

Doval, Laura; Skreta, Vasiliki

Article

Optimal mechanism for the sale of a durable good

Theoretical Economics

Provided in Cooperation with:

The Econometric Society

Suggested Citation: Doval, Laura; Skreta, Vasiliki (2024) : Optimal mechanism for the sale of a durable good, Theoretical Economics, ISSN 1555-7561, The Econometric Society, New Haven, CT, Vol. 19, Iss. 2, pp. 865-915,
<https://doi.org/10.3982/TE4485>

This Version is available at:

<https://hdl.handle.net/10419/320255>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<https://creativecommons.org/licenses/by-nc/4.0/>

Optimal mechanism for the sale of a durable good

LAURA DOVAL

Economics Division, Columbia Business School, Columbia University and CEPR

VASILIKI SKRETA

Department of Economics, University of Texas at Austin, Department of Economics, University College London, and CEPR

A buyer wishes to purchase a durable good from a seller who in each period chooses a mechanism under limited commitment. The buyer's value 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 \(2022\)](#).

KEYWORDS. Mechanism design, limited commitment, information design, public PBE, posted prices, Coase conjecture.

JEL CLASSIFICATION. D84, D86.

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 mech-

Laura Doval: laura.doval@columbia.edu

Vasiliki Skreta: vskreta@gmail.com

We thank 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 (Doval: SES-2131706; Skreta: SES-1851729).

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

anisms, 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. 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. Whereas 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 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 value being high.

Second, we argue that the seller's payoff under the posted-prices assessment is an upper bound on the seller's equilibrium payoff in the mechanism-selection game. 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 (2022). The main theorem in Doval and Skreta (2022) allows us to simplify the class of mechanisms the seller offers in any equilibrium of the game and the buyer's equilibrium behavior. This step reduces the search for the optimal sequence of mechanisms to those that satisfy, loosely speaking, a sequence of *participation* and *truth-telling* constraints, allowing us for the most part to ignore the buyer as a player. Like in standard mechanism design, the low-valuation buyer's utility and the high-valuation buyer's truth-telling constraint determine an upper bound on the revenue the seller can extract within a period (Lemma 2). Replacing this upper bound in the seller's payoff provides us with a dynamic analogue of the virtual surplus (Equation (4)), where the seller's payoff is written as a function of the allocation, but also the continuation payoffs.

We show that the value of (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. Because the auxiliary program (OPT) ignores the truth-telling constraint of the low-valuation buyer (i.e., it corresponds to the *relaxed* program in mechanism design), our result implies that the solution to (OPT) satisfies the remaining constraints and can thus be implemented as an equilibrium outcome. 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, even if in some settings the solution to the analogue of (OPT) may not deliver an implementable outcome.

Finally, following the logic in Gul, Sonnenschein, and Wilson (1986), we show that the seller's payoff in the posted-prices assessment is a lower bound on the seller's equilibrium payoff. Underlying the argument in Gul, Sonnenschein, and Wilson (1986) that a unique equilibrium payoff exists in the gap case is the property that the minimum price the seller chooses in equilibrium imposes an upper bound on the maximum payoff the buyer can obtain. Relying once again on (OPT), we establish that guarantees on the seller's equilibrium payoff translate into upper bounds on the high-valuation buyer's payoff. Armed with this result, we show that the seller can always undercut the price in the posted-prices assessment and earn close to his payoff in that assessment.

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. Because 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, e.g., the seminal work of Bester and Strausz (2001, 2007), and the discussion in the related literature), such characterization has proved elusive. In Doval and Skreta (2022), we provide a 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 we characterize is the optimal one among *all* mechanisms the seller could have offered the buyer under limited commitment.

Second, the optimality of posted prices should not be taken for granted, even if it is evocative of Skreta (2006). First, our model is not an infinite-horizon version of that in Skreta (2006), since we consider a larger class of mechanisms than Skreta (2006). Indeed, the mechanisms in Skreta (2006) presume that the seller must observe the buyer's input message (cf. Laffont and Tirole (1988), Bester and Strausz (2001)), whereas we consider mechanisms in which the seller gets to design how much he observes about the buyer's input message. In Doval and Skreta (2022), 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 (Remark 3 explains why posted prices may fail to be optimal in Doval and Skreta (2022)). 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 (2022) 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 \(2017\)](#)).² Underlying these restrictions are 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 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.³ Related to this literature is the problem of dynamic bargaining with one-sided incomplete information⁴ (e.g., [Sobel and Takahashi \(1983\)](#), [Fudenberg, Levine, and Tirole \(1985\)](#), [Ausubel and Deneckere \(1989\)](#)). In all these papers, the uninformed party's inability to commit limits his bargaining power.

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 \(2022\)](#) 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](#).

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

²[Beccuti and Möller \(2018\)](#) lies somewhat in between these two strands because they take limits of their finite-horizon results to draw conclusions about the infinite-horizon game. They do not show that this limit corresponds to the seller's revenue-maximizing equilibrium in the infinite-horizon game.

³Other 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. Relatedly, [Brzustowski, Georgiadis-Harris, and Szentes \(2023\)](#) show that smart contracts help the seller avoid the Coase conjecture.

⁴Recently, [Peski \(2022\)](#) studies alternating bargaining games where players can offer menus.

private information: before her interaction with the seller starts, she observes her value $v \in \{v_L, v_H\} \equiv V$, 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 value 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.⁵ The game ends the first time the good is sold.

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 seller's action space. In each period, the seller offers the buyer a mechanism. Following Doval and Skreta (2022), 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 $S^{\mathbf{M}} \times A$.⁶ 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 $\Delta(V)$.⁷ Denote by $\mathcal{M}_{\mathcal{I}}$ the set of all mechanisms with message sets $(M_i, S_i)_{i \in \mathcal{I}}$.⁸

Mechanism-selection game: The seller's prior μ_0 and the collection $(M_i, S_i)_{i \in \mathcal{I}}$ define a mechanism-selection game, denoted $G_{\mathcal{I}}(\mu_0)$, as follows. In each period t , as long as the good has not been sold, the game proceeds as follows. First, the seller and the buyer observe the realization of a public randomization device, $\omega \sim U[0, 1]$. Second, 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 $\varphi^{\mathbf{M}}(\cdot | m)$ and are observed by both the seller and the buyer. If the good is not sold, the game proceeds to period $t + 1$.

Histories: The game $G_{\mathcal{I}}(\mu_0)$ has two types of histories: public and private. Public histories capture what the *seller* knows through period t : the past realizations of the

⁵Even if 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.

⁶Throughout, 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 $\zeta : X \rightarrow \Delta(Y)$ 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 $x \in X$.

⁷Because V is finite, taking $M^{\mathbf{M}}$ to be finite is without loss of generality (Doval and Skreta, 2022).

⁸We 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 (2022) shows that the choice of the collection plays no further role in the analysis.

public randomization device, his past choices of mechanisms, the buyer's participation decisions, and the realized output messages and allocations. We let h^t denote a public history through period t and let H^t denote the set of all such histories. Instead, private histories capture what the *buyer* knows through period t . First, the buyer knows the public history of the game and her input messages into the mechanism (henceforth, a buyer history). Second, the buyer also knows her private information. We let h_B^t denote a buyer history through period t and let $H_B^t(h^t)$ denote the set of buyer histories consistent with public history h^t . Thus, $V \times H_B^t(h^t)$ denotes the set of private histories consistent with public history h^t .

Strategies and beliefs: A behavioral strategy for the seller is a collection of measurable mappings $\Gamma \equiv (\Gamma_t)_{t=0}^\infty$, where for each period t and each public history h^t , $\Gamma_t(h^t)$ describes the seller's (possibly random) choice of mechanism at h^t .⁹ Similarly, a behavioral strategy for the buyer is a collection of measurable mappings $(\pi_t(v, \cdot), r_t(v, \cdot))_{t=0}^\infty$, where for each period t , each private history (v, h_B^t) , and each mechanism, $\mathbf{M}_t$, $\pi_t(v, h_B^t, \mathbf{M}_t)$ describes the buyer's participation decision, whereas $r_t(v, h_B^t, \mathbf{M}_t)$ describes the buyer's choice of input messages in the mechanism, conditional on participation. 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 $\mu_t(h^t) \in \Delta(V \times H_B^t(h^t))$. 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 E contains the formal statement. For now, we note that if the seller's strategy space was finite and the mechanisms used by the seller had finite support, then this coincides with the definition in [Fudenberg and Tirole \(1991b\)](#).¹⁰

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 a PBE assessment $\langle \Gamma, (\pi_v, r_v)_{v \in V}, \mu \rangle$ exists that induces η . We denote by $\mathcal{O}_T^*(\mu_0)$ the set of PBE outcomes and by $\mathcal{E}_T^*(\mu_0) \subseteq \mathbb{R}^3$ the set of PBE payoffs of $G_T(\mu_0)$. We denote a generic element of $\mathcal{E}_T^*(\mu_0)$ 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 $\mathcal{O}_T^*(\mu_0)$ and $\mathcal{E}_T^*(\mu_0)$. 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.

⁹As we explain in Appendix E, $\mathcal{M}_T$ is a Polish space—which we endow with its Borel σ -algebra—and we can follow [Aumann \(1964\)](#) when defining the seller's strategy.

¹⁰The only difference between Bayes' rule where possible and consistency in sequential equilibrium is the following. Under PBE, the seller can assign zero probability to one of the buyer's values and then, after the buyer deviates, can assign positive probability to that same value.

2.1 Revelation principle

The game $G_{\mathcal{I}}(\mu_0)$ is not simple to analyze for at least two reasons. 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 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 $G(\mu_0)$ denote the same game in the previous section, except that in each period 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 set of buyer's values and of 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 $G(\mu_0)$. In what follows, a subset of the set of canonical mechanisms has special significance: the set of *direct Blackwell mechanisms*. A direct Blackwell mechanism is a canonical mechanism $\varphi: V \mapsto \Delta(\Delta(V) \times A)$ that can be decomposed into a Blackwell experiment, $\beta: V \mapsto \Delta(\Delta(V))$, and an *allocation rule*, $\alpha: \Delta(V) \mapsto \Delta(A)$.¹¹

Lemma 1 summarizes the key implications of Doval and Skreta (2022) for our analysis, which we explain below.

LEMMA 1 (Doval and Skreta (2022)). *For any PBE outcome of any mechanism-selection game $G_{\mathcal{I}}(\mu_0)$, an outcome-equivalent PBE of game $G(\mu_0)$ exists. That is, $\bigcup_{\mathcal{I}} \mathcal{O}_{\mathcal{I}}^*(\mu_0) = \mathcal{O}^*(\mu_0)$.*

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

- (a) *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;*
- (b) *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's value coincide with μ' ;*
- (c) *For all histories h^t , the mechanism offered by the seller at h^t is a direct Blackwell mechanism;*
- (d) *The buyer's strategy depends only on her private value and the public history.*

Lemma 1 has several implications. Part (a) of Lemma 1 implies the mechanisms chosen 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 (a) simplifies the analysis of the buyer's

¹¹That is, for all measurable subsets $U' \subset \Delta(V)$ and $A' \subset A$, we have that for all $v \in V$,

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

behavior, by reducing it to a series of constraints the seller's equilibrium offer of a mechanism must satisfy (see Equations $(PC_{h^t, v})$ and $(IC_{h^t, v, v'})$ in Section 4.1).

Part (b) 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 φ^{M_t} as a direct Blackwell mechanism described in part (c).

Part (c) implies that the choice of mechanism at history h^t can be equivalently thought of as the choice of a Blackwell experiment, β^{M_t} , and an allocation rule, α^{M_t} . Direct Blackwell mechanisms allow us to separately optimize on the allocation *given* a particular experiment, and then optimize on the experiment. As in the literature on information design, it is convenient to work with the distribution over posteriors induced by the experiment β^{M_t} , which we denote by τ^{M_t} and is defined as follows. For all Borel subsets $U' \subseteq \Delta(V)$,

$$\int_{U'} \tau^{M_t}(d\mu_{t+1}) = \sum_{v \in V} \mu_t(h^t)(v) \beta^{M_t}(U'|v), \quad (BC_{\mu_t(h^t)})$$

where $\mu_t(h^t) \in \Delta(V)$ 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 $\alpha^{M_t}(\cdot|\mu)$ as an expected payment, $x^{M_t}(\mu)$, and a probability of trade, $q^{M_t}(\mu)$.

Part (d) implies the set of PBE payoffs of $G(\mu_0)$ coincides with the set of *Public PBE* payoffs of $G(\mu_0)$ (Athey and Bagwell (2008)). Relying on Abreu, Pearce, and Stacchetti (1990), Athey and Bagwell (2008) show that Public PBE have a recursive structure and we use this property to argue the assessment we define in Section 3 is indeed a PBE assessment.

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

REMARK 1. Below, we abuse notation in the following two ways. First, because values 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 $G(\mu_0)$. To state Theorem 1, we proceed as follows. First, we informally describe the *posted-prices* assessment, $(\Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^*)$, which is singled out by the proof of Theorem 1. Second, we explain why the outcome induced by $(\Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^*)$ can be implemented via a sequence of posted prices. Finally, we state Theorem 1.

An assessment in posted prices: The posted-prices 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, Levine, and Tirole \(1985\)](#) and [Hart and Tirole \(1988\)](#). The assessment is defined by an increasing sequence of threshold beliefs $\bar{\mu}_0 = 0 < \bar{\mu}_1 < \dots < \bar{\mu}_n < \dots$ such that if the seller's belief is in $[\bar{\mu}_n, \bar{\mu}_{n+1})$, 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 and determines both the rents for v_H and the seller's revenue.

To understand how the sequence of thresholds is determined, consider first when selling the good at a price of v_L is optimal for the seller. If a seller with belief μ_0 sells the good at a price of v_L , his revenue can be written as follows:

$$\begin{aligned} v_L &= \mu_0(v_H - \Delta v) + (1 - \mu_0)v_L \\ &= \mu_0 v_H + (1 - \mu_0) \left(v_L - \frac{\mu_0}{1 - \mu_0} \Delta v \right) \equiv \mu_0 v_H + (1 - \mu_0) \hat{v}_L(\mu_0). \end{aligned} \quad (1)$$

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 sign of the virtual value $\hat{v}_L(\mu_0)$ determines whether the optimal price is v_L ($\hat{v}_L(\mu_0) > 0$) or v_H ($\hat{v}_L(\mu_0) < 0$) in the commitment solution. Because $\hat{v}_L$ is decreasing in μ_0 , a threshold belief exists such that $\hat{v}_L(\mu_0) = 0$:

$$\hat{v}_L(\mu_0) = 0 \quad \Leftrightarrow \quad \mu_0 = \frac{v_L}{v_H} \equiv \bar{\mu}_1. \quad (2)$$

At that belief, the seller is indifferent between both prices in the commitment solution. As we argue next, $\bar{\mu}_1$ also plays an important role in the posted-prices assessment.

