Optimal mechanism for the sale of a durable good*

Laura Doval† Vasiliki Skreta‡

May 28, 2021

Abstract

A buyer wishes to purchase a durable good from a seller who in each period chooses a mechanism under limited commitment. The buyer's valuation is binary and fully persistent. We show that posted prices implement all equilibrium outcomes of an infinite-horizon, mechanism selection game. Despite being able to choose mechanisms, the seller can do no better and no worse than if he chose prices in each period, so that he is subject to Coase's conjecture. Our analysis marries insights from information and mechanism design with those from the literature on durable goods. We do so by relying on the revelation principle in Doval and Skreta (2020).

KEYWORDS: *mechanism design, limited commitment, information design, self-generation, posted prices, Coase conjecture*

JEL CLASSIFICATION: D84, D86

*We thank the Editor Simon Board and three anonymous referees for excellent comments which substantially improved the paper. We would like to thank Rahul Deb, Frederic Koessler, Dan Quigley, Pablo Schenone, and especially Max Stinchcombe, as well as audiences at Cowles, SITE, and Stony Brook, for thought-provoking questions and illuminating discussions. Vasiliki Skreta is grateful for generous financial support through the ERC consolidator grant 682417 "Frontiers in design." This research was supported by grants from the National Science Foundation (SES-1851744 and SES-1851729).

†Columbia University and CEPR. E-mail: laura.doval@columbia.edu

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

1 Introduction

We characterize the equilibrium outcomes of an infinite-horizon, mechanism-selection game between a durable-good seller and a privately informed buyer under *limited commitment*, so that the seller can commit to today's mechanism, but not to the mechanism he will offer if no sale occurs. Theorem 1 shows that all equilibrium outcomes can be implemented via posted prices. We construct a Perfect Bayesian equilibrium of the mechanism-selection game, which achieves the seller's unique equilibrium payoff, and we show that it implements the *essentially* unique equilibrium outcome.[1] In this equilibrium, as long as a sale has not occurred, the seller will choose a mechanism that can be implemented as a posted price. Despite being able to choose from a rich set of mechanisms, the seller can do no better and no worse than if he could only choose prices in each period.

In our game, an uninformed seller faces a privately informed buyer, whose valuation is binary, fully persistent, and strictly above the seller's marginal cost.[2] In each period, as long as the good has not been sold, the seller offers the buyer a *mechanism*, the rules of which determine the allocation for that period. A mechanism consists of (i) a set of *input* messages for the buyer, and (ii) for each input message, a distribution over *output* messages and allocations. While the seller observes the output message and the allocation, he does not observe the input message the buyer submits to the mechanism. Thus, when designing the mechanism, the seller gets to design *how much* he observes about the buyer's choices and hence, design his beliefs about the buyer's value. The combination of mechanism design and information design elements is key to our characterization.

Our analysis bridges the literatures on mechanism design and on the durable-good monopolist, especially the work of Gul, Sonnenschein, and Wilson (1986). To see this, it is useful to review the two main steps involved in the proof of Theorem 1. First, we construct an assessment that is identical along the path to that in Hart and Tirole (1988), which we dub the *posted-prices assessment*. In this assessment, along the path of play, the seller sells the good using a decreasing sequence of prices, which reflect that conditional on the good not being sold the seller assigns less probability to the buyer's valu-

[1] Whenever the equilibrium outcome is not unique, all equilibrium outcomes are achieved also via a sequence of posted prices.

[2] In the *no gap* case, all equilibrium outcomes can still be implemented in posted prices, but equilibrium multiplicity arises as in Ausubel and Deneckere (1989) (see footnote 16).

ation being high. Using the property that beliefs go down in the posted-prices assessment, we show that the seller's payoff in this assessment is a lower bound on the seller's equilibrium payoff. The logic of the argument is analogous to the argument in Gul, Sonnenschein, and Wilson (1986) that there is a unique equilibrium in the gap case.

Second, we argue that the seller's payoff under the posted-prices assessment is an upper bound on the seller's equilibrium payoff. To do so, we rely on an *auxiliary* program, that only involves the seller (see OPT in Section 4). In this program, the seller maximizes the dynamic analogue of the *virtual surplus*, by choosing a Bayes' plausible distribution over posteriors and for each posterior (i) a probability of trade and (ii) a vector of equilibrium continuation payoffs. We arrive at the program defined in OPT by relying on the tools in our previous work, Doval and Skreta (2020). The main theorem in Doval and Skreta (2020) allows us to simplify the class of mechanisms the seller will offer in any equilibrium of the game and the buyer's equilibrium behavior. This step reduces the search of the optimal sequence of mechanisms to those that satisfy, loosely speaking, a sequence of *participation* and *truthtelling* constraints, allowing us for the most part to ignore the buyer as a player. Like in standard mechanism design, the participation constraint of the low-valuation buyer and the truthtelling constraint of the high-valuation buyer impose an upper bound on the revenue the seller can extract within a period (Section 4). Replacing this upper bound in the seller's payoff provides us with a dynamic analogue of the virtual surplus (Equation 3), where the seller's payoff is written as a function of the allocation, but also the continuation payoffs.

We show that the value of the program defined in OPT coincides with the seller's payoff in the posted-prices assessment and hence, that the seller cannot do better than in the posted-prices equilibrium. Note that the solution to the auxiliary program corresponds to the solution to the *relaxed* program in mechanism design, where only a subset of the incentive constraints have been imposed. However, as argued by Laffont and Tirole (1993) and more recently by Breig (2020), it is not immediate that the solution to the relaxed problem satisfies *all* constraints and hence, can be implemented as an equilibrium outcome. Nevertheless, as we discuss in the conclusions, we expect that in settings with transferable utility, the study of the analogous problem to OPT provides a natural benchmark to understand the properties of the principal's optimal mechanism.

The significance of our results is two-fold. First, to the best of our knowledge, this is the first paper to characterize optimal mechanisms under limited commitment and persistent private information in an infinite-horizon setting. Such characterization has proved

elusive. Indeed, the set of tools available to tackle the difficulties with the revelation principle under limited commitment do not readily apply to infinite-horizon settings (see, for instance, the seminal work of Bester and Strausz (2001, 2007), and the discussion in the related literature). In Doval and Skreta (2020), we introduce a tool akin to the revelation principle for mechanism-selection games under limited commitment that applies to a broad class of games, including infinite-horizon ones. It is the application of this tool that allows us to argue that the mechanism that we characterize is the optimal one amongst *all* mechanisms the seller could have offered the buyer under limited commitment.

Second, while the optimality of posted-prices is evocative of Skreta (2006), there are two reasons why this result should not be taken for granted. First, our model is not an infinite-horizon version of the one studied in Skreta (2006), since we allow the seller to choose from a class of mechanisms that includes those in Skreta (2006). In Doval and Skreta (2020), we study a two-period version of the model in Skreta (2006), but we allow the seller to offer mechanisms like those in this paper. We show that when the seller is sufficiently patient it is not an equilibrium for the seller to post a price in each period. It follows that we cannot take limits using the equilibrium outcome in Skreta (2006) to analyze the equilibrium outcomes of the game we study, even after showing that the game ends in finite time. Second, Breig (2020) shows in a binary-value model with a perishable good that posted-prices may not be optimal: Indeed, the seller may benefit from using random delivery contracts.

Related Literature: The paper contributes mainly to three strands of literature. The first strand, similar to this paper, derives optimal mechanisms when the designer has limited commitment. Most papers in that literature examine either finite-horizon settings (see Laffont and Tirole, 1988; Skreta, 2006, 2015; Deb and Said, 2015; Fiocco and Strausz, 2015; Beccuti and Möller, 2018), or infinite-horizon settings, imposing restrictions on the class of contracts that can be offered (e.g., Gerardi and Maestri, 2020), or on the solution concept (e.g., Acharya and Ortner, 2017).[3] Underlying these restrictions is that the results in both Bester and Strausz (2001) and Skreta (2006) do not readily extend to infinite-horizon settings. For instance, the result in Bester and Strausz (2001) applies

[3] Beccuti and Möller (2018) lie somewhat in between these two strands because they take limits of their finite-horizon results to draw conclusions about the infinite-horizon setting. However, they do not show that this limit corresponds to the seller's revenue-maximizing equilibrium in the infinite-horizon game.

only if the principal is earning his highest payoff consistent with the agent's payoff (see Lemma 1 in Bester and Strausz, 2001). Thus, implicit in their multistage extension is a restriction to equilibria of the mechanism-selection game that possess a Markov structure, which, as shown by Ausubel and Deneckere (1989), may not be enough to characterize the principal's best equilibrium payoff. The approach in Skreta (2006) has the advantage that the set of incentive feasible outcomes has a well-understood structure. It is not clear, however, how to incorporate the principal's sequential rationality constraints in infinite-horizon settings in a tractable way.

The second strand is the literature that follows the observation in Coase (1972) that the durable-good monopolist faces a time-inconsistency problem, which in turn limits his monopoly power. The papers in the durable-good monopolist literature (Stokey, 1981; Bulow, 1982; Gul, Sonnenschein, and Wilson, 1986; Sobel, 1991; Ortner, 2017) study price dynamics and establish (under some conditions) Coase's conjecture.[4] Related to this literature are the problems of dynamic bargaining with one-sided incomplete information[5] (e.g., Sobel and Takahashi, 1983; Fudenberg, Levine, and Tirole, 1985; Ausubel and Deneckere, 1989). In all these papers, the seller's inability to commit reduces monopoly profits.

Finally, as will become clear from the analysis, the paper contributes to the literature on information design (Aumann, Maschler, and Stearns (1995) and Kamenica and Gentzkow (2011)), highlighting its potential to provide tractable characterizations of equilibrium outcomes in games. Contrary to most of these papers, however, the seller aims to persuade his future self, as opposed to another player, highlighting the role of information as a commitment device (see, e.g., Carrillo and Mariotti (2000), and recently, Habibi (2020)).

Organization The rest of the paper is organized as follows. Section 2 describes the model; Section 2.1 summarizes the results in Doval and Skreta (2020) used to simplify the analysis that follows. Section 3 presents the main result of the paper, Theorem 1. Section 4 introduces the auxiliary program OPT and studies its properties. Section 5 reviews the main steps of the proof of Theorem 1. Section 6 concludes. All proofs not in the main text are in the Appendix.

[4] Other papers, like Wolinsky (1991), McAfee and Wiseman (2008) and Board and Pycia (2014) study variations on Coase's original problem and their implications for Coase's conjecture.

[5] Recently, Peski (2019) studies alternating bargaining games where players can offer menus.

2 Model

Primitives: Two players, a seller and a buyer, interact over infinitely many periods. The seller owns one unit of a durable good to which he attaches value 0. The buyer has private information: before her interaction with the seller starts, she observes her valuation $v \in \{v_L, v_H\} \equiv V$, with $0 < v_L < v_H$. Let $\Delta v \equiv v_H - v_L$ denote the difference in values. Let μ_0 denote the probability that the buyer's valuation is v_H at the beginning of the game. In what follows, we denote by $\Delta(V)$ the set of distributions on V.

An allocation in period t is a pair $(q, x) \in \{0,1\} \times \mathbb{R}$, where q indicates whether the good is traded ($q = 1$) or not ($q = 0$), and x is a payment from the buyer to the seller. Let A denote the set of allocations. The game ends the first time the good is traded.[6]

Payoffs are as follows. If in period t, the allocation is (q, x), the flow payoffs are $u_B(q, x, v) = vq - x$ and $u_S(q, x) = x$ for the buyer and the seller, respectively. The seller and the buyer maximize the expected discounted sum of flow payoffs. They share a common discount factor $\delta \in (0,1)$.

Mechanisms: To introduce the timing of the game, we first define the action space of the seller. In each period, the seller offers the buyer a mechanism. Following Myerson (1982), Bester and Strausz (2007), and Doval and Skreta (2020), we define a mechanism as follows. A mechanism, $\mathbf{M} = (M^{\mathbf{M}}, S^{\mathbf{M}}, \varphi^{\mathbf{M}})$, consists of a set of *input* messages $M^{\mathbf{M}}$, a set of *output* messages $S^{\mathbf{M}}$, and a transition probability $\varphi^{\mathbf{M}}$ from $M^{\mathbf{M}}$ to $S^{\mathbf{M}} \times A$.[7] For instance, $M^{\mathbf{M}}$ can be the set of buyer values, V, and $S^{\mathbf{M}}$ be the set of seller beliefs about the buyer's value, $\Delta(V)$. In this case, the mechanism associates to each report a distribution over beliefs and allocations. We endow the seller with a collection $(M_i, S_i)_{i \in \mathcal{I}}$ of input and output messages in which each M_i contains at least two elements, and each S_i contains $\Delta(V)$.[8] Denote by $\mathcal{M}$ the set of all mechanisms with message sets $(M_i, S_i)_{i \in \mathcal{I}}$.[9]

[6] While the set of allocations is $\{0,1\} \times \mathbb{R}$, we allow the seller to offer randomizations on A and hence induce *fractional* assignments of the good.

[7] Throughout, we assume that $M^{\mathbf{M}}$ and $S^{\mathbf{M}}$ are Polish spaces, that is, they are separable, completely metrizable topological spaces. Note that the set of allocations A is also a Polish space, and therefore $S^{\mathbf{M}} \times A$ is a Polish space. For a Polish space X, let $\Delta(X)$ denote the set of Borel measures on X. We endow $\Delta(X)$ with the weak* topology. Thus, $\Delta(X)$ is also a Polish space (Aliprantis and Border, 2013). For any two measurable spaces X and Y, a mapping $\zeta : X \mapsto \Delta(Y)$ is a *transition probability from X to Y* if for any measurable $C \subseteq Y$, $\zeta(C|x)$ is a measurable real valued function of $x \in X$.

[8] We show in Doval and Skreta (2020) that it is without loss of generality to take $M^{\mathbf{M}}$ to be finite when the set of types is finite.

[9] We restrict the seller to choose mechanisms with input and output messages in $(M_i, S_i)_{i \in \mathcal{I}}$ to have a well-defined action space for the seller. This allows us to have a well-defined set of deviations, avoiding set-theoretic issues related to self-referential sets. The analysis in Doval and Skreta (2020) shows that the

Timing: In each period t, as long as the good has not been sold, the game proceeds as follows. First, the seller offers the buyer a mechanism, $\mathbf{M}$. Observing the mechanism, the buyer decides whether to participate or not. If she does not participate in the mechanism, the good is not sold and no payments are made. If she instead chooses to participate, she sends a message $m \in M^{\mathbf{M}}$, which is unobserved by the seller. An output message and an allocation, (s, q, x), are drawn from $\varphi^{\mathbf{M}}(\cdot|m)$. The output message and the allocation are observed by both the seller and the buyer. If the good is not traded, the game proceeds to period $t+1$.

The above defines an infinite-horizon dynamic game, which we denote by $G^{\infty}_{\mathcal{M}}(\mu_0)$ to remind the reader that the seller's beliefs at the beginning of the game are given by μ_0. Public histories in this game are

$$h^t = (\mathbf{M}_0, \pi_0, s_0, (q_0, x_0), \ldots, \mathbf{M}_{t-1}, \pi_{t-1}, s_{t-1}, (q_{t-1}, x_{t-1})),$$

where $\pi_k \in \{0, 1\}$ denotes the buyer's participation decision, with the restriction that if $\pi_k = 0$, then $s_k = \emptyset$, i.e., no output message is generated, and $(q_k, x_k) = (0, 0)$.[10] Let H^t denote the set of all period t public histories; they capture what the seller knows through period t.

A *buyer* history consists of the public history of the game together with the buyer's reports into the mechanism and her private information. Formally, a buyer history is an element

$$h^t_B = (\mathbf{M}_0, \pi_0, m_0, s_0, (q_0, x_0), \ldots, \mathbf{M}_{t-1}, \pi_{t-1}, m_{t-1}, s_{t-1}, (q_{t-1}, x_{t-1})).$$

Let H^t_B denote the set of buyer histories through period t. The buyer also knows her valuation and hence, when her valuation is v, a history through period t is an element of $\{v\} \times H^t_B$.

Strategies and beliefs: Since the seller's action space, $\mathcal{M}$, is an uncountable set of functions, we model the seller's behavioral strategy $(\Gamma_t)_{t=0}^{\infty}$ following Aumann (1964). That is, endow $[0, 1]$ with the Borel σ-algebra and the Lebesgue measure. Then, Γ_t is defined as a jointly measurable function from $H^t \times [0, 1]$ to $\mathcal{M}$.[11] We denote the collection $(\Gamma_t)_{t=0}^{\infty}$ by Γ. The buyer's participation strategy is given by $\pi_t : V \times H^t_B \times \mathcal{M} \mapsto [0, 1]$. Conditional

choice of the collection plays no further role in the analysis.

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

[11] As we explain in Appendix D, the set $\mathcal{M}$ is a Polish space; we endow it with its Borel σ-algebra.

on participating in the mechanism, $\mathbf{M}_t$, her reporting strategy is a transition probability from $V \times H_B^t \times \mathcal{M} \times \{1\}$ to $\cup_{i \in \mathcal{I}} M_i$ such that $r_t(h_B^t, \mathbf{M}_t) \in \Delta(M^{\mathbf{M}_t})$. We denote the tuple $(\pi_t(v,\cdot), r_t(v,\cdot))$ by (π_{tv}, r_{tv}) and the collection $(\pi_{tv}, r_{tv})_{t=0}^{\infty}$ by (π_v, r_v).
A belief for the seller at the beginning of time t, history h^t, is a distribution $\mu_t(h^t) \in \Delta(V \times H_B^t(h^t))$, where $H_B^t(h^t)$ is the set of buyer histories consistent with the public history, h^t. The belief system, $(\mu_t)_{t=0}^{\infty}$, is denoted by μ.

Solution concept: We are interested in studying the Perfect Bayesian equilibrium (henceforth, PBE) payoffs of this game, where PBE is defined *informally* as follows. An assessment, $\langle \Gamma, (\pi_v, r_v)_{v \in V}, \mu \rangle$, is a PBE if the following hold:

1. $\langle \Gamma, (\pi_v, r_v)_{v \in V}, \mu \rangle$ satisfies sequential rationality, and
2. μ satisfies Bayes' rule where possible.

Appendix D contains the formal statement. For now, we note that if the principal's strategy space were finite and the mechanisms used by the principal had finite support, then this coincides with the definition in Fudenberg and Tirole (1991b).[12]

The prior μ_0 together with the strategy profile $(\Gamma, (\pi_v, r_v)_{v \in V})$ induce a distribution over the *terminal nodes* $V \times H_B^{\infty}$. We are interested instead on the distribution it induces over the payoff-relevant outcomes, $V \times A^{\infty}$. We say that a distribution $\eta \in \Delta(V \times A^{\infty})$ is a PBE outcome if there is a PBE assessment $\langle \Gamma, (\pi_v, r_v)_{v \in V}, \mu \rangle$ which induces η. We denote by $\mathcal{O}_{\mathcal{M}}^*(\mu_0)$ the set of PBE outcomes and by $\mathcal{E}_{\mathcal{M}}^*(\mu_0) \subseteq \mathbb{R}^3$ the set of PBE payoffs of $G_{\mathcal{M}}^{\infty}(\mu_0)$. We denote a generic element of $\mathcal{E}_{\mathcal{M}}^*(\mu_0)$ by $u \equiv (u_L, u_H, u_S)$, where u_S is the seller's payoff and u_L, u_H denote the buyer's payoff when her value is v_L, v_H, respectively.

Theorem 1 characterizes $\mathcal{O}_{\mathcal{M}}^*(\mu_0)$ and $\mathcal{E}_{\mathcal{M}}^*(\mu_0)$. In particular, we show that the *essentially* unique equilibrium outcome can be achieved via a sequence of posted prices so that the seller of a durable good can do no better and no worse than by using posted prices.

2.1 Revelation Principle

There are at least two reasons that the game $G_{\mathcal{M}}^{\infty}(\mu_0)$ is not simple to analyze. First, the seller's action space is large and, a priori, it is not clear which mechanisms could be ruled out from consideration. Second, fixing a seller's strategy, and hence a sequence of

[12] The only difference between Bayes' rule where possible and consistency in sequential equilibrium is the following. Under PBE, the principal can assign zero probability to a type and then, after the agent deviates, can assign positive probability to that same type.

mechanisms faced by the buyer, we still need to understand the buyer's best response in the game induced by the sequence of mechanisms.

