PBE such that

$$\begin{aligned} \mu(m_a) &= 1, & \mu(m_b) &= 1/2, & \mu(m_c) &= 0 \\ a^*(m_a) &= 0, & a^*(m_b) &= 1, & a^*(m_c) &= 2. \end{aligned}$$

However, the above mechanism cannot be replicated by a mechanism with two messages, when the principal observes the output messages. The reason is not that agent 1 is not indifferent between the allocations he obtains at different messages, so that he is not willing to carry out the randomization himself. Rather, under the assumptions of Bester and Strausz (2001), the mechanism can only have as many input messages as output messages. Because the agent may be of one of two types, the mechanism can only have two input messages. Therefore, the agent does not have enough room to do the mixing and generate the required posteriors for the principal.

However, if we allow the principal to offer canonical mechanisms as the ones in this paper, the following communication device implements the same allocation as the non-canonical mechanism:

$$\begin{aligned} \beta(\mu(m_a)|\underline{v}) &= 1/2, \beta(\mu(m_a)|\bar{v}) = 0, \\ \beta(\mu(m_b)|\underline{v}) &= 1/2, \beta(\mu(m_b)|\bar{v}) = 1/2, \\ \beta(\mu(m_c)|\underline{v}) &= 0, \beta(\mu(m_c)|\bar{v}) = 1/2. \end{aligned}$$

After all, the result that taking $M \simeq V$ and $S \simeq \Delta(V)$ is without loss of generality dispenses with the restriction that the number of input messages must coincide with the number of output messages.

Extending the model in Section 2 to the case of multiple agents involves dealing with a number of subtleties that merit a full separate study and are thus beyond the scope of this paper. We plan to address this important extension in future research.

A PROOF OF PRELIMINARY RESULTS

Appendix A is organized as follows:

Proposition A.1 shows we can focus without loss of generality on equilibria of the game in which the agent does not condition her strategy on the payoff-irrelevant part of her private history.

Proposition A.2 shows we can focus without loss of generality on equilibria of the game in which the agent participates with probability one.

Proposition A.3 shows we can focus without loss of generality on equilibria in which there is a one-to-one mapping between the output messages generated in the mechanism and the continuation beliefs the principal holds about the agent's type.

Because we have to deal with an abstract dynamic game, the proof is notationally involved. As a first pass to most results, except Proposition A.1, the reader is invited to first take a look at Appendix I in the supplementary material, where the constructions are performed in a two-period version of the model.

We need a few more pieces of notation and definitions.

First, as noted in footnote 12, some output messages can never be observed. Given a mechanism, \mathbf{M}_t , define $S^{*\mathbf{M}_t} = \{s \in S^{\mathbf{M}_t} : (\exists m \in M^{\mathbf{M}_t}) \beta^{\mathbf{M}_t}(s|m) > 0\}$. Since removing public histories from the tree that are consistent with mechanism \mathbf{M}_t and $s \in S^{\mathbf{M}_t} \setminus S^{*\mathbf{M}_t}$ does not change the set of equilibrium outcomes, hereafter, these histories are removed from the tree.

Second, fix a PBE of the dynamic mechanism-selection game $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$ and a public history h^t for some $t \geq 0$ (if $t = 0$, then $h^0 = \emptyset$ denotes the initial public history). For $T \geq t$, the history $h^T = (h^t, \mathbf{M}_t, p_t, s_t, a_t, \dots, \mathbf{M}_{T-1}, p_{T-1}, s_{T-1}, a_{T-1}, \omega_T)$ is on the path of the equilibrium strategy profile given h^t if for all $t \leq \tau \leq T - 1$

$$h^{\tau+1} = (h^\tau, \mathbf{M}_\tau, p_\tau, s_\tau, a_\tau, \omega_{\tau+1}),$$

where

$$\begin{aligned}
& \mathbf{M}_\tau \in \text{supp } \Gamma^*(h^\tau) \\
& \sum_{(v, h_A^\tau)} \mu^*(h^\tau)(v, h_A^\tau) \pi_v^*(h_A^\tau, \mathbf{M}_\tau)(p_\tau) > 0 \\
p_\tau = 1 \Rightarrow & \sum_{(v, h_A^\tau): \pi_v^*(h_A^\tau, \mathbf{M}_\tau) > 0} \mu^*(h^\tau)(v, h_A^\tau) \sum_{m \in M^{\mathbf{M}_\tau}} r_v^*(h_A^\tau, \mathbf{M}_\tau, 1)(m) \beta^{\mathbf{M}_\tau}(s_\tau | m) \alpha^{\mathbf{M}_\tau}(a_\tau | s_\tau) > 0 \\
p_\tau = 0 \Rightarrow & s_\tau = \emptyset, a_\tau = a^*.
\end{aligned}$$

That is, along the path from h^t to $h^{\tau+1}$, the principal made choices according to his equilibrium strategy, the agent made participation choices according to her equilibrium strategy, and the output-message labels and allocations correspond to those in the mechanism chosen by the principal. Note that we do not say anything about the reports of the agent, because they are not part of the public history.

We sometimes need to talk about the histories that are on the path of the equilibrium strategy profile given a public history h^t and a mechanism \mathbf{M}_t . The only difference with the above definition is that \mathbf{M}_t need not have positive probability according to $\Gamma^*(h^t)$.

The above notation is used as follows. Proposition A.1 shows that for any PBE, there is a payoff-equivalent PBE in which the agent does not condition her strategy on the payoff-irrelevant part of her private history. To do so, starting from any history h^t , we need to modify the strategy *for all continuation histories on the path of the strategy*. Similarly, the main theorem shows we can transform any PBE of the game into one in which the principal's strategy selects only canonical mechanisms on and off the equilibrium path. To do so, we must map the continuation strategy starting from any history h^t in the game to one in which the continuation strategy only offers canonical mechanisms. When we perform this mapping for h^0 , we are doing the transformation for the path of the equilibrium strategy; when we do this for histories that can be reached from h^t , we are doing the transformation *for the path of the continuation strategy*.

Proposition A.1. *Fix a PBE $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$ and a public history h^t . Then, there exists a continuation strategy for the agent $(\pi_v^{**}, r_v^{**})_{v \in V}$ such that:*

1. *For any mechanism \mathbf{M}_t , for all $h_A^t, h_A^{t'} \in H_A^t(h^t)$, and for all $v \in V$, $\pi_v^{**}(h_A^t, \mathbf{M}_t) = \pi_v^{**}(h_A^{t'}, \mathbf{M}_t)$ and $r_v^{**}(h_A^t, \mathbf{M}_t, 1) = r_v^{**}(h_A^{t'}, \mathbf{M}_t, 1)$.*

2. For all mechanisms \mathbf{M}_t , for $\tau \geq t + 1$, for all histories h^τ on the equilibrium path starting from (h^t, \mathbf{M}_t) , for all $h_A^\tau, h_A^{\tau'} \in H_A^\tau(h^\tau)$, for all $\mathbf{M}_\tau \in \text{supp } \Gamma^*(h^\tau)$, and for all $v \in V$ $\pi_v^{**}(h_A^\tau, \mathbf{M}_\tau) = \pi_v^{**}(h_A^{\tau'}, \mathbf{M}_\tau)$ and $r_v^{**}(h_A^\tau, \mathbf{M}_\tau, 1) = r_v^{**}(h_A^{\tau'}, \mathbf{M}_\tau, 1)$.
3. For all histories h^τ on the equilibrium path starting from h^t , the continuation payoff for the principal at $(\Gamma^*, (\pi_v^{**}, r_v^{**})_{v \in V}, \mu^*)$ is the same as at $(\Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^*)$; similarly, for each $v \in V$ and each $h_A^\tau \in H_A^\tau(h^\tau)$ the continuation payoff for the agent at (v, h_A^τ) is the same under both strategy profiles.
4. $(\Gamma^*, (\pi_v^{**}, r_v^{**})_{v \in V}, \mu^*)$ is also a PBE.

Proof. Fix a PBE $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$. Let h^t be a public history such that there exists $\mathbf{M}_t, h_A^t, h_A^{t'}$, both consistent with h^t , and $v \in V$ such that either $\pi_v^*(h_A^t, \mathbf{M}_t) \neq \pi_v^*(h_A^{t'}, \mathbf{M}_t)$ or $r_v^*(h_A^t, \mathbf{M}_t, 1) \neq r_v^*(h_A^{t'}, \mathbf{M}_t, 1)$.

Note that for each $m \in M^{\mathbf{M}_t}$, the agent's continuation payoff at (v, h_A^t) and $(v, h_A^{t'})$ must be the same: after all, the continuation strategy of (v, h_A^t) is feasible for $(v, h_A^{t'})$ and vice versa. Therefore, the agent at (v, h_A^t) is not only indifferent between all the messages in the support of $r_v^*(h_A^t, \mathbf{M}_t, 1)$, but is also indifferent between all messages in the support of $r_v^*(h_A^{t'}, \mathbf{M}_t, 1)$. Therefore, the agent at (v, h_A^t) is indifferent between $r_v^*(h_A^t, \mathbf{M}_t, 1)$ and any randomization between $r_v^*(h_A^t, \mathbf{M}_t, 1)$ and $r_v^*(h_A^{t'}, \mathbf{M}_t, 1)$.

Moreover, the above is true for any continuation public history that is reached with positive probability from h^t for the same reasons. That is, for any $\tau \geq t$ and h^τ that succeeds (h^t, \mathbf{M}_t) along which the principal follows Γ^* and the agent at (v, h_A^t) and at $(v, h_A^{t'})$ follows π_v^*, r_v^* and for any $h_A^\tau, h_A^{\tau'}$ that succeed h_A^t and $h_A^{t'}$, respectively, the agent is indifferent between her reporting strategy $r_v^*(h_A^\tau, \mathbf{M}_\tau, 1)$ and $r_v^*(h_A^{\tau'}, \mathbf{M}_\tau, 1)$ for $\mathbf{M}_\tau \in \Gamma^*(h^\tau)$.

Therefore, starting from h^t , the following is also an optimal strategy for the agent when her valuation is v . Consider first those types v for which $\sum_{h_A^t \in H_A^t(h^t)} \mu^*(h^t)(v, h_A^t) > 0$. For any $h_A^t \in H_A^t(h^t)$, she participates with probability

$$\pi_v^{**}(h_A^t, \mathbf{M}_t) = \sum_{h_A^{t'}} \frac{\mu^*(h^t)(v, h_A^{t'})}{\sum_{\tilde{h}_A^t \in H_A^t(h^t)} \mu^*(h^t)(v, \tilde{h}_A^t)} \pi_v^*(h_A^{t'}, \mathbf{M}_t),$$

and reports $m \in M^{\mathbf{M}_t}$ with probability:

$$r_v^{**}(h_A^t, \mathbf{M}_t, 1)(m) = \sum_{h_A^{t'} \in H_A^t(h^t)} \frac{\mu^*(h^t)(v, h_A^{t'}) \pi_v^*(h_A^{t'}, \mathbf{M}_t)}{\sum_{\tilde{h}_A^t \in H_A^t(h^t)} \mu^*(h^t)(v, \tilde{h}_A^t) \pi_v^*(\tilde{h}_A^t, \mathbf{M}_t)} r_v^*(h_A^{t'}, \mathbf{M}_t, 1)(m),$$

as long as $\sum_{\tilde{h}_A^t \in H_A^t(h^t)} \mu^*(h^t)(v, \tilde{h}_A^t) \pi_v^*(\tilde{h}_A^t, \mathbf{M}_t) > 0$.³¹ For each h^τ that has positive probability from (h^t, \mathbf{M}_t) and $h_A^\tau \in H_A^\tau(h^\tau)$, for each $\mathbf{M}_\tau \in \text{supp } \Gamma^*(h^\tau)$, the agent participates with probability

$$\pi_v^{**}(h_A^\tau, \mathbf{M}_\tau) = \sum_{h_A^{\tau'} \in H_A^\tau(h^\tau)} \frac{\mu^*(h^\tau)(v, h_A^{\tau'})}{\sum_{\tilde{h}_A^\tau \in H_A^\tau(h^\tau)} \mu^*(h^\tau)(v, \tilde{h}_A^\tau)} \pi_v^*(h_A^{\tau'}, \mathbf{M}_\tau), \quad (5)$$

as long as $\sum_{\tilde{h}_A^\tau \in H_A^\tau(h^\tau)} \mu^*(h^\tau)(v, \tilde{h}_A^\tau) > 0$ and reports $m \in M^{\mathbf{M}_\tau}$ with probability

$$r_v^{**}(h_A^\tau, \mathbf{M}_\tau, 1)(m) = \sum_{h_A^{\tau'} \in H_A^\tau(h^\tau)} \frac{\mu^*(h^\tau)(v, h_A^{\tau'}) \pi_v^*(h_A^{\tau'}, \mathbf{M}_\tau)}{\sum_{\tilde{h}_A^\tau \in H_A^\tau(h^\tau)} \mu^*(h^\tau)(v, \tilde{h}_A^\tau) \pi_v^*(\tilde{h}_A^\tau, \mathbf{M}_\tau)} r_v^*(h_A^{\tau'}, \mathbf{M}_\tau, 1)(m), \quad (6)$$

as long as $\sum_{\tilde{h}_A^\tau \in H_A^\tau(h^\tau)} \mu^*(h^\tau)(v, \tilde{h}_A^\tau) \pi_v^*(\tilde{h}_A^\tau, \mathbf{M}_\tau) > 0$.

Before dealing with the zero probability events, note that the above transformation can be applied to all (v, h_A^t) , $h_A^t \in H_A^t(h^t)$, regardless of whether $\mu^*(h^t)(v, h_A^t) > 0$. This is because sequential rationality applies to all information sets of the agent and, thus, to all $h_A^t \in H_A^t(h^t)$. A consequence of the above transformation is that as long as the principal assigns positive probability to the event that the agent's type is $v \in V$, the agent plays the same at all of her payoff-irrelevant histories; even at those to which the principal assigns zero probability.

Now consider those types $v \in V$ such that $\sum_{h_A^t \in H_A^t(h^t)} \mu^*(h^t)(v, h_A^t) = 0$. For any

³¹Recall we are applying this transformation only at histories that are on the equilibrium strategy given h^t . Hence, if $\sum_{\tilde{h}_A^t \in H_A^t(h^t)} \mu^*(h^t)(v, \tilde{h}_A^t) \pi_v^*(\tilde{h}_A^t, \mathbf{M}_t) = 0$, we continue doing the transformation at $h^{t+1} = (h^t, \mathbf{M}_t, 0, \emptyset, a_*)$. Remark A.1 discusses how in this case we can apply a similar construction starting from history $(h^t, \mathbf{M}_t, 1)$.

$h_A^t \in H_A^t(h^t)$, she participates with probability

$$\pi_v^{**}(h_A^t, \mathbf{M}_t) = \sum_{h_A^{t'}} \frac{\pi_v^*(h_A^{t'}, \mathbf{M}_t)}{|H_A^t(h^t)|},$$

Suppose that $(h^t, \mathbf{M}_t, 1)$ has positive probability conditional on the principal offering \mathbf{M}_t . Then, modify the agent's reporting strategy so that she reports $m \in M^{\mathbf{M}_t}$ with probability:

$$r_v^{**}(h_A^t, \mathbf{M}_t, 1)(m) = \sum_{h_A^{t'} \in H_A^t(h^t)} \frac{\pi_v^*(h_A^{t'}, \mathbf{M}_t)}{\sum_{\tilde{h}_A^t \in H_A^t(h^t)} \pi_v^*(\tilde{h}_A^t, \mathbf{M}_t)} r_v^*(h_A^{t'}, \mathbf{M}_t, 1)(m),$$

if $\sum_{\tilde{h}_A^t \in H_A^t(h^t)} \pi_v^*(\tilde{h}_A^t, \mathbf{M}_t) > 0$ and with probability

$$r_v^{**}(h_A^t, \mathbf{M}_t, 1)(m) = \sum_{h_A^{t'} \in H_A^t(h^t)} \frac{r_v^*(h_A^{t'}, \mathbf{M}_t, 1)(m)}{|H_A^t(h^t)|},$$

otherwise.