The preceding discussion implies v_L is the maximum revenue for a seller with belief $\mu_0 < \bar{\mu}_1$ in the mechanism-selection game. Instead, a seller with belief $\mu_0 > \bar{\mu}_1$ prefers waiting 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 $\mu_0 > \bar{\mu}_1$, the seller prefers to wait at least one period to sell the good to v_L . When the seller's prior belief is $\bar{\mu}_1$, 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 $\bar{\mu}_1$ reflects that the seller takes one period to trade with v_L in the posted-prices assessment. Recursively, one can define $\bar{\mu}_n$ as the belief at which the seller is indifferent between trading with v_L in n or $n - 1$ periods.

In the posted-prices assessment, play proceeds as follows. If the seller's prior is such that $\mu_0 \in [\bar{\mu}_n, \bar{\mu}_{n+1})$ and $n \geq 1$, then the seller chooses a mechanism that induces two beliefs, 1 and $\bar{\mu}_{n-1}$. At belief 1, the good is sold and the transfer is $v_L + (1 - \delta^n)\Delta v$, whereas at belief $\bar{\mu}_{n-1}$, the good is not sold and the transfer is 0 (see Figure 1(a)). Subsequently, a seller with belief $\bar{\mu}_m$ for $m \leq n - 1$, chooses a mechanism that induces two beliefs, 1 (with allocation $(1, v_L + (1 - \delta^m)\Delta v)$) and $\bar{\mu}_{m-1}$ (with allocation $(0, 0)$) (see Figure 1(b)).




FIGURE 1. Mechanism and belief dynamics under $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$ as a function of μ_0 .

Thus, starting from μ_0 , the game ends in n periods. Figure 1(c) illustrates how beliefs fall conditional on no trade.

More precisely, the assessment $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$ is as follows (the formal statement is in Appendix D.2):

1. Along the equilibrium path:

- (a) If at history h^t , the seller's beliefs, $\mu_t^*(h^t)$, are in $[\bar{\mu}_0, \bar{\mu}_1)$, he chooses a mechanism such that $(q^{\mathbf{M}^*}_t(\mu_t^*(h^t)), x^{\mathbf{M}^*}_t(\mu_t^*(h^t))) = (1, v_L)$, and the Blackwell experiment satisfies that $\beta^{\mathbf{M}^*}_t(\mu_t^*(h^t)|v) = 1$ for $v \in \{v_L, v_H\}$.
- (b) If at history h^t , the seller's beliefs, $\mu_t^*(h^t)$, are in $[\bar{\mu}_n, \bar{\mu}_{n+1})$ for $n \geq 1$, the seller's mechanism satisfies the following. First, it induces two posteriors, $\bar{\mu}_{n-1}$ 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 $(q^{\mathbf{M}^*}_t(\bar{\mu}_{n-1}), x^{\mathbf{M}^*}_t(\bar{\mu}_{n-1})) = (0, 0)$. Finally, the Blackwell experiment $\beta^{\mathbf{M}^*}_t$ maps v_L to $\bar{\mu}_{n-1}$, whereas it maps v_H to both $\bar{\mu}_{n-1}$ and 1 with positive probability. The probabilities $\beta^{\mathbf{M}^*}_t(\bar{\mu}_{n-1}|v_H)$, $\beta^{\mathbf{M}^*}_t(1|v_H)$ are chosen so that when the seller observes $\bar{\mu}_{n-1}$, his updated belief coincides with $\bar{\mu}_{n-1}$, that is,

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

2. Off the equilibrium path, the seller's strategy coincides with the above, except that when $\mu_t^*(h^t) = \bar{\mu}_n$ for some $n \geq 1$, the seller may randomize between the mechanism he offers on the path of play when his belief is $\bar{\mu}_n$, and the one he offers on the path of play when his belief is $\bar{\mu}_{n-1}$.¹²
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 value.

An implementation in posted prices: We now argue that the equilibrium outcome induced by $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$ can be implemented in posted prices. Clearly, when the seller's beliefs are below $\bar{\mu}_1$, 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 $[\bar{\mu}_1, \bar{\mu}_2)$. Note that when the realized allocation is trade, then the high-valuation buyer's payoff is $v_H - v_L - (1 - \delta)\Delta v = \delta\Delta v$. Instead, when the realized allocation is no trade, the seller's beliefs next period are $\bar{\mu}_0 = 0$, so that the high-valuation buyer's continuation payoff is $\delta\Delta v$. That is, the high-valuation buyer 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 high-valuation buyer 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 $\bar{\mu}_0$ when the allocation is $(0, 0)$. Because $\bar{\mu}_0 = 0$, it follows that when the seller's belief is in $[\bar{\mu}_1, \bar{\mu}_2)$, the high-valuation 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 $[\bar{\mu}_n, \bar{\mu}_{n+1})$, 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 $\bar{\mu}_{n-1}$ when the allocation is $(0, 0)$.

REMARK 2 (Direct vs. indirect implementation). It is interesting to contrast the implementation under the posted-prices 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 *rationes* 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 rejects the initial prices with positive probability, which like rationing, prevents the seller from becoming pessimistic too quickly about the buyer's value. 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 with a continuum of buyer values, [Denicolo and Garella \(1999\)](#) show that if the seller observes only whether

¹²The need for mixing arises for technical reasons: it ensures that the high-valuation buyer's continuation payoffs are upper semicontinuous, and thus guarantees that a best response exists after any deviation by the seller (see Appendix D.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.

trade happens, the seller may prefer to ration higher valuation buyers instead of posting a price in the first period, 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 (2022) that both posted prices and rationing as in Denicolo and Garella (1999) may be dominated in a two-period model with a continuum of buyer values by the period-1 seller offering what we dub an obfuscated nonuniform pricing mechanism. This mechanism sells the good to higher valuation buyers with probability 1, excludes lower valuation buyers with probability 1, and sells the good with probability less than 1 to middle valuation buyers. Importantly, the period-2 seller only observes whether the good is sold and middle valuation buyers report their values before knowing whether they will be rationed. That this mechanism dominates posting a price in period 1 crucially depends on the finite-horizon assumption in Doval and Skreta (2022): Because the period-2 seller can commit to exclude low-valuation buyers, the period-1 seller is willing to sacrifice trade with middle valuation buyers to induce higher period-2 prices than those that would be feasible if he instead posted a price in period 1.

Theorem 1 states that the seller's payoff in the mechanism-selection game is unique and coincides with that in the posted-prices assessment. Furthermore, the low-valuation buyer's payoff is also unique and equal to 0. Finally, except at the threshold beliefs $\{\bar{\mu}_n\}_{n \geq 1}$, the high-valuation buyer's payoff is also unique. The multiplicity of the high-valuation buyer's payoff arises because when the seller's prior belief is $\bar{\mu}_n$ 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.¹³

To state Theorem 1, let $u_H^*(\mu_0)$ and $u_S^*(\mu_0)$ denote the high-valuation buyer's 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, for $n \geq 0$ and $\mu_0 \in [\bar{\mu}_n, \bar{\mu}_{n+1})$, the set of equilibrium payoffs of $G(\mu_0)$ is the following:

$$\mathcal{E}^*(\mu_0) = \begin{cases} \{(0, u_H^*(\mu_0), u_S^*(\mu_0))\}, & n = 0, \text{ or } n \geq 1 \text{ and } \mu_0 \in (\bar{\mu}_n, \bar{\mu}_{n+1}), \\ \{(0, u_H, u_S^*(\mu_0)) : u_H \in [u_H^*(\mu_0), u_H^*(\mu_0)/\delta]\}, & n \geq 1, \mu_0 = \bar{\mu}_n. \end{cases} \quad (3)$$

¹³When the seller's prior belief coincides with $\bar{\mu}_n$ for $n \geq 1$, an alternative PBE assessment exists, which is payoff equivalent to $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$ for the seller, but delivers a higher payoff 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, $\bar{\mu}_{n-2}$ and 1. When the posterior is 1, the allocation is $(1, v_L + (1 - \delta^{n-1})\Delta v)$, whereas when the posterior is $\bar{\mu}_{n-2}$, 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 D.

Theorem 1 together with the above discussion implies the following.

COROLLARY 1. *All equilibrium outcomes of $G(\mu_0)$ can be implemented via a sequence of posted prices.*¹⁴

Section 5 reviews the main steps of the proof of Theorem 1. Section 5.1 argues that $u_S^*(\mu_0)$ is a lower bound and an upper bound on the seller's equilibrium payoff. Section 5.2 describes the arguments to show that $(\Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^*)$ is a PBE assessment. To show that $u_S^*(\mu_0)$ is the seller's unique equilibrium payoff and that $(\Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^*)$ 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 (OPT) to show both that the seller's unique equilibrium payoff corresponds to that in the posted-prices assessment and that $(\Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^*)$ is a PBE assessment. In (OPT), the seller maximizes the dynamic analogue of the virtual surplus, which we denote by $VS(\cdot)$. 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 truth-telling constraints (Section 4.1 provides the details of this derivation in our setting). As we explain below, $VS(\cdot)$ also depends on the posterior distribution induced by the mechanism's Blackwell experiment (recall Equation $(BC_{\mu_t(h')})$ on p. 872). Moreover, because of the dynamic nature of our problem, it also depends on the continuation payoffs.

Formally,

$$\begin{aligned} & \max_{(\tau_0, q_0), u} VS((\tau_0, q_0), u, \mu_0), \\ & \text{such that } \begin{cases} \tau_0 \in \Delta(\Delta(V)) \text{ is Bayes' plausible for } \mu_0, \\ q_0 : \Delta(V) \mapsto [0, 1], \\ (\forall \mu_1 \in \Delta(V)) u(\mu_1) \in \mathcal{E}^*(\mu_1), \end{cases} \end{aligned} \quad (\text{OPT})$$

where

$$\begin{aligned} & VS((\tau_0, q_0), u, \mu_0) \\ & \equiv \int_{\Delta(V)} \left[\begin{aligned} & q_0(\mu_1)(\mu_1 v_H + (1 - \mu_1) \hat{v}_L(\mu_0)) \\ & + (1 - q_0(\mu_1)) \delta \left(u_S(\mu_1) + \mu_1 u_H(\mu_1) + (1 - \mu_1) \right. \\ & \quad \left. \times \left(u_L(\mu_1) - \frac{\mu_0}{1 - \mu_0} (u_H(\mu_1) - u_L(\mu_1)) \right) \right) \end{aligned} \right] \tau_0(d\mu_1), \end{aligned} \quad (4)$$

¹⁴In the no gap case, the seller's best equilibrium outcome can also be sustained using posted prices. Indeed, 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.

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 also refer to the latter as the mechanism.

In (OPT), we not only allow the seller with prior μ_0 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, $u(\mu_1) \in \mathcal{E}^*(\mu_1)$.

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

$$q_0(\mu_1)(\mu_1 v_H + (1 - \mu_1) \hat{v}_L(\mu_0)), \quad (5)$$

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

$$(1 - q_0(\mu_1)) \delta \left(u_S(\mu_1) + \mu_1 u_H(\mu_1) + (1 - \mu_1) \left(u_L(\mu_1) - \frac{\mu_0}{1 - \mu_0} (u_H(\mu_1) - u_L(\mu_1)) \right) \right), \quad (6)$$

accounts for the rents, $u_H(\mu_1) - u_L(\mu_1)$, the high-valuation buyer 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. That rents are accounted for using the seller's prior, μ_0 , rather than his posterior belief, μ_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.

Whereas 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 truth-telling constraints would still be satisfied, whereas the seller's revenue would increase. This argument does not translate immediately to our 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 $\mathbf{M}$.

Program (OPT) allows us to circumvent this difficulty: Because in (OPT) the seller can choose both the mechanism *and* the continuation payoffs, 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 (4) 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 (4) 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, $(\Gamma, (\pi_v, r_v)_{v \in V}, \mu)$. 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 $u^{\mathbf{M}_t}$ when the seller offers mechanism $\mathbf{M}_t$ at history h^t . In what follows, we show the following.

LEMMA 2. *The virtual surplus of mechanism $\mathbf{M}_t$ is an upper bound on the sum of the seller and the low-valuation buyer's payoffs. That is,*

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

Lemma 2 is the analogue of the result in mechanism design that the mechanism's allocation together with the lowest type's utility in the mechanism pin down the seller's maximum revenue. Because the low-valuation buyer's payoff is nonnegative, Lemma 2 also implies that the virtual surplus of $\mathbf{M}_t$ is an upper bound on the seller's payoff. The inequality in Lemma 2 shows that the mechanism's virtual surplus is the maximum payoff the seller and the low-valuation buyer can share. As we explain in Section 5, once we show the seller captures the entirety of the virtual surplus, the inequality in Lemma 2 implies low-valuation buyer's payoff is 0 in any equilibrium.

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

$$U_S(h^t) = \int_{\Delta(V)} (x^{\mathbf{M}_t}(\mu_{t+1}) + (1 - q^{\mathbf{M}_t}(\mu_{t+1}))\delta u_S^{\mathbf{M}_t}(\mu_{t+1}))\tau^{\mathbf{M}_t}(d\mu_{t+1}), \quad (7)$$

where $u_S^{\mathbf{M}_t}(\mu_{t+1})$ is shorthand notation for the seller's continuation payoff when at history h^t , he offers $\mathbf{M}_t$ and the output message is μ_{t+1} .¹⁵ Equation (7) 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 $(q^{\mathbf{M}_t}, x^{\mathbf{M}_t})$.

¹⁵This continuation payoff can also depend on h^t , but we omit this dependence to simplify notation.

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:

$$\begin{aligned} \int_{\Delta(V)} (vq^{\mathbf{M}_t}(\mu_{t+1}) - x^{\mathbf{M}_t}(\mu_{t+1}) + (1 - q^{\mathbf{M}_t}(\mu_{t+1}))\delta u_v^{\mathbf{M}_t}(\mu_{t+1}))\beta^{\mathbf{M}_t}(d\mu_{t+1}|v) \\ \geq u_v^{\mathbf{M}_t}(\emptyset), \end{aligned} \quad (\text{PC}_{h^t, v})$$

where the left-hand side of Equation $(\text{PC}_{h^t, v})$ is the buyer's payoff at h^t , $U_v(h^t)$, and $u_v^{\mathbf{M}_t}(\emptyset)$ is shorthand notation for the buyer's continuation payoff 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:

$$\begin{aligned} \int_{\Delta(V)} (vq^{\mathbf{M}_t}(\mu_{t+1}) - x^{\mathbf{M}_t}(\mu_{t+1}) \\ + (1 - q^{\mathbf{M}_t}(\mu_{t+1}))\delta u_v^{\mathbf{M}_t}(\mu_{t+1}))(\beta^{\mathbf{M}_t}(d\mu_{t+1}|v) - \beta^{\mathbf{M}_t}(d\mu_{t+1}|v')) \\ \geq 0. \end{aligned} \quad (\text{IC}_{h^t, v, v'})$$

The above expressions implicitly use Lemma 1 in one more way. 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 $G(\mu_{t+1})$. Formally, $u^{\mathbf{M}_t}(\mu_{t+1}) \in \mathcal{E}^*(\mu_{t+1})$.