Let $G^{\infty}(\mu_0)$ denote the same game in the previous section, except that in each period (i) both the buyer and the seller observe a draw from a public randomization device $\omega \sim U[0,1]$, and (ii) the seller's action space is the set of *canonical mechanisms* denoted by $\mathcal{M}_C$ and defined as follows. $\mathcal{M}_C$ is the set of all mechanisms where the set of input and output messages are the buyer's values and the seller's beliefs about the buyer's value, respectively. That is, $(M,S) = (V, \Delta(V))$. Let $\mathcal{O}^*(\mu_0)$ denote the set of PBE outcomes and $\mathcal{E}^*(\mu_0)$ denote the set of PBE payoffs of $G^{\infty}(\mu_0)$.

Lemma 1 summarizes the key implications of Doval and Skreta (2020) for our analysis, which we explain below:

Lemma 1 (Doval and Skreta, 2020)**.** *$G^{\infty}_{\mathcal{M}}(\mu_0)$ and $G^{\infty}(\mu_0)$ have the same set of equilibrium outcomes, i.e., $\mathcal{O}^*_{\mathcal{M}}(\mu_0) = \mathcal{O}^*(\mu_0)$.*

Moreover, let $\eta \in \mathcal{O}^*(\mu_0)$. Then, there exists a PBE assessment $\langle \Gamma, (\pi_v, r_v)_{v \in V}, \mu \rangle$ of $G^{\infty}(\mu_0)$ that induces η and satisfies the following properties:

1. *For all histories h^t, the buyer participates in the mechanism offered by the seller at that history and truthfully reports her type, with probability* 1*,*

2. *For all histories h^t, if the mechanism offered by the seller at h^t outputs posterior μ', the seller's updated equilibrium beliefs about the buyer coincide with μ',*

3. *For all histories h^t, the mechanism offered by the seller at h^t, $\varphi^{\mathbf{M}_t} : V \mapsto \Delta(\Delta(V) \times A)$ can be decomposed into a* communication device *$\beta^{\mathbf{M}_t} : V \mapsto \Delta\Delta(V)$ and an* allocation rule *$\alpha^{\mathbf{M}_t} : \Delta(V) \mapsto \Delta(A)$,*[13]

4. *The strategy of the buyer depends only on her private valuation and the public history.*

Lemma 1 has several implications. Part 1 of Lemma 1 implies the mechanisms cho-

[13] That is, for all measurable subsets $U' \subset \Delta(V)$ and $A' \subset A$, we have that for all $v \in V$,

$$\varphi^{\mathbf{M}_t}(U' \times A' | v) = \int_{U'} \alpha^{\mathbf{M}_t}(A' | \mu) \beta^{\mathbf{M}_t}(d\mu | v).$$

sen by the seller in equilibrium must satisfy a *participation* constraint and an *incentive compatibility constraint* for each buyer value and each public history. As in the case of commitment to long-term mechanisms, part 1 simplifies the analysis of the buyer's behavior, by reducing it to a series of constraints (see Equations $\mathrm{PC}_{h^t,v}$ and $\mathrm{IC}_{h^t,v,v'}$).

Part 2 implies that the mechanism's output message encodes all of the information that the seller has in equilibrium about the buyer's value. In particular, conditional on observing the output message, the allocation carries no more information about the buyer's value. As a consequence, conditional on the output message, the allocation can be drawn independently of the buyer's report. This, in turn, delivers the decomposition of $\varphi^{\mathbf{M}_t}$ described in part 3.

Part 3 implies that the choice of mechanism at history h^t can be equivalently thought of as the choice of an experiment (the communication device, $\beta^{\mathbf{M}_t}$) and an allocation rule (the $\alpha^{\mathbf{M}_t}$). The decomposition of $\varphi^{\mathbf{M}_t}$ allows us to separately optimize on the allocation *given* a particular communication device, and then optimize on the communication device. As in the literature on information design, it is convenient to work with the distribution over posteriors induced by the experiment $\beta^{\mathbf{M}_t}$, which we denote by $\tau^{\mathbf{M}_t}$ and is defined as follows. For all Borel subsets $U' \subseteq [0,1]$,

$$\int_{U'} \tau^{\mathbf{M}_t}(d\mu_{t+1}) = \sum_{v \in V} \mu_t(h^t)(v)\beta^{\mathbf{M}_t}(U'|v), \tag{$\mathrm{BC}_{\mu_t(h^t)}$}$$

where $\mu_t(h^t)$ is the seller's belief about the buyer's value at h^t. Furthermore, as in the literature on mechanism design with quasilinear utilities, we can write $\alpha^{\mathbf{M}_t}(\cdot|\mu)$ as an expected payment, $x^{\mathbf{M}_t}(\mu)$, and a probability of trade, $q^{\mathbf{M}_t}(\mu)$.

Part 4 implies the set of PBE payoffs of $G^\infty(\mu_0)$ coincides with the set of *Public PBE* payoffs of $G^\infty(\mu_0)$ (Athey and Bagwell, 2008), allowing us to invoke the results in Abreu et al. (1990) and Athey and Bagwell (2008) to argue the assessment we construct in the next section is indeed a PBE assessment.

The rest of the paper studies the equilibrium outcomes and payoffs of $G^\infty(\mu_0)$ and when we refer to a PBE assessment, we mean one that satisfies the conditions of Lemma 1.

Remark 1. *In what follows, we abuse notation in the following two ways. First, because valuations are binary, we can think of an element in $\Delta(V)$ (a distribution over v_L and*

v_H) as an element of the interval $[0,1]$ (the probability assigned to v_H). We use the latter formulation in what follows. That is, whereas the mechanism outputs a distribution over v_L and v_H, we index this distribution by the probability of v_H. Second, even though $\beta(\cdot|v)$ is a measure over $\Delta(V)$ (in this case a c.d.f.), we sometimes write $\beta(\mu|v)$ when β has an atom at μ.

3 Main result

Section 3 contains the main result of the paper: Theorem 1 characterizes the equilibrium outcomes and payoffs of $G^{\infty}(\mu_0)$. To state Theorem 1, we proceed as follows. First, we describe a PBE assessment $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v\in V}, \mu^* \rangle$, which is singled out by the proof of Theorem 1. Second, we explain why the equilibrium outcome induced by $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v\in V}, \mu^* \rangle$ can be implemented via a sequence of posted prices. Finally, we state Theorem 1.

An assessment in posted prices: The proof of Theorem 1 singles out a particular PBE assessment, $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v\in V}, \mu^* \rangle$, that implements the seller's unique equilibrium payoff, which we now define informally (the formal statement is in Section C.2).

The assessment $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v\in V}, \mu^* \rangle$ is essentially the same along the path of play to those constructed by Fudenberg et al. (1985) and Hart and Tirole (1988). The assessment is defined by an increasing sequence of threshold beliefs $\overline{\mu}_0 = 0 < \overline{\mu}_1 < \cdots < \overline{\mu}_n < \ldots$ such that if the seller's belief is in $[\overline{\mu}_n, \overline{\mu}_{n+1})$, then it takes n periods for the good to be sold to v_L, at which point the game ends. The number of periods before v_L buys the good determines both the rents for v_H and the seller's revenue.

To understand how the sequence of thresholds is determined, note that if a seller with belief μ_0 sells the good at a price of v_L, his revenue can be written as follows:

$$v_L = \mu_0(v_H - \Delta v) + (1-\mu_0)v_L = \mu_0 v_H + (1-\mu_0)\underbrace{\left(v_L - \frac{\mu_0}{1-\mu_0}\Delta v\right)}_{\equiv \hat{v}_L(\mu_0)}. \tag{1}$$

The first equality represents revenue as the surplus extracted from each type. The second equality represents revenue as the *virtual surplus*, where the value of allocating the good to v_L is adjusted to capture that when v_L is served, so is v_H, which leaves rents to v_H. The threshold $\overline{\mu}_1 \equiv v_L/v_H$ is the belief at which the virtual value, $\hat{v}_L(\mu_0)$, equals 0. At that belief, the seller is indifferent between serving the buyer for both of her values and

excluding the low-valuation buyer.

For a seller with belief $\mu_0 < \overline{\mu}_1$, v_L is the maximum revenue that the seller can make. Instead, a seller with belief $\mu_0 > \overline{\mu}_1$ prefers to take one period to trade with v_L to trading immediately with v_L. To see this, note that if the good is sold with probability 1 in the next period, the price at that point is v_L. Thus, the maximum price that the seller can sell the good for in the first period is $v_H - \delta(v_H - v_L) = v_L + (1-\delta)\Delta v$. It is easy to check that when $\mu_0 > \overline{\mu}_1$, the seller prefers to take one period to sell the good to v_L. When the seller's prior belief is $\overline{\mu}_1$ the seller is indifferent between selling the good to v_L in one period and selling the good to v_L immediately. Indeed, the subscript 1 in $\overline{\mu}_1$ reflects that the seller takes at most one period to trade with v_L. Recursively, one can define $\overline{\mu}_n$ as the belief at which the seller is indifferent between trading with v_L in n period or $n-1$ periods.

In the assessment, play proceeds as follows. If the seller's prior is such that $\mu_0 \in [\overline{\mu}_n, \overline{\mu}_{n+1})$, then the seller chooses a mechanism that induces two beliefs, 1 and $\overline{\mu}_{n-1}$. At belief 1, the good is sold and the transfer is $v_L + (1-\delta^n)\Delta v$, whereas at belief $\overline{\mu}_{n-1}$, the good is not sold and the transfer is 0 (see Figure 1a). Subsequently, a seller with belief $\overline{\mu}_m$ for $m \leq n-1$, chooses a mechanism that induces two beliefs, 1 (with allocation $(1, v_L + (1-\delta^m)\Delta v)$) and $\overline{\mu}_{m-1}$ (with allocation (0,0)) (see Figure 1b). Thus, starting from μ_0, the game ends in n periods. Figure 1c illustrates how beliefs fall conditional on no trade.

More precisely, the PBE assessment, $\langle \Gamma^*, (\pi^*_v, r^*_v)_{v \in V}, \mu^* \rangle$, is as follows:

1. Along the equilibrium path:

   (a) If at history h^t, the seller's beliefs, $\mu^*(h^t)$, are in $[\overline{\mu}_0, \overline{\mu}_1)$, he chooses a mechanism such that $(q^{\mathbf{M}^*_t}(\mu^*(h^t)), x^{\mathbf{M}^*_t}(\mu^*(h^t))) = (1, v_L)$ and the communication device satisfies that $\beta^{\mathbf{M}^*_t}(\mu^*(h^t)|v) = 1$ for $v \in \{v_L, v_H\}$.

   (b) If at history h^t, the seller's beliefs, $\mu^*(h^t)$, are in $[\overline{\mu}_n, \overline{\mu}_{n+1})$ for $n \geq 1$, the seller's mechanism satisfies the following. First, it induces two posteriors, $\overline{\mu}_{n-1}$ and 1. Second, the allocation rule satisfies that $(q^{\mathbf{M}^*_t}(1), x^{\mathbf{M}^*_t}(1)) = (1, v_L + (1-\delta^n)\Delta v)$, whereas $(q^{\mathbf{M}^*_t}(\overline{\mu}_{n-1}), x^{\mathbf{M}^*_t}(\overline{\mu}_{n-1})) = (0,0)$. Finally, the communication device maps v_L to $\overline{\mu}_{n-1}$, whereas it maps v_H to both $\overline{\mu}_{n-1}$ and 1 with positive probability. The probabilities $\beta^{\mathbf{M}^*_t}(\overline{\mu}_{n-1}|v_H), \beta^{\mathbf{M}^*_t}(1|v_H)$ are chosen so that when the seller ob-

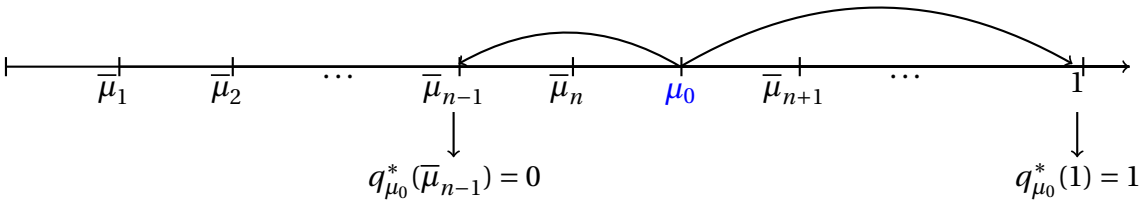

Figure (a) Mechanism for $\mu_0 \in [\overline{\mu}_n, \overline{\mu}_{n+1})$ under $\langle \Gamma^*, (\pi^*_v, r^*_v)_{v \in V}, \mu^* \rangle$

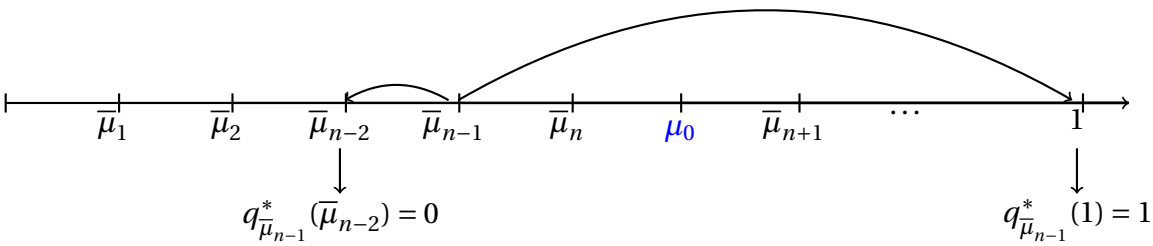

Figure (b) Mechanism for $\mu_0 = \overline{\mu}_{n-1}$ under $\langle \Gamma^*, (\pi^*_v, r^*_v)_{v \in V}, \mu^* \rangle$

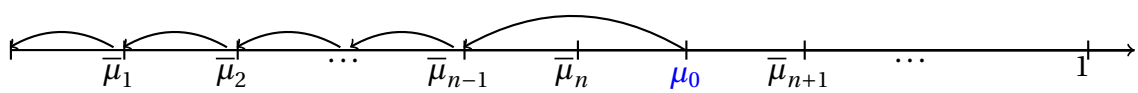

Figure (c) Belief dynamics starting from μ_0.

Figure 1: Mechanism and belief dynamics under $\langle \Gamma^*, (\pi^*_v, r^*_v)_{v \in V}, \mu^* \rangle$ as a function of μ_0.

serves $\overline{\mu}_{n-1}$, his updated belief coincides with $\overline{\mu}_{n-1}$, that is

$$\overline{\mu}_{n-1} = \frac{\mu^*(h^t)\beta^{\mathbf{M}^*_t}(\overline{\mu}_{n-1}|v_H)}{\mu^*(h^t)\beta^{\mathbf{M}^*_t}(\overline{\mu}_{n-1}|v_H) + (1 - \mu^*(h^t))}$$

2. Off the equilibrium path, the seller's strategy coincides with the above, except that when $\mu^*(h^t) = \overline{\mu}_n$ for some $n \geq 1$, the seller may randomize between the mechanism he offers on the path of play when his belief is $\overline{\mu}_n$ and the one he offers on the path of play when his belief is $\overline{\mu}_{n-1}$.[14]

3. At each history h^t, the buyer's best response to the seller's equilibrium offer at h^t is to participate in the mechanism and truthfully report her valuation.

An implementation in posted prices: We now argue that there is an implementation in posted prices of the equilibrium outcome induced by $\langle \Gamma^*, (\pi^*_v, r^*_v)_{v \in V}, \mu^* \rangle$. Clearly, when

[14]The need for mixing arises for technical reasons: it ensures that the buyer's continuation payoffs when her valuation is v_H are upper-semicontinuous and, thus, guarantees that a best response exists after any deviation by the seller (see Section C.1). Indeed, we appeal to the results in Simon and Zame (1990) to simultaneously determine the buyer's best response and the seller's mixing.

the seller's beliefs are below $\overline{\mu}_1$, the seller's mechanism corresponds to selling the good at a price of v_L. Consider then the case in which the seller's beliefs are in $[\overline{\mu}_1, \overline{\mu}_2)$. Note that when the buyer's valuation is v_H and the realized allocation is trade, then her payoff is $v_H - v_L - (1-\delta)\Delta v = \delta \Delta v$. On the other hand, when the buyer's valuation is v_H and the realized allocation is no trade, then the seller's beliefs next period are $\overline{\mu}_0 = 0$, so that the buyer's continuation payoff is $\delta \Delta v$. That is, the buyer with valuation v_H is indifferent between obtaining the good at price $v_L + (1-\delta)\Delta v$ and not obtaining the good, and paying a price of v_L in the next period. Since the buyer with valuation is v_H is indifferent between these two options, she is willing to mix between buying at price $v_L + (1-\delta)\Delta v$ and not obtaining the good. She does so in a way that the seller's belief is $\overline{\mu}_0$ when the allocation is $(0,0)$. Since $\overline{\mu}_0 = 0$, it implies that when the seller's belief is in $[\overline{\mu}_1, \overline{\mu}_2)$, the buyer buys with probability 1 at a price of $v_L + (1-\delta)\Delta v$. Working recursively through the equations, one can show that when the seller's prior is in $[\overline{\mu}_n, \overline{\mu}_{n+1})$, the mechanism is equivalent to posting a price of $v_L + (1-\delta^n)\Delta v$. In this case, the low-valuation buyer chooses the $(0,0)$ allocation, whereas the high-valuation buyer mixes so that the seller's belief is $\overline{\mu}_{n-1}$ when the allocation is $(0,0)$.

Remark 2 (Direct vs. indirect implementation)**.** *It is interesting to contrast the implementation under the PBE assessment $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$ described above with the implementation via posted prices. In the former, the buyer is truthful and the seller* rations *the high-valuation buyer, which slows down the rate at which the seller's beliefs fall conditional on the good not being sold. Instead, in the implementation via posted prices, the high-valuation buyer misreports her type with positive probability, which, like rationing, prevents the seller from becoming pessimistic too quickly about the buyer's valuation. Since both implementations are payoff equivalent, the seller cannot do better with rationing than with posted prices.*

Remark 3 (Rationing)**.** *The mechanism used by the seller in $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$ is what Denicolo and Garella (1999) dub rationing. In a two-period model, Denicolo and Garella (1999) show that if the seller observes only whether trade happens, the seller may prefer to ration high valuation buyers instead of posting a price in the first period, in order to induce a strong demand in the second period. Whereas rationing does not dominate posted prices when values are binary, we show in Doval and Skreta (2020) that both posted prices and rationing as in Denicolo and Garella (1999) may be dominated when values are*

drawn from a continuum in a two-period model.

Theorem 1 states that there is a unique equilibrium payoff for the seller and it coincides with the payoff that the seller obtains under $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v\in V}, \mu^* \rangle$. Furthermore, the low-valuation buyer's payoff is also unique and equal to 0 (see Lemma 3). Finally, except at the threshold beliefs $\{\overline{\mu}_n\}_{n\geq 1}$, the high-valuation buyer's payoff is also unique. The multiplicity of v_H's payoffs arises because when the seller's prior belief is $\overline{\mu}_n$ for $n \geq 1$, the seller is indifferent between trading with v_L in n periods and in $n-1$ periods, whereas v_H's rents are higher when v_L receives the good in $n-1$ periods. All equilibrium payoff vectors for v_H can be obtained by randomizing between trading with v_L in n and $n-1$ periods.[15]

To state Theorem 1, let $u_H^*(\mu_0)$ and $u_S^*(\mu_0)$ denote the high-valuation buyer and seller's payoff under the assessment $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v\in V}, \mu^* \rangle$ when the seller's prior belief is μ_0.