For each h^τ that has positive probability from (h^t, \mathbf{M}_t) and $h_A^\tau \in H_A^\tau(h^\tau)$, for each $\mathbf{M}_\tau \in \text{supp } \Gamma^*(h^\tau)$, the agent participates with probability

$$\pi_v^{**}(h_A^\tau, \mathbf{M}_\tau) = \sum_{h_A^{\tau'} \in H_A^\tau(h^\tau)} \frac{\pi_v^*(h_A^{\tau'}, \mathbf{M}_\tau)}{|H_A^\tau(h^\tau)|}. \quad (7)$$

If the principal assigns positive probability to $(h^\tau, \mathbf{M}_\tau, 1)$ at h^τ , then modify the agent's reporting strategy so that she reports $m \in M^{\mathbf{M}_\tau}$ with probability

$$r_v^{**}(h_A^\tau, \mathbf{M}_\tau, 1)(m) = \sum_{h_A^{\tau'} \in H_A^\tau(h^\tau)} \frac{\pi_v^*(h_A^{\tau'}, \mathbf{M}_\tau)}{\sum_{\tilde{h}_A^\tau \in H_A^\tau(h^\tau)} \pi_v^*(\tilde{h}_A^\tau, \mathbf{M}_\tau)} r_v^*(h_A^{\tau'}, \mathbf{M}_\tau, 1)(m), \quad (8)$$

if $\sum_{\tilde{h}_A^\tau \in H_A^\tau(h^\tau)} \pi_v^*(\tilde{h}_A^\tau, \mathbf{M}_\tau) > 0$ and with probability

$$r_v^{**}(h_A^\tau, \mathbf{M}_\tau, 1)(m) = \sum_{h_A^{\tau'} \in H_A^\tau(h^\tau)} \frac{r_v^*(h_A^{\tau'}, \mathbf{M}_\tau, 1)(m)}{|H_A^\tau(h^\tau)|},$$

otherwise. Thus, for those valuations $v \in V$ to which the principal assigns 0 probability—either at h^t or at a continuation history h^τ on the equilibrium path of the strategy given h^t —their strategies also do not depend on the payoff-irrelevant part of the private history.

Fix $\tau \geq t$. Under the new strategy, the principal's beliefs that the agent is of type v and her private history is $h_A^{\tau+1}$ at history $h^{\tau+1} = (h^\tau, \mathbf{M}_\tau, 1, s_\tau, a_\tau)$, $\mathbf{M}_\tau \in \text{supp } \Gamma^*(h^\tau)$ are given by:

$$\mu^{**}(h^{\tau+1})(v, h_A^{\tau+1}) = \frac{\mu^{**}(h^\tau)(v, h_A^\tau) \pi_v^*(h_A^\tau, \mathbf{M}_\tau) r_v^{**}(h_A^\tau, \mathbf{M}_\tau, 1)(m_\tau) \beta^{\mathbf{M}_\tau}(s_\tau | m_\tau)}{\sum_{\tilde{v}, \tilde{h}_A^\tau} \sum_{\tilde{m}_\tau} \mu^{**}(\tilde{v}, \tilde{h}_A^\tau) \pi_{\tilde{v}}^{**}(\tilde{h}_A^\tau, \mathbf{M}_\tau) r_{\tilde{v}}^{**}(\tilde{h}_A^\tau, \mathbf{M}_\tau, 1)(\tilde{m}_\tau) \beta^{\mathbf{M}_\tau}(s_\tau | \tilde{m}_\tau)}, \quad (9)$$

where $\mu^{**}(h^t)(v, h_A^t) = \mu^*(h^t)(v, h_A^t)$ and at history $h^{\tau+1} = (h^\tau, \mathbf{M}_\tau, 0, s_\tau, a_\tau)$, $\mathbf{M}_\tau \in \text{supp } \Gamma^*(h^\tau)$ are given by:

$$\mu^{**}(h^{\tau+1})(v, h_A^{\tau+1}) = \frac{\mu^{**}(h^\tau)(v, h_A^\tau)(1 - \pi_v^{**}(h_A^\tau, \mathbf{M}_\tau))}{\sum_{\tilde{v}, \tilde{h}_A^\tau} \mu^{**}(h^\tau)(\tilde{v}, \tilde{h}_A^\tau)(1 - \pi_{\tilde{v}}^{**}(\tilde{h}_A^\tau, \mathbf{M}_\tau))}, \quad (10)$$

We now show by induction that for any $\tau \geq t$,

$$\sum_{h_A^{\tau+1} \in H_A^{\tau+1}(h^{\tau+1})} \mu^{**}(h^{\tau+1})(v, h_A^{\tau+1}) = \sum_{h_A^{\tau+1} \in H_A^{\tau+1}(h^{\tau+1})} \mu^*(h^{\tau+1})(v, h_A^{\tau+1}).$$

We do so for histories h^τ that are consistent with the equilibrium strategy for which the agent participates; it is immediate that the same holds for those histories in which she does not.

For $\tau = t$ and $h^{t+1} = (h^t, \mathbf{M}_t, 1, s_t, a_t)$, note the denominator on the right-hand side of equation (9) can be written as:

$$\begin{aligned} & \sum_{\tilde{v}, \tilde{h}_A^t} \sum_{m \in M^{\mathbf{M}_t}} \mu^*(h^t)(\tilde{v}, \tilde{h}_A^t) \pi_{\tilde{v}}^{**}(h_A^t, \mathbf{M}_t) r_{\tilde{v}}^{**}(h_A^t, \mathbf{M}_t, 1)(m) \beta^{\mathbf{M}_t}(s | m) \\ &= \sum_{\tilde{v}} \sum_{m \in M^{\mathbf{M}_t}} \pi_{\tilde{v}}^{**}(h_A^t, \mathbf{M}_t) r_{\tilde{v}}^{**}(h_A^t, \mathbf{M}_t, 1)(m) \beta^{\mathbf{M}_t}(s | m) \sum_{h_A^t} \mu^*(h^t)(\tilde{v}, h_A^t) \\ &= \sum_{\tilde{v}, h_A^t} \sum_{m \in M^{\mathbf{M}_t}} \mu^*(h^t)(\tilde{v}, h_A^t) \pi_{\tilde{v}}^*(h_A^t, \mathbf{M}_t) r_{\tilde{v}}^*(h_A^t, \mathbf{M}_t, 1)(m) \beta^{\mathbf{M}_t}(s | m), \end{aligned}$$

where the first equality uses that π_v^{**}, r_v^{**} does not depend on h_A^t and the second equality uses the definition of π_v^{**}, r_v^{**} ; see equations (7) and (8). Note the last expression is the denominator in $\mu^*(h^{t+1})(v, h_A^{t+1})$. Therefore, for $h^{t+1} = (h^t, \mathbf{M}_t, 1, s_t, a_t)$, $h_A^{t+1} = (h_A^t, \mathbf{M}_t, 1, m_t, s_t, a_t)$, $h_A^t \in H_A^t(h^t)$

$$\begin{aligned}
\sum_{h_A^t} \mu^{**}(h^{t+1})(v, h_A^{t+1}) &= \frac{\sum_{h_A^t} \mu^*(h^t)(v, h_A^t) \pi_v^{**}(h_A^t, \mathbf{M}_t) r_v^{**}(h_A^t, \mathbf{M}_t, 1)(m_t) \beta^{\mathbf{M}_t}(s_t|m_t)}{\sum_{\tilde{v}, \tilde{h}_A^t} \sum_{m \in M^{\mathbf{M}_t}} \mu^*(h^t)(\tilde{v}, \tilde{h}_A^t) \pi_{\tilde{v}}^{**}(\tilde{h}_A^t, \mathbf{M}_t) r_{\tilde{v}}^{**}(\tilde{h}_A^t, \mathbf{M}_t, 1)(m) \beta^{\mathbf{M}_t}(s_t|m)} \\
&= \frac{\pi_v^{**}(h_A^t, \mathbf{M}_t) r_v^{**}(h_A^t, \mathbf{M}_t, 1)(m_t) \beta^{\mathbf{M}_t}(s_t|m_t) \sum_{h_A^t} \mu^*(h^t)(v, h_A^t)}{\sum_{\tilde{v}, \tilde{h}_A^t} \sum_{m \in M^{\mathbf{M}_t}} \mu^*(h^t)(\tilde{v}, \tilde{h}_A^t) \pi_{\tilde{v}}^{**}(\tilde{h}_A^t, \mathbf{M}_t, 1) r_{\tilde{v}}^{**}(\tilde{h}_A^t, \mathbf{M}_t, 1)(m) \beta^{\mathbf{M}_t}(s_t|m)} \\
&= \frac{\sum_{h_A^t} \mu^*(h^t)(v, h_A^t) \pi_v^*(h_A^t, \mathbf{M}_t) r_v^*(h_A^t, \mathbf{M}_t, 1)(m_t) \beta^{\mathbf{M}_t}(s_t|m_t)}{\sum_{\tilde{v}, \tilde{h}_A^t} \sum_{m \in M^{\mathbf{M}_t}} \mu^*(h^t)(\tilde{v}, \tilde{h}_A^t) \pi_{\tilde{v}}^*(\tilde{h}_A^t, \mathbf{M}_t) r_{\tilde{v}}^*(\tilde{h}_A^t, \mathbf{M}_t, 1)(m) \beta^{\mathbf{M}_t}(s_t|m)} \\
&= \frac{\sum_{h_A^t} \mu^*(h^t)(v, h_A^t) \pi_v^*(h_A^t, \mathbf{M}_t) r_v^*(h_A^t, \mathbf{M}_t, 1)(m_t) \beta^{\mathbf{M}_t}(s_t|m_t)}{\sum_{\tilde{v}, \tilde{h}_A^t} \sum_{m \in M^{\mathbf{M}_t}} \mu^*(h^t)(\tilde{v}, \tilde{h}_A^t) \pi_{\tilde{v}}^*(\tilde{h}_A^t, \mathbf{M}_t) r_{\tilde{v}}^*(\tilde{h}_A^t, \mathbf{M}_t, 1)(m) \beta^{\mathbf{M}_t}(s_t|m)} \\
&= \sum_{h_A^t} \mu^*(h^{t+1})(v, h_A^{t+1}),
\end{aligned}$$

where the second equality uses that π_v^{**}, r_v^{**} do not depend on h_A^t , the third equality uses the definition of π_v^{**}, r_v^{**} , the fourth equality uses the observation about the denominator, and the last equality follows by definition of $\mu^*(h^{t+1})$. Adding up both sides of the expression over m_t delivers the desired expression.

Now suppose we have established the above claim for each $\tau' < \tau$. We now show that it holds to $\tau' = \tau$. To see that it holds for $\tau' = \tau$, note the probability of $h^{\tau+1} = (h^\tau, \mathbf{M}_\tau, 1, s_\tau)$ conditional on h^τ , which is given by the denominator on the right hand

side of equation (9), can be written as:

$$\begin{aligned}
& \sum_{\tilde{v}, \tilde{h}_A^\tau} \sum_{\tilde{m}_\tau} \mu^{**}(\tilde{v}, \tilde{h}_A^\tau) \pi_{\tilde{v}}^{**}(\tilde{h}_A^\tau, \mathbf{M}_\tau) r_{\tilde{v}}^{**}(\tilde{h}_A^\tau, \mathbf{M}_\tau, 1) (\tilde{m}_\tau) \beta^{\mathbf{M}_\tau}(s_\tau | \tilde{m}_\tau) \\
&= \sum_{\tilde{v}} \sum_{\tilde{m}_\tau} \pi_{\tilde{v}}^{**}(\tilde{h}_A^\tau, \mathbf{M}_\tau) r_{\tilde{v}}^{**}(\tilde{h}_A^\tau, \mathbf{M}_\tau, 1) (\tilde{m}_\tau) \sum_{\tilde{h}_A^\tau} \mu^{**}(\tilde{v}, \tilde{h}_A^\tau) \beta^{\mathbf{M}_\tau}(s_\tau | \tilde{m}_\tau) \\
&= \sum_{\tilde{v}} \sum_{\tilde{m}_\tau} \pi_{\tilde{v}}^{**}(\tilde{h}_A^\tau, \mathbf{M}_\tau) r_{\tilde{v}}^{**}(\tilde{h}_A^\tau, \mathbf{M}_\tau, 1) (\tilde{m}_\tau) \sum_{\tilde{h}_A^\tau} \mu^*(\tilde{v}, \tilde{h}_A^\tau) \beta^{\mathbf{M}_\tau}(s_\tau | \tilde{m}_\tau) \\
&= \sum_{\tilde{v}, \tilde{h}_A^\tau} \sum_{\tilde{m}_\tau} \mu^*(\tilde{v}, \tilde{h}_A^\tau) \pi_{\tilde{v}}^*(\tilde{h}_A^\tau, \mathbf{M}_\tau) r_{\tilde{v}}^*(\tilde{h}_A^\tau, \mathbf{M}_\tau, 1) (\tilde{m}_\tau) \beta^{\mathbf{M}_\tau}(s_\tau | \tilde{m}_\tau),
\end{aligned}$$

where the second equality makes use of the inductive hypothesis for $\tau' = \tau - 1$, $\sum_{\tilde{h}_A^\tau} \mu^{**}(h^\tau)(\tilde{v}, \tilde{h}_A^\tau) = \sum_{\tilde{h}_A^\tau} \mu^*(h^\tau)(\tilde{v}, \tilde{h}_A^\tau)$, and the third equality uses the definition of the participation and reporting strategies defined in equations (7) and (8). Note the last line corresponds to the expression of the denominator of $\mu^*(h^{\tau+1})(v, h_A^{\tau+1})$ in the original PBE.