Equations $(\text{PC}_{h^t, v})$ and $(\text{IC}_{h^t, v, v'})$ 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 sidestep this challenge, we ignore the right-hand side of the low-valuation buyer's participation constraint. Instead, we use the following identity to solve for the transfers *given* the low-valuation buyer's utility in the mechanism, $U_L(h^t)$:

$$\begin{aligned} \int_{\Delta(V)} (v_L q^{\mathbf{M}_t}(\mu_{t+1}) - x^{\mathbf{M}_t}(\mu_{t+1}) \\ + (1 - q^{\mathbf{M}_t}(\mu_{t+1}))\delta u_L^{\mathbf{M}_t}(\mu_{t+1}))\beta^{\mathbf{M}_t}(d\mu_{t+1}|v_L) = U_L(h^t). \end{aligned} \quad (8)$$

Equation (8) can be used to rewrite Equation $(\text{IC}_{h^t, v, v'})$ for $v = v_H$ as follows:

$$\begin{aligned} \int_{\Delta(V)} (v_H q^{\mathbf{M}_t}(\mu_{t+1}) - x^{\mathbf{M}_t}(\mu_{t+1}) + (1 - q^{\mathbf{M}_t}(\mu_{t+1}))\delta u_H^{\mathbf{M}_t}(\mu_{t+1}))\beta^{\mathbf{M}_t}(d\mu_{t+1}|v_H) \\ \geq \int_{\Delta(V)} [\Delta v q^{\mathbf{M}_t}(\mu_{t+1}) + (1 - q^{\mathbf{M}_t}(\mu_{t+1}))\delta(u_H^{\mathbf{M}_t}(\mu_{t+1}) - u_L^{\mathbf{M}_t}(\mu_{t+1}))]\beta^{\mathbf{M}_t}(d\mu_{t+1}|v_L) \\ + U_L(h^t). \end{aligned} \quad (9)$$

Equations (8) and (9) already impose constraints on the maximum revenue the seller can make in period t when the low-valuation buyer's payoff is $U_L(h^t)$. Indeed, like in standard mechanism design, the utility of v_L and the truth-telling constraint for v_H determine the maximum expected transfer the seller can extract from the buyer in period

t . Replacing the upper bound on the expected transfers obtained from these equations in the seller's payoff (Equation (7)), 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 $VS((\tau^{\mathbf{M}_t}, q^{\mathbf{M}_t}), u^{\mathbf{M}_t}, \mu_t(h^t)) - U_L(h^t)$. Lemma 2 then follows.

Having established that (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 (OPT) that are important to prove that the seller's unique equilibrium payoff is $u_S^*(\mu_0)$. As we explain next, a key difficulty that (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, e.g., Fudenberg, Levine, and Tirole (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.¹⁶ 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. 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 (OPT) is useful to discipline belief dynamics. Whereas it may not be obvious how to rule out that an equilibrium in which the seller's beliefs may sometimes go up conditional on no trade exists, 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 $\mu_0 > \bar{\mu}_1$, conditional on selling the good with positive probability, the seller sells the good only to the high-valuation buyer.

¹⁶Whereas the truth-telling equations $(IC_{h^t, v, v'})$ can be used to derive a "monotonicity" condition analogous to that in the skimming lemma, this condition only implies that *on average* the expected probability of trade of v_H must be higher than that of v_L :

$$\int_{\Delta(V)} [\Delta v q^{\mathbf{M}_t}(\mu_{t+1}) + (1 - q^{\mathbf{M}_t}(\mu_{t+1}))\delta(u_{H'}^{\mathbf{M}_t}(\mu_{t+1}) - u_L^{\mathbf{M}_t}(\mu_{t+1}))](\beta^{\mathbf{M}_t}(d\mu_{t+1}|v_H) - \beta^{\mathbf{M}_t}(d\mu_{t+1}|v_L)) \geq 0.$$

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

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

$$\int_{[\mu_0, 1]} (1 - q_0(\mu_1)) \tau_0(d\mu_1) = 0.$$

(b) *Furthermore, if $\mu_0 > \bar{\mu}_1$ and the seller induces μ_1 and sells the good (i.e., $q_0(\mu_1) > 0$), then $\mu_1 = 1$.*

The proof is in Appendix B. In what follows, we provide intuition for Proposition 1, starting from part (a). To see why not selling the good and at the same time induce a belief $\mu_1 \geq \mu_0$ is not optimal, note the following. First, associated to any continuation payoff, $u(\mu_1)$, 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 $u(\mu_1)$ 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 (after grouping terms, the term pre-multiplying $u_H(\mu_1) - u_L(\mu_1)$ in Equation (6) 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 $u(\mu_1)$, contradicting that it is optimal to induce μ_1 and not sell the good.

Part (b) follows from the observation that a seller with prior above $\bar{\mu}_1$ prefers to trade with v_L in at least two periods (recall that $\hat{v}_L(\mu_0) < 0$ when $\mu_0 > \bar{\mu}_1$). 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), which in a slight abuse of notation, we denote by $VS(\mu_0)$:

$$\begin{aligned} VS(\mu_0) \equiv & \max_{\tau_0, u} \tau_0(\{1\})v_H + \int_{[0, \mu_0)} \delta \left[u_S(\mu_1) + u_L(\mu_1) \right. \\ & \left. + \left(\frac{\mu_1 - \mu_0}{1 - \mu_0} \right) (u_H(\mu_1) - u_L(\mu_1)) \right] \tau_0(d\mu_1). \end{aligned} \quad (10)$$

Equation (10) 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 kind 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 (10).

COROLLARY 2. *The value of (OPT), $VS(\mu_0)$, equals the value of*

$$\max_{G \in \Delta([0, \mu_0])} \max_u \int_{[0, \mu_0]} \left[\frac{\mu_0 - \mu_1}{1 - \mu_1} v_H + \frac{1 - \mu_0}{1 - \mu_1} \delta(u_S(\mu_1) + u_L(\mu_1)) + \left(\frac{\mu_1 - \mu_0}{1 - \mu_0} \right) (u_H(\mu_1) - u_L(\mu_1)) \right] G(d\mu_1). \quad (11)$$

The proof is in Appendix B and is a consequence of the constraint that τ_0 is Bayes' plausible for μ_0 . Given the preceding discussion, the term in the square brackets inside the integral in Equation (11) is 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 (10) is as if the seller were randomizing over posterior distributions that split the prior between 1 (and selling the good) and $\mu_1 < \mu_0$ (and not selling the good). In other words, 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. 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 .

Because conditional on not selling the good the seller induces beliefs below the prior, Equation (11) shows that for a fixed distribution G , the seller trades off the sum $u_S + u_L$ against the high-valuation buyer's rents, $u_H - u_L$, when choosing continuation payoffs. Indeed, if continuation payoffs $u, u' \in \mathcal{E}^*(\mu_1)$ exist such that $(u_H - u_L, u_S + u_L) \ll (u'_H - u'_L, u'_S + u'_L)$, the solution to the program in Equation (11) would choose u over u' if the value of the integrand is higher for u . That is, the seller may forgo maximizing $u_S + u_L$ —and hence, by Lemma 2 potentially forgo the maximum continuation virtual surplus—if this could lead to lower rents for the high-valuation buyer. As we explain in Section 5 below, we circumvent this trade-off in the proof of Theorem 1 by showing that for each $\mu_1 < \mu_0$ the sum $u_S + u_L$ is unique, which in turn relies on excluding the existence of payoffs like u and u' .

5. PROOF OF THEOREM 1: KEY STEPS

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

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

We show $u_S^*(\mu_0)$ is both a lower bound and an upper bound on the seller's equilibrium payoff. That is, we show 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. Program (OPT) is key to show $u_S^*(\mu_0)$ is the seller's unique equilibrium payoff. The proof that $u_S^*(\mu_0)$ is an upper bound on the seller's equilibrium payoff follows from showing that $u_S^*(\mu_0)$ is the value of (OPT). Instead, the proof that $u_S^*(\mu_0)$ is a lower bound on the seller's payoff uses a constrained version of (OPT) to then apply the logic in Gul, Sonnenschein, and Wilson (1986): If an equilibrium in which the seller earns less than $u_S^*(\mu_0)$ exists, the seller can always undercut the price in the posted-prices assessment and earn close to $u_S^*(\mu_0)$.

Using the property in the posted-prices assessment that along the path of play beliefs fall, the proof of Theorem 1 proceeds by induction on the interval the seller's prior belongs to. For each $n \geq 0$, we establish two results. First, for all $\mu_0 \in [\bar{\mu}_n, \bar{\mu}_{n+1})$, the set of equilibrium payoffs corresponds to that in Theorem 1 (cf. Equation (3)). Second, for all $\mu_0 \geq \bar{\mu}_n$, the seller can guarantee a certain payoff, denoted $u_S^*(\mu_0, n)$, which coincides with $u_S^*(\mu_0)$ for $\mu_0 \in [\bar{\mu}_n, \bar{\mu}_{n+1})$. This payoff is obtained by the seller emulating the strategy that sells the good to v_L in n periods in the posted-prices assessment (Figure 1(a)): The seller uses a mechanism that splits his beliefs between 1 and $\bar{\mu}_{n-1}$; when the belief is 1, the good is sold at a price of $v_L + (1 - \delta^n)\Delta v$, and when the belief is $\bar{\mu}_{n-1}$, the good is not sold and play proceeds according to the posted-prices assessment. As we explain below, this second result is key to establish in the $(n + 1)^{\text{th}}$ step of the induction that the seller can guarantee the payoff from the posted-prices assessment when $\mu_0 \in [\bar{\mu}_{n+1}, \bar{\mu}_{n+2})$.

To illustrate the main steps of the proof that $u_S^*(\mu_0)$ is the seller's unique equilibrium payoff, fix $n \geq 1$.¹⁷ Suppose that for all $m < n$ we have already shown that (i) Theorem 1 holds for all $\mu_0 \in [\bar{\mu}_m, \bar{\mu}_{m+1})$ and (ii) the seller can guarantee $u_S^*(\mu_0, m)$ for all $\mu_0 \geq \bar{\mu}_m$. In what follows, we argue that (i) Theorem 1 holds for $\mu_0 \in [\bar{\mu}_n, \bar{\mu}_{n+1})$ and (ii) the seller's payoff is at least $u_S^*(\mu_0, n)$ for $\mu_0 \geq \bar{\mu}_n$.

Posted prices maximize the virtual surplus: 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 $\bar{\mu}_{n-1}$ 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 that, conditional on inducing a belief $\mu_1 < \bar{\mu}_n$, the seller places weight on at most $\bar{\mu}_{n-1}$ (and on $\bar{\mu}_{n-2}$ only if $\mu_0 = \bar{\mu}_n$). We then argue it is never optimal to induce a belief $\mu_1 \in [\bar{\mu}_n, \mu_0)$.

The inductive hypothesis—which identifies the continuation payoffs below $\bar{\mu}_n$ —and the properties of the posted-prices assessment imply that conditional on inducing a belief in $[0, \bar{\mu}_n)$, the solution to the problem in Equation (11) places weight on at most $\{\bar{\mu}_{n-2}, \bar{\mu}_{n-1}\}$. First, except at the threshold beliefs $\{\bar{\mu}_m\}_{m \leq n-1}$, the inductive hypothesis pins down the value of the integrand in Equation (11) for $\mu_1 < \bar{\mu}_n$. Second, because at the threshold beliefs the seller's and the low-valuation buyer's payoff are unique, the seller then chooses the continuation payoff that minimizes the high-valuation buyer's rents, $u_H^*(\cdot)$, as is the case in the posted-prices assessment. Third, relying on the properties of the posted-prices assessment, Lemma C.2 shows that inducing beliefs in $[0, \bar{\mu}_n)$ other than $\{\bar{\mu}_{n-2}, \bar{\mu}_{n-1}\}$ is not optimal. Furthermore, $\bar{\mu}_{n-2}$ can be in the support of τ_0

¹⁷The proof also verifies that Theorem 1 holds for $\mu_0 \in [0, \bar{\mu}_1)$ (see Appendix C.2.1).

only if $\mu_0 = \bar{\mu}_n$. This is intuitive. On the one hand, it can only be optimal to induce the threshold beliefs $\{\bar{\mu}_m\}_{m \leq n-1}$: For $m \leq n-1$, both $\mu_1 \in (\bar{\mu}_m, \bar{\mu}_{m+1})$ and $\bar{\mu}_m$ imply that trade happens with v_L in m periods; however, inducing $\bar{\mu}_m$ allows the seller to trade with v_H with a higher probability. On the other hand, the indifference condition that defines the threshold beliefs implies that because $\mu_0 \geq \bar{\mu}_n$, inducing beliefs below $\bar{\mu}_{n-2}$ cannot be optimal (see Lemma C.1).

To conclude the proof that $u_S^*(\mu_0)$ is an upper bound on $VS(\mu_0)$, we show in Lemma C.3 that inducing posteriors in $[\bar{\mu}_n, \mu_0)$ is not optimal. Whereas 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 $[\bar{\mu}_n, \mu_0)$ if at the induced beliefs continuation equilibria exist where the seller somehow manages to slowly trade with v_H so as to maximally delay trade with v_L . As we show in Appendix C.2, this cannot be optimal for μ_0 close to $\bar{\mu}_n$: The closer to $\bar{\mu}_n$, 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 $\bar{\mu}_n$. It follows that μ_0^* small enough exists such that if the seller's prior μ_0 is in $[\bar{\mu}_n, \mu_0^*]$, it is better to trade with v_L in n periods in exchange of increasing the probability of trading with v_H today.

More formally, note that Lemma C.2 implies that for $\mu_0 \in [\bar{\mu}_n, \bar{\mu}_{n+1})$, the value of (OPT), $VS(\mu_0)$, is bounded above by

$$VS(\mu_0) \leq \max \left\{ u_S^*(\mu_0), \frac{\mu_0 - \bar{\mu}_n}{1 - \bar{\mu}_n} v_H + \frac{1 - \mu_0}{1 - \bar{\mu}_n} \overline{VS}_{[\bar{\mu}_n, \mu_0)} \right\}, \quad (12)$$

where $\overline{VS}_{[\bar{\mu}_n, \mu_0)}$ is the supremum of the value function $VS(\mu)$ on $[\bar{\mu}_n, \mu_0]$. To see why Equation (12) holds, note the following. First, if the solution to (OPT) places positive mass below $\bar{\mu}_n$, Lemma C.2 implies that $VS(\mu_0) = u_S^*(\mu_0)$, since the seller places weight on $\bar{\mu}_{n-1}$ (or $\bar{\mu}_{n-2}$ if $\mu_0 = \bar{\mu}_n$). The equality follows from Corollary 2 and that $u_S^*(\mu_0)$ corresponds to splitting μ_0 between 1 and $\bar{\mu}_{n-1}$. Second, if the solution to (OPT) places weight on $[\bar{\mu}_n, \mu_0)$, the second term on the right-hand side of Equation (12) is an upper bound to $VS(\mu_0)$. After all, (i) $(\mu_0 - \bar{\mu}_n)/(1 - \bar{\mu}_n)$ is the largest weight that can be assigned to v_H while still remaining on $[\bar{\mu}_n, \mu_0)$ and (ii) the remaining weight corresponds to some $\mu_1 \in [\bar{\mu}_n, \mu_0)$ with payoff

$$\begin{aligned} & \left(u_S(\mu_1) + u_L(\mu_1) + (1 - \mu_1) \left(\frac{\mu_1}{1 - \mu_1} - \frac{\mu_0}{1 - \mu_0} \right) (u_H(\mu_1) - u_L(\mu_1)) \right) \\ & \leq u_S(\mu_1) + u_L(\mu_1) \leq \overline{VS}_{[\bar{\mu}_n, \mu_0)}, \end{aligned}$$

where the first inequality follows from $\mu_1 < \mu_0$ and the second from Lemma 2 and the definition of $\overline{VS}_{[\bar{\mu}_n, \mu_0)}$ together with $\mu_1 \in [\bar{\mu}_n, \mu_0)$.