Theorem 1. *$\langle \Gamma^*, (\pi_v^*, r_v^*)_{v\in V}, \mu^* \rangle$ is a PBE assessment. Furthermore, the set of equilibrium payoffs of $G^\infty(\mu_0)$ is the following. For $\mu_0 \in [0, \overline{\mu}_1)$, $\mathcal{E}^*(\mu_0) = (0, u_H^*(\mu_0), u_S^*(\mu_0))$. Furthermore, for all $n \geq 1$, and all $\mu_0 \in [\overline{\mu}_n, \overline{\mu}_{n+1})$*

$$\mathcal{E}^*(\mu_0) = \begin{cases} (0, u_H^*(\mu_0), u_S^*(\mu_0)) & \text{if } \mu_0 > \overline{\mu}_n \\ \{0\} \times [u_H^*(\mu_0), u_H^*(\mu_0)/\delta] \times \{u_S^*(\mu_0)\} & \text{otherwise} \end{cases} . \tag{2}$$

Theorem 1 together with the above discussion implies:

Corollary 1. *All equilibrium outcomes of $G^\infty(\mu_0)$ can be implemented via a sequence of posted prices.*[16]

Section 5 reviews the main steps of the proof of Theorem 1. Section 5.1 shows that $u_S^*(\mu_0)$

[15] When the seller's *prior* belief coincides with $\overline{\mu}_n$ for $n \geq 1$, there is another PBE assessment which is payoff equivalent to $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v\in V}, \mu^* \rangle$ for the seller, but delivers a payoff of $u_H^*(\mu_0)/\delta$ to the high-valuation buyer. For $n \geq 2$, this PBE assessment is identical to $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v\in V}, \mu^* \rangle$ except that in the first period the seller offers a mechanism that induces two posteriors, $\overline{\mu}_{n-2}$ and 1. When the posterior is 1, the allocation is $(1, v_L + (1-\delta^{n-1})\Delta v)$, whereas when the posterior is $\overline{\mu}_{n-2}$, the allocation is $(0,0)$. Instead, when $n = 1$, the seller sells the good with probability 1 at a price of v_L. Verifying that this is also a PBE assessment follows immediately from the arguments in Appendix C.

[16] In the gap case, that is, when $v_L = 0$, multiplicity of the equilibrium outcomes and payoffs obtains as in Ausubel and Deneckere (1989) and yet, equilibrium outcomes can also be sustained using posted prices. For instance, the seller can obtain the monopoly profit $\mu_0 v_H$ by setting a price of v_H in period 1 and thereafter, not selling the good.

is a lower bound and an upper bound on the seller's equilibrium payoff. Section 5.2 describes the arguments to show that $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v\in V}, \mu^* \rangle$ is a PBE assessment. To show that $u_S^*(\mu_0)$ is an upper bound on the seller's equilibrium payoff and that $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v\in V}, \mu^* \rangle$ is a PBE assessment, we rely on an auxiliary program, which we denote by OPT, and define in Section 4, to which we turn next.

4 Auxiliary program: maximization of the virtual surplus

We now introduce a program, denoted by OPT, that is our main tool of analysis. We use it in Section 5.1 to show both that the seller cannot do better than in the posted-prices equilibrium and that $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v\in V}, \mu^* \rangle$ is a PBE assessment. In this program, the seller maximizes the dynamic analogue of the virtual surplus, which we denote by VS. Recall that in standard mechanism design the virtual surplus of a mechanism is an upper bound on the revenue from the mechanism that only depends on the probability of trade, since the transfers are determined by the buyer's participation and truthtelling constraints (Section 4.1 provides the details of this derivation in our setting). As we explain below, VS depends on the distribution over posteriors induced by the mechanism's communication device (recall Equation $\mathrm{BC}_{\mu_t(h^t)}$ on page 10). Moreover, because of the dynamic nature of our problem, it also depends on the continuation payoffs.

Formally,

$$\max_{(\tau_0,q_0),(u_H,u_S)} VS((\tau_0,q_0),(u_H,u_S),\mu_0), \tag{OPT}$$
$$\text{s.t.}\begin{cases} \tau_0 \in \Delta(\Delta(V)) \text{ is Bayes' plausible for } \mu_0 \\ q_0 : \Delta(V) \mapsto [0,1] \\ (\forall \mu_1 \in \Delta(V))(0, u_H(\mu_1), u_S(\mu_1)) \in \mathcal{E}^*(\mu_1) \end{cases},$$

where

$$VS((\tau_0,q_0),(u_H,u_S),\mu_0) \equiv \int_{\Delta(V)} \left[\begin{array}{c} q_0(\mu_1)(\mu_1 v_H + (1-\mu_1)\hat{v}_L(\mu_0)) + \\ (1-q_0(\mu_1))\delta\left(u_S(\mu_1) + (1-\mu_1)\left(\frac{\mu_1}{1-\mu_1} - \frac{\mu_0}{1-\mu_0}\right)u_H(\mu_1)\right) \end{array} \right] \tau_0(d\mu_1), \tag{3}$$

defines the *virtual surplus*. Note that the virtual surplus only depends on a mechanism's probability of trade, q_0, and distribution over posteriors, τ_0. In what follows, we identify a mechanism $\mathbf{M}$ with its induced $(\tau^{\mathbf{M}}, q^{\mathbf{M}})$ and refer to the latter also as the mechanism.

In OPT, we allow the seller with prior μ_0 not only to choose his most preferred mechanism (τ_0, q_0), (accruing its virtual surplus), but also his preferred continuation payoffs, subject to the constraint that these continuation payoffs are actually equilibrium payoffs. That is, $(0, u_H(\mu_1), u_S(\mu_1)) \in \mathcal{E}^*(\mu_1)$. The notation anticipates the result in Lemma 3 below that the low-valuation buyer's payoff is 0 in *any* equilibrium (see Section 4.1).

The objective in OPT, which is the virtual surplus, consists of two terms. The first term

$$q_0(\mu_1)\left(\mu_1 v_H + (1-\mu_1)\hat{v}_L(\mu_0)\right),$$

represents how much surplus the seller can extract subject to the rents he must leave to the high-valuation buyer (recall Equation 1). Indeed, whenever $\mu_1 \neq 1$ and $q_0(\mu_1) > 0$, the seller sells the good with positive probability to the low-valuation buyer. In that case, v_H gets rents equal to $q_0(\mu_1)\Delta v$, which the seller pays with probability μ_0. This, in turn, explains why the virtual value of v_L, $\hat{v}_L(\cdot)$ is evaluated at the seller's prior belief, μ_0, instead of the posterior belief, μ_1. The second term

$$(1-q_0(\mu_1))\delta\left(u_S(\mu_1) + (1-\mu_1)\left(\frac{\mu_1}{1-\mu_1} - \frac{\mu_0}{1-\mu_0}\right)u_H(\mu_1)\right),$$

accounts for the rents the v_H receives in terms of continuation payoffs, which limit how much surplus the seller can extract in period 0 whenever he induces no trade with positive probability. The term pre-multiplying v_H's continuation payoffs, $u_H(\mu_1)$, reflects that conditional on no trade, the seller's beliefs about the buyer's value may change so that the optimal mechanism from period 1 onward may not coincide with what is optimal from period 0 onward. The *wedge* between the likelihood ratio of μ_1 and μ_0 reflects this disagreement and it is the source of the time inconsistency of the commitment solution.

While the construction of the upper bound on the seller's payoff is reminiscent of standard observations in mechanism design, it is important to note one challenge relative to this literature, which explains why we allow the seller to select continuation payoffs in OPT. The standard argument for why a revenue-maximizing seller can secure the virtual surplus of a mechanism is that if he were not, then there would be a way to increase transfers so that the buyer's participation and truthtelling constraints would still be satisfied, whereas the seller's revenue would increase. This argument does not translate

immediately to a game. Because continuation play may depend on the mechanism $\mathbf{M}$ chosen by the seller, it may not be possible for him to secure the mechanism's virtual surplus. To do so, the seller would need to offer a mechanism $\mathbf{M}'$ that coincides with $\mathbf{M}$, except that transfers are higher. However, even if $\mathbf{M}'$ induces the same distribution over allocations and seller's beliefs as $\mathbf{M}$, this could trigger lower continuation payoffs for the seller, preventing him from obtaining the virtual surplus of $\mathbf{M}$. The program defined in OPT allows us to circumvent this difficulty: Because the seller can choose both the mechanism *and* the continuation payoffs in OPT, the seller does not need to consider how his choice of mechanism may adversely affect his continuation payoffs.

Section 4.1 contains the details of the construction of the virtual surplus defined in Equation 3 and verifies that it is an upper bound on the seller's equilibrium payoff (Lemma 2). The reader interested in the properties of the solution to OPT and how it is used in the proof of Theorem 1 can proceed to Section 4.2 with little loss of continuity.

4.1 Derivation of the virtual surplus

We now derive the virtual surplus defined in Equation 3 by relying on Lemma 1 and show that it is an upper bound on the seller's equilibrium payoff. To do so, consider a PBE assessment, $\langle \Gamma, (\pi_v, r_v)_{v \in V}, \mu \rangle$. Fix a history h^t and let $\mathbf{M}_t$ denote the mechanism offered by the seller at h^t under the assessment. Let $(\tau^{\mathbf{M}_t}, q^{\mathbf{M}_t})$ denote the distribution over posteriors and the probability of trade associated with $\mathbf{M}_t$. Furthermore, the PBE assessment specifies continuation payoffs $(u_H^{\mathbf{M}_t}, u_S^{\mathbf{M}_t})$ when the seller offers mechanism $\mathbf{M}_t$ at history h^t. In what follows, we show that the seller's equilibrium payoff at h^t satisfies the following:

Lemma 2. *The seller's payoff at h^t, $U_S(h^t)$ satisfies*

$$U_S(h^t) \leq VS((\tau^{\mathbf{M}_t}, q^{\mathbf{M}_t}), (u_H^{\mathbf{M}_t}, u_S^{\mathbf{M}_t}), \mu_t(h^t)).$$

To see why Lemma 2 holds, note that Lemma 1 implies that the seller's equilibrium payoff at h^t, $U_S(h^t)$, can be written as:

$$U_S(h^t) = \int_{\Delta(V)} \left(x^{\mathbf{M}_t}(\mu_{t+1}) + (1 - q^{\mathbf{M}_t}(\mu_{t+1})) u_S^{\mathbf{M}_t}(\mu_{t+1}) \right) \tau^{\mathbf{M}_t}(d\mu_{t+1}), \tag{4}$$

where $u_S^{\mathbf{M}_t}(\mu_{t+1})$ is short-hand notation for the continuation payoff of the seller when

at history h^t, he offers $\mathbf{M}_t$ and the output message is μ_{t+1}.[17] Equation 4 uses Lemma 1 as follows. First, the seller's payoff from offering $\mathbf{M}_t$ is written under the assumption that the buyer participates in the mechanism and truthfully reports her value. Second, it uses Lemma 1 to write the mechanism in terms of the distribution over posteriors $\tau^{\mathbf{M}_t}$ and the allocation $(q^{\mathbf{M}_t}, x^{\mathbf{M}_t})$.

In particular, the mechanism $\mathbf{M}_t$ together with the continuation payoffs $(u_L^{\mathbf{M}_t}, u_H^{\mathbf{M}_t})$ satisfy the following constraints. First, the buyer prefers to participate in the mechanism for both her values, that is for $v \in \{v_L, v_H\}$ the following holds:

$$\int_{\Delta(V)} \left(v q^{\mathbf{M}_t}(\mu_{t+1}) - x^{\mathbf{M}_t}(\mu_{t+1}) + (1 - q^{\mathbf{M}_t}(\mu_{t+1}))\delta u_v^{\mathbf{M}_t}(\mu_{t+1}) \right) \beta^{\mathbf{M}_t}(d\mu_{t+1}|v) \geq u_v^{\mathbf{M}_t}(\emptyset), \qquad (\mathrm{PC}_{h^t,v})$$

where $u_v^{\mathbf{M}_t}(\emptyset)$ is short-hand notation for the continuation payoff of the buyer when at history h^t, the seller offers $\mathbf{M}_t$ and the buyer rejects. Also, the buyer prefers to truthfully report her value to the mechanism, that is, for $v \in \{v_L, v_H\}$ and $v' \neq v$, the following holds:

$$\int_{\Delta(V)} \left(v q^{\mathbf{M}_t}(\mu_{t+1}) - x^{\mathbf{M}_t}(\mu_{t+1}) + (1 - q^{\mathbf{M}_t}(\mu_{t+1}))\delta u_v^{\mathbf{M}_t}(\mu_{t+1}) \right) (\beta^{\mathbf{M}_t}(d\mu_{t+1}|v) - \beta^{\mathbf{M}_t}(d\mu_{t+1}|v')) \geq 0. \qquad (\mathrm{IC}_{h^t,v,v'})$$

There is one more way in which the above expressions implicitly use Lemma 1. By Lemma 1, the assessment $\langle \Gamma, (\pi_v, r_v)_{v \in V}, \mu \rangle$ is a Public PBE, so that the continuation payoff vector $u^{\mathbf{M}_t}(\mu_{t+1}) \equiv (u_L^{\mathbf{M}_t}(\mu_{t+1}), u_H^{\mathbf{M}_t}(\mu_{t+1}), u_S^{\mathbf{M}_t}(\mu_{t+1}))$ is an equilibrium payoff vector of $G^\infty(\mu_{t+1})$. Formally, $u^{\mathbf{M}_t}(\mu_{t+1}) \in \mathcal{E}^*(\mu_{t+1})$.

Equations $\mathrm{PC}_{h^t,v}$ and $\mathrm{IC}_{h^t,v,v'}$ are analogous to the participation and incentive compatibility constraints one would obtain in mechanism design except for the following: The participation constraint is potentially type-dependent (the right hand side is $u_v^{\mathbf{M}_t}(\emptyset)$). To obtain a participation constraint that more closely resembles that of the commitment case, we show that the low-valuation buyer's equilibrium payoff is 0 in *any* equilibrium of $G^\infty(\mu_0)$:

Lemma 3. *For all $\mu_0 \in \Delta(V)$, if $u \in \mathcal{E}^*(\mu_0)$, then $u_L = 0$.*

Lemma 3 implies that when the buyer's value is v_L, her utility from participation is 0.

[17] This continuation payoff can also depend on h^t, but we omit this dependence to simplify notation.

Thus, we can write Equation $\text{PC}_{h^t,v}$ when $v = v_L$ as follows:

$$\int_{\Delta(V)} \left(v_L q^{\mathbf{M}_t}(\mu_{t+1}) - x^{\mathbf{M}_t}(\mu_{t+1})\right) \beta^{\mathbf{M}_t}(d\mu_{t+1}|v_L) = 0, \tag{5}$$

where to obtain Equation 5, we have used Lemma 3 twice in Equation $\text{PC}_{h^t,v}$. First, to replace the continuation payoffs by 0 both after participation and after rejection of $\mathbf{M}_t$. Second, to change the inequality in Equation $\text{PC}_{h^t,v}$ to an equality. After all, the buyer's utility when she participates in the mechanism $\mathbf{M}_t$ is her equilibrium payoff at h^t, and hence an element of $\mathcal{E}^*(\mu_t(h^t))$. Equation 5 can be used to rewrite Equation $\text{IC}_{h^t,v,v'}$ for $v = v_H$ as follows:

$$\begin{aligned}&\int_{\Delta(V)} \left(v_H q^{\mathbf{M}_t}(\mu_{t+1}) - x^{\mathbf{M}_t}(\mu_{t+1}) + (1 - q^{\mathbf{M}_t}(\mu_{t+1}))\delta u_H^{\mathbf{M}_t}(\mu_{t+1})\right) \beta^{\mathbf{M}_t}(d\mu_{t+1}|v_H) \geq \\ &\int_{\Delta(V)} \left(\Delta v q^{\mathbf{M}_t}(\mu_{t+1}) + (1 - q^{\mathbf{M}_t}(\mu_{t+1}))\delta u_H^{\mathbf{M}_t}(\mu_{t+1})\right) \beta^{\mathbf{M}_t}(d\mu_{t+1}|v_L).\end{aligned} \tag{6}$$

Equations 5 and 6 already impose constraints on the maximum revenue the seller can make in period t. Indeed, like in standard mechanism design, the participation constraint for v_L and the truthtelling constraint for v_H determine the maximum expected payment the seller can extract from the buyer in period t. Replacing the upper bound on the expected payments obtained from these equations in the seller's payoff (Equation 4), we obtain an upper bound on the seller's revenue at h^t when he offers mechanism $\mathbf{M}_t$. It is immediate to check that this upper bound corresponds to $VS((\tau^{\mathbf{M}_t}, q^{\mathbf{M}_t}), (u_H^{\mathbf{M}_t}, u_S^{\mathbf{M}_t}), \mu(h^t))$. Lemma 2 then follows.

Having established why the optimization program in OPT is an upper bound on the seller's equilibrium payoff, Section 4.2 studies properties of its solution which are important for the proof of Theorem 1.

4.2 Properties of the solution to OPT

The main result in this section, Proposition 1, shows two properties of the solution to the program defined in OPT that are important to prove that the seller's maximum equilibrium payoff is $u_S^*(\mu_0)$. As we explain next, a key difficulty that the program in OPT allows us to circumvent is the analysis of the *belief dynamics* in the game.

An important result in the bargaining literature is the *skimming lemma*, which states that incentive compatibility of the buyer's behavior implies that the expected discounted

probability of trade of v_H is higher than that of v_L (see, for instance, Fudenberg et al., 1985). This property immediately implies that along the path of play the seller's beliefs fall conditional on no trade. As a consequence, prices must fall along the path of play.

Contrast this to the game we analyze in which the seller offers mechanisms that enable him to design how much he observes about the buyer's choices. In particular, the seller can choose how fast he learns about the buyer's value conditional on no trade; he could even choose to become more optimistic about the buyer's value conditional on no trade.[18] On the one hand, this would allow the seller to avoid the belief dynamics associated with posted-prices and hence, avoid the temptation to trade more often with v_L in future rounds. On the other hand, this comes at a cost. Recall that Lemma 1 implies that the mechanism's allocation is measurable with respect to the information generated by the mechanism and this information is, in turn, subject to the Bayes' plausibility constraint. This implies that for the seller's beliefs conditional on no trade to fall slowly (or not fall at all), it must be that the seller is selling the good to v_H with small probability.

It turns out that the program defined in OPT is useful to discipline belief dynamics. While it may not be obvious how to rule out that there is an equilibrium in which the seller's beliefs may sometimes go up conditional on no trade, it turns out that this is never the case in a solution to OPT. Indeed, as we establish in Proposition 1 below, it is never optimal to not sell the good and induce a belief above the prior. Furthermore, whenever $\mu_0 > \overline{\mu}_1$, conditional on selling the good with positive probability, the seller sells the good only to the high-valuation buyer.

Proposition 1. *Suppose that $\mu_0 \geq \overline{\mu}_1$ and let (τ_0, q_0),(u_H, u_S) denote a solution to OPT. Then, the following hold:*

1. *It is never optimal to induce a belief $\mu_1 \geq \mu_0$ and not sell the good. That is, letting*

[18] While the truthtelling equations $\mathrm{IC}_{h^t,v,v'}$ can be used to derive a "monotonicity" condition analogous to that in the skimming lemma, it only implies that on average the expected probability of trade of v_H must be higher than that of v_L:

$$\int_{\Delta(V)} \left(\Delta v q^{\mathbf{M}_t}(\mu_{t+1}) + (1 - q^{\mathbf{M}_t}(\mu_{t+1})) u_H^{\mathbf{M}_t}(\mu_{t+1})\right) \left(\beta^{\mathbf{M}_t}(d\mu_{t+1}|v_H) - \beta^{\mathbf{M}_t}(d\mu_{t+1}|v_L)\right) \geq 0.$$

$A = [\mu_0, 1)$, *then*

$$\int_A (1 - q_0(\mu_1))\tau_0(d\mu_1) = 0.$$

2. *Furthermore, if* $\mu_0 > \overline{\mu}_1$ *and the seller induces* μ_1 *and sells the good (that is,* $q_0(\mu_1) > 0$*), then* $\mu_1 = 1$.

The proof is in Section A.2. In what follows, we provide intuition for Proposition 1, starting from part 1. To see why it is not optimal to not sell the good while at the same time inducing a belief $\mu_1 \geq \mu_0$, note the following. First, associated to any continuation payoff, $(u_H(\mu_1), u_S(\mu_1))$, there is a mechanism chosen by the seller when his belief is μ_1, and continuation payoffs for the seller and the buyer in the event that the good is not sold. This implies that, conditional on inducing a belief μ_1, the seller with belief μ_0 could always choose today the mechanism and the continuation payoffs associated with $(u_H(\mu_1), u_S(\mu_1))$ in a solution to OPT. Second, the seller with belief μ_0 below μ_1 pays rents to v_H with lower probability than the seller with belief μ_1 (the term premultiplying u_H in Equation 3 is positive). It follows that the seller with belief μ_0 prefers to accrue today the payoff from the mechanism (and continuation payoffs) that induce $(u_H(\mu_1), u_S(\mu_1))$, contradicting that it is optimal to induce μ_1 and not sell the good.