Therefore,

$$\begin{aligned}
\sum_{h_A^\tau} \mu^{**}(h^{\tau+1})(v, h_A^{\tau+1}) &= \frac{\sum_{h_A^\tau} \mu^{**}(h^\tau)(v, h_A^\tau) \pi_v^{**}(h_A^\tau, \mathbf{M}_\tau) r_v^{**}(h_A^\tau, \mathbf{M}_\tau, 1) (m_\tau) \beta^{\mathbf{M}_\tau}(s_\tau | m_\tau)}{\sum_{\tilde{v}, \tilde{h}_A^\tau} \sum_{\tilde{m}_\tau} \mu^*(h^\tau)(\tilde{v}, \tilde{h}_A^\tau) \pi_{\tilde{v}}^{**}(\tilde{h}_A^\tau, \mathbf{M}_\tau) r_{\tilde{v}}^*(\tilde{h}_A^\tau, \mathbf{M}_\tau, 1) (\tilde{m}_\tau) \beta^{\mathbf{M}_\tau}(s_\tau | \tilde{m}_\tau)} \\
&= \frac{\pi_v^{**}(h_A^\tau, \mathbf{M}_\tau) r_v^{**}(h_A^\tau, \mathbf{M}_\tau, 1) (m_\tau) \beta^{\mathbf{M}_\tau}(s_\tau | m_\tau) \sum_{h_A^\tau} \mu^{**}(h^\tau)(v, h_A^\tau)}{\sum_{\tilde{v}, \tilde{h}_A^\tau} \sum_{\tilde{m}_\tau} \mu^*(h^\tau)(\tilde{v}, \tilde{h}_A^\tau) \pi_{\tilde{v}}^*(\tilde{h}_A^\tau, \mathbf{M}_\tau) r_{\tilde{v}}^*(\tilde{h}_A^\tau, \mathbf{M}_\tau, 1) (\tilde{m}_\tau) \beta^{\mathbf{M}_\tau}(s_\tau | \tilde{m}_\tau)} \\
&= \frac{\pi_v^{**}(h_A^\tau, \mathbf{M}_\tau) r_v^{**}(h_A^\tau, \mathbf{M}_\tau, 1) (m_\tau) \beta^{\mathbf{M}_\tau}(s_\tau | m_\tau) \sum_{h_A^\tau} \mu^*(h^\tau)(v, h_A^\tau)}{\sum_{\tilde{v}, \tilde{h}_A^\tau} \sum_{\tilde{m}_\tau} \mu^*(h^\tau)(\tilde{v}, \tilde{h}_A^\tau) \pi_{\tilde{v}}^{**}(\tilde{h}_A^\tau, \mathbf{M}_\tau) r_{\tilde{v}}^*(\tilde{h}_A^\tau, \mathbf{M}_\tau, 1) (\tilde{m}_\tau) \beta^{\mathbf{M}_\tau}(s_\tau | \tilde{m}_\tau)} \\
&= \frac{\sum_{h_A^\tau} \mu^*(h^\tau)(v, h_A^\tau) \pi_v^*(h_A^\tau, \mathbf{M}_\tau) r_v^*(h_A^\tau, \mathbf{M}_\tau, 1) (m_\tau) \beta^{\mathbf{M}_\tau}(s_\tau | m_\tau)}{\sum_{\tilde{v}, \tilde{h}_A^\tau} \sum_{\tilde{m}_\tau} \mu^*(h^\tau)(\tilde{v}, \tilde{h}_A^\tau) \pi_{\tilde{v}}^*(\tilde{h}_A^\tau, \mathbf{M}_\tau) r_{\tilde{v}}^*(\tilde{h}_A^\tau, \mathbf{M}_\tau, 1) (\tilde{m}_\tau) \beta^{\mathbf{M}_\tau}(s_\tau | \tilde{m}_\tau)} \\
&= \sum_{h_A^\tau} \mu^*(h^{\tau+1})(v, h_A^{\tau+1}),
\end{aligned}$$

where the first equality makes use of our conclusion for the denominator, the second equality uses that π_v^{**}, r_v^{**} do not depend on h_A^τ , the third equality makes use of the inductive hypothesis, the fourth equality makes use of equations (7) and (8), and the fifth equality follows from the definition of the beliefs via Bayes' rule. Adding up the

above expression over m_τ delivers the desired conclusion.

We now use the above properties to show the payoffs of the principal do not change along the path that starts from h^t . Fix any history h^τ that is on the path of the equilibrium strategy starting from h^t . The principal's payoffs are given by:

$$\begin{aligned}
& \sum_{(v, h_A^\tau)} \mu^{**}(h^\tau)(v, h_A^\tau) \left\{ (1 - \pi_v^{**}(v, h_A^\tau)) \mathbb{E}^{\Gamma^*, \pi_v^{**}, r_v^{**}} [W(a(h^\tau), a^*, a^{\geq \tau+1}) | v, h_A^\tau, \mathbf{M}_\tau, 0] + \pi_v^{**}(h_A^\tau, \mathbf{M}_\tau) \times \right. \\
& \left. \sum_{m \in M^{\mathbf{M}_\tau}} r_v^{**}(h_A^\tau, \mathbf{M}_\tau, 1)(m) \sum_{s \in S^{\mathbf{M}_\tau}} \beta^{\mathbf{M}_\tau}(s|m) \mathbb{E}^{\Gamma^*, \pi_v^{**}, r_v^{**}} [W(a(h^\tau), a_s, a^{\geq \tau+1}, v) | v, h_A^\tau, 1, m, s] \right\} \\
& = \sum_v \left\{ (1 - \pi_v^{**}(h_A^\tau, \mathbf{M}_\tau)) \sum_{h_A^\tau} \mu^{**}(h^\tau)(v, h_A^\tau) \mathbb{E}^{\Gamma^*, \pi_v^{**}, r_v^{**}} [W(a(h^\tau), a^*, a^{\geq \tau+1}) | v, h_A^\tau, \mathbf{M}_\tau, 0] + \pi_v^{**}(h_A^\tau, \mathbf{M}_\tau) \times \right. \\
& \left. \sum_{m \in M^{\mathbf{M}_\tau}} r_v^{**}(h_A^\tau, \mathbf{M}_\tau, 1)(m) \sum_{h_A^\tau} \mu^{**}(h^\tau)(v, h_A^\tau) \sum_{s \in S^{\mathbf{M}_\tau}} \beta^{\mathbf{M}_\tau}(s|m) \mathbb{E}^{\Gamma^*, \pi_v^{**}, r_v^{**}} [W(a(h^\tau), a_s, a^{\geq \tau+1}, v) | v, h_A^\tau, 1, m, s] \right\} \\
& = \sum_v \left\{ (1 - \pi_v^*(h_A^\tau, \mathbf{M}_\tau)) \sum_{h_A^\tau} \mu^*(h^\tau)(v, h_A^\tau) \mathbb{E}^{\Gamma^*, \pi_v^*, r_v^*} [W(a(h^\tau), a^*, a^{\geq \tau+1}) | v, h_A^\tau, \mathbf{M}_\tau, 0] + \pi_v^*(h_A^\tau, \mathbf{M}_\tau) \times \right. \\
& \left. \sum_{m \in M^{\mathbf{M}_\tau}} r_v^{**}(h_A^\tau, \mathbf{M}_\tau, 1)(m) \sum_{h_A^\tau} \mu^*(h^\tau)(v, h_A^\tau) \sum_{s \in S^{\mathbf{M}_\tau}} \beta^{\mathbf{M}_\tau}(s|m) \mathbb{E}^{\Gamma^*, \pi_v^{**}, r_v^{**}} [W(a(h^\tau), a_s, a^{\geq \tau+1}, v) | v, h_A^\tau, 1, m, s] \right\} \\
& = \sum_{(v, h_A^\tau)} \mu^*(h^\tau)(v, h_A^\tau) \left\{ (1 - \pi_v^*(v, h_A^\tau)) \mathbb{E}^{\Gamma^*, \pi_v^*, r_v^*} [W(a(h^\tau), a^*, a^{\geq \tau+1}) | v, h_A^\tau, \mathbf{M}_\tau, 0] \right. \\
& \left. + \pi_v^*(h_A^\tau, \mathbf{M}_\tau) \times \sum_{m \in M^{\mathbf{M}_\tau}} r_v^*(h_A^\tau, \mathbf{M}_\tau, 1)(m) \sum_{s \in S^{\mathbf{M}_\tau}} \beta^{\mathbf{M}_\tau}(s|m) \mathbb{E}^{\Gamma^*, \pi_v^*, r_v^*} [W(a(h^\tau), a_s, a^{\geq \tau+1}, v) | v, h_A^\tau, 1, m, s] \right\} \\
& = \sum_{(v, h_A^\tau)} \mu^*(h^\tau)(v, h_A^\tau) \left\{ (1 - \pi_v^*(v, h_A^\tau)) \mathbb{E}^{\Gamma^*, \pi_v^*, r_v^*} [W(a(h^\tau), a^*, a^{\geq \tau+1}) | v, h_A^\tau, \mathbf{M}_\tau, 0] \right. \\
& \left. + \pi_v^*(h_A^\tau, \mathbf{M}_\tau) \times \sum_{m \in M^{\mathbf{M}_\tau}} r_v^*(h_A^\tau, \mathbf{M}_\tau, 1)(m) \sum_{s \in S^{\mathbf{M}_\tau}} \beta^{\mathbf{M}_\tau}(s|m) \mathbb{E}^{\Gamma^*, \pi_v^*, r_v^*} [W(a(h^\tau), a_s, a^{\geq \tau+1}, v) | v, h_A^\tau, 1, m, s] \right\},
\end{aligned}$$

where the first equality follows from noting that the agent's strategy does not depend on h_A^τ , the second equality follows from the previous result and noting that under π_v^{**}, r_v^{**} , the continuation strategy does not depend on h_A^τ , the third equality follows from the definition of the strategy, and the last equality follows from noting this equality holds after every history on the path of the equilibrium strategy starting from h^t . \square

Remark A.1. Suppose that at history h^t and after offering \mathbf{M}_t , the principal assigns probability 0 to the agent participating. In that case, his beliefs $\mu^*(h^t, \mathbf{M}_t, 1) \in \Delta(V \times H_A^t(h^t, \mathbf{M}_t, 1))$ are not determined by Bayes' rule. It is immediate to extend the proof of Proposition A.1 to show that starting from $(h^t, \mathbf{M}_t, 1)$, we can modify the agent's strategy along the path of the equilibrium strategy so that she does not condition on her payoff-irrelevant private history.

Remark A.2. The payoff-equivalent PBE assessment one obtains from Proposition A.1 satisfies the following property. On the equilibrium path, the principal's beliefs over the agent's payoff-relevant type, $v \in V$, do not depend on her payoff-irrelevant history, h_A^t . That is, for any public history on the equilibrium path of the strategy profile given the initial history, for any $v \in V$, $h_A^t, h_A^{t'} \in H_A^t(h^t)$ such that $\mu^*(h^t)(v, h_A^t), \mu^*(h^t)(v, h_A^{t'}) > 0$, we have $\mu^*(h^t)(v, h_A^t) = \mu^*(h^t)(v, h_A^{t'})$.

However, at a public history h^t reached after a deviation by the agent, either because she changes her participation strategy in a detectable way or because she triggers an output message that was not supposed to be triggered according to the equilibrium strategy, the requirements of PBE do not rule out that the principal's *updated* beliefs depend non-trivially on both v and h_A^t .

However, it follows from Proposition A.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 . After all, the agent's behavior after the deviation does not depend on his payoff irrelevant private history and the principal cannot offer mechanisms as a function of h_A^t . We record this in Corollary A.1 below and prove it in Section III of the supplementary material.

Corollary A.1. *Fix a PBE, $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$. Then, without loss of generality, for any t , public history h^t , $v \in V$ and private histories $h_A^t, h_A^{t'} \in H_A^t(h^t)$ such that $\mu^*(h^t)(v, h_A^t), \mu^*(h^t)(v, h_A^{t'}) > 0$,*

$$\mu^*(h^t)(v, h_A^t) = \mu^*(h^t)(v, h_A^{t'}).$$

Remark A.3. Note that the corollary states that equality holds only for $(v, h_A^t), (v, h_A^{t'})$ that have positive probability given the equilibrium beliefs, because the agent's strategy

may assign probability 0 to some input messages and hence some h_A^t 's cannot be given positive probability.

Hereafter, we focus on equilibria in which the agent's strategy does not depend on the payoff-irrelevant part of her private history.

We introduce one final piece of notation. Given a strategy profile $(\Gamma^*, (\pi_v^*, r_v^*)_{v \in V})$ and a history h^t , denote the continuation strategy starting at h^t implied by this profile as $(\Gamma^*, (\pi_v^*, r_v^*)_{v \in V})|_{h^t}$.

The next result shows that without loss of generality, we can focus on equilibria in which the agent participates with probability 1:

Proposition A.2. *Fix a PBE, $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$. Then, there is a PBE, $\langle \Gamma^{**}, (\pi_v^{**}, r_v^{**})_{v \in V}, \mu^{**} \rangle$ where*

1. *For every $t \geq 0$, for every $v \in V$, for every h_A^t , $\pi_v^{**}(h_A^t, \mathbf{M}_t) = 1$ for all $\mathbf{M}_t \in \text{supp } \Gamma^{**}(h^t)$.*
2. *The principal and the agent's payoffs are the same after every history as in $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$.*
3. *For every t and public history h^t , the distribution over allocations is the same as in $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$.*

Proof. Fix t and h^t such that there exists $\mathbf{M}_t \in \text{supp } \Gamma^*(h^t)$ such that $\pi_v^*(h_A^t, \mathbf{M}_t) < 1$. Recall $M^{\mathbf{M}_t}$ is a finite set and $S^{\mathbf{M}_t}$ contains $\Delta(M^{\mathbf{M}_t})$. Recall that for all $m \in M^{\mathbf{M}_t}$, $\beta^{\mathbf{M}_t}(\cdot|m) \in \Delta^*(S^{\mathbf{M}_t})$ has finite support. Then there exists $s^* \in S^{\mathbf{M}_t}$ such that $\beta^{\mathbf{M}_t}(s^*|m) = 0$ for all $m \in M^{\mathbf{M}_t}$.

Let $V_1 = \{v \in V : \pi_v^*(h_A^t, \mathbf{M}_t) > 0\} = \{v_1, \dots, v_{|V_1|}\}$. Since $|V| \leq |M^{\mathbf{M}_t}|$, we can label the latter set $M^{\mathbf{M}_t} = \{m_1, \dots, m_{|V_1|}, \dots, m_{|M^{\mathbf{M}_t}|}\}$. Modify $\beta^{\mathbf{M}_t}$ as follows. For $i = 1, \dots, |V_1|$, let

$$\tilde{\beta}(s|m_i) = \sum_{m \in M^{\mathbf{M}_t}} \beta^{\mathbf{M}_t}(s|m) r_{v_i}^*(h_A^t, \mathbf{M}_t, 1)(m).$$

Note that $\tilde{\beta}$ does not depend on h_A^t since $r_{v_i}^*$ does not depend on h_A^t .

If $|V_1| < |M^{\mathbf{M}_t}|$, let $\tilde{\beta}(s^*|m_i) = 1$ for all $i > |V_1|$ and let $\tilde{\alpha}(s^*) = \delta_{a^*}$. Modify the

strategies so that the principal, instead of offering \mathbf{M}_t , offers $\widetilde{\mathbf{M}}_t = \{\langle M^{\mathbf{M}_t}, \widetilde{\beta}, S^{\mathbf{M}_t} \rangle, \widetilde{\alpha}\}$

$$r_{v_i}^{**}(h_A^t, \mathbf{M}_t, 1)(m) = \begin{cases} \pi_{v_i}^*(h_A^t, \mathbf{M}_t) & \text{if } m = m_i \\ (1 - \pi_{v_i}^*(h_A^t, \mathbf{M}_t)) & \text{if } m = m_{|V_1|+1} \\ 0 & \text{otherwise} \end{cases}$$

and let $(\Gamma^{**}, (\pi_v^{**}, r_v^{**})_{v \in V})|_{(h^t, \widetilde{\mathbf{M}}_t, 1, s^*)} = (\Gamma^*, (\pi_v^*, r_v^*)_{v \in V})|_{(h^t, \mathbf{M}_t, 0)}$, for all other $s \in S^{\mathbf{M}_t}$, let $(\Gamma^{**}, (\pi_v^{**}, r_v^{**})_{v \in V})|_{(h^t, \widetilde{\mathbf{M}}_t, 1, s)} = (\Gamma^*, (\pi_v^*, r_v^*)_{v \in V})|_{(h^t, \mathbf{M}_t, 1, s)}$.

If $|V_1| = |M^{\mathbf{M}_t}|$ (which implies that $V_1 = V$), modify $\widetilde{\beta}$ once more so that:

$$\widetilde{\beta}(s|m_i) = \begin{cases} \pi_{v_i}^*(h_A^t, \mathbf{M}_t) \widetilde{\beta}(s|m_i) & \text{if } s \neq s^* \\ (1 - \pi_{v_i}^*(h_A^t, \mathbf{M}_t)) & \text{otherwise} \end{cases},$$

and let $\widetilde{\alpha}(s^*) = \delta_{a^*}$ as before. Modify the strategies so that the principal, instead of offering \mathbf{M}_t , offers $\widetilde{\mathbf{M}}_t = \{\langle M^{\mathbf{M}_t}, \widetilde{\beta}, S^{\mathbf{M}_t} \rangle, \widetilde{\alpha}\}$

$$r_{v_i}^{**}(h_A^t, \mathbf{M}_t, 1)(m) = \mathbb{1}[m = m_i],$$

and let $(\Gamma^{**}, (\pi_v^{**}, r_v^{**})_{v \in V})|_{(h^t, \widetilde{\mathbf{M}}_t, 1, s^*)} = (\Gamma^*, (\pi_v^*, r_v^*)_{v \in V})|_{(h^t, \mathbf{M}_t, 0)}$, for all other $s \in S^{\mathbf{M}_t}$, let $(\Gamma^{**}, (\pi_v^{**}, r_v^{**})_{v \in V})|_{(h^t, \widetilde{\mathbf{M}}_t, 1, s)} = (\Gamma^*, (\pi_v^*, r_v^*)_{v \in V})|_{(h^t, \mathbf{M}_t, 1, s)}$.