For μ_0^* close to $\bar{\mu}_n$, the probability of trading with v_H today and at the same time remaining above $\bar{\mu}_n$ is small and $u_S^*(\mu_0)$ attains the maximum on the right-hand side of Equation (12) for $\mu_0 \in [\bar{\mu}_n, \mu_0^*]$ (see Lemma C.3 for details). In other words, μ_0^* small enough exists such that for all $\mu_0 \in [\bar{\mu}_n, \mu_0^*]$, the value of (OPT), $VS(\mu_0)$, coincides with the seller's payoff in the posted-prices assessment, $u_S^*(\mu_0)$. Replacing $\bar{\mu}_n$ with μ_0^* in Equation (12), one can argue that for beliefs μ_0 close enough to μ_0^* , $u_S^*(\mu_0)$ is also an

upper bound on the seller's payoff. Proceeding this way, one establishes that $u_S^*(\mu_0)$ is an upper bound on the value of (OPT) for all beliefs $\mu_0 \in [\bar{\mu}_n, \bar{\mu}_{n+1})$.

Seller can guarantee the payoff from selling to v_L in n periods: We show in Proposition C.2 that the seller can guarantee the payoff $u_S^*(\mu_0, n)$ for $\mu_0 \geq \bar{\mu}_n$, and hence the payoff of the posted-prices assessment for $\mu_0 \in [\bar{\mu}_n, \bar{\mu}_{n+1})$. The logic is similar to that in Gul, Sonnenschein, and Wilson (1986): We show that the seller can always undercut the price in the posted-prices assessment—say, by offering a mechanism that sells the good for $v_L + (1 - \delta^n)\Delta v - \delta F$ for some small $F > 0$ —and earn close to $u_S^*(\mu_0, n)$.

Provided the high-valuation buyer participates in the mechanism with positive probability (note that the low-valuation buyer *always* rejects), we show that F small enough can be chosen so that the seller's beliefs conditional on rejection are $\bar{\mu}_{n-1}$. That beliefs conditional on rejection are $\bar{\mu}_{n-1}$, in turn, implies the seller earns a payoff close to $u_S^*(\mu_0, n)$: First, it implies that the high-valuation buyer's acceptance probability coincides with the probability of selling the good in the posted-prices assessment. Second, because $\bar{\mu}_{n-1} < \bar{\mu}_n$, the inductive hypothesis implies the seller's continuation payoffs coincide with those in the posted-prices assessment. Thus, it suffices to show that the seller can guarantee that the high-valuation buyer participates in the mechanism.

Proposition 2 is key to showing that the high-valuation buyer does not reject the mechanism with probability 1.

PROPOSITION 2 (Buyer's maximal rents for $\mu_0 \geq \bar{\mu}_n$). *For all $\mu_0 \geq \bar{\mu}_n$, $u_H \leq \delta^{n-1}\Delta v$.*

The proof is in Appendix C.3. Proposition 2 implies the high-valuation buyer *cannot* reject the price of $v_L + (1 - \delta^n)\Delta v - \delta F$ with probability 1, as doing so can yield a payoff of at most $\delta\delta^{n-1}\Delta v$. Together with the argument in the preceding paragraph, we conclude that as F becomes small the seller can secure $u_S^*(\mu_0, n)$ for $\mu_0 \geq \bar{\mu}_n$.

To prove Proposition 2, we show that if the high-valuation buyer makes at least $\delta^{n-1}\Delta v$, the seller makes at most $u_S^*(\mu_0, n - 1)$ (Lemma C.5). Because by the inductive hypothesis, the seller can guarantee $u_S^*(\mu_0, n - 1)$, and we conclude the high-valuation buyer can make at most $\delta^{n-1}\Delta v$. To show Lemma C.5, we first argue that whenever the high-valuation buyer makes at least $\delta^{n-1}\Delta v$, the seller's payoff is bounded above by the value of a *constrained* version of (OPT) stated in Lemma C.4. In this program, the seller maximizes the virtual surplus subject to the constraint that the high-valuation buyer obtains at least $\delta^{n-1}\Delta v$. Relying on Proposition 1 and that the value of (OPT) at $\bar{\mu}_n$ is the seller's payoff in the posted-prices assessment, Lemma C.5 shows that the value of this constrained program is exactly $u_S^*(\mu_0, n - 1)$.

The upper and lower bound results for $\mu_0 \in [\bar{\mu}_n, \bar{\mu}_{n+1})$ imply that $u_S^*(\mu_0)$ is the seller's unique equilibrium payoff. This is key to show that the buyer's payoff is as in Theorem 1.

Buyer's payoff for $\mu_0 \in [\bar{\mu}_n, \bar{\mu}_{n+1})$: We first argue that the low-valuation buyer's payoff is 0 in any equilibrium.

PROPOSITION 3 ($u_L = 0$ for $\mu_0 \in [\bar{\mu}_n, \bar{\mu}_{n+1})$). *Let $\mu_0 \in [\bar{\mu}_n, \bar{\mu}_{n+1})$ and suppose the following hold. First, $u_S \geq u_S^*(\mu_0)$ for all $u \in \mathcal{E}^*(\mu_0)$. Second, $u_S^*(\mu_0) = VS(\mu_0)$. Then, for all $u \in \mathcal{E}^*(\mu_0)$, $u_L = 0$.*

In other words, when, as we have argued above, the seller can capture the entirety of the maximum virtual surplus, there is nothing left for the low-valuation buyer.

PROOF. We have that for all $u \in \mathcal{E}^*(\mu_0)$, the following holds:

$$u_S^*(\mu_0) + u_L \leq u_S + u_L \leq u_S^*(\mu_0),$$

where the first inequality follows from $u_S \geq u_S^*(\mu_0)$, and the second inequality follows from Lemma 2 and that $u_S^*(\mu_0)$ is the value of (OPT). It follows that $u_L = 0$ for all $u \in \mathcal{E}^*(\mu_0)$. $\square$

As we show in Appendix C.2, that the seller can capture the maximum virtual surplus also implies the uniqueness of the high-valuation buyer's payoff except at $\bar{\mu}_n$, as the solution to (OPT) is unique except at $\bar{\mu}_n$.

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 D). To do this, we first complete the equilibrium assessment by specifying the seller's and the buyer's strategy after every history (see Appendices D.1–D.2). We then show that given beliefs and continuation payoffs, neither the buyer nor the seller have a one shot deviation (Appendix D.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 $\{\bar{\mu}_n\}_{n \geq 1}$, 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 $\{\bar{\mu}_n\}_{n \geq 1}$ needs to be determined jointly with the buyer's strategy, to which we turn next.

Buyer's strategy (Appendix D.1) To complete the buyer's strategy, we first classify mechanisms according to whether they satisfy the participation and truth-telling 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 payoff 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 her continuation payoff. 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 her continuation payoff. It remains to show that the seller has no one-shot deviations.

The seller has no one-shot deviations (Appendix D.3): To show that the seller does not have one-shot deviations, we rely once again on (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 Appendix D.3 that the payoff from any deviation at history h^t is bounded above by the value of (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 $(0, u_H^*(\mu_{t+1}), u_S^*(\mu_{t+1}))$. 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 $\bar{\mu}_n$ 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, $u_H^*(\bar{\mu}_n)$. 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 buyer's participation and truth-telling constraint *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 (2022), we endow the seller with a 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 (2022) also allows us to simplify the buyer's equilibrium behavior, so that for the most part we were able to focus on the strategic considerations that pertain to the seller.

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. Program (OPT) is exactly like the relaxed program in standard mechanism design: If a solution to (OPT) satisfies the ignored constraints, then a PBE of the mechanism-selection game

exists that implements the solution to (OPT). However, as the analysis of Laffont and Tirole (1993) underscores, it is not immediate that the solution to the relaxed problem satisfies *all* constraints, and hence can be implemented as an equilibrium outcome.¹⁸ 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 with little knowledge of how continuation equilibria look like, which reinforces its applicability (recall the results in Proposition 1).

It is an open question whether posted prices are optimal outside of the binary-values case. In particular, we expect that belief dynamics will be more difficult to analyze. To see this, note that incentive compatibility implies that the probability of trade is increasing in the buyer's value only on average (cf. footnote 16). Thus, in contrast to the analysis of the posted-prices game, we lose the result that conditional on no trade the seller's beliefs are a truncation of the prior. With binary types, we circumvented this issue with Proposition 1, (a), which shows that the solution to (OPT) never induces beliefs above the prior. We expect that arguments similar to those in the proof of Proposition 1 can be used to rule out beliefs that dominate the prior in the hazard rate order. However, outside the binary-types case, this is not enough to restrict continuation beliefs to be truncations of the prior.

Whereas some of these difficulties could be circumvented by taking advantage that (OPT) is in essence an information design problem, this would require establishing further results. For instance, with finitely many types, information design results bound the number of beliefs induced at a solution of (OPT) by the number of types (Rockafellar, 1970). However, the optimality of posted prices implicitly requires that the seller uses at most two beliefs. Furthermore, with a continuum of types, 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 Dworzak 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.

APPENDIX A: AUXILIARY RESULTS

Appendix A collects a series of lemmas that we use in the proofs, starting from the following inequality, which we state as a lemma for easy reference.

LEMMA A.1 (Maximum surplus). *For any μ_0 and any payoff $u \in \mathcal{E}^*(\mu_0)$, the following holds:*

$$u_S + \mu_0 u_H + (1 - \mu_0) u_L \leq \mu_0 v_H + (1 - \mu_0) v_L. \quad (\text{A.1})$$

LEMMA A.2 (Equilibrium payoffs for $\mu \in \{0, 1\}$). *The following hold:*

¹⁸In their binary-type regulation model, Laffont and Tirole (1993) show that the solution to the relaxed program violates the analogue to the low-valuation buyer's truth-telling constraint.

(a) For $\mu_0 = 0$, $\mathcal{E}^*(\mu_0) = \{(0, \Delta v, v_L)\}$.

(b) For $\mu_0 = 1$, $\mathcal{E}^*(\mu_0) = \{(0, 0, v_H)\}$.

PROOF OF LEMMA A.2. We first prove part (a). Let $\bar{u}_L(0)$ denote the low-valuation buyer's highest equilibrium payoff at $\mu_0 = 0$. We claim that $\bar{u}_L(0) = 0$. Toward a contradiction, assume not and consider a PBE assessment in which $u_L \geq \bar{u}_L(0) - \epsilon > 0$ for some $\epsilon > 0$. Lemma A.1 implies $u_S + u_L \leq v_L$, so that $u_S \leq v_L - (\bar{u}_L(0) - \epsilon)$. Consider the following deviation for the seller: the seller offers the buyer a mechanism that sells the good with probability 1 at $v_L - \delta \bar{u}_L(0) - \eta$ for some small $\eta > 0$. The low-valuation buyer must accept the mechanism: Indeed, if the low-valuation buyer rejects the mechanism with positive probability, the seller's beliefs upon rejection are 0, in which case the low-valuation buyer obtains at most $\delta \bar{u}_L(0)$, a contradiction. It follows that the seller has a deviation: The seller's payoff under the new mechanism is $v_L - \delta \bar{u}_L(0) - \eta$, which for $\eta < (1 - \delta)\bar{u}_L(0) + \epsilon$ is larger than $v_L - (\bar{u}_L(0) - \epsilon)$. This contradicts that an assessment exists in which $u_L > 0$, and hence $\bar{u}_L(0) = 0$. Note that the same proof shows that the seller obtains at least v_L in any PBE assessment. Finally, $u_S \geq v_L$ and Lemma A.1 imply that $v_L \leq u_S + u_L \leq v_L$. Thus, the seller's unique equilibrium payoff is v_L .¹⁹ That $u_H = \Delta v$ follows from noting that the seller trades with v_L immediately with probability 1 and incentive compatibility. Indeed, given the seller's belief μ_0 and an equilibrium payoff, u , let $Q_L(\mu_0, u)$, $Q_H(\mu_0, u)$ denote the total discounted probability of trade with v_L and v_H , respectively. Then the equilibrium payoffs satisfy that $u_L + \Delta v Q_L(\mu_0, u) \leq u_H \leq u_L + \Delta v Q_H(\mu_0, u)$. That $u_H = \Delta v$ follows from noting that $Q_L(\mu_0, u) = 1$ implies $Q_H(\mu_0, u) = 1$, and $u_L = 0$.

We now prove part (b). Similar steps as those used to establish that $\bar{u}_L(0) = 0$ imply that $\bar{u}_H(1)$ —the high-valuation buyer's maximum equilibrium payoff at $\mu_0 = 1$ —is 0 and that the seller's equilibrium payoff is at least v_H . Furthermore, because $u_S + u_H \leq v_H$ (cf. Lemma A.1) and $u_H = 0$, the seller's unique equilibrium payoff is v_H . Finally, it is immediate that $u_H = 0$ implies that $u_L = 0$. $\square$

LEMMA A.3 (Buyer's rents vs. seller's payoff). Let $\mu_0 \in \Delta(V)$ and $u \in \mathcal{E}^*(\mu_0)$ be such that $u_L \geq F$ and $u_H > F + \Delta v$. Then $u_S < v_L - F$.

PROOF OF LEMMA A.3. Lemma A.1 and the conditions on $u \in \mathcal{E}^*(\mu_0)$ imply $u_S \leq \mu_0 v_H + (1 - \mu_0)v_L - \mu_0 u_H - (1 - \mu_0)u_L < v_L - F$. $\square$

LEMMA A.4 (Beliefs fall after price rejection). Let $\mu_0 \in \Delta(V)$ and consider the mechanism that sells the good with probability 1 at a price of p . That is, $\beta(\cdot | v_L) = \beta(\cdot | v_H)$ and $(q(\cdot), x(\cdot)) = (1, p)$. Then, if this mechanism is rejected with positive probability in a PBE assessment, then the seller's beliefs conditional on rejection, μ_R , satisfy that $\mu_R \leq \mu_0$.