That part 2 holds follows from the observation that a seller with prior above $\overline{\mu}_1$ prefers to trade with v_L in at least two periods (recall that $\hat{v}_L(\mu_0) < 0$ when $\mu_0 > \overline{\mu}_1$). Thus, it is never optimal to sell the good to v_L with positive probability today. It follows that if $q_0(\mu_1) > 0$, then the seller must assign the good to v_H, and hence $\mu_1 = 1$.

Proposition 1 implies that a solution to OPT never induces posteriors in $[\mu_0, 1)$. This, in turn, delivers the following expression for the value of OPT, where $VS(\mu_0)$ denotes the value function of the program in OPT:

$$VS(\mu_0) \equiv \max_{\tau_0, (u_H, u_S)} \tau_0(\{1\})v_H + \int_{[0,\mu_0)} \delta\left(u_S(\mu_1) + \left(\frac{\mu_1}{1-\mu_1} - \frac{\mu_0}{1-\mu_0}\right)u_H(\mu_1)\right)\tau_0(d\mu_1). \quad (7)$$

Equation 7 simply states that the solution to OPT can be described by the probability of selling to v_H today (the probability of inducing a belief $\mu_1 = 1$) and the probability with which the good is not sold and a belief below the prior is induced. One distribution over posteriors is of particular interest in what follows: the one that splits μ_0 between 1 (with

$q_0(1) = 1$) and $\mu_1 < \mu_0$ (with $q_0(\mu_1) = 0$). Bayes' plausibility implies that the weights on 1 and μ_1 are

$$\frac{\mu_0 - \mu_1}{1 - \mu_1} \quad \text{and} \quad \frac{1 - \mu_0}{1 - \mu_1},$$

respectively. Note that this is precisely the type of mechanism that the seller uses in the posted-prices assessment. Corollary 2 shows that these are essentially the distributions over posteriors that solve the problem in Equation 7:

Corollary 2. *The value of the program in Equation 7 equals the value of*

$$\max_{G \in \Delta([0,\mu_0))} \max_{(u_H, u_S)} \int_{[0,\mu_0)} \left[\frac{\mu_0 - \mu_1}{1 - \mu_1} v_H + \frac{1 - \mu_0}{1 - \mu_1} \delta \left(u_S(\mu_1) + \left(\frac{\mu_1}{1 - \mu_1} - \frac{\mu_0}{1 - \mu_0} \right) u_H(\mu_1) \right) \right] G(d\mu_1). \tag{8}$$

The proof is in Section A.2 and is a consequence of the constraint that τ_0 is Bayes' plausible for μ_0. Corollary 2 implies that if posteriors μ_1 and μ_1' are on the support of τ_0, then the seller is indifferent between splitting μ_0 between μ_1 and 1 and splitting μ_0 between μ_1' and 1. Indeed, given the preceding discussion, the term in the square brackets inside the integral in Equation 8 represents the payoff from splitting μ_0 between 1 and $\mu_1 < \mu_0$. Corollary 2 implies that, for a fixed choice of continuation payoffs, the solution to the problem in Equation 7 is as if the seller were randomizing between splitting the prior between 1 (and selling the good) and $\mu_1 < \mu_0$ (and not selling the good). As a consequence, to determine the optimal τ_0 it is enough to compare the payoffs of the splittings of μ_0 between μ_1 and 1 for different μ_1.

5 Proof of Theorem 1: key steps

Section 5 overviews the main steps of the proof of Theorem 1. Taking as given that $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$ is a PBE assessment, Section 5.1 shows that it achieves the unique equilibrium payoff for the seller and that, except for the threshold beliefs, $\{\overline{\mu}_n\}_{n \geq 1}$, there is a unique equilibrium outcome and hence, a unique equilibrium payoff for the high-valuation buyer as well. Section 5.2 reviews the main steps to show the first statement of Theorem 1; namely, that $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$ is a PBE assessment.

5.1 Characterization of the equilibrium payoffs of $G^\infty(\mu_0)$

We show that $u_S^*(\mu_0)$ is both a lower bound and an upper bound on the seller's equilibrium payoff. That is, we show that the seller can never do better nor worse than if

he were limited to choose prices in each period so that having access to a richer action space does not increase nor decrease the seller's payoff. The proof that $u_S^*(\mu_0)$ is a lower bound is similar to the logic in Gul, Sonnenschein, and Wilson (1986): if there is an equilibrium where the seller earns less than $u_S^*(\mu_0)$, then he can always *undercut* the price in the posted-prices assessment and earn close to $u_S^*(\mu_0)$. The only difference is that here this undercutting is done using mechanisms. Instead, the proof that $u_S^*(\mu_0)$ is an upper bound on the seller's equilibrium payoff takes advantage of the program in OPT.

To illustrate the main steps of the proof that $u_S^*(\mu_0)$ is both an upper bound and a lower bound on the seller's equilibrium payoff, fix $\mu_0 \in [\overline{\mu}_n, \overline{\mu}_{n+1})$ for some $n \geq 1$ and suppose that we have already shown that Theorem 1 holds for all $m < n$ and all $\mu \in [\overline{\mu}_m, \overline{\mu}_{m+1})$.[19] In what follows, we show that Theorem 1 holds for $\mu_0 \in [\overline{\mu}_n, \overline{\mu}_{n+1})$.

$u_S^*(\mu_0)$ is an upper bound on the seller's equilibrium payoff: To show that the value of OPT corresponds to the seller's payoff in the posted-prices assessment, we argue that splitting the prior between 1 and $\overline{\mu}_{n-1}$ dominates all other splittings. Corollary 2 implies that this is enough to show that $u_S^*(\mu_0) = VS(\mu_0)$. In what follows, we first argue why, conditional on inducing a belief $\mu_1 < \overline{\mu}_n$, the seller places weight on at most $\overline{\mu}_{n-1}$ (and on $\overline{\mu}_{n-2}$ only if $\mu_0 = \overline{\mu}_n$). We then argue why it is never optimal to induce a belief $\mu_1 \in [\overline{\mu}_n, \mu_0)$.

To see why conditional on inducing a belief in $[0, \overline{\mu}_n)$, the solution to the problem in Equation 8 places at most weight on $\{\overline{\mu}_{n-2}, \overline{\mu}_{n-1}\}$, recall that by the inductive hypothesis, for all $\mu_1 < \overline{\mu}_n$, the seller's continuation payoff is as described in Theorem 1. Furthermore, note the seller with prior belief μ_0 prefers the continuation payoff that minimizes the high-valuation buyer's rents whenever, conditional on no trade, he induces a posterior $\mu_1 < \mu_0$.[20] This means that at the threshold beliefs, the seller prefers to induce a continuation payoff of $u_H^*(\cdot)$ for the high-valuation buyer, as is the case in the posted-prices assessment.

Lemma B.2 shows that if the seller places positive probability on $[0, \overline{\mu}_n)$, then this can happen at most at $\{\overline{\mu}_{n-2}, \overline{\mu}_{n-1}\}$. Furthermore, $\overline{\mu}_{n-2}$ can be in the support of τ_0 only if $\mu_0 = \overline{\mu}_n$. This is intuitive: Beliefs in $[0, \overline{\mu}_n)$ can be classified in terms of the time it takes

[19] We verify that Theorem 1 holds for $\mu_0 \in [0, \overline{\mu}_1)$ in Section B.2.

[20] This can be easily seen from noting that when $\mu_1 < \mu_0$, the term pre-multiplying $u_H(\cdot)$ in Equation 3 is negative.

to trade with the low valuation buyer, where if $\mu_1 \in [\overline{\mu}_m, \overline{\mu}_{m+1})$, then trade happens with v_L in m periods. It can never be optimal to induce a belief $\mu_1 \in (\overline{\mu}_m, \overline{\mu}_{m+1})$ for $m \leq n-1$: Both μ_1 and $\overline{\mu}_m$ imply that trade happens with v_L in m periods; however, inducing $\overline{\mu}_m$ allows the seller to trade with v_H with a higher probability. It follows that conditional on inducing beliefs in $[0, \overline{\mu}_n)$, the solution to OPT induces at most beliefs in $\{\overline{\mu}_m\}_{m \leq n-1}$. The indifference condition that defines the threshold beliefs implies that since $\mu_0 \geq \overline{\mu}_n$, then it cannot be optimal to induce beliefs below $\overline{\mu}_{n-2}$ (see Lemma B.1).

To conclude the proof that $u_S^*(\mu_0)$ is an upper bound on $VS(\mu_0)$, it remains to show that it is not optimal to induce posteriors in $[\overline{\mu}_n, \mu_0)$. While Proposition 1 implies that at the solution to OPT the seller's beliefs go down conditional on no trade, it does not say *how fast* they go down. Indeed, it could be optimal to induce beliefs in $[\overline{\mu}_n, \mu_0)$ if there exist continuation equilibria where the seller at the induced beliefs somehow manages to slowly trade with v_H so as to maximally delay trade with v_L. As we show next, this cannot be optimal for μ_0 close to $\overline{\mu}_n$: The closer to $\overline{\mu}_n$, the smaller the probability that the seller with belief μ_0 can trade with v_H if conditional on no trade, his beliefs must remain above $\overline{\mu}_n$. Thus, there exists μ_0 small enough for which it is better to trade with v_L in n periods in exchange of increasing the probability of trading with v_H today.

Formally, let

$$F = \left\{\mu_0 \in [\overline{\mu}_n, \overline{\mu}_{n+1}) : \text{Equation 2 does not hold}\right\}, \tag{9}$$

and let $\mu^\star = \inf F$. Note that $\mu^\star \in [\overline{\mu}_n, \overline{\mu}_{n+1})$. Then, for all $\epsilon > 0$, there exists $\mu_0 \in [\mu^\star, \mu^\star + \epsilon)$ such that Equation 2 does not hold for μ_0. For each $\epsilon \in (0, \overline{\mu}_{n+1} - \overline{\mu}_n)$, let $\overline{VS}_\epsilon$ denote the supremum of VS over $\mu_0 \in [\mu^\star, \mu^\star + \epsilon)$. For all $\mu_0 \in [\mu^\star, \mu^\star + \epsilon)$, the following holds:

$$VS(\mu_0) \leq \max\left\{u_S^*(\mu_0), \frac{\mu_0 - \mu^\star}{1 - \mu^\star} v_H + \frac{1 - \mu_0}{1 - \mu^\star} \delta \overline{VS}_\epsilon\right\}. \tag{10}$$

To see why Equation 10 holds note the following. First, if the solution to OPT places positive mass below $\mu^\star$, Lemma B.2 implies that $VS(\mu_0) = u_S^*(\mu_0)$, since the seller places weight on $\overline{\mu}_{n-1}$ (or $\overline{\mu}_{n-2}$ if $\mu_0 = \overline{\mu}_n$). The equality then follows because Corollary 2 implies that if μ_1 is in the support of τ_0, then the seller's payoff coincides with the payoff he would obtain from splitting μ_0 between 1 and μ_1. The payoff $u_S^*(\mu_0)$ corresponds to splitting μ_0 between 1 and $\overline{\mu}_{n-1}$. Second, if the solution to OPT places weight on $[\mu^\star, \mu_0)$,

then the second term on the RHS of Equation 10 is an upper bound to $VS(\mu_0)$. After all, (i) $(\mu_0-\mu^\star)/(1-\mu^\star)$ is the largest weight that can be assigned to v_H while still remaining on $[\mu^\star,\mu_0)$ and (ii) the remaining weight corresponds to some $\mu_1 \in [\mu^\star,\mu_0)$ with payoff,

$$\left(u_S(\mu_1)+(1-\mu_1)\left(\frac{\mu_1}{1-\mu_1}-\frac{\mu_0}{1-\mu_0}\right)u_H(\mu_1)\right) \le u_S(\mu_1) \le \overline{VS}_\epsilon,$$

where the first inequality follows from $\mu_1<\mu_0$ and the second from the definition of $\overline{VS}_\epsilon$ together with $\mu_1 \in [\mu^\star,\mu_0) \subset [\mu^\star,\mu^\star+\epsilon)$. Taking the supremum over $\mu_0 \in [\mu^\star,\mu^\star+\epsilon)$ on both sides of Equation 10, we obtain

$$\overline{VS}_\epsilon \le \max\left\{\overline{u}_{S\epsilon}, \frac{\epsilon}{1-\mu^\star}(v_H-\delta\overline{VS}_\epsilon)+\delta\overline{VS}_\epsilon\right\}, \tag{11}$$

where $\overline{u}_{S\epsilon}$ is the supremum of $u_S^*(\mu_0)$ over $\mu_0 \in [\mu^\star,\mu^\star+\epsilon)$, and the expression in the second term follows from noting that $\mu_0-\mu^\star<\epsilon$ and $v_H>\delta\overline{VS}_\epsilon$.

We claim that there exists $\overline{\epsilon}>0$ such that for all $\epsilon \in (0,\overline{\epsilon})$,

$$\frac{\epsilon}{1-\mu^\star}(v_H-\delta\overline{VS}_\epsilon)+\delta\overline{VS}_\epsilon \le \overline{u}_{S\epsilon}.$$

Toward a contradiction, suppose not. Then, for all $\epsilon \in (0,\overline{\mu}_{n+1}-\overline{\mu}_n)$, there exists $f(\epsilon) \in (0,\epsilon)$ such that

$$\overline{u}_{Sf(\epsilon)} < \frac{f(\epsilon)}{1-\mu^\star}(v_H-\delta\overline{VS}_{f(\epsilon)})+\delta\overline{VS}_{f(\epsilon)}. \tag{12}$$

Since $f(\epsilon)<\epsilon$, then $f(\epsilon)\to 0$ as $\epsilon\to 0$. Equation 11 together with Equation 12 imply that:

$$\overline{VS}_{f(\epsilon)} \le \frac{f(\epsilon)}{1-\mu^\star}(v_H-\delta\overline{VS}_{f(\epsilon)})+\delta\overline{VS}_{f(\epsilon)},$$

so that as $\epsilon\to 0$,

$$\lim_{\epsilon\to 0}\overline{VS}_{f(\epsilon)} \le \delta\lim_{\epsilon\to 0}\overline{VS}_{f(\epsilon)},$$

a contradiction, since for all ϵ, $\overline{VS}_\epsilon \ge v_L>0$.[21] It follows that there exists $\overline{\epsilon}>0$ such that $\forall\epsilon\in(0,\overline{\epsilon})$, $\overline{VS}_\epsilon=\overline{u}_{S\epsilon}$.

[21] Note that the limit $\lim_{\epsilon\to 0}\overline{VS}_{f(\epsilon)}$ exists up to a convergent subnet because $\overline{VS}_{f(\epsilon)}$ is bounded.

We now claim that $VS(\mu_0) = u_S^*(\mu_0)$ for all $\mu_0 \in [\mu^\star, \mu^\star + \overline{\epsilon})$. Toward a contradiction, suppose that there exists $\mu_0 \in [\mu^\star, \mu^\star + \overline{\epsilon})$ such that $VS(\mu_0) > u_S^*(\mu_0)$. By continuity of u_S^* on $[\overline{\mu}_n, \overline{\mu}_{n+1})$ (see Equation 19), there exists $\gamma > 0$ such that letting $\epsilon = \mu_0 + \gamma - \mu^\star$, we have

$$u_S^*(\mu_0 + \gamma) = \overline{u}_{S\epsilon} < VS(\mu_0) \leq \overline{VS}_\epsilon,$$

where the first equality follows from u_S^* being increasing on $[\overline{\mu}_n, \overline{\mu}_{n+1})$. This is a contradiction.

Thus, for all $\mu_0 \in [\mu^\star, \mu^\star + \overline{\epsilon})$ we have that $u_S^*(\mu_0)$ is an upper bound on the seller's payoff.[22] By definition of $\mu^\star$, then there exists $\mu_0 \in [\mu^\star, \mu^\star + \overline{\epsilon})$ such that either (i) $u_S^*(\mu_0)$ is not the seller's unique equilibrium payoff, or (ii) there exists $u' \in \mathcal{E}^*(\mu_0)$ such that $u'_S = u_S^*(\mu_0)$, but $u'_H \neq u_H^*(\mu_0)$. Otherwise, we would have that there exists $\epsilon (= \overline{\epsilon})$ such that for all $\mu_0 \in [\mu^\star, \mu^\star + \epsilon)$, Equation 2 holds, which contradicts the definition of $\mu^\star$.

In what follows, we first show that $u_S^*(\mu_0)$ is a lower bound on the seller's equilibrium payoff, and hence the seller's unique equilibrium payoff. We then show that there is a unique equilibrium payoff for the high-valuation buyer, except if $\mu_0 = \overline{\mu}_n$.

$u_S^*(\mu_0)$ is a lower bound on the seller's equilibrium payoff (Proposition B.1): Suppose there is an equilibrium of $G^\infty(\mu_0)$ in which the seller earns a payoff $\underline{u}_S < u_S^*(\mu_0)$. Note that in the posted-prices equilibrium, the seller's mechanism induces posterior beliefs below $\overline{\mu}_n$, where by assumption the unique continuation payoff for the seller coincides with that of the posted-prices equilibrium. Then, if the seller were to deviate to the mechanism he offers in the posted-prices equilibrium *and* we could guarantee that the buyer participates and tells the truth, the seller's continuation payoff will coincide with those in the posted-prices equilibrium. Thus, it suffices to show that the seller can guarantee that the buyer participates and tells the truth when he attempts to deviate to such a mechanism at the beginning of the game.

Key to showing that the seller can ensure such a deviation is the observation that whenever the low-valuation buyer has strict incentives to participate, equilibrium considerations imply that the high-valuation buyer also participates with probability 1. If not,

[22] The argument above is reminiscent to that in Fudenberg and Tirole (1991a)'s treatment of the equilibrium in the posted-prices game with binary values.

Bayes' rule would imply that the seller assigns probability 1 to the high-valuation buyer after non-participation and thus, the high-valuation buyer earns a continuation payoff of 0 (Lemma A.1). Instead, since the low-valuation buyer strictly prefers to participate, it follows that the high-valuation buyer can ensure a strictly positive payoff from participation. Thus, the high-valuation buyer participates with probability 1 whenever the low-valuation buyer strictly prefers to participate.

Using the above observation, we construct a mechanism that is "close" to the one the seller offers under $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$, but modified so that (i) the low-valuation buyer gets paid for participating, (ii) the high-valuation buyer gets paid for telling the truth, and (iii) instead of generating a belief of $\overline{\mu}_{n-1}$, the seller generates a belief slightly above $\overline{\mu}_{n-1}$. The property in (iii) ensures that the seller with belief μ_0 cannot be "punished" by the seller with belief $\overline{\mu}_{n-1}$, who could choose to trade with v_L in $n-2$ periods.

The proof so far implies that for all $\mu_0 \in [\mu^\star, \mu^\star + \overline{\epsilon})$, there is a unique equilibrium payoff for the seller and it coincides with that of the posted-prices equilibrium. It remains to show that the high-valuation buyer's payoff coincides with $u_H^*(\mu_0)$, unless μ_0=$\overline{\mu}_n$.

High-valuation buyer's payoff: Suppose there exists $u' \in \mathcal{E}^*(\mu_0)$ such that $u'_S = u_S^*(\mu_0)$, but $u'_H \neq u_H^*(\mu_0)$. This equilibrium payoff, u', is associated to a mechanism in period 0, (τ'_0, q'_0), and continuation payoffs, (u'_H, u'_S), for the buyer and the seller. We argue that the virtual surplus associated to $((\tau'_0, q'_0),(u'_H, u'_S))$ must equal $u_S^*(\mu_0)$. To see this, note that

$$VS((\tau'_0, q'_0), (u'_H, u'_S), \mu_0) \geq u_S^*(\mu_0).$$

Otherwise, we have that

$$u'_S \leq VS((\tau'_0, q'_0), (u'_H, u'_S), \mu_0) < u_S^*(\mu_0),$$

where the first inequality follows from Lemma 2 and the second inequality is by assumption. Since $u'_S = u_S^*(\mu_0)$, we have a contradiction.

Furthermore, it cannot be the case that

$$VS((\tau'_0, q'_0), (u'_H, u'_S), \mu_0) > u_S^*(\mu_0),$$

since (τ_0', q_0') together with the continuation payoffs (u_H', u_S') are feasible choices in OPT. It follows that

$$VS((\tau_0', q_0'), (u_H', u_S'), \mu_0) = u_S^*(\mu_0).$$

Thus,

$$((\tau_0', q_0'), (u_H', u_S')) \in \arg\max_{(\tau_0, q_0), (u_H, u_S)} VS((\tau_0, q_0), (u_H, u_S), \mu_0).$$

However, the proof that $u_S^*(\mu_0)$ is the value of OPT implies that the solution to OPT is unique, unless $\mu_0 = \overline{\mu}_n$. It then follows that $\mu_0 = \overline{\mu}_n$, so that there is a continuum of solutions to OPT, with v_H's payoff ranging from $u_H^*(\mu_0)$ to $u_H^*(\mu_0)/\delta$. This completes the proof.