It follows immediately that the principal's and the agent's payoffs remain the same and we have not changed the distribution over allocations at any history starting from h^t . \square

Remark A.4. To keep things simple, the proof of Proposition A.2 uses the restriction that $\beta^{\mathbf{M}_t}$ has finite support and $S^{\mathbf{M}_t}$ is a *large* set to add an output message that allows the principal to

1. replicate the agent's non-participation decision and,
2. make it incentive compatible for him to offer the same continuation upon observing a^* as he was offering before.

One can write an albeit more notationally involved proof that (i) does not rely on the existence of an output message that is never sent and (ii) respects the one-to-one mapping between posteriors and output messages. This alternative proof is available from the authors upon request.

Proposition A.3. Fix a PBE, $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$ that satisfies the properties of Propositions A.1 and A.2. Then, without loss of generality, there is a one-to-one map between output messages and continuation beliefs. That is, for every t , public history h^t , $\mathbf{M}_t \in \text{supp } \Gamma^*(h^t)$, if $s_t, s'_t \in S^{*\mathbf{M}_t}$ is such that $s_t \neq s'_t$, then $\mu^*(h^t, \mathbf{M}_t, 1, s_t) \neq \mu^*(h^t, \mathbf{M}_t, 1, s'_t)$.

Lemma A.1 is used to prove Proposition A.3:

Lemma A.1. Fix a PBE assessment, $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$, that satisfies the properties of Proposition A.1 and A.2. Then, there is another assessment $\langle \Gamma^{**}, (\pi_v^{**}, r_v^{**})_{v \in V}, \mu^{**} \rangle$ that satisfies the properties of Proposition A.1 and the following holds:

1. For all h^t , for all $\mathbf{M}_t \in \text{supp } \Gamma^{**}(h^t)$, $\pi_v^{**}(h_A^t, \mathbf{M}_t) = 1$ for all $v \in V$ such that $\sum_{h_A^t} \mu^*(h^t)(v, h_A^t) > 0$.
2. For all h^t , for all $\mathbf{M}_t \in \text{supp } \Gamma^{**}(h^t)$, if $s \in S^{*\mathbf{M}_t}$, then

$$\sum_{(v, h_A^t), m \in M^{\mathbf{M}_t}} \mu^{**}(h^t)(v, h_A^t) r_v^{**}(h_A^t, \mathbf{M}_t)(m) \beta^{\mathbf{M}_t}(s|m) > 0.$$

3. For all h^t , the principal's continuation payoff remains the same and he faces the same distribution over allocations at each continuation history on the path of the equilibrium strategy given h^t . The same holds for the agent for each of her types $v \in V$ which have positive probability at h^t .

Among other things, Lemma A.1 guarantees that if the principal's strategy specifies that mechanism \mathbf{M}_t is played at history h^t , then any output message $s \in S^{*\mathbf{M}_t} (\equiv \{s \in S^{\mathbf{M}_t} : (\exists m \in M^{\mathbf{M}_t}) \beta^{\mathbf{M}_t}(s|m) > 0\})$ has positive probability under the equilibrium strategy profile. Thus, the principal is never *surprised* by the output messages he observes.

Proof of Lemma A.1. Consider a PBE assessment, $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$, that satisfies the properties of Propositions A.1 and A.2. Suppose there exists a history h^t and a type $v \in V$ to which the principal assigns probability 0. That is, $\sum_{h_A^t \in H_A^t(h^t)} \mu^*(h^t)(v, h_A^t) = 0$. Let $\mathbf{M}_t \in \text{supp } \Gamma^*(h^t)$. Let $M^{+\mathbf{M}_t} = \{m \in M^{\mathbf{M}_t} : \sum_{(\tilde{v}, \tilde{h}_A^t) \in V \times H_A^t(h^t)} \mu^*(h^t)(\tilde{v}, \tilde{h}_A^t) r_{\tilde{v}}(\tilde{h}_A^t, \mathbf{M}_t, 1)(m) > 0\}$. If $M^{\mathbf{M}_t} \setminus M^{+\mathbf{M}_t} \neq \emptyset$, note that we can do the following transformation without upsetting the equilibrium:

First, replace \mathbf{M}_t by $\mathbf{M}'_t = (\langle \beta^{\mathbf{M}'_t}, M^{\mathbf{M}'_t}, S^{\mathbf{M}'_t} \rangle, \alpha^{\mathbf{M}'_t})$ where $M^{\mathbf{M}'_t} = M^{\mathbf{M}_t}$, $S^{\mathbf{M}'_t} =$

$S^{\mathbf{M}_t}, \alpha^{\mathbf{M}'_t} = \alpha^{\mathbf{M}_t}$ and $\beta^{\mathbf{M}'_t}(\cdot|m) = \beta^{\mathbf{M}_t}(\cdot|m)$ for $m \in M^{+\mathbf{M}_t}$ and otherwise, let $\beta^{\mathbf{M}'_t}(\cdot|m) = \beta^{\mathbf{M}_t}(\cdot|m^+)$ for some $m^+ \in M^{+\mathbf{M}_t}$. Modify the principal's strategy at h^t so that instead of offering \mathbf{M}_t , he offers \mathbf{M}'_t ; that is, let $\Gamma^{**}(h^t)(\mathbf{M}'_t) = \Gamma^*(h^t)(\mathbf{M}_t)$

Second, modify continuation strategies so that

$$(\Gamma^{**}, (\pi_v^{**}, r_v^{**})_{v \in V})|_{(h^t, \mathbf{M}'_t, 1, s_t, a_t)} = (\Gamma^*, (\pi_v^*, r_v^*)_{v \in V})|_{(h^t, \mathbf{M}_t, 1, s_t, a_t)}$$

for those output messages and allocations consistent with \mathbf{M}'_t .

Third, modify the agent's strategy as follows. For $v' \in V$ such that

$\sum_{h_A^t \in H_A^t(h^t)} \mu^*(h^t)(v', h_A^t) > 0$, let $\pi_{v'}^{**}(h_A^t, \mathbf{M}'_t) = \pi_{v'}^*(h_A^t, \mathbf{M}_t) = 1$ and $r_{v'}^{**}(h_A^t, \mathbf{M}'_t, 1) = r_{v'}^*(h_A^t, \mathbf{M}_t, 1)$. Note that for these types we have not really modified the mechanism—in effect, we have removed the choices they were not making and, hence, removed possible deviations for them.

Consider now $v \in V$ such that $\sum_{h_A^t \in H_A^t(h^t)} \mu^*(h^t)(v, h_A^t) = 0$. Set $r_v^{**}(h_A^t, \mathbf{M}'_t, 1)$ so that $r_v^{**}(h_A^t, \mathbf{M}'_t, 1)(m) > 0$ if and only if m solves

$$\max_{m \in M^{+\mathbf{M}_t}} \sum_{s \in S^{\mathbf{M}_t}} \beta^{\mathbf{M}_t}(s|m) \sum_{a \in A} \alpha^{\mathbf{M}_t}(m|s) \mathbb{E}^{\Gamma^*, \pi^*, r^*} [U(a(h^t), a, a^{\geq t+1}, v) | h_A^t, \mathbf{M}_t, 1, m, s, a],$$

where we are using that \mathbf{M}'_t is only relabeling the input messages and the continuation histories remain the same as before so that for all $m \in M_t^{+\mathbf{M}_t}$:

$$\begin{aligned} & \sum_{s \in S^{\mathbf{M}'_t}} \beta^{\mathbf{M}'_t}(s|m) \sum_{a \in A} \alpha^{\mathbf{M}'_t}(m|s) \mathbb{E}^{\Gamma^{**}, \pi^{**}, r^{**}} [U(a(h^t), a, a^{\geq t+1}, v) | h_A^t, \mathbf{M}'_t, 1, m, s, a] \\ &= \sum_{s \in S^{\mathbf{M}_t}} \beta^{\mathbf{M}_t}(s|m) \sum_{a \in A} \alpha^{\mathbf{M}_t}(m|s) \mathbb{E}^{\Gamma^*, \pi^*, r^*} [U(a(h^t), a, a^{\geq t+1}, v) | h_A^t, \mathbf{M}_t, 1, m, s, a]. \end{aligned}$$

Note that the PBE assessment already specified what the agent would have done when her type is v after she reported $m \in M^{+\mathbf{M}_t}$, so we only need to choose her strategy at $(h_A^t, \mathbf{M}'_t, 1)$.

For such a type $v \in V$, however, it may no longer be optimal to participate in the mechanism when the principal offers \mathbf{M}'_t . Thus, set $\pi_v^*(h_A^t, \mathbf{M}'_t) = 1$ only if the agent's payoff from participating is at least the payoff from not participating. Note that since we

only made worse the mechanism at h^t for the agent when her type has zero probability at h^t , this does not affect her incentives at earlier histories. Hence, this modification does not alter the PBE, nor the payoffs or the distribution over allocations at each continuation history from the perspective of the principal and those types that have positive probability at h^t . It does alter the payoff and the distribution over allocations for the agent when her type has zero probability at h^t ; however, this only happens at an event that has zero probability for her given her type. \square

Proof of Proposition A.3. Take any $h^t, \mathbf{M}_t \in \text{supp } \Gamma^*(h^t)$ such that there exists $s_t, s'_t \in S^{\mathbf{M}_t}$ with $\mu^*(h^t, \mathbf{M}_t, 1, s_t) = \mu^*(h^t, \mathbf{M}_t, 1, s'_t)$, where $\mu^*(h^t, \mathbf{M}_t, 1, s_t) \in \Delta(V \times H_A^{t+1}(h^t, \mathbf{M}_t, 1, s_t))$. Note that by construction, the belief does not depend on the agent's private history. In what follows, we abuse notation and denote by μ the marginal distribution on the agent's type.

The finite support assumption implies that there is $K \geq 1$ such that we can index the principal's posteriors at history $(h^t, \mathbf{M}_t, 1, \cdot)$ as follows $\{\mu_1, \dots, \mu_K\}$. Partition $S^{*\mathbf{M}_t}$ as follows:

$$S^{*\mathbf{M}_t} = \bigcup_{l=1}^K \{s_t \in S^{*\mathbf{M}_t} : \mu^*(h^t, \mathbf{M}_t, 1, s_t) = \mu_l\} = \bigcup_{l=1}^K S^{*\mathbf{M}_t}(\mu_l).$$

Item 2 in Lemma A.1 implies that all the output messages in $S^{*\mathbf{M}_t}$ are generated with positive probability (from the point of view of the principal).

For each $l \in \{1, \dots, K\}$, let $S^{\mathbf{M}'_t}(\mu_l) = \{s_{t,1}^{\mu_l}, \dots, s_{t,H_l}^{\mu_l}\}$.

Consider the following mechanism: $\mathbf{M}'_t = (\langle \beta^{\mathbf{M}'_t}, M^{\mathbf{M}'_t}, S^{\mathbf{M}'_t} \rangle, \alpha^{\mathbf{M}'_t})$, where $M^{\mathbf{M}'_t} = M^{\mathbf{M}_t}$, $S^{\mathbf{M}'_t} = S^{\mathbf{M}_t}$. For each $l \in \{1, \dots, K\}$, let

$$\begin{aligned} \beta^{\mathbf{M}'_t}(s_{t,1}^{\mu_l} | m) &= \sum_{h=1}^{H_l} \beta^{\mathbf{M}_t}(s | m), & \beta^{\mathbf{M}'_t}(s_{t,h}^{\mu_l} | m) &= 0 \quad h \in \{2, \dots, H_l\} \\ \alpha^{\mathbf{M}'_t}(\cdot | s_{t,1}^{\mu_l}) &= \sum_{h=1}^{H_l} \frac{Pr_{\mu^*, \Gamma^*, \pi^*, r^*}(s_{t,h}^{\mu_l})}{\sum_{h'=1}^{H_l} Pr_{\mu^*, \Gamma^*, \pi^*, r^*}(s_{t,h'}^{\mu_l})} \alpha^{\mathbf{M}_t}(\cdot | s_{t,h}^{\mu_l}), \end{aligned}$$

where for $h \in \{1, \dots, H_l\}$,

$$Pr_{\mu^*, \Gamma^*, \pi^*, r^*}(s_{t,h}^{\mu_l}) = \sum_{(v, h_A^t)} \mu^*(h^t)(v, h_A^t) \sum_{m \in M^{\mathbf{M}_t}} r_v^*(h_A^t, \mathbf{M}_t, 1)(m) \beta^{\mathbf{M}_t}(s_{t,h}^{\mu_l} | m),$$

where we are using that $\mu^*(h^t)(v, h_A^t) > 0$ implies that the agent participates with probability 1.

Modify the continuation strategies as follows:

First, for those types $v \in V$ such that $\sum_{h_A^t} \mu^*(h^t)(v, h_A^t) > 0$, let $\pi_v^*(h_A^t, \mathbf{M}_t) = \pi_v^*(h_A^t, \mathbf{M}_t) = 1$ and $r_v^*(h_A^t, \mathbf{M}_t, 1) = r_v^*(h_A^t, \mathbf{M}_t, 1)$. Because the original strategies do not depend on h_A^t beyond h^t , the new strategies inherit this feature. We modify the participation and reporting strategy of the types that have zero probability at h^t at the end since their strategies do not matter for the principal's incentives.

Second, for each $l \in \{1, \dots, K\}$ and each $a \in \mathcal{A}(h^t)$ such that $\sum_{h=1}^{H_l} \alpha^{\mathbf{M}_t}(a | s_{t,h}^{\mu_l}) > 0$, partition $[0, 1] = \cup_{h=0}^{H_l-1} [\omega_h^a, \omega_{h+1}^a)$, where $\omega_0^a = 0, \omega_{H_l}^a = 1$ and for $h = 1, \dots, H_l$:

$$\omega_h^a - \omega_{h-1}^a = \frac{Pr_{\mu^*, \Gamma^*, \pi^*, r^*}(s_{t,h}^{\mu_l}) \alpha^{\mathbf{M}_t}(a | s_{t,h}^{\mu_l})}{\sum_{h'=1}^{H_l} Pr_{\mu^*, \Gamma^*, \pi^*, r^*}(s_{t,h'}^{\mu_l}) \alpha^{\mathbf{M}_t}(a | s_{t,h'}^{\mu_l})}.$$

Fix any $a \in \text{supp } \alpha^{\mathbf{M}_t}(\cdot | s_{t,\mu}^1)$. Then, for $\omega \in [\omega_{h-1}^a, \omega_h^a)$, $h \geq 1$, let

$$(\Gamma^*, (\pi_v^*, r_v^*)_{v \in V}) \Big|_{(h^t, \mathbf{M}_t, s_{t,1}^{\mu_l}, a, \omega)} = (\Gamma^*, (\pi_v^*, r_v^*)_{v \in V}) \Big|_{(h^t, \mathbf{M}_t, s_{t,h}^{\mu_l}, a, \frac{\omega - \omega_{h-1}^a}{\omega_h^a - \omega_{h-1}^a)}.$$

That is we append to history $(h^t, \mathbf{M}_t, 1, s_{t,1}^{\mu_l}, a, \omega)$, the continuation strategy that corresponds to $(h^t, \mathbf{M}_t, 1, s_{t,h}^{\mu_l}, a, \frac{\omega - \omega_{h-1}^a}{\omega_h^a - \omega_{h-1}^a})$.