PROOF OF LEMMA A.4. Toward a contradiction, suppose that $\mu_R \in (\mu_0, 1]$. Note it cannot be that $\mu_R = 1$ as this would imply that $v_L - p \geq 0 \geq v_H - p$ (cf. Lemma A.2, (b)),

¹⁹That a PBE assessment exists that delivers the equilibrium payoff $(0, \Delta v, v_L)$ follows from the construction in Appendix D.

a contradiction. Then $\mu_R \in (\mu_0, 1)$ so that the low-valuation buyer rejects with positive probability. Then the following must hold:

$$v_H - p \leq \delta u_H(\mu_R), \quad v_L - p = \delta u_L(\mu_R).$$

Because the low-valuation buyer can always imitate the high-valuation buyer's strategy, the continuation payoffs satisfy $\delta u_H(\mu_R) \leq \delta u_L(\mu_R) + \delta \Delta v$. This is a contradiction as $\Delta v > 0$ implies that we cannot have

$$v_H - p \leq \delta u_H(\mu_R) \leq v_L - p + \delta \Delta v.$$

It then follows that $\mu_R \in [0, \mu_0]$. □

APPENDIX B: OMITTED PROOFS FROM SECTION 4.2

PROOF OF PROPOSITION 1. To see that part (a) holds, let $A = [\mu_0, 1]$. Toward a contradiction, suppose that

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

Associated to any $u(\mu_1) \in \mathcal{E}^*(\mu_1)$ chosen by the policy that solves (OPT) there is a mechanism $\mathbf{M}(\mu_1)$ and continuation payoffs, $(u_L(\mu_2), u_H(\mu_2), u_S(\mu_2)) \in \mathcal{E}^*(\mu_2)$ for the buyer and the seller that are feasible in equilibrium.²⁰ Define $u_{H|L}(\mu_1)$ as

$$\begin{aligned} u_{H|L}(\mu_1) &= \int_{\Delta(V)} [\Delta v q^{\mathbf{M}(\mu_1)}(\mu_2) \\ &\quad + (1 - q^{\mathbf{M}(\mu_1)}(\mu_2)) \delta (u_H(\mu_2) - u_L(\mu_2))] \frac{1 - \mu_2}{1 - \mu_1} \tau^{\mathbf{M}(\mu_1)}(d\mu_2), \end{aligned}$$

and note that

$$\begin{aligned} VS((\tau^{\mathbf{M}(\mu_1)}, q^{\mathbf{M}(\mu_1)}), u^{\mathbf{M}(\mu_1)}, \mu_1) \\ = u_S(\mu_1) + u_L(\mu_1) + \mu_1(u_H(\mu_1) - u_{H|L}(\mu_1) - u_L(\mu_1)). \end{aligned}$$

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

$$\begin{aligned} VS((\tau^{\mathbf{M}(\mu_1)}, q^{\mathbf{M}(\mu_1)}), u^{\mathbf{M}(\mu_1)}, \mu_1) \\ + (1 - \mu_1) \left(\frac{\mu_1}{1 - \mu_1} - \frac{\mu_0}{1 - \mu_0} \right) (u_{H|L}(\mu_1) + u_L(\mu_1) - u_L(\mu_1)) \\ = u_S(\mu_1) + u_L(\mu_1) + \mu_1(u_H(\mu_1) - u_{H|L}(\mu_1) - u_L(\mu_1)) \\ + (1 - \mu_1) \left(\frac{\mu_1}{1 - \mu_1} - \frac{\mu_0}{1 - \mu_0} \right) (u_{H|L}(\mu_1) + u_L(\mu_1) - u_L(\mu_1)) \end{aligned}$$

²⁰Lemma D.1 in Doval and Skreta (2022) shows that restricting attention to PBE assessments in which the seller plays a pure strategy is without loss of generality.

$$\begin{aligned}
&= u_S(\mu_1) + u_L(\mu_1) + (1 - \mu_1) \left(\frac{\mu_1}{1 - \mu_1} - \frac{\mu_0}{1 - \mu_0} \right) (u_H(\mu_1) - u_L(\mu_1)) \\
&\quad + (1 - \mu_1) \frac{\mu_0}{1 - \mu_0} (u_H(\mu_1) - (u_{H|L}(\mu_1) + u_L(\mu_1))) \\
&\geq u_S(\mu_1) + u_L(\mu_1) + (1 - \mu_1) \left(\frac{\mu_1}{1 - \mu_1} - \frac{\mu_0}{1 - \mu_0} \right) (u_H(\mu_1) - u_L(\mu_1)), \quad (\text{B.1})
\end{aligned}$$

with the inequality being strict whenever $u_H(\mu_1) > u_L(\mu_1) + u_{H|L}(\mu_1)$.²¹

Consider the mechanism (τ'_0, q'_0) , which coincides with (τ_0, q_0) except that whenever $\mu_1 \in \mathcal{A}$ 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 (B.1), the payoff difference between the two policies is

$$\begin{aligned}
&VS((\tau'_0, q'_0), u', \mu_0) - VS((\tau_0, q_0), u, \mu_0) \\
&= (1 - \delta) \int_{\mathcal{A}} (u_S(\mu_1) + u_L(\mu_1) \\
&\quad + (1 - \mu_1) \left(\frac{\mu_1}{1 - \mu_1} - \frac{\mu_0}{1 - \mu_0} \right) (u_H(\mu_1) - u_L(\mu_1))) \tau_0(d\mu_1) \\
&\quad + \int_{\mathcal{A}} (1 - \mu_1) \frac{\mu_0}{1 - \mu_0} (u_H(\mu_1) - u_{H|L}(\mu_1) - u_L(\mu_1)) \tau_0(d\mu_1) > 0,
\end{aligned}$$

contradicting the optimality of the original policy.

To see that part (b) in Proposition 1 holds, assume toward a contradiction that $\mu_1 < 1$ exists 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

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

since $\delta < 1$ and $\hat{v}_L(\mu_0) < 0$, where we implicitly use Lemma A.2, (a). This is a contradiction as long as $\int_{[0,1]} q_0(\mu_1) \tau_0(d\mu_1) > 0$. $\square$

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

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

²¹One interpretation of Equation (B.1) is the following. In the equilibrium of $G(\mu_1)$ with payoff vector $(u_L(\mu_1), u_H(\mu_1), u_S(\mu_1))$, 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 high-valuation 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.

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

where the second equality follows from the constraint that τ_0 is Bayes' plausible for μ_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

$$\tau_0(\{1\}) = \int_0^{\mu_0} \frac{\mu_0 - \mu_1}{1 - \mu_1} G(d\mu_1) \quad \text{and} \quad \tau_0(U') = \int_{U'} \frac{1 - \mu_0}{1 - \mu_1} G(d\mu_1), \quad U' \subset [0, \mu_0).$$

The result follows. $\square$

APPENDIX C: OMITTED PROOFS FROM SECTION 5.1

Appendix C is organized as follows. Appendix C.1 recursively constructs 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 $\{\bar{\mu}_n\}_{n \geq 0}$ is well-defined. Second, taking as given that $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$ is a PBE assessment and, therefore, that the buyer's and the seller's payoffs under $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$ are equilibrium payoffs (i.e., $(0, u_H^*(\mu_0), u_S^*(\mu_0)) \in \mathcal{E}^*(\mu_0)$), Appendix C.2 provides the proof that the set of equilibrium payoffs is as in the statement of Theorem 1.

C.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 payoff: Whereas in the posted-prices assessment the low-valuation buyer's payoff is 0, the high-valuation buyer's payoff depends on the seller's prior and we denote it by $u_H^*(\mu_0)$. Inductively, we show that for $n \geq 0$ and $\mu_0 \in [\bar{\mu}_n, \bar{\mu}_{n+1})$,

$$u_H^*(\mu_0) = \delta^n \Delta v. \quad (\text{C.1})$$

First, fix $\mu_0 \in [0, \bar{\mu}_1)$. 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 $\mu_0 \in [\bar{\mu}_n, \bar{\mu}_{n+1})$. Suppose we have already shown that $u_H^*(\bar{\mu}_m) =$

$\delta^m \Delta v$ for $m \leq n-1$. Then the high-valuation buyer's payoff under the equilibrium strategy satisfies

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

Note that the high-valuation buyer is indifferent between reporting v_L and reporting v_H along the path of play. This is immediate for $\mu_0 \in [0, \bar{\mu}_1)$. When $\mu_0 \in [\bar{\mu}_n, \bar{\mu}_{n+1})$, the buyer's payoff from reporting v_L at the beginning of the game and then following her equilibrium strategy is $\delta u_H^*(\bar{\mu}_{n-1}) \equiv \delta^n \Delta v$. 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 payoff: Consider a seller with belief $\mu_0 \in [\bar{\mu}_n, \bar{\mu}_{n+1})$. The mechanism used by the seller determines $(\tau_{\mu_0}^*, q_{\mu_0}^*)$ defined as follows. First, if $\mu_0 \in [0, \bar{\mu}_1)$, we have that $(\tau_{\mu_0}^*(\mu_0), q_{\mu_0}^*(\mu_0)) = (1, 1)$. Second, for $n \geq 1$ and $\mu_0 \in [\bar{\mu}_n, \bar{\mu}_{n+1})$,

$$\tau_{\mu_0}^*(1) = \frac{\mu_0 - \bar{\mu}_{n-1}}{1 - \bar{\mu}_{n-1}} = 1 - \tau_{\mu_0}^*(\bar{\mu}_{n-1}), \quad q_{\mu_0}^*(1) = 1 = 1 - q_{\mu_0}^*(\bar{\mu}_{n-1}).$$

To define the seller's payoff, we first define the threshold beliefs $\{\bar{\mu}_n\}_{n \geq 0}$ and the seller's payoff in the equilibrium when his belief is $\bar{\mu}_n$. Letting $\bar{\mu}_0 = 0$, define $u_S^*(0) = v_L$. Define $\bar{\mu}_1$ as the belief $\mu \in [0, 1]$ such that

$$\begin{aligned} & \mu v_H + (1 - \mu) \delta \left[u_S^*(0) + (1 - 0) \left(\frac{0}{1 - 0} - \frac{\mu}{1 - \mu} \right) u_H^*(0) \right] \\ & = \mu v_H + (1 - \mu) \hat{v}_L(\mu), \end{aligned} \quad (\text{C.2})$$

where the RHS of Equation (C.2) uses Equation (1) to express the revenue from selling at v_L in terms of the virtual surplus. Simple algebra shows that $\bar{\mu}_1 = v_L/v_H$ (cf. Equation (2)). Let $u_S^*(\bar{\mu}_1)$ denote the value of the expression in Equation (C.2) evaluated at $\mu = \bar{\mu}_1$. Suppose we have defined $u_S^*(\bar{\mu}_m)$ for $1 \leq m \leq n-1$. Define $\bar{\mu}_n$ to be the belief $\mu \in [0, 1]$ such that

$$\begin{aligned} & \frac{\mu - \bar{\mu}_{n-1}}{1 - \bar{\mu}_{n-1}} v_H + \frac{1 - \mu}{1 - \bar{\mu}_{n-1}} \delta \left[u_S^*(\bar{\mu}_{n-1}) + (1 - \bar{\mu}_{n-1}) \left(\frac{\bar{\mu}_{n-1}}{1 - \bar{\mu}_{n-1}} - \frac{\mu}{1 - \mu} \right) u_H^*(\bar{\mu}_{n-1}) \right] \\ & = \frac{\mu - \bar{\mu}_{n-2}}{1 - \bar{\mu}_{n-2}} v_H + \frac{1 - \mu}{1 - \bar{\mu}_{n-2}} \delta \left[u_S^*(\bar{\mu}_{n-2}) \right. \\ & \quad \left. + (1 - \bar{\mu}_{n-2}) \left(\frac{\bar{\mu}_{n-2}}{1 - \bar{\mu}_{n-2}} - \frac{\mu}{1 - \mu} \right) u_H^*(\bar{\mu}_{n-2}) \right], \end{aligned} \quad (\text{C.3})$$

and define $u_S^*(\bar{\mu}_n)$ to be the value of the expression in Equation (C.3) evaluated at $\mu = \bar{\mu}_n$.

Lemma C.1 shows the threshold beliefs determined by Equation (C.3) are well-defined and satisfy $\bar{\mu}_0 = 0 < \bar{\mu}_1 < \dots < \bar{\mu}_n < \bar{\mu}_{n+1} < \dots$

LEMMA C.1. *The sequence of thresholds defined by Equation (C.3) is increasing.*

PROOF OF LEMMA C.1. Recall that $\bar{\mu}_0 = 0$ and as argued below Equation (C.2), $\bar{\mu}_1 = v_L/v_H$. Suppose we have shown that $\bar{\mu}_{n-1} > \bar{\mu}_{n-2}$. We show that $\bar{\mu}_n > \bar{\mu}_{n-1}$.

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

$$\begin{aligned} \Delta_n(\mu; \bar{\mu}_{n-1}, \bar{\mu}_{n-2}) &= \frac{\mu - \bar{\mu}_{n-1}}{1 - \bar{\mu}_{n-1}} v_H \\ &+ \frac{1 - \mu}{1 - \bar{\mu}_{n-1}} \delta \left[u_S^*(\bar{\mu}_{n-1}) + (1 - \bar{\mu}_{n-1}) \left(\frac{\bar{\mu}_{n-1}}{1 - \bar{\mu}_{n-1}} - \frac{\mu}{1 - \mu} \right) u_H^*(\bar{\mu}_{n-1}) \right] \\ &- \left\{ \frac{\mu - \bar{\mu}_{n-2}}{1 - \bar{\mu}_{n-2}} v_H + \frac{1 - \mu}{1 - \bar{\mu}_{n-2}} \delta \left[u_S^*(\bar{\mu}_{n-2}) \right. \right. \\ &\quad \left. \left. + (1 - \bar{\mu}_{n-2}) \left(\frac{\bar{\mu}_{n-2}}{1 - \bar{\mu}_{n-2}} - \frac{\mu}{1 - \mu} \right) u_H^*(\bar{\mu}_{n-2}) \right] \right\}, \end{aligned} \quad (\text{C.4})$$

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

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

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

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

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

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

$$\begin{aligned} U_S^*(h^t) &= \frac{\mu_t^*(h^t) - \bar{\mu}_{n-1}}{1 - \bar{\mu}_{n-1}} v_H + \frac{1 - \mu_t^*(h^t)}{1 - \bar{\mu}_{n-1}} \delta \left[u_S^*(\bar{\mu}_{n-1}) \right. \\ &\quad \left. + (1 - \bar{\mu}_{n-1}) \left(\frac{\bar{\mu}_{n-1}}{1 - \bar{\mu}_{n-1}} - \frac{\mu_t^*(h^t)}{1 - \mu_t^*(h^t)} \right) u_H^*(\bar{\mu}_{n-1}) \right] \\ &\equiv u_S^*(\mu_t^*(h^t)). \end{aligned} \quad (\text{C.5})$$

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

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

and note that $R^{(\tau^*, q^*)}(\mu_0, \mu_0) = u_S^*(\mu_0)$.