5.2 $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$ is an equilibrium assessment

The analysis so far has relied on the observation that $(0, u_H^*(\mu_0), u_S^*(\mu_0))$ is an equilibrium payoff. The rest of the proof of Theorem 1 shows that $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$ is an equilibrium assessment (see Appendix C). To do this, we first complete the equilibrium assessment by specifying the seller's and the buyer's strategy after every history (see Section C.1-C.2). We then show that given beliefs and continuation payoffs, neither the buyer nor the seller have a one shot deviation (Section C.3). The results in Athey and Bagwell (2008) imply that this is enough to conclude that $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$ is an equilibrium assessment.

Seller's strategy: Except for the cutoff beliefs $\{\overline{\mu}_n\}_{n \geq 1}$, we specify that the seller plays the mechanism described in the posted-prices assessment on and off the path of play. Instead, the seller's strategy off the path of play when his beliefs are in $\{\overline{\mu}_n\}_{n \geq 1}$ needs to be determined *jointly* with the buyer's strategy, to which we turn next.

Buyer's strategy (Section C.1): To complete the buyer's strategy, we first classify mechanisms according to whether they satisfy the participation and truthtelling constraints for the buyer given the continuation payoffs under $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$. For mechanisms that satisfy these constraints, we specify that the buyer indeed participates and truthfully reports her value to the mechanism.

To specify the buyer's strategy for mechanisms that fail to satisfy either constraint, one needs to determine simultaneously the buyer's best response and the seller's beliefs conditional on observing either the buyer reject the mechanism, or the buyer accept the

mechanism and the output message that results from the buyer's report. On the one hand, the buyer's continuation payoffs depend on the seller's beliefs, which are determined by the buyer's strategy. On the other hand, whether the buyer's strategy is a best response depends on the continuation payoffs. We use the results in Simon and Zame (1990) to solve for this fixed-point. It is at this point where the possibility that the seller randomizes when indifferent between trading with v_L in n or $n-1$ periods arises to ensure that the buyer's best response is well-defined.

By specifying the buyer's strategy in the way described above, we ensure that the buyer is best responding to the seller's strategy given the continuation payoffs. It remains to show that the seller has no one-shot deviations:

The seller has no one-shot deviations (Section C.3): To show that the seller does not have one-shot deviations, we rely once again on the optimization problem defined by OPT. For concreteness, suppose we are at a history h^t and let μ_t denote the seller's belief at h^t.

We show in Section C.3 that the payoff from any deviation at history h^t is bounded above by the value of the program in OPT evaluated at $\mu_0 = \mu_t$ subject to the constraint that for posterior beliefs μ_{t+1} below μ_t, the continuation payoffs for the buyer and the seller are given by $(u_H^*(\mu_{t+1}), u_S^*(\mu_{t+1}))$. The results in Section 5.1 imply that the value of this program coincides with $u_S^*(\mu_t)$. Since this is the seller's payoff under the equilibrium strategy at h^t, it follows that the seller has no one-shot deviations at h^t.

Two observations are key to show that this upper bound holds. The first is one we have already used in the solution to OPT: For threshold beliefs $\overline{\mu}_n$ below μ_t, the seller with prior belief μ_t prefers the continuation payoff vector in which the high-valuation buyer receives her lowest equilibrium payoff, $u_H^*(\overline{\mu}_n)$. Thus, by selecting continuation payoffs in this way, we exaggerate the payoff that the seller can guarantee from a deviation. The second is that *any* mechanism $\mathbf{M}_t$ together with the buyer's best response to $\mathbf{M}_t$ define a new mechanism, $\mathbf{M}_t'$, that satisfies the participation and truthtelling constraint of the buyer *given* the continuation payoffs associated to $\mathbf{M}_t$. Thus, we can use mechanism $\mathbf{M}_t'$ together with the continuation payoffs specified by the assessment when $\mathbf{M}_t$ is offered to bound the revenue from $\mathbf{M}_t$ by its virtual surplus. This completes the description of the main steps in the proof of Theorem 1.

6 Conclusions

This is the first paper to characterize all equilibrium outcomes in an infinite-horizon, mechanism-selection game between an uninformed designer and a privately informed agent with persistent private information under limited commitment. We do so by marrying insights from the literatures on bargaining and mechanism design. Following the results in our previous work, Doval and Skreta (2020), we endow the seller with a larger class of mechanisms that enables the seller to design his posterior beliefs about the buyer's value. The combination of mechanism design and information design elements was key in obtaining a tractable characterization. The revelation principle in Doval and Skreta (2020) also allows us to simplify the equilibrium behavior of the buyer, so that for the most part we were able to focus on the strategic considerations that pertain the seller.

The study of the properties of the solution to OPT can provide a useful benchmark for the analysis of the designer's best equilibrium payoff in other settings with quasilinear utility. In such settings, the virtual surplus is an upper bound on the designer's equilibrium payoff. Furthermore, when there is a continuum of types, the application of the envelope theorem implies that the designer's payoff can be represented as the virtual surplus. The program in OPT is exactly like the relaxed problem in standard mechanism design: If a solution to OPT satisfies the ignored constraints, then there is a PBE of the mechanism-selection game that implements the solution to OPT. However, as the analysis of Laffont and Tirole (1993) and, more recently, by Breig (2020), underscores it is not immediate that the solution to the relaxed problem satisfies *all* constraints and hence, can be implemented as an equilibrium outcome. Nevertheless, like the solution to the relaxed problem in standard mechanism design, the solution to OPT provides a natural benchmark to study the properties of the designer's optimal mechanism. As the analysis in this paper shows, the properties of the solution to OPT can be derived even without much knowledge of how continuation equilibria look like, which reenforces its applicability (recall the results in Proposition 1).

Still, unresolved issues remain for settings with a continuum of types.[23] The most im-

[23] In settings with finitely many types the study of OPT, which is an information design problem, would reveal information about the cardinality of posteriors. For instance, Carathéodory's theorem (Rockafellar, 1970) suggests that with three valuations the optimality of posted prices might be overturned because a solution to OPT might use three beliefs.

portant is that the solution to OPT involves solving an information design problem with a continuum of states for which the existing tools do not readily apply. Indeed, the tools in Gentzkow and Kamenica (2016), Kolotilin (2018) and Dworczak and Martini (2019) apply to settings where the objective function depends only on the posterior mean. Instead, the virtual surplus may depend on the *entire* shape of the distribution, and not just the posterior mean. The characterization of equilibrium outcomes of our game with a continuum of values remains an open question left for future research.

References

Abreu, D., D. Pearce, and E. Stacchetti (1990): "Toward a theory of discounted repeated games with imperfect monitoring," *Econometrica*, 58, 1041–1063.

Acharya, A. and J. Ortner (2017): "Progressive learning," *Econometrica*, 85, 1965–1990.

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

Athey, S. and K. Bagwell (2008): "Collusion with persistent cost shocks," *Econometrica*, 76, 493–540.

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

Aumann, R. J., M. Maschler, and R. E. Stearns (1995): *Repeated games with incomplete information*, MIT press.

Ausubel, L. M. and R. J. Deneckere (1989): "Reputation in bargaining and durable goods monopoly," *Econometrica*, 511–531.

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

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

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

Board, S. and M. Pycia (2014): "Outside options and the failure of the Coase conjecture," *American Economic Review*, 104, 656–71.

Breig, Z. (2020): "Repeated Contracting without Commitment," *Unpublished Manuscript*.

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

CARRILLO, J. D. AND T. MARIOTTI (2000): "Strategic ignorance as a self-disciplining device," *The Review of Economic Studies*, 67, 529–544.

COASE, R. H. (1972): "Durability and monopoly," *The Journal of Law and Economics*, 15, 143–149.

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

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

DOVAL, L. AND V. SKRETA (2020): "Mechanism Design with Limited Commitment," *arXiv preprint arXiv:1811.03579*.

DWORCZAK, P. AND G. MARTINI (2019): "The simple economics of optimal persuasion," *Journal of Political Economy*, 127, 1993–2048.

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

FUDENBERG, D., D. LEVINE, AND J. TIROLE (1985): "Infinite horizon models of bargaining with one-sided uncertainty," in *Game Theoretic Models of Bargaining*, Cambridge University Press, vol. 73, 79.

FUDENBERG, D. AND J. TIROLE (1991a): *Game theory*, MIT Press.

——— (1991b): "Perfect Bayesian equilibrium and sequential equilibrium," *journal of Economic Theory*, 53, 236–260.

GENTZKOW, M. AND E. KAMENICA (2016): "A Rothschild-Stiglitz approach to Bayesian persuasion," *American Economic Review*, 106, 597–601.

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

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

HABIBI, A. (2020): "Motivation and information design," *Journal of Economic Behavior & Organization*, 169, 1–18.

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

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

KOLOTILIN, A. (2018): "Optimal information disclosure: A linear programming approach," *Theoretical Economics*, 13, 607–635.

LAFFONT, J.-J. AND J. TIROLE (1988): "The dynamics of incentive contracts," *Econometrica*, 1153–1175.

——— (1993): *A theory of incentives in procurement and regulation*, MIT press.

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

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

ORTNER, J. (2017): "Durable goods monopoly with stochastic costs," *Theoretical Economics*, 12, 817–861.

PESKI, M. (2019): "Alternating-offer bargaining with incomplete information and mechanisms," .

ROCKAFELLAR, R. T. (1970): *Convex analysis*, Princeton University Press.

SIMON, L. K. AND W. R. ZAME (1990): "Discontinuous games and endogenous sharing rules," *Econometrica*, 861–872.

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.

SOBEL, J. (1991): "Durable goods monopoly with entry of new consumers," *Econometrica*, 59, 1455–1485.

SOBEL, J. AND I. TAKAHASHI (1983): "A multistage model of bargaining," *The Review of Economic Studies*, 50, 411–426.

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

WOLINSKY, A. (1991): "Durable-good monopoly with inventories," *Economics Letters*, 37, 339–343.

A Omitted proofs from Section 4

Appendix A is organized as follows. Section A.1 introduces two auxiliary lemmas regarding the equilibrium payoffs of v_L (Lemma 3) and v_H (Lemma A.1). Section A.2 presents the proof of Proposition 1.

A.1 Auxiliary results

Lemma A.1. *When $\mu_0 = 1$, in any PBE of $G^\infty(\mu_0)$, the buyer's payoff is 0 when her valuation is v_H.*

Proof. To see that Lemma A.1 holds, we proceed as follows. Let $u_H^\star$ denote the highest equilibrium payoff that the buyer when her valuation is v_H can obtain in a PBE when μ_0 equals 1. Towards a contradiction, suppose that $u_H^\star > 0$. In particular, this means that $\delta u_H^\star < u_H^\star$. Let $\mathbf{M}$ denote the mechanism offered by the seller in $t = 0$.

Consider now a mechanism $\mathbf{M}'$ that coincides with $\mathbf{M}$ except that transfers are raised by $\epsilon > 0$, where ϵ is chosen so that $u_H^\star - \epsilon > \delta u_H^\star$. We argue that the buyer must accept this mechanism with probability 1.

Suppose that in equilibrium the buyer rejected such a mechanism. Then, Bayes' rule where possible pins down the seller's beliefs after the rejection. In particular, the buyer when her valuation is v_H can make at most $\delta u_H^\star$ in the continuation. This contradicts that it is sequentially rational to reject the mechanism. The same argument also implies that it cannot be that the buyer is indifferent between accepting and rejecting the mechanism. Thus, any such offer must be accepted with probability 1.

This is a contradiction because the seller's revenue under $\mathbf{M}'$ is higher than that under $\mathbf{M}$, so that the seller has a deviation. We conclude that $u_H^\star = 0$. □

Proof of Lemma 3.

Proof. We establish that for any $\mu_0 \in \Delta(V)$, $u \in \mathcal{E}^*(\mu_0)$ implies that $u_L = 0$. Toward a contradiction, let $\overline{u}_L$ denote

$$\sup\{u_L : (\exists \mu_0 \in \Delta(V))(\exists (u_S, u_H)) : (u_S, u_H, u_L) \in \mathcal{E}^*(\mu_0)\}.$$

By assumption $\overline{u}_L > 0$. Then, there exists a belief μ_0 and an equilibrium payoff vector

$(u_L, u_H, u_S) \in \mathcal{E}^*(\mu_0)$ such that $u_L > \delta \overline{u}_L$, by definition of the supremum. Let $\langle \Gamma, (\pi_v, r_v)_{v \in V}, \mu \rangle$ denote the assessment that achieves that payoff. Let $\epsilon > 0$ be such that $u_L - \epsilon > \delta \overline{u}_L$. Consider the mechanism that the seller offers, $\mathbf{M}_0$ and consider a deviation by the seller such that he raises all payments by ϵ. Since $u_L - \epsilon > \delta \overline{u}_L$, the best response for the low-valuation buyer is to accept this mechanism. Given this, a rejection of the mechanism reveals that the buyer's valuation is v_H. Thus, by rejecting the new mechanism, a buyer with valuation v_H earns at most 0 by Lemma A.1. Therefore, the buyer also accepts the mechanism when her valuation is v_H. It follows that the seller has a profitable deviation, a contradiction. □

A.2 Omitted proofs from Section 4.2

Proof of Proposition 1. To see that part 1 holds, recall that $A = [\mu_0, 1)$. Toward a contradiction, suppose that

$$\int_A (1 - q_0(\mu_1))\tau_0(d\mu_1) > 0.$$

Note that associated to any $(0, u_H(\mu_1), u_S(\mu_1)) \in \mathcal{E}^*(\mu_1)$ chosen by the policy that solves OPT there is an associated mechanism $\mathbf{M}(\mu_1)$ and continuation payoffs, $(0, u_H(\mu_2), u_S(\mu_2)) \in \mathcal{E}^*(\mu_2)$ for the buyer and the seller that are feasible in equilibrium.[24] Furthermore,

$$\begin{aligned}
u_S(\mu_1) &= \int_{\Delta(V)} \left(x^{\mathbf{M}(\mu_1)}(\mu_2) + (1 - q^{\mathbf{M}(\mu_1)}(\mu_2))\delta u_S(\mu_2) \right) \tau^{\mathbf{M}(\mu_1)}(d\mu_2) \leq \\
&\int_{\Delta(V)} \left[\begin{array}{l} q^{\mathbf{M}(\mu_1)}(\mu_2)(\mu_2 v_H + (1-\mu_2)\hat{v}_L(\mu_1)) + \\ (1 - q^{\mathbf{M}(\mu_1)}(\mu_2))\delta \left(u_S(\mu_2) + (1-\mu_2)\left(\frac{\mu_2}{1-\mu_2} - \frac{\mu_1}{1-\mu_1} \right) u_H(\mu_2) \right) \end{array} \right] \tau^{\mathbf{M}(\mu_1)}(d\mu_2) \\
&= u_S(\mu_1) + \mu_1(u_H(\mu_1) - u_{H|L}(\mu_1)),
\end{aligned}$$

where the inequality follows from Lemma 2 and $u_{H|L}(\mu_1)$ is defined as:

$$u_{H|L}(\mu_1) = \int_{\Delta(V)} \left[\Delta v q^{\mathbf{M}(\mu_1)}(\mu_2) + (1 - q^{\mathbf{M}(\mu_1)}(\mu_2))\delta u_H(\mu_2) \right] \frac{1-\mu_2}{1-\mu_1} \tau^{\mathbf{M}(\mu_1)}(d\mu_2).$$

[24] Since the set of mechanisms, $\mathcal{M}_C$, is convex, it is without loss of generality to consider pure strategies for the seller which do not condition on the realization of the public randomization device. However, as in the proof that $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$ is an equilibrium assessment, it is sometimes convenient to specify the seller's strategy as if the seller were mixing.

It can be verified that the payoff from implementing $\mathbf{M}(\mu_1)$ in $t = 0$ and choosing its associated continuation payoffs for $t = 1$ yields a payoff equal to:

$$
\begin{aligned}
&u_S(\mu_1) + \mu_1 \left(u_H(\mu_1) - u_{H|L}(\mu_1)\right) + (1-\mu_1)\left(\frac{\mu_1}{1-\mu_1} - \frac{\mu_0}{1-\mu_0}\right)(u_H(\mu_1) - (u_H(\mu_1) - u_{H|L}(\mu_1))) \\
&= u_S(\mu_1) + (1-\mu_1)\left(\frac{\mu_1}{1-\mu_1} - \frac{\mu_0}{1-\mu_0}\right) u_H(\mu_1) + (1-\mu_1)\frac{\mu_0}{1-\mu_0}(u_H(\mu_1) - u_{H|L}(\mu_1)) \\
&\geq u_S(\mu_1) + (1-\mu_1)\left(\frac{\mu_1}{1-\mu_1} - \frac{\mu_0}{1-\mu_0}\right) u_H(\mu_1), \qquad (13)
\end{aligned}
$$

with the inequality being strict whenever $u_H(\mu_1) > u_{H|L}(\mu_1)$.[25]

Consider the mechanism (τ_0', q_0'), which coincides with (τ_0, q_0) except that whenever $\mu_1 \in A$ is induced, we implement $\mathbf{M}(\mu_1)$ in period 0 and the continuation payoffs associated with $\mathbf{M}(\mu_1)$ in period 1. Since (τ_0', q_0') is feasible, it must be that it is less preferred to (τ_0, q_0). Using Equation 13, we have that the payoff difference between the two policies is:

$$
\begin{aligned}
&VS((\tau_0', q_0'), (u_S', u_H'), \mu_0) - VS((\tau_0, q_0), (u_H, u_S), \mu_0) = \\
&= (1-\delta)\int_A \left(u_S(\mu_1) + (1-\mu_1)\left(\frac{\mu_1}{1-\mu_1} - \frac{\mu_0}{1-\mu_0}\right) u_H(\mu_1)\right)\tau_0(d\mu_1) \\
&+ \int_A (1-\mu_1)\frac{\mu_0}{1-\mu_0}(u_H(\mu_1) - u_{H|L}(\mu_1))\tau_0(d\mu_1) > 0,
\end{aligned}
$$

contradicting the optimality of the original policy.