This clearly guarantees that the principal's payoff is the same, that he updates to μ_l after observing $s_{t,1}^{\mu_l}$ for $l \in \{1, \dots, K\}$, and that the continuation strategies are indeed sequentially rational for the principal. Next we show that payoffs remain the same for those types of the agent to which the principal assigns positive probability, so that the above strategies are sequentially rational for them.

Fix $l \in \{1, \dots, K\}$. For $h \in \{2, \dots, H_l\}$, let k_{1h}^l denote the following ratio:

$$k_{1h}^l = \frac{Pr_{\mu^*, \Gamma^*, \pi^*, r^*}(s_{t,h}^{\mu_l})}{Pr_{\mu^*, \Gamma^*, \pi^*, r^*}(s_{t,1}^{\mu_l})}$$

and let $k_{11}^l = 1$. Because the principal updates to the same belief about the agent's type after each $s \in S^{\mathbf{M}_t}(\mu)$, we have that for all $v_j \in \text{supp } \mu$ and for all $h \in \{1, \dots, |S^{\mathbf{M}_t}(\mu)|\}$,

$$\sum_{m \in M^{\mathbf{M}_t}} r_v^*(h_A^t, \mathbf{M}_t, 1)(m) \beta(s_{t,h}^{\mu_l} | m) = k_{1h}^l \left[\sum_{m \in M^{\mathbf{M}_t}} r_v^*(h_A^t, \mathbf{M}_t, 1)(m) \beta(s_{t,1}^{\mu_l} | m) \right].$$

The above expression implies that we can write

$$\begin{aligned} \alpha^{\mathbf{M}_t}(\cdot | s_{t,1}^{\mu_l}) &= \sum_{h=1}^{H_l} \frac{Pr_{\mu^*, \Gamma^*, \pi^*, r^*}(s_{t,h}^{\mu_l})}{\sum_{h'=1}^{H_l} Pr_{\mu^*, \Gamma^*, \pi^*, r^*}(s_{t,h'}^{\mu_l})} \alpha^{\mathbf{M}_t}(\cdot | s_{t,h}^{\mu_l}) = \sum_{h=1}^{H_l} \frac{k_{1h}^l}{\sum_{h'=1}^{H_l} k_{1h'}^l} \alpha^{\mathbf{M}_t}(a | s_{t,h}^{\mu_l}), \\ \omega_h^a - \omega_{h-1}^a &= \frac{Pr_{\mu^*, \Gamma^*, \pi^*, r^*}(s_{t,h}^{\mu_l}) \alpha^{\mathbf{M}_t}(a | s_{t,h}^{\mu_l})}{\sum_{h'=1}^{H_l} Pr_{\mu^*, \Gamma^*, \pi^*, r^*}(s_{t,h'}^{\mu_l}) \alpha^{\mathbf{M}_t}(a | s_{t,h'}^{\mu_l})} = \frac{k_{1h}^l \alpha^{\mathbf{M}_t}(a | s_{t,h}^{\mu_l})}{\sum_{h'=1}^{H_l} k_{1h'}^l \alpha^{\mathbf{M}_t}(a | s_{t,h'}^{\mu_l})}. \end{aligned}$$

Moreover, we also have the following:

$$\begin{aligned} &\sum_{h=1}^{H_l} \frac{\sum_{m \in M^{\mathbf{M}_t}} r_v^*(h_A^t, \mathbf{M}_t, 1)(m) \beta^{\mathbf{M}_t}(s_{t,h}^{\mu_l} | m)}{\sum_{h'=1}^{H_l} \sum_{m' \in M^{\mathbf{M}_t}} r_v^*(h_A^t, \mathbf{M}_t, 1)(m') \beta^{\mathbf{M}_t}(s_{t,h'}^{\mu_l} | m')} \alpha^{\mathbf{M}_t}(a | s_{t,h}^{\mu_l}) \\ &= \sum_{h=1}^{H_l} \frac{k_{1h}^l}{\sum_{h'=1}^{H_l} k_{1,h'}^l} \alpha^{\mathbf{M}_t}(a | s_{t,h}^{\mu_l}) = \alpha^{\mathbf{M}_t}(a | s_{t,1}^{\mu_l}), \end{aligned}$$

and

$$\frac{\sum_{m \in M^{\mathbf{M}_t}} r_v^*(h_A^t, \mathbf{M}_t, 1)(m) \beta^{\mathbf{M}_t}(s_{t,h}^{\mu_l} | m) \alpha^{\mathbf{M}_t}(a | s_{t,h}^{\mu_l})}{\sum_{\tilde{h}=1}^{H_l} \sum_{\tilde{m} \in M^{\mathbf{M}_t}} r_v^*(h_A^t, \mathbf{M}_t, 1)(\tilde{m}) \beta^{\mathbf{M}_t}(s_{t,\tilde{h}}^{\mu_l} | \tilde{m}) \alpha^{\mathbf{M}_t}(a | s_{t,\tilde{h}}^{\mu_l})} = \frac{k_{1h}^l \alpha^{\mathbf{M}_t}(a | s_{t,h}^{\mu_l})}{\sum_{h'=1}^{H_l} k_{1h'}^l \alpha^{\mathbf{M}_t}(a | s_{t,h'}^{\mu_l})} = \omega_h^a - \omega_{h-1}^a.$$

Thus, we can write the agent's payoff in the original mechanism \mathbf{M}_t as follows:

$$\begin{aligned}
& \sum_{m \in M^{\mathbf{M}_t}} r_v^*(h_A^t, \mathbf{M}_t, 1)(m) \sum_{s \in S^{\mathbf{M}_t}} \beta^{\mathbf{M}_t}(s|m) \sum_{a \in \mathcal{A}(h^t)} \alpha^{\mathbf{M}_t}(a|s) \mathbb{E}[U(a(h^t), a, a^{\geq t+1}, v)|h_A^t, \mathbf{M}_t, 1, m, s, a] = \\
& = \sum_{m \in M^{\mathbf{M}_t}} r_v^*(h_A^t, \mathbf{M}_t, 1)(m) \sum_{l=1}^K \sum_{h=1}^{H_l} \beta^{\mathbf{M}_t}(s_{t,h}^{\mu_l}|m) \sum_{a \in \mathcal{A}(h^t)} \alpha^{\mathbf{M}_t}(a|s_{t,h}^{\mu_l}) \mathbb{E}[U(a(h^t), a, a^{\geq t+1}, v)|h_A^t, \mathbf{M}_t, 1, m, s_{t,h}^{\mu_l}, a] \\
& = \sum_{l=1}^K \sum_{a \in \mathcal{A}(h^t)} \left[\sum_{h=1}^{H_l} \sum_{m \in M^{\mathbf{M}_t}} r_v^*(h_A^t, \mathbf{M}_t, 1)(m) \beta^{\mathbf{M}_t}(s_{t,h}^{\mu_l}|m) \alpha^{\mathbf{M}_t}(a|s_{t,h}^{\mu_l}) \mathbb{E}[U(a(h^t), a, a^{\geq t+1}, v)|h_A^t, \mathbf{M}_t, 1, m, s_{t,h}^{\mu_l}, a] \right] \\
& = \sum_{l=1}^K \sum_{a \in \mathcal{A}(h^t)} \left(\sum_{\tilde{h}=1}^{H_l} \sum_{\tilde{m} \in M^{\mathbf{M}_t}} r_v^*(h_A^t, \mathbf{M}_t, 1)(\tilde{m}) \beta^{\mathbf{M}_t}(s_{t,\tilde{h}}^{\mu_l}|\tilde{m}) \alpha^{\mathbf{M}_t}(a|s_{t,\tilde{h}}^{\mu_l}) \right) \times \\
& \sum_{h=1}^{H_l} \frac{\sum_{m \in M^{\mathbf{M}_t}} r_v^*(h_A^t, \mathbf{M}_t, 1)(m) \beta^{\mathbf{M}_t}(s_{t,h}^{\mu_l}|m) \alpha^{\mathbf{M}_t}(a|s_{t,h}^{\mu_l})}{\sum_{\tilde{h}=1}^{H_l} \sum_{\tilde{m} \in M^{\mathbf{M}_t}} r_v^*(h_A^t, \mathbf{M}_t, 1)(\tilde{m}) \beta^{\mathbf{M}_t}(s_{t,\tilde{h}}^{\mu_l}|\tilde{m}) \alpha^{\mathbf{M}_t}(a|s_{t,\tilde{h}}^{\mu_l})} \mathbb{E}[U(a(h^t), a, a^{\geq t+1}, v)|h_A^t, \mathbf{M}_t, 1, m, s_{t,h}^{\mu_l}, a] \\
& = \sum_{l=1}^K \left(\sum_{h'=1}^{H_l} \sum_{m' \in M^{\mathbf{M}_t}} r_v^*(h_A^t, \mathbf{M}_t, 1)(m') \beta^*(s_{t,h'}^{\mu_l}|m') \right) \sum_{a \in \mathcal{A}(h^t)} \alpha^{\mathbf{M}_t}(a|s_{t,1}^{\mu_l}) \times \\
& \sum_{h=1}^{H_l} \frac{\sum_{m \in M^{\mathbf{M}_t}} r_v^*(h_A^t, \mathbf{M}_t, 1)(m) \beta^{\mathbf{M}_t}(s_{t,h}^{\mu_l}|m) \alpha^{\mathbf{M}_t}(a|s_{t,h}^{\mu_l})}{\sum_{\tilde{h}=1}^{H_l} \sum_{\tilde{m} \in M^{\mathbf{M}_t}} r_v^*(h_A^t, \mathbf{M}_t, 1)(\tilde{m}) \beta^{\mathbf{M}_t}(s_{t,\tilde{h}}^{\mu_l}|\tilde{m}) \alpha^{\mathbf{M}_t}(a|s_{t,\tilde{h}}^{\mu_l})} \mathbb{E}[U(a(h^t), a, a^{\geq t+1}, v)|h_A^t, \mathbf{M}_t, 1, m, s_{t,h}^{\mu_l}, a] \\
& = \sum_{m \in M^{\mathbf{M}_t}} r_v^*(h_A^t, \mathbf{M}'_t, 1)(m) \sum_{l=1}^K \beta^{\mathbf{M}'_t}(s_{t,1}^{\mu_l}|m) \sum_{a \in \mathcal{A}(h^t)} \alpha^{\mathbf{M}'_t}(a|s_{t,1}^{\mu_l}) \times \mathbb{E}[U(a(h^t), a, a^{\geq t+1}, v)|h_A^t, \mathbf{M}'_t, 1, m, s_{t,1}^{\mu_l}, a],
\end{aligned}$$

where the first equality uses the labeling of the posteriors we have used throughout the proof, the second equality is obtained by changing the order of summation, the third equality is obtained by multiplying and dividing by the probability that, conditional on the belief being μ_l , allocation a was obtained, the fourth equality is obtained by using the definition of $\alpha^{\mathbf{M}'_t}$ and grouping the terms that represent the total probability that the output message corresponds to belief μ_l , and the final equality is obtained by realizing this rewriting corresponds to the payoff the agent obtains under mechanism \mathbf{M}'_t .

Therefore, the agent's incentives remain the same when her type has positive probability at h^t .

Finally, for those $v \in V$ such that $\sum_{h_A^t \in H_A^t(h^t)} \mu^*(h^t)(v, h_A^t) = 0$, choose $r_v^*(h_A^t, \mathbf{M}'_t, 1)$ to

solve:

$$\max_{m \in M^{\mathbf{M}'_t}} \sum_{l=1}^K \beta^{\mathbf{M}'_t}(s_{t,1}^{\mu_l} | m) \sum_{a \in \mathcal{A}(h^t)} \alpha^{\mathbf{M}'_t}(a | s_{t,1}^{\mu_l}) \times \mathbb{E}^{\Gamma^*, \pi^*, r^*} [U(a(h^t), a, a^{\geq t+1}, v) | h_A^t, \mathbf{M}'_t, 1, m, s_{t,1}^{\mu_l}, a],$$

and use the payoff of this to calculate $\pi_v^*(h_A^t, \mathbf{M}'_t)$.

It is immediate that with all these modifications the assessment remains a PBE. \square

B PROOF OF THEOREM 3.1

Proof. Let $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$ denote a PBE that satisfies the properties of Propositions A.1-A.3. That is, the agent's strategy only depends on her payoff-relevant type, $v \in V$, the agent participates with probability 1 when her type has positive probability, and each output message corresponds to exactly one posterior belief.

Fix $t \geq 0$, a history h^t , and a history h^τ on the path given h^t for some $\tau \geq t$. For each $\mathbf{M}_\tau \in \text{supp } \Gamma^*(h^\tau)$, define the injective mapping:

$$\begin{aligned} \sigma(h^\tau, \mathbf{M}_\tau) : S^{\mathbf{M}_\tau} &\mapsto \Delta(V) \\ \sigma(h^\tau, \mathbf{M}_\tau)(s) &= \sum_{h_A^\tau \in H_A^\tau, m \in M^{\mathbf{M}_\tau}} \mu^*(h^\tau, \mathbf{M}_\tau, 1, s_\tau)(\cdot, h_A^\tau, m), \end{aligned} \quad (11)$$

which is well-defined by Propositions A.1 and A.3

Using this, we can define the corresponding canonical mechanism \mathbf{M}_τ^C as follows:

$$\mathbf{M}_\tau^C = \{(V, \beta^{\mathbf{M}_\tau^C}, \Delta(V)), \alpha^{\mathbf{M}_\tau^C}\}, \quad (12)$$

where³²

$$\beta^{\mathbf{M}_\tau^C}(\mu|v) = \sum_{m \in \mathbf{M}_\tau} \beta^{\mathbf{M}_\tau}(\sigma^{-1}(h^\tau, \mathbf{M}_\tau)(\mu)|m)r_v^*(h_A^\tau, \mathbf{M}_\tau, 1)(m), \quad (13)$$

$$\alpha^{\mathbf{M}_\tau^C}(\mu) = \alpha^{\mathbf{M}_\tau}(\sigma^{-1}(h^\tau, \mathbf{M}_\tau)(\mu)). \quad (14)$$

Note the construction of $\beta^{\mathbf{M}_\tau^C}$ uses the fact that the agent's reporting strategy only depends on her private type and not on the payoff-irrelevant part of her private history.

Having done this transformation for $t \leq \tau' \leq \tau$, we can map any history

$$h^\tau = (h^t, \mathbf{M}_t, 1, s_t, \dots, \mathbf{M}_\tau, 1, s_\tau, a_\tau, \omega_\tau),$$

on the path of the equilibrium strategy starting from h^t to:³³

$$h^{\tau^C} = (h^t, \mathbf{M}_t^C, 1, \sigma(h^t, \mathbf{M}_t)(s_t), a_t, \dots, \mathbf{M}_\tau^C, 1, \sigma(h^\tau, \mathbf{M}_\tau)(s_\tau), a_\tau, \omega_\tau).$$

Thus, we can define the principal's strategy so that $\Gamma'(h^{\tau^C})(\mathbf{M}_\tau^C) = \Gamma^*(h^\tau)(\mathbf{M}_\tau)$.

Given $h^\tau = (h^t, \mathbf{M}_t, 1, s_t, a_t, \dots, \mathbf{M}_\tau, 1, s_\tau, a_\tau, \omega_\tau)$ and the corresponding h^{τ^C} , the set of agent histories that is consistent with h^{τ^C} is given by:

$$\begin{aligned} H_A^{\tau^C}(h^{\tau^C}) &= \\ &= \left\{ \begin{array}{l} (h_A^t, \mathbf{M}_t^C, 1, v_t, \sigma(h^t, \mathbf{M}_t)(s_t), a_t, \dots, \mathbf{M}_{\tau-1}^C, 1, v_{\tau-1}, \sigma_{\tau-1}(h^{\tau-1}, \mathbf{M}_{\tau-1})(s_{\tau-1}), a_{\tau-1}, \omega_\tau) : \\ h_A^t \in H_A^t(h^t), (v_t, \dots, v_{\tau-1}) \in V^\tau \end{array} \right\}. \end{aligned}$$

Let $\pi'_v(h_A^{\tau^C}, \mathbf{M}_\tau^C) = \pi_v^*(h_A^\tau, \mathbf{M}_\tau) = 1$ and $r'_v(h_A^{\tau^C}, \mathbf{M}_\tau^C, 1) = \delta_v$.