Equation (C.3) implies that under the specification of equilibrium play, when the seller's belief is $\bar{\mu}_n$, 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 high-valuation buyer 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's and the seller's strategies off the path of play. For future reference, let $\mathcal{U}_H^*$ denote the following correspondence:

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

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

C.2 Omitted proofs from Section 5.1

Appendix C.2 completes the steps to show that the set of equilibrium payoffs of $G(\mu_0)$ is as described in Equation (3). In what follows, we first state the inductive hypothesis. We then prove the base case, which establishes Equation (3) for $\mu_0 \in [0, \bar{\mu}_1)$ (Appendix C.2.1). We then prove the inductive step in Appendix C.2.2. Along the way, we provide the omitted proofs of the statements in Section 5.1 regarding the properties of the solution to (OPT) (Lemmas C.2 and C.3). In what follows, $\bar{u}_S(\mu_0)$, $\underline{u}_S(\mu_0)$ denote the seller's maximum and minimum equilibrium payoffs when his belief is μ_0 . Moreover, for $n \geq 0$ and $\mu_0 \geq \bar{\mu}_n$, we let $u_S^*(\mu_0, n)$ denote the seller's payoff from "emulating" the strategy in the posted-prices assessment that sells the good to the low-valuation buyer in n periods from now. Formally,

$$u_S^*(\mu_0, n) = \frac{\mu_0 - \bar{\mu}_{n-1}}{1 - \bar{\mu}_{n-1}} v_H + \frac{1 - \mu_0}{1 - \bar{\mu}_{n-1}} \delta R^{(\tau^*, q^*)}(\bar{\mu}_{n-1}, \mu_0). \quad (\text{C.6})$$

Inductive hypothesis: Fix $n \in \mathbb{N}_0$. The inductive hypothesis $P(n)$ is given by $P(n)$. The following hold:

- (n.1) For all $\mu_0 \geq \bar{\mu}_n$, $\underline{u}_S(\mu_0) \geq u_S^*(\mu_0, n)$.
- (n.2) For all $\mu_0 \in [\bar{\mu}_n, \bar{\mu}_{n+1})$, Equation (3) in Theorem 1 holds.

C.2.1 Base case We first show that $P(0) = 1$.

PROPOSITION C.1 (Seller payoff guarantee for $\mu_0 \geq \bar{\mu}_0$). *For all $\mu_0 \geq \bar{\mu}_0$, $\underline{u}_S(\mu_0) \geq u_S^*(\mu_0, 0)$.*

PROOF OF PROPOSITION C.1. Define $M_S = \{\mu_0 \in \Delta(V) : \underline{u}_S(\mu_0) < v_L\}$. Toward a contradiction, suppose that M_S is nonempty and let $\underline{\mu}_S = \inf M_S$. We consider two cases.

Case 1: $\underline{\mu}_S \in M_S$. Let $\bar{u}_L(\underline{\mu}_S) = \sup\{u_L : u \in \mathcal{E}^*(\underline{\mu}_S)\}$. Suppose the seller with prior $\underline{\mu}_S$ offers mechanism $\mathbf{M}_F$ that sells the good at price $v_L - \delta F$ for $F > 0$. We argue that if $\mathbf{M}_F$ is rejected with positive probability in a PBE assessment, then the seller's beliefs

conditional on rejection, μ_R , coincide with $\underline{\mu}_S$. Suppose that $\mathbf{M}_F$ is rejected with positive probability. Then Lemma A.4 implies that $\mu_R \leq \underline{\mu}_S$. Furthermore, it cannot be that $\mu_R = 0$ (Lemma A.2, (a) implies that $\underline{\mu}_S > 0$) because $\delta u_L(\mu_R) = 0 < \delta F$. Thus, the high-valuation buyer must reject the mechanism with positive probability. We finally rule out that $\mu_R \in (0, \underline{\mu}_S)$. In that case, we would need that $\delta F + \Delta v \leq \delta u_H(\mu_R)$ and $\delta F \leq \delta u_L(\mu_R)$, contradicting that $\underline{u}_S(\mu_R) \geq v_L$ for $\mu_R < \underline{\mu}_S$ (Lemma A.3).

Thus, $\mu_R = \underline{\mu}_S$, so that for $F = \bar{u}_L(\underline{\mu}_S) + \epsilon/\delta$, the low-valuation buyer accepts $\mathbf{M}_F$ with probability 1 and so does the high-valuation buyer (cf. Lemma A.2, (b)). Then

$$\underline{u}_S(\underline{\mu}_S) \geq v_L - \delta \bar{u}_L(\underline{\mu}_S). \quad (\text{C.7})$$

If $\bar{u}_L(\underline{\mu}_S) = 0$, we arrive at a contradiction. Suppose then that $\bar{u}_L(\underline{\mu}_S) > 0$. We now argue that if the seller with prior $\underline{\mu}_S$ offers $\mathbf{M}_F$ for $F = \delta \bar{u}_L(\underline{\mu}_S) + \epsilon/\delta$, then the buyer accepts this mechanism with probability 1. By the same logic as above, if $\mathbf{M}_F$ is rejected with positive probability, then $\mu_R = \underline{\mu}_S$ and

$$\delta F + \Delta v \leq \delta u_H(\mu_R), \quad \delta F \leq \delta u_L(\mu_R). \quad (\text{C.8})$$

Lemma A.3 implies that $\underline{u}_S(\underline{\mu}_S) < v_L - \delta \bar{u}_L(\underline{\mu}_S)$, a contradiction to Equation (C.7). Proceeding iteratively, we conclude that for all n , $\underline{u}_S(\underline{\mu}_S) \geq v_L - \delta^n \bar{u}_L(\underline{\mu}_S)$, so that as $n \rightarrow \infty$ we have that $\underline{u}_S(\underline{\mu}_S) \geq v_L$, contradicting the definition of $\underline{\mu}_S$.

Case 2: $\underline{\mu}_S \notin M_S$. Fix $\eta > 0$ and let $\bar{u}_L(\eta) = \sup\{u_L : u \in \mathcal{E}^*(\mu_0), \mu_0 \in (\underline{\mu}_S, \underline{\mu}_S + \eta)\}$. Fix $\mu_0 \in (\underline{\mu}_S, \underline{\mu}_S + \eta)$. Suppose the seller with prior μ_0 offers $\mathbf{M}_F$ that sells the good at price $v_L - \delta F$, $F > 0$. Similar logic to Case 1 implies that if $\mathbf{M}_F$ is rejected with positive probability in a PBE assessment, then $\mu_R \in (\underline{\mu}_S, \mu_0]$. In particular, Lemma A.2, (a) implies that $[0, \underline{\mu}_S]$ is nonempty, and hence we can rule out that only the low-valuation buyer rejects $\mathbf{M}_F$.

We conclude that if $F = \bar{u}_L(\eta) + \epsilon/\delta$, the low-valuation buyer accepts with probability 1 and so does the high-valuation buyer. We conclude that for all $\mu_0 \in (\underline{\mu}_S, \underline{\mu}_S + \eta)$,

$$\underline{u}_S(\mu_0) \geq v_L - \delta \bar{u}_L(\eta). \quad (\text{C.9})$$

If $\bar{u}_L(\eta) = 0$, we arrive at a contradiction since Equation (C.9) holds for all $\mu_0 \in (\underline{\mu}_S, \underline{\mu}_S + \eta)$ contradicting the definition of $\underline{\mu}_S$.

Suppose then that $\bar{u}_L(\eta) > 0$. We now argue that for $\mu_0 \in (\underline{\mu}_S, \underline{\mu}_S + \eta)$ the mechanism $\mathbf{M}_F$ must be accepted with probability 1 for $F = \delta \bar{u}_L(\eta) + \epsilon/\delta$. To see this, note that if $\mathbf{M}_F$ is rejected with positive probability, by the same logic as above $\mu_R \in (\underline{\mu}_S, \mu_0]$, so that Equation (C.8) holds. Lemma A.3 implies that $\mu_R \in (\underline{\mu}_S, \mu_0]$ exists such that $u_S < v_L - \delta \bar{u}_L(\eta)$, a contradiction to Equation (C.9). Proceeding iteratively, we conclude that for all n and all $\mu_0 \in (\underline{\mu}_S, \underline{\mu}_S + \eta)$,

$$\underline{u}_S(\mu_0) \geq v_L - \delta^n \bar{u}_L(\eta), \quad (\text{C.10})$$

so that as $n \rightarrow \infty$ we have that $\underline{u}_S(\mu_0) \geq v_L$ for all $\mu_0 \in (\underline{\mu}_S, \underline{\mu}_S + \eta)$, contradicting the definition of $\underline{\mu}_S$. It follows that for all $\mu_0 \in \Delta(V)$, $\underline{u}_S(\mu_0) \geq v_L = u_S^*(\mu_0, 0)$.

We conclude that part (n.1) holds. $\square$

We now show that part (n.2) holds for $\mu_0 \in [\bar{\mu}_0, \bar{\mu}_1)$, starting from the seller's payoff:

$u_S \leq u_S^*(\mu_0)$ for $\mu_0 \in [\bar{\mu}_0, \bar{\mu}_1)$: Note that $\bar{u}_S(\mu_0) \leq v_L$ as v_L is the seller's payoff in the commitment solution for $\mu_0 \in [0, \bar{\mu}_1)$. Since v_L is the seller's payoff in the posted-prices assessment, then $\bar{u}_S(\mu_0) = u_S^*(\mu_0) = v_L$. Together with Proposition C.1, this implies that the seller's equilibrium payoff is unique and coincides with v_L .

$u_L = 0$ in $[\bar{\mu}_0, \bar{\mu}_1)$: Proposition C.1 and the above argument imply that the assumptions in Proposition 3 hold for $\mu_0 \in [\bar{\mu}_0, \bar{\mu}_1)$. Hence, $u_L = 0$ for all $u \in \mathcal{E}^*(\mu_0)$

$u_H = \Delta v$ in $[\bar{\mu}_0, \bar{\mu}_1)$: It remains to show that the high-valuation buyer's payoff is unique in $[\bar{\mu}_0, \bar{\mu}_1)$. That the seller's payoff is at least v_L for $\mu_0 \in [\bar{\mu}_0, \bar{\mu}_1)$ and Lemma A.3 imply that $u_H \leq \Delta v$. Moreover, that the seller's payoff is at most v_L implies that $u_H = \Delta v$.

C.2.2 Inductive step Fix $n \geq 1$ and suppose that $P(k) = 1$ for all $k \in \{0, \dots, n-1\}$.

Part (n.1): We now show that $P(n) = 1$, starting from the lower bound on the seller's payoff.

PROPOSITION C.2 (Seller payoff guarantee for $\mu_0 \geq \bar{\mu}_n$). *For all $\mu_0 \geq \bar{\mu}_n$, $u_S(\mu_0) \geq u_S^*(\mu_0, n)$.*

The proof of Proposition C.2 relies on Proposition 2, the proof of which is in Appendix C.3.

PROOF OF PROPOSITION C.2. Fix $\mu_0 \geq \bar{\mu}_n$ and assume $u \in \mathcal{E}^*(\mu_0)$ exists such that $u_S < u_S^*(\mu_0, n)$. Let $(\Gamma, (\pi_v, r_v)_{v \in V}, \mu)$ denote a PBE assessment with payoff u and consider the following deviation for the seller. The seller offers a mechanism that sells the good at price $v_L + (1 - \delta^n)\Delta v - \delta F$, that is, for $v \in \{v_L, v_H\}$, $\beta(\mu_0|v) = 1$, $(q(\mu_0), x(\mu_0)) = (1, v_L + (1 - \delta^n)\Delta v - \delta F)$, where $F > 0$ satisfies that for $n \geq 2$, $\delta^{n-1}\Delta v < \delta^{n-1}\Delta v + F < \delta^{n-2}\Delta v$.

Because $-\Delta v(1 - \delta^n) + \delta F < -\Delta v(1 - \delta^{n-1}) < 0$, the low-valuation buyer rejects the mechanism with probability 1, so that the seller's beliefs upon rejection μ_R are determined by Bayes' rule and satisfy that $\mu_R \in [0, \mu_0]$.

For $n = 1$, the high-valuation buyer must accept the mechanism with probability 1: Proposition 2 implies $\delta\Delta v$ is the largest payoff from rejection and the high-valuation buyer gets strictly more by accepting the mechanism. The seller's payoff is then $\mu_0(v_H - \delta\Delta v - \delta F) + (1 - \mu_0)\delta v_L = u_S^*(\mu_0, 1) - \delta F$, where we use Lemma A.2, (a) to determine the seller's payoff conditional on rejection. Letting $F \rightarrow 0$, proves Proposition C.2 for $n = 1$.

For $n \geq 2$, we argue the high-valuation buyer must randomize between accepting and rejecting the mechanism. It cannot be the case that the high-valuation buyer rejects the mechanism with probability 1: Proposition 2 implies that $\delta u_H(\mu_R) \leq \delta\delta^{n-1}\Delta v < \delta(\delta^{n-1}\Delta v + F)$, where $\mu_R = \mu_0$. Similarly, it cannot be the case that the high-valuation buyer accepts the mechanism with probability 1: In that case, rejection reveals the low-valuation buyer (i.e., $\mu_R = 0$) and the high-valuation buyer's continuation payoff is $\delta\Delta v$, which is strictly larger than $\delta^n\Delta v + \delta F$ (cf. Lemma A.2, (a)).

It follows that the high-valuation buyer must be indifferent between accepting and rejecting so that the continuation payoffs after rejection satisfy that $\delta^{n-1}\Delta v + F = u_H(\mu_R)$, which can only be the case if $\mu_R = \bar{\mu}_{n-1}$. Indeed, $\mu_R \in (\bar{\mu}_{n-1}, \mu_0)$ would yield

a payoff of at most $\delta^{n-1}\Delta v$ (cf. Proposition 2) and $\mu_R < \bar{\mu}_{n-1}$ would yield a payoff of at least $\delta^{n-2}\Delta v$. This pins down the high-valuation buyer's acceptance probability to be $(\bar{\mu}_{n-1}/\mu_0)(\mu_0 - \bar{\mu}_{n-1}/(1 - \bar{\mu}_{n-1}))$. Simple algebra shows that the seller's payoff from offering the above mechanism is

$$\begin{aligned} & \frac{\mu_0 - \bar{\mu}_{n-1}}{1 - \bar{\mu}_{n-1}} v_H + \delta \frac{1 - \mu_0}{1 - \bar{\mu}_{n-1}} \left(u_S^*(\bar{\mu}_{n-1}) + \left(\frac{\bar{\mu}_{n-1} - \mu_0}{1 - \mu_0} \right) (u_H^*(\bar{\mu}_{n-1}) + F) \right) \\ &= u_S^*(\mu_0, n) - \delta \frac{\mu_0 - \bar{\mu}_{n-1}}{1 - \bar{\mu}_{n-1}} F. \end{aligned}$$

Letting $F \rightarrow 0$ completes the proof of Proposition C.2 for $n \geq 2$. $\square$

Part (n.2): We now prove that part (n.2) holds for $\mu_0 \in [\bar{\mu}_n, \bar{\mu}_{n+1})$. Because we have already established that the seller can guarantee the payoff from the posted-prices assessment for $\mu_0 \in [\bar{\mu}_n, \bar{\mu}_{n+1})$, it remains to show that $u_S^*(\mu_0)$ is an upper bound on the seller's payoff.

Define $M_n = \{\mu_0 \in [\bar{\mu}_n, \bar{\mu}_{n+1}) : \text{Equation (3) does not hold}\}$. Toward a contradiction, suppose $M_n \neq \emptyset$ and let $\mu_* = \inf M_n$. Then, for all $\epsilon > 0$, $\mu_0^\epsilon \in [\mu_*, \mu_* + \epsilon)$ exists such that either $\bar{u}_S(\mu_0^\epsilon) > u_S^*(\mu_0^\epsilon)$ or the buyer's payoff is not as in Equation (3).