To see that part 2 in Proposition 1 holds, assume towards a contradiction that there exists $\mu_1 < 1$ such that $q_0(\mu_1) > 0$. If $\mu_1 > 0$, consider an alternative policy which splits the weight on μ_1 conditional on $q_0(\mu_1) > 0$ between 0 and 1, and sets $q_0(0) = 0$ and $q_0(1) = 1$. Instead, if $\mu_1 = 0$, set $q_0(0) = 0$. This leads to a change in payoffs of

$$
\int_{[0,1)} q_0(\mu_1)\left[(\mu_1 v_H + (1-\mu_1)\delta\,\hat{v}_L(\mu_0)) - (\mu_1 v_H + (1-\mu_1)\,\hat{v}_L(\mu_0))\right]\tau_0(d\mu_1) > 0
$$

[25] One interpretation of Equation 13 is the following. In the equilibrium of $G^\infty(\mu_1)$ with payoff vector $(0, u_H(\mu_1), u_S(\mu_1))$, the incentive compatibility constraint for the high valuation buyer may not have been binding. When the new policy implements the virtual surplus of the mechanism associated with μ_1, two things happen. First, we increase the seller's payoff by the same amount that we decrease the buyer's payoff. Second, from today's perspective there is an additional gain: by decreasing v_H's continuation payoffs, we decrease the rents today and hence increase the virtual surplus.

since $\delta < 1$ and $\hat{v}_L(\mu_0) < 0$. This is a contradiction as long as $\int_{[0,1)} q_0(\mu_1)\tau_0(d\mu_1) > 0$. □

Proposition 1 implies that if $((\tau_0, q_0),(u_H, u_S))$ is a maximizer of $VS(\cdot, \mu_0)$, then

$$\begin{aligned}
&VS((\tau_0, q_0), (u_H, u_S), \mu_0) = \\
&\tau_0(\{1\})v_H + \int_{[0,\mu_0)} \delta \left[\frac{\mu_0 - \mu_1}{1-\mu_1} v_H + \frac{1-\mu_0}{1-\mu_1}\delta \left(u_S(\mu_1) + \left(\frac{\mu_1}{1-\mu_1} - \frac{\mu_0}{1-\mu_0} \right) u_H(\mu_1) \right) \right] \tau_0(d\mu_1) \\
&= \int_{[0,\mu_0)} \left[\frac{\mu_0 - \mu_1}{1-\mu_1} v_H + \frac{1-\mu_0}{1-\mu_1}\delta \left(u_S(\mu_1) + \left(\frac{\mu_1}{1-\mu_1} - \frac{\mu_0}{1-\mu_0} \right) u_H(\mu_1) \right) \right] \frac{1-\mu_1}{1-\mu_0} \tau_0(d\mu_1) \\
&= \int_{[0,\mu_0)} \left[\frac{\mu_0 - \mu_1}{1-\mu_1} v_H + \frac{1-\mu_0}{1-\mu_1}\delta \left(u_S(\mu_1) + \left(\frac{\mu_1}{1-\mu_1} - \frac{\mu_0}{1-\mu_0} \right) u_H(\mu_1) \right) \right] G(d\mu_1)
\end{aligned}$$

where the second equality follows from the constraint that τ_0 is Bayes' plausible for μ_0 and the third equality from noting that $(1-\mu_1)/(1-\mu_0)\tau_0$ defines a measure $G \in \Delta([0, \mu_0])$. We now prove Corollary 2:

Proof of Corollary 2. Any distribution over posteriors $\tau \in \Delta\Delta(V)$ with support on $[0, \mu_0] \cup \{1\}$ induces $G \in \Delta([0, \mu_0])$ via

$$G(U') = \int_{U'} \frac{1-\mu_1}{1-\mu_0} \tau(d\mu_1).$$

Furthermore, any distribution $G \in \Delta([0, \mu_0])$ induces a Bayes' plausible distribution over posteriors τ_0 with support on $[0, \mu_0] \cup \{1\}$, where

$$\begin{aligned}
\tau_0(U') &= \int_{U'} \frac{1-\mu_0}{1-\mu_1} G(d\mu_1), U' \subset [0, \mu_0] \\
\tau_0(\{1\}) &= \int_0^{\mu_0} \frac{\mu_0 - \mu_1}{1-\mu_1} G(d\mu_1).
\end{aligned}$$

The result follows. □

B Omitted proofs from Section 5.1

Appendix B is organized as follows. Section B.1 constructs recursively the buyer's and seller's payoffs under $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$. In particular, we verify that the sequence $\{\overline{\mu}_n\}$ is well-defined. Finally, taking as given that $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$ is a PBE assessment and therefore, that the buyer and the seller's payoff under $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$ is an equilibrium payoff (that is,$(0, u_H^*(\mu_0), u_S^*(\mu_0)) \in \mathcal{E}^*(\mu_0)$), Section B.2 completes the proof that

Equation 2 holds by showing that for all $\mu_0 \in \Delta(V)$ (i) $u_S^*(\mu_0)$ is the seller's unique equilibrium payoff for the seller, and (ii) except for the threshold beliefs, $\{\overline{\mu}_n\}_{n\geq 1}$, there is a unique equilibrium payoff for the buyer.

B.1 Payoffs under $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v\in V}, \mu^* \rangle$

Notation: In what follows, we denote by $(\beta_\mu^*, q_\mu^*, x_\mu^*)$ the mechanism used by the seller with belief μ in the assessment $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v\in V}, \mu^* \rangle$.

Buyer's payoffs: While the buyer's payoff is 0 when her valuation is v_L, the buyer's payoff when her valuation is v_H depends on the seller's prior. Let $u_H^*(\mu_0)$ denote the buyer's equilibrium payoff in the posted prices equilibrium. We characterize this payoff inductively.

First, fix $\mu_0 \in [0, \overline{\mu}_1)$. Then, the seller sells the good at a price of v_L, so that

$$u_H^*(\mu_0) = \Delta v.$$

Second, fix $n \geq 1$ and let $\mu_0 \in [\overline{\mu}_n, \overline{\mu}_{n+1})$. Suppose we have already shown that $u_H^*(\overline{\mu}_m) = \delta^m \Delta v$ for $m \leq n-1$. Then, the buyer's payoff when her valuation is v_H under the equilibrium strategy is the following:

$$\beta_{\mu_0}^*(1|v_H)(v_H - (v_H - \delta^n \Delta v)) + (1 - \beta_{\mu_0}^*(1|v_H))\delta u_H^*(\overline{\mu}_{n-1}) = \delta u_H^*(\overline{\mu}_{n-1}) = \delta^n \Delta v.$$

Note that the buyer is indifferent between reporting v_L and reporting v_H along the path of play. This is immediate for $\mu_0 \in [0, \overline{\mu}_1)$. When $\mu_0 \in [\overline{\mu}_n, \overline{\mu}_{n+1})$, the buyer's payoff from reporting v_L at the beginning of the game and then following her equilibrium strategy is $\delta u_H^*(\overline{\mu}_{n-1}) \equiv \delta^n \Delta v$. Thus, along the path of play, the low-valuation buyer is indifferent between participating in the mechanism and not, whereas the high-valuation buyer is indifferent between reporting the truth and not.

Seller's payoffs: Consider a seller with belief $\mu_0 \in [\overline{\mu}_n, \overline{\mu}_{n+1})$. The mechanism used by the seller determines $(\tau_{\mu_0}^*, q_{\mu_0}^*)$ defined as follows. First, if $\mu_0 \in [0, \overline{\mu}_1)$, we have that

$(\tau^*_{\mu_0}(\{\mu_0\}), q^*_{\mu_0}(\mu_0)) = (1,1)$. Second, if $\mu_0 \in [\overline{\mu}_n, \overline{\mu}_{n+1})$ for $n \geq 1$

$$\tau^*_{\mu_0}(1) = \frac{\mu_0 - \overline{\mu}_{n-1}}{1 - \overline{\mu}_{n-1}} = 1 - \tau^*_{\mu_0}(\overline{\mu}_{n-1}) \tag{14}$$

$$q^*_{\mu_0}(1) = 1 - q^*_{\mu_0}(\overline{\mu}_{n-1}). \tag{15}$$

To define the seller's payoff, we first define the cutoffs $\{\overline{\mu}_n\}_{n\geq 0}$ and the seller's payoff in the equilibrium when his belief is $\overline{\mu}_n$.

For $n = 0$, let

$$u^*_S(\mu_0) = v_L = \mu_0 v_H + (1 - \mu_0)\hat{v}_L(\mu_0).$$

Define $\overline{\mu}_1$ as the belief $\mu \in [0,1]$ such that

$$\overline{\mu}_1 v_H + (1 - \overline{\mu}_1)\delta \left[u^*_S(0) + (1-0) \times (\frac{0}{1-0} - \frac{\overline{\mu}_1}{1-\overline{\mu}_1}) u^*_H(0) \right] = \overline{\mu}_1 v_H + (1 - \overline{\mu}_1)\hat{v}_L(\overline{\mu}_1) \tag{16}$$

and let $u^*_S(\overline{\mu}_1)$ denote the value of the expression in Equation 16. Recursively, suppose we have defined $u^*_S(\overline{\mu}_m)$ for $1 \leq m \leq n-1$. Define $\overline{\mu}_n$ to be the belief $\mu \in [0,1]$ such that

$$\begin{aligned} &\frac{\overline{\mu}_n - \overline{\mu}_{n-1}}{1-\overline{\mu}_{n-1}} v_H + \frac{1-\overline{\mu}_n}{1-\overline{\mu}_{n-1}}\delta \left[u^*_S(\overline{\mu}_{n-1}) + (1-\overline{\mu}_{n-1}) \left(\frac{\overline{\mu}_{n-1}}{1-\overline{\mu}_{n-1}} - \frac{\overline{\mu}_n}{1-\overline{\mu}_n} \right) u^*_H(\overline{\mu}_{n-1}) \right] \\ &= \frac{\overline{\mu}_n - \overline{\mu}_{n-2}}{1-\overline{\mu}_{n-2}} v_H + \frac{1-\overline{\mu}_n}{1-\overline{\mu}_{n-2}}\delta \left[u^*_S(\overline{\mu}_{n-2}) + (1-\overline{\mu}_{n-2}) \left(\frac{\overline{\mu}_{n-2}}{1-\overline{\mu}_{n-2}} - \frac{\overline{\mu}_n}{1-\overline{\mu}_n} \right) u^*_H(\overline{\mu}_{n-2}) \right], \end{aligned} \tag{17}$$

and define $u^*_S(\overline{\mu}_n)$ to be the LHS of Equation 17.

Lemma B.1 implies that the cutoffs determined by Equation 17 are well-defined and satisfy $0 < \overline{\mu}_1 < \cdots < \overline{\mu}_n < \overline{\mu}_{n+1} < \ldots$:

Lemma B.1. *The sequence of thresholds defined by Equation 17 is increasing.*

Proof of Lemma B.1. Recall that we are defining $\overline{\mu}_0 = 0$ and $\overline{\mu}_1 = v_L/v_H$. Suppose we have shown that $\overline{\mu}_{n-1} > \overline{\mu}_{n-2}$. We show that $\overline{\mu}_n > \overline{\mu}_{n-1}$.

To do so, fix $\mu \geq \overline{\mu}_{n-1}$. We claim that the difference:

$$\Delta_n(\mu;\overline{\mu}_{n-1},\overline{\mu}_{n-2}) = \frac{\mu-\overline{\mu}_{n-1}}{1-\overline{\mu}_{n-1}} v_H + \frac{1-\mu}{1-\overline{\mu}_{n-1}}\delta\left[u_S^*(\overline{\mu}_{n-1}) + (1-\overline{\mu}_{n-1})\left(\frac{\overline{\mu}_{n-1}}{1-\overline{\mu}_{n-1}} - \frac{\mu}{1-\mu}\right)u_H^*(\overline{\mu}_{n-1})\right]$$
$$-\left(\frac{\mu-\overline{\mu}_{n-2}}{1-\overline{\mu}_{n-2}} v_H + \frac{1-\mu}{1-\overline{\mu}_{n-2}}\delta\left[u_S^*(\overline{\mu}_{n-2}) + (1-\overline{\mu}_{n-2})\left(\frac{\overline{\mu}_{n-2}}{1-\overline{\mu}_{n-2}} - \frac{\mu}{1-\mu}\right)u_H^*(\overline{\mu}_{n-2})\right]\right) \tag{18}$$

is increasing in μ. Note that Δ_n is differentiable in μ, and

$$\frac{\partial}{\partial\mu}\Delta_n(\mu;\overline{\mu}_{n-1},\overline{\mu}_{n-2}) = \frac{v_H(\overline{\mu}_{n-1}-\overline{\mu}_{n-2})}{(1-\overline{\mu}_{n-1})(1-\overline{\mu}_{n-2})} - \delta v_H \frac{\overline{\mu}_{n-1}-\overline{\mu}_{n-2}}{(1-\overline{\mu}_{n-1})(1-\overline{\mu}_{n-2})} + \delta(\delta^{n-2}-\delta^{n-1})v_H > 0.$$

The cutoff $\overline{\mu}_n$ is defined by $\Delta_n(\overline{\mu}_n;\overline{\mu}_{n-1},\overline{\mu}_{n-2}) = 0$. Note that $\overline{\mu}_n \neq \overline{\mu}_{n-1}$ if $n \geq 1$. If $\overline{\mu}_n = \overline{\mu}_{n-1}$, then

$$0 = \Delta_n(\overline{\mu}_n;\overline{\mu}_{n-1},\overline{\mu}_{n-2}) = \delta u_S^*(\overline{\mu}_{n-1}) - u_S^*(\overline{\mu}_{n-1}) < 0,$$

since $\delta < 1$. Because $\Delta_n(\cdot;\overline{\mu}_{n-1},\overline{\mu}_{n-2})$ is increasing, we conclude that $\overline{\mu}_n > \overline{\mu}_{n-1}$. □

Having established this, we can now define the seller's payoff at any history h^t under $\langle\Gamma^*,(\pi_v^*,r_v^*)_{v\in V},\mu^*\rangle$. If $\mu^*(h^t) \in [\overline{\mu}_n,\overline{\mu}_{n+1})$, then

$$U_S^*(h^t) = \frac{\mu^*(h^t)-\overline{\mu}_{n-1}}{1-\overline{\mu}_{n-1}} v_H + \frac{1-\mu^*(h^t)}{1-\overline{\mu}_{n-1}}\delta\left[u_S^*(\overline{\mu}_{n-1}) + (1-\overline{\mu}_{n-1})\left(\frac{\overline{\mu}_{n-1}}{1-\overline{\mu}_{n-1}} - \frac{\mu^*(h^t)}{1-\mu^*(h^t)}\right)u_H^*(\overline{\mu}_{n-1})\right]$$
$$\equiv u_S^*(\mu^*(h^t)) \tag{19}$$

In what follows, we simplify notation by denoting for any $\mu_0,\mu_1 \in \Delta(V)$:

$$R^{(\tau^*,q^*)}(\mu_1,\mu_0) = u_S^*(\mu_1) + (1-\mu_1)\left(\frac{\mu_1}{1-\mu_1} - \frac{\mu_0}{1-\mu_0}\right)u_H^*(\mu_1), \tag{20}$$

and note that $R^{(\tau^*,q^*)}(\mu_0,\mu_0) = u_S^*(\mu_0)$. Furthermore, if $\mu_0 \in [\overline{\mu}_n,\overline{\mu}_{n+1})$, we can write $R^{(\tau^*,q^*)}(\mu_0,\mu_0)$ as

$$R^{(\tau^*,q^*)}(\mu_0,\mu_0) = \frac{\mu_0-\overline{\mu}_{n-1}}{1-\overline{\mu}_{n-1}} v_H + \frac{1-\mu_0}{1-\overline{\mu}_{n-1}}\delta R^{(\tau^*,q^*)}(\overline{\mu}_{n-1},\mu_0).$$

Equation 17 implies that under the specification of equilibrium play, when the seller's belief is $\overline{\mu}_n$, he is indifferent between a posted price of $v_L + (1-\delta^n)\Delta v$ and a price of

$v_L + (1 - \delta^{n-1}\Delta v)$. This, in turn, implies that the buyer when her payoff is v_H may obtain any payoff in $[\delta^n \Delta v, \delta^{n-1}\Delta v]$, if the seller were to randomize between these two posted prices. This randomization is important for the specification of the buyer and the seller's strategies off the path of play. For future reference, let $\mathcal{U}_H^*$ denote the following correspondence:

$$\mathcal{U}_H^*(\mu_0) = \begin{cases} u_H^*(\mu_0) & \text{if } \mu_0 \neq \overline{\mu}_n, n \geq 1 \\ \left[\delta^n \Delta v, \delta^{n-1}\Delta v\right] & \text{if } \mu_0 = \overline{\mu}_n \text{ for some } n \geq 1 \end{cases}. \tag{21}$$

Note $\mathcal{U}_H^*$ is upper-hemicontinuous, convex-valued, and compact-valued.

B.2 Omitted proofs from Section 5.1

Section B.2 completes the steps to show that the set of equilibrium payoffs of $G^\infty(\mu_0)$ is as described in Equation 2. In what follows, we first establish that Equation 2 holds for $\mu_0 < \overline{\mu}_1$. We then establish that $u_S^*(\mu_0)$ is a lower bound on the seller's equilibrium payoff (Proposition B.1). Finally, we provide the omitted proofs of the statements in Section 5.1 regarding the properties of the solutions to OPT (Lemma B.2).

Equation 2 holds for $\mu_0 < \overline{\mu}_1$: This case is simple. First we need to show that the seller's highest equilibrium payoff is v_L for $\mu_0 < \overline{\mu}_1$ which is immediate since this is the seller's payoff in the commitment solution and there exists a PBE assessment $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$ which achieves it. Very similar arguments to those in Proposition B.1 imply that the seller's payoff cannot be below $u_S^*(\mu_0)$.

Equation 2 holds for $\mu_0 \geq \overline{\mu}_1$: In the rest of this section, we focus on the case in which $\mu_0 \geq \overline{\mu}_1$. As in the main text, let $n \geq 1$ denote the smallest m such that there exists $\mu_0 \in [\overline{\mu}_m, \overline{\mu}_{m+1})$ such that Equation 2 does not hold. To complete the proof in Section 5.1, we first show that $u_S^*(\mu_0)$ is a lower bound on the seller's equilibrium payoff (Proposition B.1). We then complete the steps to show that when $\mu_0 \in [\overline{\mu}_n, \overline{\mu}_{n+1})$, the solution to OPT places weight at most on $\{\overline{\mu}_{n-2}, \overline{\mu}_{n-1}\}$ (Lemma B.2):

Proposition B.1 shows that the seller can always guarantee $u_S^*(\mu_0)$:

Proposition B.1 ($u_S^*(\mu_0)$ is a lower bound)**.** *Let $n \geq 1$ denote the minimum $m \geq 1$ such that there exists $\mu_0 \in [\overline{\mu}_m, \overline{\mu}_{m+1})$ such that Equation 2 fails to hold. Let $\mu_0 \in [\overline{\mu}_n, \overline{\mu}_{n+1})$ and let $u \in \mathcal{E}^*(\mu_0)$. Then, $u_S \geq u_S^*(\mu_0)$.*

Proof. Toward a contradiction, let $\underline{u} \equiv (0, \underline{u}_H, \underline{u}_S) \in \mathcal{E}^*(\mu_0)$ be such that $\underline{u}_S < u_S^*(\mu_0)$. Fix $\epsilon, \gamma > 0$ and consider the following mechanism, $\mathbf{M}_0'$:

$$\beta^{\mathbf{M}_0'}(\overline{\mu}_{n-1}+\gamma|v_H) = \frac{\overline{\mu}_{n-1}+\gamma}{\mu_0}\frac{1-\mu_0}{1-(\overline{\mu}_{n-1}+\gamma)}, \beta^{\mathbf{M}_0'}(1|v_H) = \frac{1}{\mu_0}\frac{\mu_0-(\overline{\mu}_{n-1}+\gamma)}{1-(\overline{\mu}_{n-1}+\gamma)}$$
$$\beta^{\mathbf{M}_0'}(\overline{\mu}_{n-1}+\gamma|v_L) = 1$$
$$q^{\mathbf{M}_0'}(1) = 1, x^{\mathbf{M}_0'}(1) = v_L + (1-\delta^n)\Delta v - \epsilon,$$
$$q^{\mathbf{M}_0'}(\overline{\mu}_{n-1}+\gamma) = 0, x^{\mathbf{M}_0'}(\overline{\mu}_{n-1}+\gamma) = -\epsilon/2.$$

Pick γ so that $\overline{\mu}_{n-1}+\gamma \in [\overline{\mu}_{n-1}, \overline{\mu}_n)$. Then, the seller's continuation payoff is $u_S^*(\overline{\mu}_{n-1}+\gamma)$.

We claim that equilibrium considerations imply that the buyer has to accept $\mathbf{M}_0'$ with probability one for both of her values. By Lemma 3, the low-valuation buyer obtains a payoff of 0 by rejecting the mechanism. Instead, by participating the buyer can guarantee a payoff of $\epsilon/2$. Thus, sequential rationality of the buyer's strategy when her value is v_L implies that $\pi_{v_L}(\mathbf{M}_0') = 1$. Suppose then that v_H in equilibrium rejects such an offer. In this case, Bayes' rule where possible implies that the seller must assign probability 1 to v_H, in which case the buyer's continuation value is 0 (Lemma A.1). Then, it cannot be a best response to reject the mechanism with positive probability since $\epsilon/2 > 0$ is a lower bound on the buyer's payoff from participating in the mechanism. Thus, in the PBE assessment that sustains $\underline{u}$ the buyer must accept *any* such mechanism and beliefs must be specified so that participation is indeed a best response. Note that the latter is always feasible since assigning probability one to the rejection coming from v_H justifies 0 continuation values.