Let $[V \times H_A^{\tau^C}(h^{\tau^C})]^*$ denote the set of truthful histories starting from h^t , i.e., those that have the agent of type v report v throughout $t, \dots, \tau - 1$ conditional on her participating

³²Suppose that $v \in V$ has probability 0 at history h^τ . We can use the agent's strategy profile to construct $\beta^{\mathbf{M}_\tau^C}(\cdot|v)$. A consequence of Lemma A.1 and Proposition A.3 is that the principal assigns probability 0 to such v for all $s \in S^{*\mathbf{M}_\tau}$ and hence at all $\mu \in S^{*\mathbf{M}_\tau^C}$. In other words, $\beta^{\mathbf{M}_\tau^C}(\mu|v) > 0$ implies that $\mu(v) = 0$.

³³Note that the agent always participates on the path of the strategy.

in the mechanism. With this notation at hand, let:

$$\mu'(h^{\tau^C})(v, h_A^{\tau^C}) = \mu^*(h^\tau)(v) \mathbb{1}[(v, h_A^{\tau^C}) \in [V \times H_A^{\tau^C}(h^{\tau^C})]^*]$$

It remains to check that at history h^{τ^C} , when the principal offers $\mathbf{M}_\tau^C \in \text{supp } \Gamma'(h^{\tau^C})$ and the output message is ν , his beliefs are

$$\sum_{h_A^\tau, m \in M^{\mathbf{M}_\tau}} \mu^*(h^\tau, \mathbf{M}_\tau, 1, \sigma^{-1}(h^\tau, \mathbf{M}_\tau)(\nu))(\cdot, h_A^\tau, m) = \nu.$$

Suppose that we have shown this for h^τ and now we show it for $h^{\tau+1} = (h^\tau, \mathbf{M}_\tau^C, 1, \nu)$. Note that the probability that the agent is of type v and reports v to the mechanism is:

$$\begin{aligned} \mu'(h^{\tau^C}, \mathbf{M}_\tau^C, 1, \nu)(v, h_A^\tau, \mathbf{M}_\tau^C, 1, v, \nu) &= \frac{\mu^*(h^\tau)(v) \beta^{\mathbf{M}_\tau^C}(\nu|v)}{\sum_{\tilde{v}} \mu^*(h^\tau)(\tilde{v}) \beta^{\mathbf{M}_\tau^C}(\nu|\tilde{v})} \\ &= \frac{\mu^*(h^\tau)(v) \sum_{h_A^\tau \in H_A^\tau(h^\tau)} \sum_{m \in M^{\mathbf{M}_\tau}} r_v^*(h_A^\tau, \mathbf{M}_\tau, 1)(m) \beta^{\mathbf{M}_\tau}(\sigma^{-1}(h^\tau, \mathbf{M}_\tau)(\nu)|m)}{\sum_{\tilde{v}} \mu^*(h^\tau)(\tilde{v}) \sum_{h_A^{\tau'} \in H_A^{\tau'}(h^\tau)} \sum_{m' \in M^{\mathbf{M}_\tau}} r_{\tilde{v}}^*(h_A^{\tau'}, \mathbf{M}_\tau, 1)(m') \beta^{\mathbf{M}_\tau}(\sigma^{-1}(h^\tau, \mathbf{M}_\tau)(\nu)|m')} \\ &= \sum_{h_A^\tau \in H_A^\tau(h^\tau), m \in M^{\mathbf{M}_\tau}} \frac{\mu^*(h^\tau)(v) r_v^*(h_A^\tau, \mathbf{M}_\tau, 1)(m) \beta^{\mathbf{M}_\tau}(\sigma^{-1}(h^\tau, \mathbf{M}_\tau)(\nu)|m)}{\sum_{\tilde{v}} \mu^*(h^\tau)(\tilde{v}) \sum_{h_A^{\tau'} \in H_A^{\tau'}(h^\tau)} \sum_{m' \in M^{\mathbf{M}_\tau}} r_{\tilde{v}}^*(h_A^{\tau'}, \mathbf{M}_\tau, 1)(m') \beta^{\mathbf{M}_\tau}(\sigma^{-1}(h^\tau, \mathbf{M}_\tau)(\nu)|m')} \\ &= \sum_{m \in M^{\mathbf{M}_\tau}} \sum_{h_A^\tau} \mu^*(h^\tau, \mathbf{M}_\tau, 1, \sigma^{-1}(h^\tau, \mathbf{M}_\tau)(\nu))(v, h_A^\tau, m) = \mu^*(h^\tau, \mathbf{M}_\tau, \sigma^{-1}(h^\tau, \mathbf{M}_\tau)(\nu))(v) \\ &= \nu(v) \end{aligned}$$

where the first equality uses that the agent participates with probability one and reports her type truthfully, the second equality uses the definition of $\beta^{\mathbf{M}_\tau^C}$ in equation (13), the third line is a rewriting of the second by taking the summation over $h_A^\tau \in H_A^\tau, m \in M^{\mathbf{M}_\tau}$ outside, the fourth is obtained by recognizing the expression within the summation is the principal's belief at history h^τ that the agent is of type v is at history h_A^τ and submitted message m , conditional on the output message being $\sigma^{-1}(h^\tau, \mathbf{M}_\tau)(\nu)$, and the final line uses the definition of σ to arrive to the desired expression.

For any h^t , any $\tau \geq t$, and any h^τ that is on the path given h^t , and the corresponding

h^{τ^C} , if $\mathbf{M}_\tau \notin \text{supp } \Gamma'(h^{\tau^C})$, then

$$\begin{aligned}\pi'_v(h_A^{\tau^C}, \mathbf{M}_\tau) &= \pi_v^*(h_A^\tau, \mathbf{M}_\tau) \\ r'_v(h_A^{\tau^C}, \mathbf{M}_\tau, 1) &= r_v^*(h_A^\tau, \mathbf{M}_\tau, 1)\end{aligned}$$

Note that for any $m \in M^{\mathbf{M}_\tau}$, $s \in S^{\mathbf{M}_\tau}$, the previous transformation will take the continuation strategy that follows $(h^\tau, \mathbf{M}_\tau, s_\tau, a_\tau, \omega_{\tau+1}) = h^{\tau+1}$ to one in which the principal offers canonical mechanisms.

For any other histories, specify the strategies as in the original game.

Note that we have not modified the outcome of the game after any history h^t ; in particular, the new strategy profile implements the path of the original profile. Moreover, the agent does not have an incentive to lie; otherwise, she would have had a deviation in the original profile. Additionally, the principal also has no deviations; otherwise, he would have had an incentive to deviate in the original profile. This completes the proof of Theorem 3.1. \square

C PROPERTIES OF THE CANONICAL GAME

Proposition C.1. *Fix a canonical PBE of the mechanism-selection game $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$. Then, without loss of generality, for any public history h^t , there exists a canonical mechanism \mathbf{M}_t^C such that*

1. \mathbf{M}_t^C maximizes the principal's payoff from a deviation at h^t ,
2. $\pi_v^*(h_A^t, \mathbf{M}_t^C) = 1$,
3. $r_v^*(h_A^t, \mathbf{M}_t^C, 1) = \delta_v$.

Proof. Fix a history h^t and suppose that there exists a non-canonical mechanism \mathbf{M}_t^* that maximizes the principal's payoff from a deviation. Let $\pi_v^*(h_A^t, \mathbf{M}_t^*), r_v^*(h_A^t, \mathbf{M}_t^*, 1)$ denote the agent's participation and reporting strategy upon observing the principal's choice of \mathbf{M}_t^* .

We make three observations:

1. In a canonical PBE the continuation strategy for the agent for any $h_A^{t+1} = (h_A^t, \mathbf{M}_t^*, \dots)$ does not depend on h_A^{t+1} . (Recall the proof of Proposition A.1 does not rely on the public history h^t being on the path.)
2. Therefore, we can use the same construction as in Propositions A.2 and A.3 to transform the mechanism and the continuation strategy to guarantee that the agent participates with probability 1 after observing \mathbf{M}_t^* and each output message maps exactly to one continuation belief. Denote by \mathbf{M}_t^{**} the transformed mechanism.
3. Finally, we can use the transformation in Theorem 3.1 to construct from \mathbf{M}_t^{**} a *canonical* mechanism, $\mathbf{M}_t^{C^*}$, in which the agent reports truthfully and the recommended beliefs for the principal are indeed the beliefs obtained via Bayesian updating.

Note that $\mathbf{M}_t^{C^*}$ is an available choice for the principal. It follows from the previous observations that in the original strategy profile, we can replace the best response for the agent and the continuation strategy after the principal offers $\mathbf{M}_t^{C^*}$ by those obtained in transforming \mathbf{M}_t^* to $\mathbf{M}_t^{C^*}$. In the new strategy profile, the principal is now indifferent between deviating to \mathbf{M}_t^* and to $\mathbf{M}_t^{C^*}$. \square

Corollary C.1. *If $\langle \Gamma^{*C}, (\pi_v^{*C}, r_v^{*C})_{v \in V}, \mu^{*C} \rangle$ is a canonical PBE of the canonical game, then there is an equilibrium of the mechanism-selection game $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$ that implements the same choices by the principal and the agent on the equilibrium path.*

D PROOFS OF SECTION 4

D.1 Proof of Proposition 4.1

Compare the solution to (\mathcal{P}) to the solution to the following program:

$$\begin{aligned}
& \max_{\beta, a, y, t} \sum_{i=1}^N \mu_{0,i} \sum_{h=1}^M \beta_{i,h} \sum_{q \in Q} \alpha_h(q) [w_i(q_h, y_h(q)) + t_h] & (\mathcal{A}) \\
& \text{s.t.} \begin{cases} \sum \beta_{1,h} [u_{1h} - t_h] \geq 0 \\ \sum_h (\beta_{i,h} - \beta_{k,h}) [u_{i,h} - t_h] \geq 0, (\forall i)(\forall k \in \{i-1, i+1\}) \\ y_h(q_h) \in \arg \max \sum_{i=1}^N \mu_{0,i} \beta_{i,h} w_i(q_h, y) \end{cases}
\end{aligned}$$

where $u_{i,h}$ is shorthand for $\sum_{q \in Q} \alpha_h(q) u_i(q_h, y_h(q))$.

Theorem D.1. If u_i satisfies Definition 4.1, then the values of (\mathcal{P}) and (\mathcal{A}) coincide.

Proof. We show that the solution to (\mathcal{A}) satisfies all the constraints of (\mathcal{P}) . Note first that the solution to (\mathcal{A}) satisfies that for all $i \geq 2$,

$$\begin{aligned} \sum_h \beta_{i,h} [u_{i,h} - t_h] &\geq \sum_h \beta_{i-1,h} [u_{i,h} - t_h] \\ \sum_h \beta_{i-1,h} [u_{i-1,h} - t_h] &\geq \sum_h \beta_{i,h} [u_{i-1,h} - t_h], \end{aligned}$$

so that for all $i \geq 2$, we have

$$\sum_h (\beta_{i,h} - \beta_{i-1,h}) (u_{i,h} - u_{i-1,h}) \geq 0. \quad (15)$$

We note two implications of equation (15). First, monotonic expectational differences together with equation (15) implies that if $k < i$, then it cannot be the case that

$$\sum_h (\beta_{k,h} - \beta_{k-1,h}) u_{k,h} > \sum_h (\beta_{k,h} - \beta_{k-1,h}) u_{i,h}.$$

Hence, we must have $\sum_h (\beta_{k,h} - \beta_{k-1,h}) u_{i,h} \geq \sum_h (\beta_{k,h} - \beta_{k-1,h}) u_{k,h}$, when $k < i$. Second, if $k > i$, (15) evaluated at k together with monotonic expectational differences implies that $\sum_h (\beta_{k,h} - \beta_{k-1,h}) u_{k,h} \geq \sum_h (\beta_{k,h} - \beta_{k-1,h}) u_{i,h}$. We use these two implications in what follows.

To show that the statement of the theorem holds, consider i and $j < i - 1$. The solution to (\mathcal{A}) satisfies

$$\begin{aligned} \sum_h \beta_{i,h} [u_{i,h} - t_h] &\geq \sum_h \beta_{i-1,h} [u_{i,h} - t_h] \\ \sum_h \beta_{i-1,h} [u_{i-1,h} - t_h] &\geq \sum_h \beta_{i-2,h} [u_{i-1,h} - t_h] \\ &\dots \\ \sum_h \beta_{j+1,h} [u_{j+1,h} - t_h] &\geq \sum_h \beta_{j,h} [u_{j+1,h} - t_h]. \end{aligned}$$

Adding up, we obtain

$$\sum_{k=j+1}^i \sum_{h=1}^M (\beta_{k,h} - \beta_{k-1,h}) u_{k,h} \geq \sum_h (\beta_{i,h} - \beta_{j,h}) t_h. \quad (16)$$

As discussed above, monotonic expectational differences together with equation (15) implies the left-hand side is bounded above by

$$\sum_{k=j+1}^i \sum_{h=1}^M (\beta_{k,h} - \beta_{k-1,h}) u_{i,h} = \sum_{h=1}^M (\beta_{i,h} - \beta_{j,h}) u_{i,h}. \quad (17)$$

Equations (16) and (17) imply

$$\sum_h \beta_{i,h} [u_{i,h} - t_h] \geq \sum_h \beta_{j,h} [u_{j,h} - t_h].$$

Therefore, the constraint that i does not report $j < i - 1$ holds.

Similarly, consider i and $j > i + 1$. The solution to (A) satisfies

$$\begin{aligned} \sum_h \beta_{i,h} [u_{i,h} - t_h] &\geq \sum_h \beta_{i+1,h} [u_{i,h} - t_h] \\ \sum_h \beta_{i+1,h} [u_{i+1,h} - t_h] &\geq \sum_h \beta_{i+2,h} [u_{i+1,h} - t_h] \\ &\dots \\ \sum_h \beta_{j-1,h} [u_{j-1,h} - t_h] &\geq \sum_h \beta_{j,h} [u_{j-1,h} - t_h]. \end{aligned}$$

Adding up, we obtain

$$\sum_{k=i}^{j-1} \sum_{h=1}^M (\beta_{k,h} - \beta_{k+1,h}) u_{k,h} \geq \sum_h (\beta_{i,h} - \beta_{j,h}) t_h. \quad (18)$$

As discussed above, monotonic expectational differences together with equation (15) im-

ply that the left-hand side is bounded above by

$$\sum_{k=i}^{j-1} \sum_{h=1}^M (\beta_{k,h} - \beta_{k+1,h}) u_{i,h} = \sum_{h=1}^M (\beta_{i,h} - \beta_{j,h}) u_{i,h}. \quad (19)$$

Equation (19) follows because equation (15) implies $\sum_{h=1}^M (\beta_{k,h} - \beta_{k+1,h}) u_{k,h}$ is decreasing in k .

Equations (18) and (19) imply

$$\sum_h \beta_{i,h} [u_{i,h} - t_h] \geq \sum_h \beta_{j,h} [u_{j,h} - t_h].$$

Therefore, the incentive constraint that i does not report j , $j > i + 1$ holds.