Lemmas C.2 and C.3 below deliver that the seller's payoff in the posted-prices assessment is the seller's maximal equilibrium payoff in $[\bar{\mu}_n, \bar{\mu}_{n+1})$. Fix $\mu_0 \in [\mu_*, \bar{\mu}_{n+1})$. By definition, for $\mu_1 \in [0, \mu_*)$, the seller's and the low-valuation buyer's payoffs are given by $u_S^*(\mu_1)$ and 0, respectively. Furthermore, since $\mu_0 \geq \bar{\mu}_n$, the seller prefers to minimize the high-valuation buyer's continuation payoff at $\{\bar{\mu}_m\}_{m \leq n-1}$. Hence, for $\mu_0 \in [\mu_*, \bar{\mu}_{n+1})$, the objective function in Equation (11) equals

$$\begin{aligned} & \int_{[0, \mu_*)} \left(\frac{\mu_0 - \mu_1}{1 - \mu_1} v_H \right. \\ & \quad \left. + \frac{1 - \mu_0}{1 - \mu_1} \delta \underbrace{\left(u_S^*(\mu_1) + (1 - \mu_1) \left(\frac{\mu_1}{1 - \mu_1} - \frac{\mu_0}{1 - \mu_0} \right) u_H^*(\mu_1) \right)}_{R(\tau^*, q^*)(\mu_1, \mu_0)} \right) G(d\mu_1) \\ & \quad + \int_{[\mu_*, \mu_0)} \left(\frac{\mu_0 - \mu_1}{1 - \mu_1} v_H + \frac{1 - \mu_0}{1 - \mu_1} \delta \left(u_S(\mu_1) + u_L(\mu_1) \right. \right. \\ & \quad \left. \left. + (1 - \mu_1) \left(\frac{\mu_1}{1 - \mu_1} - \frac{\mu_0}{1 - \mu_0} \right) (u_H(\mu_1) - u_L(\mu_1)) \right) \right) G(d\mu_1). \end{aligned} \quad (\text{C.11})$$

LEMMA C.2. Fix $\mu_0 \in [\mu_*, \bar{\mu}_{n+1})$ and suppose G maximizes the expression in Equation (C.11) and is such that $G([0, \mu_*)) > 0$. Then G places positive probability on at most $\{\bar{\mu}_{n-2}, \bar{\mu}_{n-1}\}$.

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

$$\frac{\mu_0 - \bar{\mu}_m}{1 - \bar{\mu}_m} v_H + \frac{1 - \mu_0}{1 - \bar{\mu}_m} \delta R(\tau^*, q^*)(\bar{\mu}_m, \mu_0) - \left(\frac{\mu_0 - \mu_1}{1 - \mu_1} v_H + \frac{1 - \mu_0}{1 - \mu_1} \delta R(\tau^*, q^*)(\mu_1, \mu_0) \right)$$

$$= v_H \left[\frac{\mu_0 - \bar{\mu}_m}{1 - \bar{\mu}_m} + \delta \frac{1 - \mu_0}{1 - \bar{\mu}_m} \frac{\bar{\mu}_m - \bar{\mu}_{m-1}}{1 - \bar{\mu}_{m-1}} - \frac{\mu_0 - \mu_1}{1 - \mu_1} - \delta \frac{1 - \mu_0}{1 - \mu_1} \frac{\mu_1 - \bar{\mu}_{m-1}}{1 - \bar{\mu}_{m-1}} \right] > 0.$$

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

$$\max_{G \in \Delta(\{\bar{\mu}_0, \dots, \bar{\mu}_{n-1}\})} \sum_{m=0}^{n-1} \left[\frac{\mu_0 - \bar{\mu}_m}{1 - \bar{\mu}_m} v_H + \frac{1 - \mu_0}{1 - \bar{\mu}_m} \delta R^{(\tau^*, q^*)}(\bar{\mu}_m, \mu_0) \right] G(\{\bar{\mu}_m\}). \quad (\text{C.12})$$

However, for $m \leq n - 2$, Lemma C.1 implies that

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

since $\mu_0 > \bar{\mu}_{n-1}$. Thus, any solution to the problem in Equation (C.12) satisfies that $\alpha \in [0, 1]$ exists such that the value of this problem is given by

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

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

LEMMA C.3. *A real number $\bar{\epsilon} > 0$ exists such that for all $\mu_0 \in [\mu_\star, \mu_\star + \bar{\epsilon})$, $VS(\mu_0) = u_S^*(\mu_0)$.*

PROOF. For $\epsilon \in (0, \bar{\mu}_{n+1} - \mu_\star)$, let $\overline{VS}_\epsilon$ denote the supremum of $VS(\cdot)$ over $[\mu_\star, \mu_\star + \epsilon)$. The same arguments as the ones after Equation (12) imply that for all $\mu_0 \in [\mu_\star, \mu_\star + \epsilon)$, the following holds:

$$VS(\mu_0) \leq \max \left\{ u_S^*(\mu_0), \frac{\mu_0 - \mu_\star}{1 - \mu_\star} v_H + \frac{1 - \mu_0}{1 - \mu_\star} \delta \overline{VS}_\epsilon \right\}. \quad (\text{C.13})$$

Taking the supremum over $\mu_0 \in [\mu_\star, \mu_\star + \epsilon)$ on both sides of Equation (C.13), we obtain

$$\overline{VS}_\epsilon \leq \max \left\{ \overline{u}_{S\epsilon}, \frac{\epsilon}{1 - \mu_\star} (v_H - \delta \overline{VS}_\epsilon) + \delta \overline{VS}_\epsilon \right\}, \quad (\text{C.14})$$

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

We claim $\bar{\epsilon} > 0$ exists such that for all $\epsilon \in (0, \bar{\epsilon})$, the right-hand side of Equation (C.14) equals $\bar{u}_{S\epsilon}$. Toward a contradiction, suppose not. Then, for all $\epsilon \in (0, \bar{\mu}_{n+1} - \mu_*)$, $f(\epsilon) \in (0, \epsilon)$ exists such that

$$\bar{u}_{Sf(\epsilon)} < \frac{f(\epsilon)}{1 - \mu_*} (v_H - \delta \bar{V}S_{f(\epsilon)}) + \delta \bar{V}S_{f(\epsilon)}. \quad (\text{C.15})$$

Since $f(\epsilon) < \epsilon$, then $f(\epsilon) \rightarrow 0$ as $\epsilon \rightarrow 0$. Equation (C.14) together with Equation (C.15) imply

$$\bar{V}S_{f(\epsilon)} \leq \frac{f(\epsilon)}{1 - \mu_*} (v_H - \delta \bar{V}S_{f(\epsilon)}) + \delta \bar{V}S_{f(\epsilon)} \Rightarrow \lim_{\epsilon \rightarrow 0} \bar{V}S_{f(\epsilon)} \leq \delta \lim_{\epsilon \rightarrow 0} \bar{V}S_{f(\epsilon)},$$

a contradiction, since for all ϵ , $\bar{V}S_\epsilon \geq v_L > 0$.²² It follows that $\bar{\epsilon} > 0$ exists such that $\forall \epsilon \in (0, \bar{\epsilon})$, $\bar{V}S_\epsilon = \bar{u}_{S\epsilon}$.

We now claim that $VS(\mu_0) = u_S^*(\mu_0)$ for all $\mu_0 \in [\mu_*, \mu_* + \bar{\epsilon}]$. Toward a contradiction, suppose $\mu_0 \in [\mu_*, \mu_* + \bar{\epsilon}]$ exists such that $VS(\mu_0) > u_S^*(\mu_0)$. By continuity of u_S^* on $[\bar{\mu}_n, \bar{\mu}_{n+1})$ (see Equation (C.5)), $\eta > 0$ exists such that letting $\epsilon = \mu_0 + \eta - \mu_*$, we have

$$u_S^*(\mu_0 + \eta) = \bar{u}_{S\epsilon} < VS(\mu_0) \leq \bar{V}S_\epsilon,$$

where the first equality follows from u_S^* being increasing on $[\bar{\mu}_n, \bar{\mu}_{n+1})$ (cf. Equation (C.5)). This is a contradiction. Thus, for all $\mu_0 \in [\mu_*, \mu_* + \bar{\epsilon})$ we have that $u_S^*(\mu_0)$ is an upper bound on the seller's payoff.²³ $\square$

Lemmas C.2 and C.3 imply that $\bar{u}_S(\mu_0) = u_S^*(\mu_0)$ for all $\mu_0 \in [\mu_*, \mu_* + \bar{\epsilon}]$. By definition of μ_* , $\mu_0 \in [\mu_*, \mu_* + \bar{\epsilon})$ exists such that $\mathcal{E}^*(\mu_0)$ fails to satisfy Equation (3) because of the buyer's payoff. In what follows, we rule this out by showing that for all $\mu_0 \in [\mu_*, \mu_* + \bar{\epsilon})$, the buyer's payoff is as in Theorem 1 for $\mu_0 \in [\mu_*, \mu_* + \bar{\epsilon})$.

Low-valuation buyer's payoff is 0 in $[\mu_, \mu_* + \bar{\epsilon})$:* Proposition C.2 and $VS(\mu_0) = u_S^*(\mu_0)$ for $\mu_0 \in [\mu_*, \mu_* + \bar{\epsilon})$ imply that the assumptions in Proposition 3 hold for $\mu_0 \in [\mu_*, \mu_* + \bar{\epsilon})$. Hence, $u_L = 0$ for all $u \in \mathcal{E}^*(\mu_0)$ and all $\mu_0 \in [\mu_*, \mu_* + \bar{\epsilon})$.

High-valuation buyer's payoff in $[\mu_, \mu_* + \bar{\epsilon})$:* For any such μ_0 , suppose $u' \in \mathcal{E}^*(\mu_0)$ exists 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' . We argue that

$$u_S^*(\mu_0) = VS((\tau'_0, q'_0), u', \mu_0). \quad (\text{C.16})$$

To see this, note that we must have $VS((\tau'_0, q'_0), u', \mu_0) \geq u_S^*(\mu_0)$; otherwise $u'_S \leq VS((\tau'_0, q'_0), u', \mu_0) < u_S^*(\mu_0)$, where the first inequality follows from Lemma 2 and the

²²Note that the limit $\lim_{\epsilon \rightarrow 0} \bar{V}S_{f(\epsilon)}$ exists up to a convergent subnet because $\bar{V}S_{f(\epsilon)}$ is bounded.

²³The argument above is reminiscent to that in Fudenberg and Tirole (1991a)'s treatment of the equilibrium in the posted-prices game with binary values.

second inequality is by assumption. Since $u'_S = u^*_S(\mu_0)$, we have a contradiction. Furthermore, it cannot be the case that $VS((\tau'_0, q'_0), u', \mu_0) > u^*_S(\mu_0)$, because (τ'_0, q'_0) together with the continuation payoffs u' are feasible choices in (OPT). Equation (C.16) then follows. Thus, $((\tau'_0, q'_0), u')$ is also a solution to (OPT).

However, the proof that $u^*_S(\mu_0)$ is the value of (OPT) implies that the solution to (OPT) is unique, unless $\mu_0 = \bar{\mu}_n$. It then follows that $\mu_0 = \bar{\mu}_n$, 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 that u_H is as in Equation (3) for all $\mu_0 \in [\mu_*, \mu_* + \bar{\epsilon}]$.

We conclude that Equation (3) holds for all $\mu_0 \in [\mu_*, \mu_* + \bar{\epsilon}]$, contradicting the definition of μ_* . It follows that M_n is empty, and hence part (n.2) of the inductive statement holds.

C.3 Proof of Proposition 2

We now prove Proposition 2, which provides the bound on the high-valuation buyer's payoff. For $n = 1$, the result follows from Proposition C.1 and Lemma A.3, as $\underline{u}_S(\mu_0) \geq v_L$ implies that $u_H \leq \Delta v$. It remains to establish Proposition 2 for $n \geq 2$. Thus, fix $n \geq 2$ and assume that for $k \leq n - 1$, $P(k) = 1$. The proof of Proposition 2 relies on Lemmas C.4 and C.5:

LEMMA C.4 (Maximal seller's payoff given high-valuation buyer's rents). *Fix $n \geq 2$ and $\mu_0 \geq \bar{\mu}_n$. Let $\epsilon > 0$ be such that an equilibrium payoff $u \in \mathcal{E}^*(\mu_0)$ exists such that $u_H \geq \delta^{n-1}\Delta v + \delta\epsilon$. Then the seller's payoff is bounded above by $\underline{VS}(\mu_0, n - 1, \epsilon)$, where*

$$\underline{VS}(\mu_0, n - 1, \epsilon) \equiv \max_{(q, \tau, u)} VS((\tau, q), u, \mu_0) \quad (\text{OPT}(\mu_0, n - 1, \epsilon))$$

such that

$$\begin{aligned} & \int_{\Delta(V)} [\Delta v q(\mu) + (1 - q(\mu))\delta(u_H(\mu) - u_L(\mu))] \beta(d\mu|v_L) \\ & \geq \delta^{n-1}\Delta v + \delta\epsilon, \end{aligned} \quad (\text{R}(n - 1))$$

$$\mathbb{E}_\tau[\mu] = \mu_0, \quad (\text{BP})$$

$$(\forall \mu \in \Delta(V)) u(\mu) \in \mathcal{E}^*(\mu_0). \quad (\text{Eqbm})$$

Note that $(\text{OPT}(\mu_0, n - 1, \epsilon))$ does not reduce to (OPT) when $\epsilon = 0$ and $\mu_0 \in (\bar{\mu}_n, \bar{\mu}_{n+1})$. Indeed, for $\mu_0 \in (\bar{\mu}_n, \bar{\mu}_{n+1})$, the solution to (OPT) delivers rents $\delta^n\Delta v$ to the high-valuation buyer, which violates $(\text{R}(n - 1))$. Instead, when $\mu_0 = \bar{\mu}_n$, a solution to (OPT) exists that delivers rents $\delta^{n-1}\Delta v$ to the high-valuation buyer.

PROOF OF LEMMA C.4. Let $\mu_0 \geq \bar{\mu}_n$. Let $u \in \mathcal{E}^*(\mu_0)$ be such that $u_H \geq \delta^{n-1}\Delta v + \delta\epsilon$. Letting $\mathbf{M}$ denote the mechanism offered by the seller in the first period, we have that $\mathbf{M}$ satisfies at least the following constraints:

$$\int_{\Delta(V)} [v_H q^{\mathbf{M}}(\mu) + (1 - q^{\mathbf{M}}(\mu))\delta u_H^{\mathbf{M}}(\mu)] \beta^{\mathbf{M}}(d\mu|v_H) - x_H$$

$$\geq \delta^{n-1} \Delta v + \delta \epsilon, \quad (\text{PCH})$$

$$\int_{\Delta(V)} [v_H q^{\mathbf{M}}(\mu) + (1 - q^{\mathbf{M}}(\mu)) \delta u_H^{\mathbf{M}}(\mu)] (\beta^{\mathbf{M}}(d\mu|v_H) - \beta^{\mathbf{M}}(d\mu|v_L)) \geq x_H - x_L, \quad (\text{ICH})$$

$$\int_{\Delta(V)} [v_L q^{\mathbf{M}}(\mu) + (1 - q^{\mathbf{M}}(\mu)) \delta u_L^{\mathbf{M}}(\mu)] \beta^{\mathbf{M}}(d\mu|v_L) \geq x_L, \quad (\text{PCL})$$

where $x_v = \int_{\Delta(V)} x^{\mathbf{M}}(\mu) \beta^{\mathbf{M}}(d\mu|v)$. Equations (PCH), (ICH), and (PCL) are only a subset of the constraints the mechanism $\mathbf{M}$ must satisfy. Indeed, we are ignoring the low-valuation buyer's truth-telling constraint and (PCL) is only a necessary condition for the low-valuation buyer to participate in the mechanism.