Furthermore, as long as $\epsilon < 2(1-\delta^n)\Delta v$, the above mechanism gives the buyer a strict incentive to tell the truth. Thus, the seller's payoff from a deviation is lower bounded by

$$-\epsilon + \frac{\mu_0-(\overline{\mu}_{n-1}+\gamma)}{1-(\overline{\mu}_{n-1}+\gamma)}(v_L+(1-\delta^n)\Delta v) + \frac{1-\mu_0}{1-(\overline{\mu}_{n-1}+\gamma)}\delta u_S^*(\overline{\mu}_{n-1}+\gamma) =$$
$$-\epsilon + \frac{\mu_0-(\overline{\mu}_{n-1}+\gamma)}{1-(\overline{\mu}_{n-1}+\gamma)}v_H + \frac{1-\mu_0}{1-(\overline{\mu}_{n-1}+\gamma)}\times$$
$$\delta\left[u_S^*(\overline{\mu}_{n-1}+\gamma) + (1-(\overline{\mu}_{n-1}+\gamma))\left(\frac{\overline{\mu}_{n-1}+\gamma}{1-(\overline{\mu}_{n-1}+\gamma)} - \frac{\mu_0}{1-\mu_0}\right)\delta^{n-1}\Delta v\right]$$
$$= -\epsilon + [\frac{\mu_0-(\overline{\mu}_{n-1}+\gamma)}{1-(\overline{\mu}_{n-1}+\gamma)}v_H + \frac{1-\mu_0}{1-(\overline{\mu}_{n-1}+\gamma)}\delta R^{(\tau^*,q^*)}(\overline{\mu}_{n-1}+\gamma,\mu_0)].$$

Note that at $\gamma = 0$, the above is precisely $u_S^*(\mu_0) - \epsilon$ (recall Equation 20), so that for ϵ, γ small enough we have that

$$-\epsilon + [\frac{\mu_0 - (\overline{\mu}_{n-1} + \gamma)}{1 - (\overline{\mu}_{n-1} + \gamma)} v_H + \frac{1 - \mu_0}{1 - (\overline{\mu}_{n-1} + \gamma)} \delta R^{(\tau^*, q^*)}(\overline{\mu}_{n-1} + \gamma, \mu_0)] > \underline{u}_S.$$

Thus, the seller has a profitable deviation. This contradicts that there is a PBE assessment where the seller earns less than $u_S^*(\mu_0)$. □

To conclude the proof that $u_S^*(\mu_0)$ is an upper bound on the seller's payoff, we prove Lemma B.2. As in the main text, let $\mu^\star = \inf\{\mu_0 \in [\overline{\mu}_n, \overline{\mu}_{n+1}) : \text{Equation 2 does not hold}\}$. Thus, for $\mu_1 \in [0, \mu^\star)$, the seller's payoffs are given by $u_S^*(\mu_1)$. Furthermore, since $\mu_0 \geq \overline{\mu}_n$, the seller prefers to minimize the buyer's continuation payoff at $\{\overline{\mu}_m\}_{m \leq n}$. This implies that the objective function in Equation 8 can be written as:

$$\int_{[0,\mu^\star)} \left[\frac{\mu_0 - \mu_1}{1 - \mu_1} v_H + \frac{1 - \mu_0}{1 - \mu_1} \delta \underbrace{\left(u_S^*(\mu_1) + \left(\frac{\mu_1}{1 - \mu_1} - \frac{\mu_0}{1 - \mu_0} \right) u_H^*(\mu_1) \right)}_{R^{(\tau^*, q^*)}(\mu_1, \mu_0)} \right] G(d\mu_1) \tag{22}$$
$$+ \int_{[\mu^\star, \mu_0)} \left[\frac{\mu_0 - \mu_1}{1 - \mu_1} v_H + \frac{1 - \mu_0}{1 - \mu_1} \delta \left(u_S(\mu_1) + \left(\frac{\mu_1}{1 - \mu_1} - \frac{\mu_0}{1 - \mu_0} \right) u_H(\mu_1) \right) \right] G(d\mu_1)$$

Lemma B.2. *Suppose G maximizes the expression in Equation 22 and is such that $G([0, \mu^\star)) > 0$. Then, G places positive probability on at most $\{\overline{\mu}_{n-2}, \overline{\mu}_{n-1}\}$.*

Proof. It is immediate to see that no $\mu_1 \in \cup_{m=0}^{n-1} (\overline{\mu}_m, \overline{\mu}_{m+1})$ can be on the support of G: if $\mu_1 \in [\overline{\mu}_m, \overline{\mu}_{m+1})$ for $m \leq n-1$, this is dominated by choosing $\overline{\mu}_m$:[26]

$$\frac{\mu_0 - \overline{\mu}_m}{1 - \overline{\mu}_m} v_H + \frac{1 - \mu_0}{1 - \overline{\mu}_m} \delta R^{(\tau^*, q^*)}(\overline{\mu}_m, \mu_0) - \left(\frac{\mu_0 - \mu_1}{1 - \mu_1} v_H + \frac{1 - \mu_0}{1 - \mu_1} \delta R^{(\tau^*, q^*)}(\mu_1, \mu_0) \right)$$
$$= v_H \left[\frac{\mu_0 - \overline{\mu}_m}{1 - \overline{\mu}_m} + \delta \frac{1 - \mu_0}{1 - \overline{\mu}_m} \frac{\overline{\mu}_m - \overline{\mu}_{m-1}}{1 - \overline{\mu}_{m-1}} - \frac{\mu_0 - \mu_1}{1 - \mu_1} - \delta \frac{1 - \mu_0}{1 - \mu_1} \frac{\mu_1 - \overline{\mu}_{m-1}}{1 - \overline{\mu}_{m-1}} \right] > 0$$

A similar argument implies that $\mu_1 \in [\overline{\mu}_n, \mu^\star)$ is dominated by choosing $\overline{\mu}_{n-1}$. Thus, if it is optimal to set $G([0, \mu^\star)) > 0$, we can reduce the problem of finding the optimal such G

[26] While the algebra is tedious, the intuition is immediate: μ_1 in $(\overline{\mu}_m, \overline{\mu}_{m+1})$ generate the same probability of trading with v_L as $\overline{\mu}_m$, but a lower probability of trading with v_H.

to

$$\max_{G\in\Delta(\{\overline{\mu}_0,\ldots,\overline{\mu}_{n-1}\})} \sum_{m=0}^{n-1} \left[\frac{\mu_0-\overline{\mu}_m}{1-\overline{\mu}_m} v_H + \frac{1-\mu_0}{1-\overline{\mu}_m} \delta R^{(\tau^*,q^*)}(\overline{\mu}_m,\mu_0) \right] G(\overline{\mu}_m) \tag{23}$$

However, for $m \le n-2$, Lemma B.1 implies that

$$\begin{aligned}
&\Delta_{m+1}(\mu_0;\overline{\mu}_m,\overline{\mu}_{m-1}) \\
&= \left[\frac{\mu_0-\overline{\mu}_m}{1-\overline{\mu}_m} v_H + \frac{1-\mu_0}{1-\overline{\mu}_m} \delta R^{(\tau^*,q^*)}(\overline{\mu}_m,\mu_0) \right] - \left[\frac{\overline{\mu}_{m-1}-\mu_0}{1-\overline{\mu}_{m-1}} v_H + \frac{1-\mu_0}{1-\overline{\mu}_{m-1}} \delta R^{(\tau^*,q^*)}(\overline{\mu}_{m-1},\mu_0) \right] > 0,
\end{aligned}$$

since $\mu_0 > \overline{\mu}_{n-1}$. Thus, any solution to the problem in Equation 23 satisfies that there exists some $\alpha \in [0,1]$ such that the value of this problem is given by:

$$\begin{aligned}
&= \alpha \left[\frac{\mu_0-\overline{\mu}_{n-1}}{1-\overline{\mu}_{n-1}} v_H + \frac{1-\mu_0}{1-\overline{\mu}_{n-1}} \delta R^{(\tau^*,q^*)}(\overline{\mu}_{n-1},\mu_0) \right] + (1-\alpha) \left[\frac{\mu_0-\overline{\mu}_{n-2}}{1-\overline{\mu}_{n-2}} v_H + \frac{1-\mu_0}{1-\overline{\mu}_{n-2}} \delta R^{(\tau^*,q^*)}(\overline{\mu}_{n-2},\mu_0) \right] \\
&= \alpha \Delta_n(\mu_0;\overline{\mu}_{n-1},\overline{\mu}_{n-2}) + \left[\frac{\mu_0-\overline{\mu}_{n-2}}{1-\overline{\mu}_{n-2}} v_H + \frac{1-\mu_0}{1-\overline{\mu}_{n-2}} \delta R^{(\tau^*,q^*)}(\overline{\mu}_{n-2},\mu_0) \right]
\end{aligned}$$

so that unless $\mu_0 = \overline{\mu}_n$ it is not optimal to set $\alpha < 1$. Instead, for $\mu_0 = \overline{\mu}_n$ any $\alpha \in [0,1]$ is a maximizer. □

The argument in Section 5.1 implies that if $\mu_0 \in [\overline{\mu}_n, \overline{\mu}_{n+1})$, then any G that maximizes the expression in Equation 22 is such that $G([\overline{\mu}_n, \mu_0)) = 0$. Thus, the value of OPT is $u_S^*(\mu_0)$. This completes the proof that the seller's equilibrium payoff is unique and coincides with that under $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v\in V}, \mu^* \rangle$.

C $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v\in V}, \mu^* \rangle$ IS AN EQUILIBRIUM ASSESSMENT

To complete the proof of Theorem 1, it remains to show that $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v\in V}, \mu^* \rangle$ is an equilibrium assessment. First, we finalize the construction of the assessment $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v\in V}, \mu^* \rangle$, by constructing the buyer's (Section C.1) and seller's strategy profile and system of beliefs (Section C.2) after every history. Section C.3 shows that neither the seller nor the buyer have one-shot deviations from the equilibrium strategy, given the continuation values implied by $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v\in V}, \mu^* \rangle$. The results in Athey and Bagwell (2008) imply that this is enough to conclude we have indeed constructed a PBE of $G^\infty(\mu_0)$.

In what follows, to simplify notation we denote by $\gamma^*(\mu)$ the mechanism $(\beta_\mu^*, q_\mu^*, x_\mu^*)$ that the seller uses in the posted-prices assessment.

C.1 Completing the buyer's strategy

To complete the buyer's strategy, for a mechanism $\mathbf{M}$ let

$$U_{1v}(\mathbf{M}) = \max_{\rho \in \Delta(V)} \sum_{v' \in V} \rho(v') \int_{\Delta(V) \times A} (vq - x)\varphi^{\mathbf{M}}(d(\mu, q, x)|v'), \tag{24}$$

denote the buyer's maximum payoff when her value is v from participating in mechanism $\mathbf{M}$, without taking into account her continuation payoffs.

Fix a public history h^t and let μ_t denote the seller's beliefs at that public history.[27] To complete the buyer's strategy, we classify mechanisms, $\mathbf{M}$, in four categories:[28]

0. Mechanisms that given the continuation values, $u_H^*(\cdot)$, satisfy participation and truthtelling. Denote the set of these mechanisms $\mathcal{M}_C^0$.

1. Mechanisms $\mathbf{M}$ not in $\mathcal{M}_C^0$ such that $U_{1,v_H}(\mathbf{M}) < 0$. Let $\mathcal{M}_C^1$ denote the set of these mechanisms.

2. Mechanisms not in $\mathcal{M}_C^0$ such that $U_{1v_H}(\mathbf{M}) \geq 0 > U_{1v_L}(\mathbf{M})$. Let $\mathcal{M}_C^2$ denote the set of these mechanisms.

3. Mechanisms not in $\mathcal{M}_C^0$ such that $U_{1v}(\mathbf{M}) \geq 0$ for $v \in \{v_L, v_H\}$. Let $\mathcal{M}_C^3$ denote the set of these mechanisms.

If $\mathbf{M} \in \mathcal{M}_C^1$, specify that the buyer rejects the mechanism for both of her values. Hence, under this strategy, the seller does not update his beliefs after observing a rejection. If, however, the buyer accepts, the seller believes $v = v_H$. Note that, in this case, continuation payoffs for the buyer are 0 from then on, regardless of her value. For each type v, let $r_v^*(\mathbf{M}, 1)$ denote a maximizer of Equation 24.[29]

If $\mathbf{M} \in \mathcal{M}_C^2$, specify that the buyer rejects $\mathbf{M}$ when her value is v_L. Hence, without loss of generality, we can specify that if the seller observes that the buyer accepts the mechanism, the buyer's value is v_H. For v_H, let $r_{v_H}^*(\mathbf{M}, 1)$ denote a maximizer of Equation 24 for $v = v_H$. Note we are using that the buyer's continuation payoffs are 0 conditional on her accepting the mechanism, so that $U_{1v_H}(\mathbf{M})$ is indeed the utility of the buyer when

[27] The belief, μ_t, is an equilibrium object, but we suppress this from the notation to keep things simple.

[28] Gerardi and Maestri (2020) use a similar trick to complete the worker's strategy in their paper.

[29] Even if the buyer does not participate on the equilibrium path, we still need to guarantee the reporting strategy is sequentially rational.

her value is v_H and the seller offers mechanism $\mathbf{M}$ in the game. Then, v_H's payoff from participating in the mechanism $\mathbf{M}$ is given by $U_{1v_H}(\mathbf{M})$ whereas the payoff from rejecting is $U_{0,v_H}(\pi_{v_H}, f) = \delta f(\nu_2(\mu_t, \pi_{v_H}))$, where

$$\nu_2(\mu_t, \pi_{v_H}) = \frac{\mu_t(1-\pi_{v_H})}{\mu_t(1-\pi_{v_H}) + 1 - \mu_t}.$$

is the seller's belief that the buyer's value is v_H when observing a rejection, according to Bayes' rule, and $f(\nu_2(\mu_t, \pi_{v_H}))$ is a measurable selection from $\mathcal{U}_H^*(\nu_2(\mu_t, \pi_{v_H}))$. Note the payoff from rejecting is specified under the assumption that in the continuation, the equilibrium path coincides with that of the posted prices equilibrium when beliefs are $\nu_2(\mu_t, \pi_{v_H})$. We use, however, a selection from $\mathcal{U}_H^*$ to ensure that if needed, the seller randomizes between the posted prices when indifferent to help make the buyer's continuation problem well-behaved.

Now, $(\pi^*_{v_H}(\mathbf{M}), f_2^*(\mathbf{M}))$ are chosen so that

$$\pi^*_{v_H} \in \arg\max_{p\in[0,1]} (1-p)U_{0,v_H}(\pi^*_{v_H}, f_2^*(\mathbf{M})) + pU_{1,v_H}(\mathbf{M})$$

The main result in Simon and Zame (1990) implies a solution to the above problem exists, given the properties of $\mathcal{U}_H^*$ and the linearity in p of the objective. Let $\pi^*_{v_H}(\mathbf{M})$ denote this fixed point. Now, if $\pi^*_{v_H}(\mathbf{M}) < 1$ and $\nu_2(\mu_t, \pi^*_{v_H}(\mathbf{M})) = \overline{\mu}_i$ for some $i \geq 1$, then there exists a weight $\phi_2(\overline{\mu}_i, \mathbf{M}) \in [0,1]$ that solves the following:

$$f_2^*(\mathbf{M}, \overline{\mu}_i) = (1-\phi)\delta^{i-1}\Delta v + \phi\delta^i \Delta v. \tag{25}$$

This weight captures the probability with which the seller, when his belief is $\overline{\mu}_i$, mixes between $\gamma^*(\overline{\mu}_i)$ and a mechanism that splits the prior $\overline{\mu}_i$ between 1 and $\overline{\mu}_{i-2}$, which we denote by $\gamma^{**}(\overline{\mu}_i)$.[30] For a mechanism in $\mathcal{M}_C^2$, let $r^*_{v_L}(\mathbf{M}, 1)$ denote a solution to Equation 24 for $v = v_L$.

Finally, if $\mathbf{M} \in \mathcal{M}_C^3$, specify that the buyer participates for both her values. If the seller observes that the buyer rejects the mechanism, he assigns probability 1 to the buyer's valuation being v_H. Thus, upon rejection, the buyer's continuation payoff is 0 indepen-

[30] Formally, if $i = 1$, $\gamma^{**}(\overline{\mu}_i)$ is the mechanism that sells the good with probability 1 at a price of v_L, that is, $\beta^{\mathbf{M}}(\overline{\mu}_i | v_L) = \beta^{\mathbf{M}}(\overline{\mu}_i | v_H) = 1$ and $(q^{\mathbf{M}}(\overline{\mu}_i), x^{\mathbf{M}}(\overline{\mu}_i) = (1, v_L)$. For $i \geq 2$, $\gamma^{**}(\overline{\mu}_i)$ is the mechanism such that $\beta^{\mathbf{M}}(\overline{\mu}_{i-2} | v_L) = 1$, $\beta^{\mathbf{M}}(\overline{\mu}_{i-2} | v_H) = (\overline{\mu}_{i-2}/\overline{\mu}_i)(1-\overline{\mu}_i/1-\overline{\mu}_{i-2}) = 1 - \beta^{\mathbf{M}}(1|v_H)$ and sets $(q^{\mathbf{M}}(\overline{\mu}_{i-2}), x^{\mathbf{M}}(\overline{\mu}_{i-2})) = (0,0)$ and $(q^{\mathbf{M}}(1), x^{\mathbf{M}}(1)) = (1, v_L + (1-\delta^{i-1})\Delta v)$.

dently of her value. Note that mechanisms in $\mathcal{M}_C^3$ satisfy the participation constraint given the continuation values of the posted-prices assessment. Now, let m_L^* satisfy[31]

$$\int_{\Delta(V)\times A}(v_L q - x)\varphi^{\mathbf{M}}(d(\mu_{t+1}, q, x)|m_L^*) \geq \int_{\Delta(V)\times A}(v_L q - x)\varphi^{\mathbf{M}}(d(\mu_{t+1}, q, x)|v),$$

for all $v \in V$. Set $r_{v_L}^*(\mathbf{M}, 1)(m_L^*) = 1$. Let $\{m_H^*\} = V \setminus \{m_L^*\}$ and define $\Delta(V)^H \equiv \text{supp}\ \varphi^{\mathbf{M}}(\cdot \times A|m_H^*)$, $\Delta(V)^L \equiv \text{supp}\ \varphi^{\mathbf{M}}(\cdot \times A|m_L^*)$. Let $r \in [0,1]$ denote the weight the buyer assigns to m_L^* when her valuation is v_H. Given this notation, let $\nu_3(\mu_t, \mu_{t+1}, r)$ denote the seller's belief that the buyer's valuation is v_H when he observes output message μ_{t+1}. We assume that if μ_{t+1} is not consistent with m_L^*, i.e., $\mu_{t+1} \in \Delta(V)^H \setminus \Delta(V)^L$, then $\nu_3(\mu_t, \mu_{t+1}, r) = 1$. This specification of beliefs does not conflict with Bayes' rule where possible: either m_H^* has positive probability in the optimal reporting strategy of the buyer when her valuation is v_H, in which case, ν_3 would be consistent with Bayes' rule, or it does not, in which case, Bayes' rule where possible places no restrictions on $\nu_3(\mu_t, \mu_{t+1}, \cdot)$ for $\mu_{t+1} \in \Delta(V)^H \setminus \Delta(V)^L$.

Given ν_3, the buyer when her valuation is v_H obtains a payoff of

$$\hat{U}_{1,v_H}(r, m, f) = \int_{\Delta(V)\times A}(v_H q - x + \delta(1-q) f(\nu_3(\mu_t, \mu_{t+1}, r)))\varphi^{\mathbf{M}}(d(\mu_{t+1}, q, x)|m),$$

when she reports $m \in V$, where f is a selection from $\mathcal{U}_H^*$. We want to find $(r_{v_H}^*, f_3^*(\mathbf{M}))$ so that

$$r_{v_H}^* \in \arg\max_{r\in[0,1]} r\hat{U}_{1,v_H}(r_{v_H}^*, m_L^*, f_3^*(\mathbf{M})) + (1-r)\hat{U}_{1,v_H}(r_{v_H}^*, m_H^*, f_3^*(\mathbf{M})). \tag{26}$$

The main theorem in Simon and Zame (1990) implies the existence of such an $(r_{v_H}^*, f_3^*(\mathbf{M}))$. Set $r_{v_H}^*(\mathbf{M}, 1)(m_L^*) = r_{v_H}^*$. As we did before, whenever $\nu_3(\mu_t, \mu_{t+1}, r_{v_H}^*) = \overline{\mu}_i$ for $i \geq 1$, we can define $\phi_3(\mathbf{M}, \overline{\mu}_i)$ as the weight on $\gamma^*(\overline{\mu}_i)$ implied by $f_3^*(\mathbf{M}, \overline{\mu}_i)$.