Finally, because we have all incentive compatibility constraints, it follows that, when u_i satisfies Definition 4.1, the participation constraints for $i \geq 2$ are implied by the participation constraint for $i = 1$.³⁴

To see this, note the following. First, because all incentive compatibility constraints are satisfied, we have that for all $i \geq 2$,

$$\sum_h \beta_{i,h} (u_{i,h} - t_h) \geq \sum_h \beta_{1,h} (u_{i,h} - t_h). \quad (20)$$

We can write the right-hand side of equation (20) as

$$\sum_h \beta_{1,h} (u_{1,h} - t_h + u_{i,h} - u_{1,h}) \quad (21)$$

Now, since the family $\{u_i\}_{i=1}^N$ satisfies monotonic expectational differences, we have that

$$u_{i,h} - u_{1,h} = (f(i) - f(1))g_1(a_h, y_h) + c(i) - c(1).$$

Moreover, recall from footnote 21, that we assume that $u_i(a^*, y^*) = 0$ for all $i \in \{1, \dots, N\}$.

³⁴We thank an anonymous referee for making us notice that we omitted an assumption in a previous version, which can be accessed [here](#).

Hence, we can rewrite the above as:

$$u_{i,h} - u_{1,h} = (f(i) - f(1))g_1(a_h, y_h) + c(i) - c(1) = (f(i) - f(1))(g_1(a_h, y_h) - g_1(a^*, y^*)) \geq 0$$

since $f(\cdot)$ is increasing in i and g_1 is minimized at (a^*, y^*) by the assumption in footnote 21. \square

Proposition D.1. *The participation constraint for $i = 1$ binds in the solution to (A).*

Proof. Otherwise, let $\epsilon = \beta_1 \cdot (u_1 - t)$ and consider the mechanism that charges $\tilde{t}_h = t_h + \epsilon$. All incentive constraints continue to be satisfied, the participation constraint for $i = 1$ holds, and revenue increases. \square

D.2 Proof of Propositions 4.2 and 4.3

We consider program (A) but with the following modifications:

1. The participation constraint binds for $i = 1$.
2. We write everything in terms of distribution over posteriors as opposed to communication devices.
3. We replace the principal's sequential rationality constraint by the correspondence $y_\mu^*(q) \equiv \arg \max_{y \in Y(q)} \sum_{i=1}^N \mu_i w_i(q, y)$ and the Bayesian plausibility constraint.

Therefore, we obtain

$$\begin{aligned} \max_{\tau, \mathbf{a}, t} \quad & \sum_{h=1}^M \tau(\mu_h) \sum_{i=1}^N \mu_{h,i} \sum_{q \in Q} \alpha_h(q) [w_i(q, y_h^*(q)) + t_h] \\ \text{s.t.} \quad & \begin{cases} \sum_{h=1}^M \tau(\mu_h) \frac{\mu_{h,1}}{\mu_{0,1}} [u_{1h} - t_h] = 0 \\ \sum_{h=1}^M \tau(\mu_h) \left(\frac{\mu_{h,i}}{\mu_{0,i}} - \frac{\mu_{h,k}}{\mu_{0,k}} \right) [u_{i,h} - t_h] \geq 0, (\forall i)(\forall k \in \{i-1, i+1\}) \\ \sum_{h=1}^M \tau(\mu_h) \mu_{h,i} = \mu_{0,i}, i \in \{1, \dots, N\} \end{cases} \end{aligned} \quad (\mathcal{A}')$$

where $u_{i,h}$ is shorthand for $\sum_{q \in Q} \alpha_h(q) u_i(q, y_h^*(q))$.

Now fix an allocation $\mathbf{a} = (\alpha, t, y)$ where $\alpha : \Delta(V) \mapsto \Delta^*(Q)$, $t : \Delta(V) \mapsto \mathbb{R}$ and $y : Q \times \Delta(V) \mapsto \Delta^*(\cup_{q,\mu} y^*(q, \mu))$, $y^*(q, \mu) \equiv \arg \max_{y \in Y(q)} \sum \mu_i w_i(q, y)$, $\text{supp } y(q, \mu) \subseteq y^*(q, \mu)$.

Consider the program

$$\begin{aligned} \max_{\tau} \quad & \sum_{h=1}^M \tau(\mu_h) \sum_{i=1}^N \mu_{h,i} \sum_{q \in Q} \alpha_h(q) [w_i(q, y_h(q)) + t_h] \\ \text{s.t.} \quad & \begin{cases} \sum_{h=1}^M \tau(\mu_h) \frac{\mu_{h,1}}{\mu_{0,1}} [u_{1h} - t_h] = 0 \\ \sum_{h=1}^M \tau(\mu_h) \left(\frac{\mu_{h,i}}{\mu_{0,i}} - \frac{\mu_{h,k}}{\mu_{0,k}} \right) [u_{i,h} - t_h] \geq 0, (\forall i)(\forall k \in \{i-1, i+1\}) \\ \sum_{h=1}^M \tau(\mu_h) \mu_{h,i} = \mu_{0,i}, i \in \{1, \dots, N\} \end{cases} \end{aligned} \quad (\mathcal{A}_{\mathbf{a}})$$

Note that not all allocations \mathbf{a} can be made incentive compatible. To address this issue, let $C_0^{\mathbf{a}}$ denote the policies τ that satisfy the constraints in $(\mathcal{A}_{\mathbf{a}})$. Letting $f_0^{\mathbf{a}}(\tau) = \sum_{h=1}^M \tau(\mu_h) \sum_{i=1}^N \mu_{h,i} [w_i(q_h, y_h(q_h)) + t_h]$, consider the modified objective function

$$f^{\mathbf{a}}(\tau) = \begin{cases} f_0^{\mathbf{a}}(\tau) & \text{if } \tau \in C_0^{\mathbf{a}} \\ -\infty & \text{otherwise} \end{cases}.$$

In what follows, $f^{\mathbf{a}}(\tau)$ is the objective function under consideration.

In [Doval and Skreta \(2018a\)](#), we extend the results in [Le Treust and Tomala \(2017\)](#) to show that given a constrained maximization problem,³⁵

$$\text{cav}_{g_1, \dots, g_K} f(\mu, \gamma_1, \dots, \gamma_K) := \sup \left\{ \sum_m \lambda_m f(\mu_m) : \begin{array}{l} \sum_m \lambda_m \mu_m = \mu, \\ \sum_m \lambda_m g_l(\mu_m) \geq \gamma_l, l \in \{1, \dots, r\}, \\ \sum_m \lambda_m g_l(\mu_m) = \gamma_l, l \in \{r+1, \dots, K\} \end{array} \right\}, \quad (22)$$

where $f, g_1, \dots, g_r, g_{r+1}, \dots, g_K : \Delta(V) \mapsto \mathbb{R} \cup \{-\infty\}$ is a tuple of functions defined on $\Delta(V)$, it follows that

$$\text{cav}_{g_1, \dots, g_K} f(\mu, \gamma_1, \dots, \gamma_K) = \text{cav} f^{g_1, \dots, g_K}(\mu, \gamma_1, \dots, \gamma_K)$$

³⁵We extend the construction in their paper for completeness given that our problem includes multiple inequality constraints and equality constraints.

where $f^{g_1, \dots, g_K} : \Delta(V) \times \mathbb{R}^K \mapsto \mathbb{R} \cup \{-\infty\}$ is such that

$$f^{g_1, \dots, g_K}(\mu, \gamma_1, \dots, \gamma_K) = \begin{cases} f(\mu) & \text{if } \gamma_i \leq g_i(\mu), i \in \{1, \dots, r\} \wedge \gamma_i = g_i(\mu), i \in \{r+1, \dots, K\} \\ -\infty & \text{otherwise} \end{cases} \quad (23)$$

That is, the constrained *Bayesian persuasion* problem with r inequality constraints and $K - r$ equality constraints in (22) can be thought of a Bayesian persuasion problem in which the objective has domain $\Delta(V) \times \mathbb{R}^K$ as defined in (23). We use this to derive properties about the number of posteriors used in the optimal solution.

Note that (\mathcal{A}_a) is a version of this program with $r = 2N - 2$ and

$$\begin{aligned} g_i(\mu) &= \left[\frac{\mu_i}{\mu_{0,i}} - \frac{\mu_{i+1}}{\mu_{0,i+1}} \right] \sum_{q \in Q} \alpha(\mu)(q) [u_i(q, y_\mu(q)) - t(\mu)], i \in \{1, \dots, N-1\} \\ g_{N-2+i}(\mu) &= \left[\frac{\mu_i}{\mu_{0,i}} - \frac{\mu_{i-1}}{\mu_{0,i-1}} \right] \sum_{q \in Q} \alpha(\mu)(q) [u_i(q, y_\mu(q)) - t(\mu)], i \in \{2, \dots, N\} \\ g_{2N-1}(\mu) &= \frac{\mu_1}{\mu_{0,1}} \sum_{q \in Q} \alpha(\mu)(q) [u_1(q, y_\mu(q)) - t(\mu)] \end{aligned}$$

and $\gamma_i = 0$ for all i . We then have the following:

Corollary D.1. *Suppose the value of (\mathcal{A}_a) is finite. Then, τ puts positive probability in at most $3N - 1$ beliefs.*

Proof. This follows from Proposition 3.1 in Doval and Skreta (2018a) and Carathéodory's theorem (see, e.g., Rockafellar (1970)). \square

Similarly, we can construct a program (\mathcal{M}_a)

$$\begin{aligned} \max_{\tau} \quad & \sum_{h=1}^M \tau(\mu_h) \sum_{i=1}^N \mu_{h,i} \sum_{q \in Q} \alpha_h(q) [w_i(q, y_h(q)) + u_i(q, y_h(q)) - \sum_{i+1 \leq l} \frac{\mu_{0,l}}{\mu_{0,i}} (u_{i+1,h} - u_{i,h})] \\ \text{s.t.} \quad & \begin{cases} \sum_{h=1}^M \tau(\mu_h) \left[\frac{\mu_{h,i+1}}{\mu_{0,i+1}} - \frac{\mu_{h,i}}{\mu_{0,i}} \right] (u_{i+1,h} - u_{i,h}) \geq 0, & i \in \{1, \dots, N-1\} \\ \sum_{h=1}^M \tau(\mu_h) \mu_h = \mu_0 \end{cases} \quad (\mathcal{M}_a) \end{aligned}$$

Corollary D.2. *Suppose the value of (\mathcal{M}_a) is finite. Then τ puts positive probability*

on at most $2N - 1$ beliefs.

Another immediate corollary is that dropping constraints from a program lowers the upper bound on the number of beliefs in the support of the solution to the program:

Corollary D.3. *Suppose the value of (\mathcal{M}_a) is finite and only M constraints bind. Then, τ puts positive probability on at most $N + M$ beliefs.*

Proof. See Corollary 3.2 in [Doval and Skreta \(2018a\)](#). □

D.3 Example 2

The next example illustrates that even if the solution to the relaxed program, (\mathcal{R}) , satisfies the monotonicity constraints (\mathbf{M}) , it may not be a solution to the original problem.

Example 2. Consider the sale of a durable good example in Section 2.1, but with three types $V \equiv \{v_L, v_M, v_H\}$. We provide a parametrization of the problem such that the solution to the relaxed program (\mathcal{R}) has the following properties:

1. In period 1, v_H buys with probability 1, v_L buys with probability 0, and v_M buys with positive probability (but bounded away from one).
2. The allocation satisfies the monotonicity constraints, (\mathbf{M}) .
3. The communication device generates two posteriors, μ^{HM}, μ^{ML} , where³⁶

$$\begin{aligned}\mu^{HM}(v_H) &= \frac{v_M}{v_H}, \mu^{HM}(v_L) = 0 \\ \mu^{ML}(v_M) &= \frac{v_L}{v_M}, \mu^{ML}(v_H) = 0.\end{aligned}$$

4. However, it is not possible to find two transfers, $t(\mu^{HM}), t(\mu^{ML})$, that satisfy that (i) v_L 's participation constraint binds, and (ii) both v_M and v_H 's downward-looking incentive constraints bind.

³⁶The reader may recognize μ^{HM}, μ^{LM} as two of the *extremal beliefs* in [Bergemann et al. \(2015\)](#): μ^{HM} makes the principal indifferent between setting a price of v_M or v_H , whereas μ^{ML} makes the principal indifferent between setting a price of v_M or v_L . Indeed, the optimal mechanism for the principal need only put weight on the extremal beliefs. Details are available from the authors upon request.

The parameters are as follows.³⁷ First, the prior is given by

$$\mu_0(v_H) = 0.4637, \mu_0(v_L) = 0.1194, \mu_0(v_M) = 0.4169,$$

and is chosen so that it is a convex combination of μ^{HM}, μ^{ML} . The values for the types are

$$v_H = 4.8385, v_M = 2.5528, v_L = 0.0357,$$

and are chosen so that v_L 's virtual valuation is negative, whereas v_M 's virtual valuation is positive. Also, we set $\delta = 0.95$. With these values, we have that

$$\begin{aligned}\mu^{HM}(v_H) &= 0.5276 \\ \mu^{ML}(v_M) &= 0.0140.\end{aligned}$$

and the communication device satisfies:

$$\begin{aligned}\beta(\mu^{HM}|v_H) &= 1 \\ \beta(\mu^{ML}|v_L) &= 1 \\ \beta(\mu^{ML}|v_M) &= \frac{\mu^{ML}(v_M)\tau(\mu^{ML})}{\mu_0(v_M)} = \frac{0.0140 \times 0.8789}{0.4169}\end{aligned}$$

Because v_L never buys (the monopolist in period 1 recommends a price of v_M in period 2), it has to be that $t(\mu^{ML}) = 0$. To determine $t(\mu^{HM})$, we have the following two equations:

$$\begin{aligned}v_H - t(\mu^{HM}) &= \beta(\mu^{HM}|v_M)(v_H - t(\mu^{HM})) + \beta(\mu^{ML}|v_M)\delta \times (v_H - v_M) \\ \beta(\mu^{HM}|v_M)(v_M - t(\mu^{HM})) + \beta(\mu^{ML}|v_M)\delta \times (v_M - v_M) &= -t(\mu^{ML})\delta(v_M - v_M) = 0.\end{aligned}$$

The first equality implies $t(\mu^{HM}) = v_H - \delta(v_H - v_M)$, whereas the second implies $t(\mu^{HM}) = v_M$. Hence, it is not possible to find two transfers, $t(\mu^{HM}), t(\mu^{ML})$ that satisfy that the downward looking constraints bind and implements the solution to the relaxed program.

³⁷The Matlab code, which implements the linear program used to derive the example, is available upon request.

E OUTPUT MESSAGES AS RECOMMENDATIONS

Proof of Proposition 5.1. We prove 2 implies 3. That 1 implies 2 follows from the results of the paper. Because a straightforward mechanism is a particular case of a mechanism, it follows immediately that 3 implies 1, thus completing the proof.