Summing up, the buyer's strategy at any history h_B^t is given by

$$\pi_{tv}^*(h_B^t, \mathbf{M}) = \begin{cases} 1 & \text{if } \mathbf{M} \in \mathcal{M}_C^0 \cup \mathcal{M}_C^3 \\ \pi_{v_H}^*(\mathbf{M}) & \text{if } \mathbf{M} \in \mathcal{M}_C^2 \text{ and } v = v_H \\ 0 & \text{otherwise} \end{cases} \tag{27}$$

[31] We cannot ensure that truthtelling will hold for mechanisms in $\mathcal{M}_C^3$, which is why we need the extra piece of notation.

and

$$r^*_{tv}(h^t_B, \mathbf{M}, 1) = \begin{cases} \delta_v & \text{if } \mathbf{M} \in \mathcal{M}^0_C \\ r^*_v(\mathbf{M}, 1) & \text{otherwise} \end{cases} . \tag{28}$$

C.2 Full specification of the PBE assessment

To complete the PBE assessment for $G^\infty(\mu_0)$, we now specify the seller's strategy and the belief system. We introduce two pieces of notation: the first one allows us to keep track of the last payoff-relevant event; the second one allows us to keep track of how the seller's beliefs evolve given the buyer's strategy (Equation 29).

From the beginning of period t of the game until the end of that period, the following things are determined: (i) the seller's choice of mechanism, $\mathbf{M}$; (ii) the buyer's participation decision, $\hat{\pi} \in \{0, 1\}$; and (iii) if the buyer participates in the mechanism, the allocation and the output message. Let $z_\emptyset(\mathbf{M}_t)$ denote the tuple $(\mathbf{M}_t, 0, \emptyset, (0, 0))$ and $z_{(\mu', (q,x))}(\mathbf{M}_t)$ denote the tuple $(\mathbf{M}_t, 1, \mu', (q, x))$. Of course, if $q = 1$, the game ends. Note that any public history h^t can be written as (h^{t-1}, z) for some z as defined above. Given h^t, let $z(h^t)$ denote the corresponding outcome z. Given any prior μ_0, let $T(\mu_0, z)$ denote the map that associates to each prior belief μ_0 and each outcome z, a seller's belief that the buyer's valuation is v_H. Formally,

$$T(\mu_0, z) = \begin{cases} \mu' & \text{if } z = z_{(\mu',\cdot)}(\mathbf{M}) \text{ and } \mathbf{M} \in \mathcal{M}^0_C \\ 1 & \begin{array}{c} \text{if } z = z_\emptyset(\mathbf{M}) \text{ and } \mathbf{M} \in \mathcal{M}^0_C \cup \mathcal{M}^3_C \\ \text{or} \\ z = z_{(\mu',\cdot)}(\mathbf{M}) \text{ and } \mathbf{M} \in \mathcal{M}^1_C \cup \mathcal{M}^2_C \end{array} \\ \mu_0 & \text{if } z = z_\emptyset(\mathbf{M}) \text{ and } \mathbf{M} \in \mathcal{M}^1_C \\ \nu_2(\mu_0, \pi^*_{v_H}(\mathbf{M})) & \text{if } z = z_{(\emptyset)}(\mathbf{M}) \text{ and } \mathbf{M} \in \mathcal{M}^2_C \\ \nu_3(\mu_0, \mu', r^*_{v_H}(\mathbf{M}, 1), r^*_{v_L}(\mathbf{M}, 1)) & \text{if } z = z_{(\mu',\cdot)}(\mathbf{M}) \text{ and } \mathbf{M} \in \mathcal{M}^3_C \end{cases} . \tag{29}$$

Let μ_0 denote the seller's prior in $G^\infty(\mu_0)$. Define $\mu^*(\emptyset) = \mu_0$ and $\Gamma^*(\emptyset) = \gamma^*(\mu_0)$, where γ^* is the strategy in the posted-prices assessment $\langle \Gamma^*, (\pi^*_v, r^*_v)_{v \in V}, \mu^* \rangle$. For any public history h^t, define

$$\mu^*_t(h^t) = T(\mu^*_t(h^{t-1}), z(h^t))). \tag{30}$$

If either (i) $z = z_{(\mu',(0,\cdot))}(\mathbf{M})$ and $\mathbf{M} \in \mathcal{M}_C^0 \cup \mathcal{M}_C^1 \cup \mathcal{M}_C^2$, (ii) $z = z_{\emptyset}(\mathbf{M})$, $\mu^*(h^t) \notin \{\overline{\mu}_i\}_{i\geq 1}$ and $\mathbf{M} \in \mathcal{M}_C^2$, (iii) $z = z_{(\mu',(0,\cdot))}(\mathbf{M})$, $\mu^*(h^t) \notin \{\overline{\mu}_i\}_{i\geq 1}$ and $\mathbf{M} \in \mathcal{M}_C^3$, or (iv) $z = z_{\emptyset}(\mathbf{M})$ and $\mathbf{M} \in \mathcal{M}_C^0 \cup \mathcal{M}_C^1 \cup \mathcal{M}_C^3$, define

$$\Gamma_t^*(h^t)(\mathbf{M}) = \mathbb{1}[\mathbf{M} = \gamma^*(T(\mu_t^*(h^{t-1}), z(h^t))))]. \tag{31}$$

If $z(h^t) = z_{\emptyset}(\mathbf{M})$ for $\mathbf{M} \in \mathcal{M}_C^2$ and $\mu_t^*(h^t) = \overline{\mu}_i$, $i \geq 1$, define

$$\Gamma_t^*(h^t)(\mathbf{M}) = \begin{cases} \phi_2(\mathbf{M}, \overline{\mu}_i) & \mathbf{M} = \gamma^*(\overline{\mu}_i) \\ 1 - \phi_2(\mathbf{M}, \overline{\mu}_i) & \mathbf{M} = \gamma^{**}(\overline{\mu}_i) \\ 0 & \text{otherwise} \end{cases}. \tag{32}$$

Finally, if $z(h^t) = z_{(\mu',(0,\cdot))}(\mathbf{M})$ for $\mathbf{M} \in \mathcal{M}_C^3$ and $\mu_t^*(h^t) = \overline{\mu}_i$, $i \geq 1$, define

$$\Gamma_t^*(h^t)(\mathbf{M}) = \begin{cases} \phi_3(\mathbf{M}, \overline{\mu}_i) & \mathbf{M} = \gamma^*(\overline{\mu}_i) \\ 1 - \phi_3(\mathbf{M}, \overline{\mu}_i) & \mathbf{M} = \gamma^{**}(\overline{\mu}_i) \\ 0 & \text{otherwise} \end{cases}. \tag{33}$$

Equations 27-28, together with equations 30- 33, define a PBE assessment where the system of beliefs is derived from the strategy profile where possible.

C.3 Seller's and buyer's sequential rationality

We now check that at all histories, neither the seller nor the buyer have a one-shot deviation from the prescribed strategy profile, given the continuation values constructed in Section B.1 and Section C.1. That this is true for the buyer given the seller's strategy follows from the construction in Section C.1.

To verify that sequential rationality holds for the seller we proceed in two steps. First, we show that for every $\mathbf{M}_t$ not in $\mathcal{M}_C^0$, there exists a $\mathbf{M}_t' \in \mathcal{M}_C^0$ such that the revenue of offering $\mathbf{M}_t$ (under the buyer's strategy) coincides with the revenue from offering $\mathbf{M}_t'$. Second, we use this property to relate the seller's payoff from offering $\mathbf{M}_t$ to the virtual surplus of $\mathbf{M}_t'$ to show that the revenue from $\mathbf{M}_t$ is bounded above by $u_S^*(\mu^*(h^t))$.

Step 1: We now show how to obtain $\mathbf{M}_t'$ from $\mathbf{M}_t$ for the case in which $\mathbf{M}_t$ is in $\mathcal{M}_C^2$. The construction for $\mathbf{M}_t \in \mathcal{M}_C^1 \cup \mathcal{M}_C^3$ follows similar steps. Suppose that the seller with belief $\mu^*(h^t)$ offers a mechanism $\mathbf{M}_t$ in $\mathcal{M}_C^2$. Let $\pi_{v_H}^*(h_B^t, \mathbf{M}_t), r_{v_H}^*(h_B^t, \mathbf{M}_t, 1)$ denote the buyer's

best responses as constructed in Section C.1. Denote by $\nu_2 \equiv \nu_2(\mu(h^t), \pi^*_{v_H}(h^t_B, \mathbf{M}))$ the seller's belief that the buyer's value is v_H when he observes non-participation. The seller's payoff is then

$$\mu_0 \pi^*_{tv_H}(h^t_B, \mathbf{M}_t) \int_{\Delta(V)\times A} [x + \delta(1-q)v_H] \left(\sum_{v\in V} \varphi^{\mathbf{M}_t}(d(\mu,q,x)|v) r^*_{tv_H}(h^t_B, \mathbf{M}_t)(v) \right)$$
$$+ (1 - \mu^*(h^t)\pi^*_{tv_H}(h^t_B, \mathbf{M}_t))\delta u^*_S(\nu_2), \tag{34}$$

where the continuation values after rejection are constructed using the equilibrium strategy when the seller has posterior ν_2.[32] Now, consider the following alternative mechanism, $\mathbf{M}'_t$:

$$\begin{aligned}
\beta^{\mathbf{M}'}(1|v_H) &= \pi^*_{tv_H}(h^t_B, \mathbf{M}_t) \\
\beta^{\mathbf{M}'}(\nu_2|v_H) &= (1 - \pi^*_{tv_H}(h^t_B, \mathbf{M}_t)), \quad \beta^{\mathbf{M}'}(\nu_2|v_L) = 1 \\
q^{\mathbf{M}'_t}(\nu_2) &= x^{\mathbf{M}'_t}(\nu_2) = 0 \\
x^{\mathbf{M}'_t}(1) &= \int_{\Delta(V)\times A} x \left(\sum_{v\in V} \varphi^{\mathbf{M}_t}(d(\mu,q,x)|v) r^*_{tv_H}(h^t_B, \mathbf{M}_t)(v) \right) \\
q^{\mathbf{M}'_t}(1) &= \int_{\Delta(V)\times A} q \left(\sum_{v\in V} \varphi^{\mathbf{M}_t}(d(\mu,q,x)|v) r^*_{tv_H}(h^t_B, \mathbf{M}_t)(v) \right),
\end{aligned}$$

Note that if the buyer participates and truthfully reports her type, $\mathbf{M}'_t$ gives the seller a payoff equal to the expression in Equation 34. Then, we can write the payoff to mechanism, $\mathbf{M}_t$, as

$$U_S(h^t, \mathbf{M}_t) \equiv \sum_{\mu_{t+1}\in\{\nu_2,1\}} \tau^{\mathbf{M}'_t}(\mu_{t+1})[x^{\mathbf{M}'_t}(\mu_{t+1}) + \delta(1 - q^{\mathbf{M}'_t}(\mu_{t+1}))u^*_S(\mu_{t+1})].$$

Step 2: We now show that $U_S(h^t, \mathbf{M}_t)$ is bounded above by $u^*_S(\mu^*(h^t))$. Indeed, we have:

$$\begin{aligned}
U_S(h^t, \mathbf{M}_t) &= \int_{\Delta(V)} [x^{\mathbf{M}'_t}(\mu_{t+1}) + (1 - q^{\mathbf{M}'_t}(\mu_{t+1}))\delta u^*_S(\mu_{t+1})]\tau^{\mathbf{M}'_t}(d\mu_{t+1}) \\
&\le VS(\tau^{\mathbf{M}'_t}, q^{\mathbf{M}'_t}, (u^{\mathbf{M}}_H, u^*_S), \mu(h^t)) \\
&\le VS(\tau^{\mathbf{M}'_t}, q^{\mathbf{M}'_t}, (u^*_H, u^*_S)_{\mu_{t+1}<\mu(h^t)}, (u^{\mathbf{M}}_H, u^*_S)_{\mu_{t+1}\ge\mu(h^t)}), \mu(h^t)) \\
&\le \max_{(\tau,q),(u_H,u_S)_{\mu_{t+1}>\mu(h^t)}} VS(\tau^{\mathbf{M}'_t}, q^{\mathbf{M}'_t}, (u^*_H, u^*_S)_{\mu_{t+1}<\mu(h^t)}, (u_H, u_S)_{\mu_{t+1}\ge\mu(h^t)}), \mu(h^t)) = u^*_S(\mu(h^t))
\end{aligned}$$

[32] Recall that the seller's continuation payoff only depends on his belief about the buyer's valuation being v_H; it is only the high-valuation buyer's payoff which may be different from u^*_H off the path of play.

The first equality represents the payoff from offering $\mathbf{M}_t$ as the payoff from offering the mechanism $\mathbf{M}'_t$ that satisfies the incentive compatibility and participation constraints (note that the seller's continuation payoffs are those of the assessment $\langle \Gamma^*, (\pi^*_v, r^*_v)_{v\in V}, \mu^* \rangle$ even after a deviation). The first inequality follows from Lemma 2. Now, while the seller's continuation payoffs are given by u^*_S, the buyer's continuation payoffs need not be exactly those of $\langle \Gamma^*, (\pi^*_v, r^*_v)_{v\in V}, \mu^* \rangle$ when the seller's posterior belief is one of the threshold beliefs. The second inequality follows from noting that when $\mu_{t+1} < \mu(h^t)$, the term multiplying u_H in the virtual surplus is negative. Thus, the seller with belief $\mu(h^t)$ prefers that whenever his posterior belief μ_{t+1} is below $\mu(h^t)$, the high-valuation buyer's continuation payoff is that which minimizes her rents, i.e., $u^*_H(\mu_{t+1})$. The last inequality follows by definition. The proofs in Sections 4.2, 5.1, and B.2 imply the last equality. We conclude that the seller has no one-shot deviations.

D Perfect Bayesian equilibrium: formal statement

We introduce in this section the necessary formalisms to define PBE. To simplify notation, we assume in what follows that $\mathcal{M}$ is such that M_i is finite for all $i \in \mathcal{I}$. When V is finite, it follows from Doval and Skreta (2020) that this is without loss of generality. It affords two important simplifications. First, $\mathcal{M}$ is itself a Polish space, which means that we can condition the buyer's strategy directly on the mechanism, $\mathbf{M}$, chosen by the seller, as opposed to the realization of the randomization device in the construction in Aumann (1964). Second, given a public history h^t, the set of buyer histories consistent with h^t, $H^t_B(h^t)$, is finite and therefore the support of $\mu_t(h^t) \in \Delta(V \times H^t_B(h^t))$ is finite.
To define the seller and the buyer's payoffs, we introduce the following notation. Given a period t and a mechanism $\mathbf{M}_t$, we let $z_{(s_t,(q_t,x_t))}(\mathbf{M}_t)$ denote the tuple $(\mathbf{M}_t, 1, s_t, (q_t, x_t))$ and $z_\emptyset(\mathbf{M}_t)$ denote the tuple $(\mathbf{M}_t, 0, \emptyset, (0,0))$. Furthermore, given the buyer's participation and reporting strategy and a mechanism $\mathbf{M}_t$ define a new distribution over $\Delta(S^{\mathbf{M}_t} \times \{0,1\} \times \mathbb{R})$ so that

$$\rho^{(\pi,r)}(S' \times A' | v, h^t_B, \mathbf{M}_t) = \pi_{tv}(h^t_B, \mathbf{M}_t) \sum_{m \in M^{\mathbf{M}_t}} \varphi^{\mathbf{M}_t}(S' \times A' | m) r_{tv}(h^t_B, \mathbf{M}_t, 1)(m)$$

for all measurable subsets $S' \times A' \subseteq S^{\mathbf{M}_t} \times \{0,1\} \times \mathbb{R}$.
Fix an assessment $\langle \Gamma, (\pi_v, r_v)_{v\in V}, \mu \rangle$, a public history h^t and a mechanism, $\mathbf{M}_t$. The

seller's payoffs are given by

$$U_S(\Gamma,(\pi_v,r_v)_{v\in V},\mu|h^t,\mathbf{M}_t) = \sum_{(v,h_B^t)} \mu_t(h^t)(v,h_B^t)(1-\pi_{tv}(h_B^t,\mathbf{M}_t))\delta\mathbb{E}_\Gamma\left[U_S(\Gamma,(\pi_v,r_v)_{v\in V},\mu|h^t,z_\emptyset(\mathbf{M}_t),\cdot)\right]$$
$$+\sum_{(v,h_B^t)} \mu_t(h^t)(v,h_B^t)\int_{S^{\mathbf{M}_t}\times\{0,1\}\times\mathbb{R}}\left(x+(1-q)\delta\mathbb{E}_\Gamma\left[U_S(\Gamma,(\pi_v,r_v)_{v\in V},\mu|h^t,z_{(s_t,(0,x))}(\mathbf{M}_t),\cdot)\right]\right)\rho^{(\pi,r)}(d(s_t,q,x)|v,h_B^t,\mathbf{M}_t).$$

Similarly, the buyer's payoffs when her valuation is v, the history is h_B^t, and the seller offers mechanism $\mathbf{M}_t$ are given by

$$U_v(\Gamma,(\pi_v,r_v),\mu|h_B^t,\mathbf{M}_t) = (1-\pi_{tv}(h_B^t,\mathbf{M}_t))\delta\mathbb{E}_\Gamma\left[U_v(\Gamma,(\pi_v,r_v),\mu|h_B^t,z_\emptyset(\mathbf{M}_t),\cdot)\right]+\pi_{tv}(h_B^t,\mathbf{M}_t)\times$$
$$\sum_{m\in M^{\mathbf{M}_t}} r_{tv}(h_B^t,\mathbf{M}_t,1)(m)\int_{S^{\mathbf{M}_t}\times A}\left(vq-x+(1-q)\delta\mathbb{E}_\Gamma\left[U_v(\Gamma,(\pi_v,r_v)_{v\in V},\mu|h_B^t,m,z_{(s_t,(q,x))}(\mathbf{M}_t),\cdot)\right]\right)\varphi^{\mathbf{M}_t}(d(s_t,q,x)|m)$$

Definition 1. *The assessment $\langle\Gamma,(\pi_v,r_v)_{v\in V},\mu\rangle$ satisfies sequential rationality if for all periods t, and all public histories h^t, we have*

1. *For all mechanisms $\mathbf{M}_t$ in the support of $\Gamma_t(h^t)$, $U_S(\Gamma,(\pi_v,r_v)_{v\in V},\mu|h^t,\mathbf{M}_t)\geq U_S(\Gamma,(\pi_v,r_v)_{v\in V},\mu|h^t,\mathbf{M}'_t)$ for all $\mathbf{M}'_t\neq\mathbf{M}_t$,*
2. *For all $v\in V$, all buyer histories $h_B^t\in H_B^t(h^t)$, and all mechanisms $\mathbf{M}_t$, $U_v(\Gamma,(\pi_v,r_v),\mu|h_B^t,\mathbf{M}_t)\geq U_v(\Gamma,(\pi'_v,r'_v),\mu|h_B^t,\mathbf{M}_t)$ for all alternative strategies (π'_v,r'_v).*

Definition 2. *The belief system μ satisfies Bayes' rule where possible if for all public histories h^t, and mechanisms $\mathbf{M}_t$, the following hold,*

$$\mu_{t+1}(h^t,z_\emptyset(\mathbf{M}_t))(v_H,h_B^t,z_\emptyset(\mathbf{M}_t))\left[\sum_{v\in V,h_B^t\in H_B^t(h^t)}\mu_t(h^t)(v,h_B^t)(1-\pi_{tv}(h_B^t,\mathbf{M}_t))\right]=\mu_t(h^t)(v_H,h_B^t)(1-\pi_{tv_H}(h_B^t,\mathbf{M}_t)),$$

and, for all measurable subsets $S'\times A'$ of $S^{\mathbf{M}_t}\times A$

$$\sum_{(v,h_B^t)}\mu_t(h^t)(v,h_B^t)\int_{S'\times A'}\mu_{t+1}(h^t,\cdot)(\overline{v},\overline{h}_B^t,z_{(s_t,(q_t,x_t))}(\mathbf{M}_t),\overline{m})\rho^{(\pi,r)}(d(s_t,q,x)|v,h_B^t)$$
$$=\mu_t(h^t)(\overline{v},\overline{h}_B^t)\pi_{t\overline{v}}(\overline{h}_B^t,\mathbf{M}_t)r_{t\overline{v}}(\overline{h}_B^t,\mathbf{M}_t,1)(\overline{m})\varphi^{\mathbf{M}_t}(S'\times A'|\overline{m}).$$

Definition 3. *An assessment $\langle\Gamma,(\pi_v,r_v)_{v\in V},\mu\rangle$ is a Perfect Bayesian equilibrium if it is sequentially rational and satisfies Bayes' rule where possible.*