Thus, consider $\langle (V, \beta, \Delta(V)), \alpha \rangle$ and $y : \Delta(V) \times Q \mapsto \cup \Delta(Y(q))$ that solves (R), or equivalently (3). The finite support assumption implies we can label $\{\mu_1, \dots, \mu_H\}$ the posteriors that are induced with positive probability by the mechanism. Given $h \in \{1, \dots, H\}$, let $(QY)_h = \{(q, y) : \alpha_h(q) \times y_h(q)(y) > 0\}$ denote the pairs (q, y) that are implemented when the belief is μ_h . Let $(QY)^* = \cup_{h \in H} (QY)_h$. Consider now the following mechanism $\langle (V, \beta', (QY)^*), \alpha' \rangle$ and the continuation strategy $y' : (QY)^* \times Q \mapsto \cup \Delta(Y(q))$ such that $\alpha'_{(q,y)}(q') = \mathbb{1}[q' = q]$ and $y'_{(q,y)}(q)(y') = \mathbb{1}[y' = y]$. Moreover, let

$$\beta'((q, y)|v_i) = \sum_{h=1}^H \beta(\mu_h|v_i) \alpha_h(q) y_h(q)(y).$$

Clearly, this mechanism delivers the same payoff to the principal and the agent. We now verify that it remains incentive compatible for the principal to follow the recommendations. Fix $(q, y) \in (QY)^*$. The principal's belief upon observing the output (q, y) is

$$\begin{aligned} \mu_{(q,y)}(v_i) &= \frac{\mu_i^0 \beta'((q, y)|v_i)}{\sum_j \mu_j^0 \beta'((q, y)|v_j)} \\ &= \frac{\mu_i^0 \sum_{h=1}^H \beta(\mu_h|v_i) \alpha_h(q) y_h(q)(y)}{\sum_j \mu_j^0 \sum_{h=1}^H \beta(\mu_h|v_j) \alpha_h(q) y_h(q)(y)} \\ &= \sum_{h \in H: (q,y) \in (QY)_h} \frac{\mu_i^0 \beta(\mu_h|v_i)}{\sum_{j'} \mu_{j'}^0 \beta(\mu_h|v_{j'})} \frac{\sum_{j'} \mu_{j'}^0 \beta(\mu_h|v_{j'})}{\sum_{h=1}^H \sum_j \mu_j^0 \beta(\mu_h|v_j) \alpha_h(q) y_h(q)(y)} \\ &= \sum_{h \in H: (q,y) \in (QY)_h} \mu_{h,i} \frac{\tau(\mu_h)}{\sum_{h': (q,y) \in (QY)_{h'}} \tau(\mu_{h'})}, \end{aligned}$$

where recall that $\tau(\mu_h) = \sum_{j=1}^N \mu_j^0 \beta(\mu_j|v_j)$. Then the payoff of the second-period principal

when he observes (q, y) and chooses $y' \in Y(q)$ can be written as

$$\sum_{h \in H: (q, y) \in (QY)_h} \frac{\tau(\mu_h)}{\sum_{h': (q, y) \in (QY)_{h'}} \tau(\mu_{h'})} \sum_{i=1}^N \mu_{h,i} w_i(q, y') \leq \sum_{h \in H: (q, y) \in (QY)_h} \frac{\tau(\mu_h)}{\sum_{h': (q, y) \in (QY)_{h'}} \tau(\mu_{h'})} \sum_{i=1}^N \mu_{h,i} w_i(q, y),$$

where the inequality follows from knowing that $y_h(q)(y) > 0$ for all h such that $(q, y) \in (QY)_h$.

Similarly, using the expression in (3), we can write the principal's payoff from the new mechanism conditional on the output being (q, y) as:

$$\sum_{h \in H: (q, y) \in (QY)_h} \frac{\tau(\mu_h)}{\sum_{h': (q, y) \in (QY)_{h'}} \tau(\mu_{h'})} \sum_{i=1}^N \mu_{h,i} [w_i(q, y; \mu^0) + \hat{u}_i(q, y; \mu_0)],$$

and note he has no incentive to choose another q , because for each h such that $(q, y) \in (QY)_h$, we have that q is in the set of maximizers of $\sum_{i=1}^N \mu_{h,i} [w_i(\cdot, y_{\mu_h}; \mu^0) + \hat{u}_i(\cdot, y_{\mu_h}; \mu_0)]$. \square

F IMPLEMENTATION VIA CONTRACTS

Proof of Proposition 5.2.

Necessity: Assume there exists t' such that (β, q, y) satisfies (DIC-P). Consider $i < j$ and μ, μ' such that $\mu(v_i)\mu'(v_j) > 0$. Then, the following must hold:

$$\begin{aligned} u_i(q(\mu), y(\mu)) - t'(\mu) &\geq u_i(q(\mu'), y(\mu')) - t'(\mu') \\ u_j(q(\mu'), y(\mu')) - t'(\mu') &\geq u_j(q(\mu), y(\mu)) - t'(\mu), \end{aligned}$$

which implies that

$$u_j(q(\mu'), y(\mu')) - u_j(q(\mu), y(\mu)) \geq u_i(q(\mu'), y(\mu')) - u_i(q(\mu), y(\mu)) \quad (24)$$

That is, letting

$$D_i(\mu', \mu) = u_i(q(\mu'), y(\mu')) - u_i(q(\mu), y(\mu)), \quad (25)$$

we need that

$$D_j(\mu', \mu) \geq D_i(\mu', \mu), \text{ whenever } \mu'(v_j)\mu(v_i) > 0. \quad (\text{DIC-M})$$

Note that Assumption 1 implies $D_i(\mu, \mu')$ is strictly increasing in i . Thus, (DIC-M) holds with strict inequality when $i < j$.

To derive the necessary conditions for the communication device, β , suppose now that $\mu(v_i)\mu'(v_i) > 0$ for $\beta(\mu|v_i)\beta(\mu'|v_i) > 0$. Because (β, q, y) satisfies (DIC-P) for t' ,

$$t'(\mu) - t'(\mu') = D_i(\mu, \mu').$$

Because under Assumption 1 $D_i(\mu, \mu')$ is strictly increasing in i , for all $j > i$, it has to be the case that $\mu'(v_j) = 0$, and for all $j < i$, it has to be the case that $\mu(v_j) = 0$. To see this, note that if $j > i$, then $D_j(\mu, \mu') > t(\mu) - t(\mu')$, and hence $v_j > v_i$ can never select the allocation at μ' . Likewise, if $j < i$, then $D_j(\mu, \mu') < t(\mu) - t(\mu')$, and hence $v_j < v_i$ can never select the allocation at μ .

Moreover, if there are three beliefs μ, μ', μ'' such that $\mu(v_i)\mu'(v_i)\mu''(v_i) > 0$ such that $u_i(q(\mu), y(\mu)) \geq u_i(q(\mu'), y(\mu')) \geq u_i(q(\mu''), y(\mu''))$ and $D_i(\mu, \mu')$ and $D_i(\mu', \mu'')$ are strictly increasing in i , then it has to be the case that: (i) $j > i$, then $\mu'(v_j) = \mu''(v_j) = 0$, and (ii) $j < i$, then $\mu'(v_j) = \mu(v_j) = 0$. Then, $\mu'(v_i) = 1$. It follows then that there are at most three beliefs at which v_i has positive probability – if we had four or more, the ones that give intermediate utility to v_i assign probability one to v_i . Hence, they must correspond to the same belief.

Finally, suppose $\mu(v_i)\mu'(v_i) > 0$, $D_i(\mu, \mu') > 0$ and $\mu(v_j) > 0$ for $j > i$. We now show that for all $l \in \{i + 1, \dots, j - 1\}$, it has to be the case that $\mu(v_l) > 0$. Towards a contradiction, assume there exists $v_l, l \in \{i + 1, \dots, j - 1\}$ such that $\mu(v_l) = 0$. Because all types have positive probability, there exists $\mu' : \beta(\mu'|v_l) > 0$. Because under t' , (β, q, y) satisfies (DIC-P), it follows that

$$u_l(q(\mu'), y(\mu')) - t'(\mu') \geq u_l(q(\mu), y(\mu)) - t'(\mu).$$

Because $\mu(v_j) > 0$, we have that

$$u_j(q(\mu), y(\mu)) - t'(\mu) \geq u_j(q(\mu'), y(\mu')) - t'(\mu').$$

It follows that

$$D_l(\mu', \mu) \geq D_j(\mu', \mu),$$

and under our assumption, this inequality is strict. Monotonic expectational differences implies that

$$D_i(\mu', \mu) \geq D_l(\mu', \mu) \geq t'(\mu') - t'(\mu).$$

The above expression contradicts that (β, q, y) satisfies (DIC-P) under t' , because v_i would strictly prefer to select the allocation in μ' to the allocation in μ .

Sufficiency: Suppose (β, q, y) satisfies the assumptions in the statement of the proposition. Then, it is possible to label the beliefs μ^1, \dots, μ^M so that $k < l$ implies that $\bar{v}(\mu^k) \equiv \max\{v : \mu^k(v) > 0\} \leq \underline{v}(\mu^l) \equiv \min\{v : \mu^l(v) > 0\}$.

Set $t'(\mu^1) = u_1(q(\mu^1), y(\mu^1))$. Note that, by definition, $v_1 = \underline{v}(\mu^1)$. For $n > 1$, define recursively

$$t'(\mu^n) = u_{\underline{v}(\mu^n)}(q(\mu^n), y(\mu^n)) - (u_{\underline{v}(\mu^n)}(q(\mu^{n-1}), y(\mu^{n-1})) - t'(\mu^{n-1})). \quad (26)$$

We now verify that (β, q, y) together with t' satisfies (DIC-P). We first check that v_i is indifferent between μ and μ' whenever $\mu(v_i)\mu'(v_i) > 0$. Monotonicity of the information structure induced by β implies that, without loss of generality, $v_i = \underline{v}(\mu) = \bar{v}(\mu')$; moreover, $\mu = \mu^l, \mu' = \mu^k$, with $k \in \{l+1, l+2\}$ and if $k = l+2$, then $v_i = \underline{v}(\mu^{l+1}) = \bar{v}(\mu^{l+1})$.

Consider first the case in which $k = l+1$. Then, it follows from equation (26) that

$$t'(\mu^{l+1}) - t'(\mu^l) = u_i(q(\mu^{l+1}), y(\mu^{l+1})) - u_i(q(\mu^l), y(\mu^l)),$$

so v_i is indeed indifferent. Now consider the case in which $k = l+2$, then recalling that

$\underline{v}(\mu^{l+1}) = \underline{v}(\mu^{l+2}) = v_i$, we have

$$\begin{aligned} t'(\mu^{l+1}) - t'(\mu^l) &= u_i(q(\mu^{l+1}), y(\mu^{l+1})) - u_i(q(\mu^l), y(\mu^l)), \\ t'(\mu^{l+2}) - t'(\mu^{l+1}) &= u_i(q(\mu^{l+2}), y(\mu^{l+2})) - u_i(q(\mu^{l+1}), y(\mu^{l+1})), \end{aligned}$$

so that v_i is indifferent between selecting the outcome that corresponds to either of the three beliefs, $\mu^l, \mu^{l+1}, \mu^{l+2}$. Finally, we show that when the agent is of type v_i , she does not want to announce any other belief μ^k such that $\mu^k(v_i) = 0$. To see this, let $\mu^l(v_i) > 0$ and consider the case in which $l < k$. Note first that

$$t'(\mu^l) = u_1(q(\mu^1), y(\mu^1)) + \sum_{n=2}^l (u_{\underline{v}(\mu^n)}(q(\mu^n), y(\mu^n)) - u_{\underline{v}(\mu^{n-1})}(q(\mu^{n-1}), y(\mu^{n-1}))),$$

so that

$$\begin{aligned} t'(\mu^k) - t'(\mu^l) &= \sum_{n=l+1}^k (u_{\underline{v}(\mu^n)}(q(\mu^n), y(\mu^n)) - u_{\underline{v}(\mu^{n-1})}(q(\mu^{n-1}), y(\mu^{n-1}))) \\ &= \sum_{n=l+1}^k D_{\underline{v}(\mu^n)}(\mu^n, \mu^{n-1}). \end{aligned}$$

Then,

$$\begin{aligned} &u_i(q(\mu^l), y(\mu^l)) - u_i(q(\mu^k), y(\mu^k)) + t'(\mu^k) - t'(\mu^l) = \\ &= u_i(q(\mu^l), y(\mu^l)) - u_i(q(\mu^k), y(\mu^k)) + \sum_{n=l+1}^k D_{\underline{v}(\mu^n)}(\mu^n, \mu^{n-1}) \\ &\geq u_i(q(\mu^l), y(\mu^l)) - u_i(q(\mu^k), y(\mu^k)) + \sum_{n=l+1}^k D_i(\mu^n, \mu^{n-1}) = 0, \end{aligned}$$

where the inequality follows from (DIC-M). A similar argument shows that the same holds for $l > k$.

□

REFERENCES

- AKBARPOUR, M. AND S. LI (2018): “Credible mechanisms,” in *Proceedings of the 2018 ACM Conference on Economics and Computation*, ACM, 371–371.
- BERGEMANN, D., B. BROOKS, AND S. MORRIS (2015): “The limits of price discrimination,” *American Economic Review*, 105, 921–57.
- BEST, J. AND D. QUIGLEY (2017): “Persuasion for the long run,” *Working Paper*.
- BESTER, H. AND R. STRAUZ (2000): “Imperfect commitment and the revelation principle: the multi-agent case,” *Economics Letters*, 69, 165–171.
- (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.
- BOLESLAVSKY, R. AND K. KIM (2018): “Bayesian persuasion and moral hazard,” Tech. rep.
- 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.
- CRAWFORD, V. P. AND J. SOBEL (1982): “Strategic information transmission,” *Econometrica: Journal of the Econometric Society*, 1431–1451.
- 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.

- DOVAL, L. AND V. SKRETA (2018a): “Constrained information design: Toolkit,” Tech. rep.
- (2018b): “Sequentially Optimal Mechanisms: Infinite horizon,” Tech. rep.
- FREIXAS, X., R. GUESNERIE, AND J. TIROLE (1985): “Planning under incomplete information and the ratchet effect,” *The review of economic studies*, 52, 173–191.
- GEORGIADIS, G. AND B. SZENTES (2018): “Optimal Monitoring Design,” Tech. rep., mimeo.
- GERARDI, D. AND L. MAESTRI (2018): “Dynamic Contracting with Limited Commitment and the Ratchet Effect,” Tech. rep.
- GERSHKOV, A., J. K. GOEREE, A. KUSHNIR, B. MOLDOVANU, AND X. SHI (2013): “On the equivalence of Bayesian and dominant strategy implementation,” *Econometrica*, 81, 197–220.
- GIBBARD, A. (1973): “Manipulation of voting schemes: a general result,” *Econometrica: journal of the Econometric Society*, 587–601.
- GOLOSOV, M. AND L. IOVINO (2016): “Social Insurance, Information Revelation, and Lack of Commitment,” *Working Paper*.
- GUL, F., H. SONNENSCHNEIN, 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.
- HENDON, E., H. J. JACOBSEN, AND B. SLOTH (1996): “The one-shot-deviation principle for sequential rationality,” *Games and Economic Behavior*, 12, 274–282.
- KAMENICA, E. AND M. GENTZKOW (2011): “Bayesian persuasion,” *American Economic Review*, 101, 2590–2615.
- KARTIK, N., S. LEE, AND D. RAPPOPORT (2017): “Single-crossing differences on distributions,” Tech. rep., Working paper.

- 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: Journal of the Econometric Society*, 1153–1175.
- LE TREUST, M. AND T. TOMALA (2017): “Persuasion with limited communication capacity,” *Working Paper*.
- LIPNOWSKI, E. AND D. RAVID (2017): “Cheap talk with transparent motives,” *Working Paper*.
- LIU, Q., K. MIERENDORFF, X. SHI, AND W. ZHONG (2018): “Auctions with limited commitment,” *American Economic Review*, *forthcoming*.
- MANELLI, A. M. AND D. R. VINCENT (2010): “Bayesian and dominant-strategy implementation in the independent private-values model,” *Econometrica*, 78, 1905–1938.
- 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.
- MOOKHERJEE, D. AND S. REICHELSTEIN (1992): “Dominant strategy implementation of Bayesian incentive compatible allocation rules,” *Journal of Economic Theory*, 56, 378 – 399.
- MYERSON, R. B. (1979): “Incentive compatibility and the bargaining problem,” *Econometrica: journal of the Econometric Society*, 61–73.
- (1982): “Optimal coordination mechanisms in generalized principal–agent problems,” *Journal of Mathematical Economics*, 10, 67–81.
- (1983): “Mechanism design by an informed principal,” *Econometrica: Journal of the Econometric Society*, 1767–1797.
- ROCKAFELLAR, R. T. (1970): *Convex analysis*, Princeton University Press.
- SALAMANCA, A. (2016): “The value of mediated communication,” *Working Paper*.

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