Program ($\star$).

The seller's payoff, u_S , is bounded above by the solution to the following program: Maximize the seller's payoff by choosing (i) transfers x_H, x_L , (ii) trade probabilities $q : \Delta(V) \mapsto [0, 1]$, (iii) a Bayes' plausible posterior distribution τ , and (iv) continuation payoffs $u(\cdot) \in \mathcal{E}^*(\cdot)$, subject to the constraints (PCH), (ICH), and (PCL). Because we allow the seller to choose x_L and x_H instead of $x(\mu)$, we give the seller more degrees of freedom than in the game. In what follows, we argue that the value of ($\star$) equals $\underline{VS}(\mu_0, n-1, \epsilon)$.

It is immediate to see that x_L is chosen so that (PCL) binds. We can then write (ICH) as follows:

$$\begin{aligned} & \int_{\Delta(V)} [v_H q(\mu)(1 - q(\mu)) \delta u_H(\mu)] \beta(d\mu|v_H) - x_H \\ & \geq \int_{\Delta(V)} [\Delta v q(\mu) + (1 - q(\mu)) \delta (u_H(\mu) - u_L(\mu))] \beta(d\mu|v_L). \end{aligned} \quad (\text{ICH}')$$

We now show that (ICH') must bind at the solution to ($\star$), which in turn implies that (R($n-1$)) holds. Toward a contradiction, suppose that (ICH') does not bind. Then x_H must be chosen so that (PCH) binds. The binding constraints (PCH) and (PCL) imply that the value of ($\star$) obtains from maximizing

$$\int_{\Delta(V)} \left[q(\mu) \left(\mu v_H + (1 - \mu) \left(v_L - \frac{\mu_0}{1 - \mu_0} (\delta^{n-1} \Delta v + \delta \epsilon) \right) \right) + (1 - q(\mu)) \delta \left(u_S(\mu) + \mu u_H(\mu) + (1 - \mu) \left(u_L(\mu) - \frac{\mu_0}{1 - \mu_0} (\delta^{n-2} \Delta v + \epsilon) \right) \right) \right] \tau(d\mu), \quad (\text{C.17})$$

subject to (BP) and (Eqbm).

We argue that the solution to the above problem is not feasible for ($\star$). Indeed, the optimal value of the objective in Equation (C.17) is

$$\mu_0 v_H + (1 - \mu_0) \left(v_L - \frac{\mu_0}{1 - \mu_0} (\Delta v \delta^{n-1} + \delta \epsilon) \right).$$

To see this, note that Lemma A.1 and $\delta < 1$ imply that, for all μ , the following holds:

$$\begin{aligned} & \delta \left(u_S(\mu) + \mu u_H(\mu) + (1 - \mu) \left(u_L(\mu) - \frac{\mu_0}{1 - \mu_0} (\Delta v \delta^{n-2} + \epsilon) \right) \right) \\ & \leq \delta \left(\mu v_H + (1 - \mu) \left(v_L - \frac{\mu_0}{1 - \mu_0} (\Delta v \delta^{n-2} + \epsilon) \right) \right) \\ & < \mu v_H + (1 - \mu) \left(v_L - \frac{\mu_0}{1 - \mu_0} (\delta^{n-1} \Delta v + \delta \epsilon) \right). \end{aligned}$$

Thus, for all μ , setting $q(\mu) = 1$ is preferred to setting $q(\mu) = 0$. The linearity in μ of the term associated to q implies the result. However, this is not feasible since $q(\mu) = 1$ for all μ and the binding (PCH) violates (ICH'). It follows that (ICH') must hold with equality at a solution to $(\star)$.

Replacing the binding (PCL) and (ICH') in the seller's payoff yields that the objective in $(\star)$ is $VS((\tau, q), u, \mu_0)$. Moreover, replacing the binding (ICH') in (PCH) yields Equation (R($n - 1$)). We conclude that $(\star)$ coincides with $(\text{OPT}(\mu_0, n - 1, \epsilon))$. $\square$

LEMMA C.5 (Value of $(\text{OPT}(\mu_0, n - 1, \epsilon))$ at $\epsilon = 0$). Fix $n \geq 2$ and $\mu_0 \geq \bar{\mu}_n$. Then $VS(\mu_0, n - 1, 0) = u_S^*(\mu_0, n - 1)$.

PROOF OF LEMMA C.5. Let λ denote the Lagrange multiplier on the constraint (R($n - 1$))) in the program $(\text{OPT}(\mu_0, n - 1, \epsilon))$ for $\epsilon = 0$. In a slight abuse of notation, define $\mu_0(\lambda) = \frac{\mu_0 - \lambda}{1 - \lambda}$. The Lagrangian is given by

$$L((\tau, q), u; \lambda) = VS((\tau, q), u, \mu_0(\lambda)) - \lambda \delta^{n-1} \Delta v. \quad (\text{L})$$

That is, up to a constant, the Lagrangian corresponds to the virtual surplus of a seller with belief $\mu_0(\lambda)$. Let λ^* be such that $\mu_0(\lambda^*) = \bar{\mu}_n$. By the upper bound proof, we know that one of the solutions to maximizing the virtual surplus for a seller with belief $\bar{\mu}_n$ delivers rents $\delta^{n-1} \Delta v$. However, this solution satisfies the Bayes' plausibility constraint of a seller with belief $\bar{\mu}_n$, whereas $(\text{OPT}(\mu_0, n - 1, \epsilon))$ requires that the distribution over posteriors averages to μ_0 . Note that the policy that delivers a payoff of $u_S^*(\mu_0, n - 1)$ to the seller gives rents $\delta^{n-1} \Delta v$ to the high-valuation buyer and satisfies the Bayes' plausibility constraint at μ_0 . In what follows, we describe the policy and argue that it is a solution to $(\text{OPT}(\mu_0, n - 1, \epsilon))$. Important for this argument is that the seller with belief $\bar{\mu}_n$ finds it optimal to give the high-valuation buyer rents equal to $\delta^{n-1} \Delta v$.

Let $(\hat{\tau}, \hat{q})$ denote the following mechanism:

$$\hat{\tau}(1) = \frac{\mu_0 - \bar{\mu}_{n-2}}{1 - \bar{\mu}_{n-2}} = 1 - \hat{\tau}(\bar{\mu}_{n-2}), \quad \hat{q}(1) = 1 = 1 - \hat{q}(\bar{\mu}_{n-2}),$$

and continuation payoffs $\hat{u}(\bar{\mu}_{n-2}) = (0, \delta^{n-2} \Delta v, u_S^*(\bar{\mu}_{n-2}))$. The mechanism $(\hat{\tau}, \hat{q})$ together with the equilibrium continuation payoffs $\hat{u}(\bar{\mu}_{n-2})$ satisfies the constraints. Furthermore, we now argue that that for all $((\tau, q), u)$ that satisfy the constraints in $(\text{OPT}(\mu_0, n - 1, \epsilon))$,

$$L((\tau, q), u; \lambda^*) \leq L((\hat{\tau}, \hat{q}), \hat{u}; \lambda^*). \quad (\text{C.18})$$

To show that Equation (C.18) holds note that Proposition 1 implies that any maximizer of the LHS of Equation (C.18) satisfies the following. First, $q(\mu) = 1$ only if $\mu = 1$ and the support of τ is included in $[0, \mu(\lambda^*)) \cup \{1\}$. Second, steps similar to those in the proof of Corollary 2 imply that maximizing $L(\cdot; \lambda^*)$ subject to the remaining constraints (namely, (BP) and (Eqbm)) is equivalent to choosing $G \in \Delta([0, \bar{\mu}_n])$ to maximize

$$\begin{aligned} & \frac{\mu_0 - \bar{\mu}_n}{1 - \bar{\mu}_n} v_H + \frac{1 - \mu_0}{1 - \bar{\mu}_n} \int_{[0, \bar{\mu}_n]} \left[\frac{\bar{\mu}_n - \mu_1}{1 - \mu_1} v_H \right. \\ & \quad \left. + \frac{1 - \bar{\mu}_n}{1 - \mu_1} \delta \left(u_S^*(\mu_1) + \left(\frac{\mu_1 - \bar{\mu}_n}{1 - \bar{\mu}_n} \right) u_H^*(\mu_1) \right) \right] G(d\mu_1), \end{aligned}$$

where by the inductive hypothesis, continuation payoffs for $\mu_1 < \bar{\mu}_n$ are those in the posted-prices assessment. Replacing μ_* for $\bar{\mu}_n$ in the proof of Lemma C.2 implies that G placing an atom at $\bar{\mu}_{n-2}$ is a solution to the above expression, completing the proof. $\square$

PROOF OF PROPOSITION 2. Lemma C.4 implies that $u_S \leq \underline{VS}(\mu_0, n-1, \epsilon) \leq \overline{VS}(\mu_0, n-1, 0) = u_S^*(\mu_0, n-1)$, where the second inequality follows from noting that $\overline{VS}(\mu_0, n-1, \epsilon)$ is decreasing in ϵ . Furthermore, because the solution at $\epsilon = 0$ is unique and the constraint binds, it must be that the second inequality is strict.²⁴ The inductive hypothesis at stage $n-1$ implies that $\underline{u}_S(\mu_0) \geq u_S^*(\mu_0, n-1)$ for $\mu_0 \geq \bar{\mu}_n$, and hence $u_H \leq \delta^{n-1} \Delta v$. $\square$

APPENDIX D: $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$ 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 posted-prices assessment, by constructing the buyer's (Appendix D.1) and seller's strategy profile and system of beliefs (Appendix D.2) after every history. Appendix D.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 $G(\mu_0)$.

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.

D.1 Completing the buyer's strategy

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

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

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

²⁴Whereas multiple policies exist that maximize the virtual surplus of the seller with prior belief $\bar{\mu}_n$, only one of these policies satisfies the constraint on the high-valuation buyer's rents.

Fix a public history h^t and let μ_t denote the seller's beliefs at that public history.²⁵ To complete the buyer's strategy, we classify mechanisms, $\mathbf{M}$, in four categories:²⁶

0. Mechanisms that given the continuation values, $(0, u_H^*(\cdot))$, satisfy participation and truth-telling. 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 $\mathbf{M}$ not in $\mathcal{M}_C^0$ such that $U_{1,v_H}(\mathbf{M}) \geq 0 > U_{1,v_L}(\mathbf{M})$. Let $\mathcal{M}_C^2$ denote the set of these mechanisms.
3. Mechanisms $\mathbf{M}$ not in $\mathcal{M}_C^0$ such that $U_{1,v}(\mathbf{M}) \geq 0$ 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 $r_v^*(\mathbf{M})$ denote a maximizer of Equation (D.1).²⁷

If $\mathbf{M} \in \mathcal{M}_C^2$, specify that the low-valuation buyer rejects $\mathbf{M}$. 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 $r_{v_H}^*(\mathbf{M})$ denote a maximizer of Equation (D.1) for $v = v_H$. Because the high-valuation buyer's payoff is 0 conditional on her accepting the mechanism (Lemma A.2, (b)), $U_{1,v_H}(\mathbf{M})$ is indeed the utility of the high-valuation buyer when the seller offers mechanism $\mathbf{M}$ in the game. Then v_H 's payoff from participating in the mechanism $\mathbf{M}$ is given by $U_{1,v_H}(\mathbf{M})$, whereas the payoff from rejecting is $U_{0,v_H}(\pi_{v_H}, f) = \delta f(v_2(\mu_t, \pi_{v_H}))$, where

$$v_2(\mu_t, \pi_{v_H}) = \frac{\mu_t(1 - \pi_{v_H})}{\mu_t(1 - \pi_{v_H}) + 1 - \mu_t},$$

is the seller's belief that the buyer's value is v_H when observing a rejection, according to Bayes' rule, and $f(v_2(\mu_t, \pi_{v_H}))$ is a measurable selection from $\mathcal{U}_H^*(v_2(\mu_t, \pi_{v_H}))$. 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 $v_2(\mu_t, \pi_{v_H})$. 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

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

²⁵The belief, μ_t , is an equilibrium object, but we suppress this from the notation to keep things simple.

²⁶Gerardi and Maestri (2020) use a similar trick to complete the worker's strategy in their paper.

²⁷Even if the buyer does not participate on the equilibrium path, we still need to guarantee the reporting strategy is sequentially rational.

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 $v_2(\mu_t, \pi_{v_H}^*(\mathbf{M})) = \bar{\mu}_n$ for some $n \geq 1$, then a weight $\phi_2(\bar{\mu}_n, \mathbf{M}) \in [0, 1]$ exists that solves the following:

$$f_2^*(\mathbf{M}, \bar{\mu}_n) = (1 - \phi)\delta^{n-1}\Delta v + \phi\delta^n\Delta v. \quad (\text{D.2})$$

This weight captures the probability with which the seller, when his belief is $\bar{\mu}_n$, mixes between $\gamma^*(\bar{\mu}_n)$ and a mechanism that splits the prior $\bar{\mu}_n$ between 1 and $\bar{\mu}_{n-2}$, which we denote by $\gamma^{**}(\bar{\mu}_n)$.²⁸ For a mechanism in $\mathcal{M}_C^2$, let $r_{v_L}^*(\mathbf{M})$ denote a solution to Equation (D.1) 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 value being v_H . Thus, upon rejection, the buyer's continuation payoff is 0 for both her values. 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^* satisfy²⁹

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

for all $v \in V$. Set $r_{v_L}^*(\mathbf{M})(m_L^*) = 1$. 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 high-valuation buyer assigns to m_L^* . Given this notation, let $v_3(\mu_t, \mu_{t+1}, r)$ denote the seller's belief that the buyer's value is v_H when he observes output message μ_{t+1} . We assume that if μ_{t+1} is not consistent with m_L^* , that is, $\mu_{t+1} \in \Delta(V)^H \setminus \Delta(V)^L$, then $v_3(\mu_t, \mu_{t+1}, r) = 1$. 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 high-valuation buyer, in which case, v_3 would be consistent with Bayes' rule, or it does not, in which case, Bayes' rule where possible places no restrictions on $v_3(\mu_t, \mu_{t+1}, \cdot)$ for $\mu_{t+1} \in \Delta(V)^H \setminus \Delta(V)^L$.

Given v_3 , the high-valuation buyer obtains a payoff of

$$\hat{U}_{1,v_H}(r, m, f) = \int_{\Delta(V) \times A} (v_H q - x + \delta(1 - q)f(v_3(\mu_t, \mu_{t+1}, r))) \varphi^{\mathbf{M}}(d(\mu_{t+1}, q, x) | m),$$

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

$$r_{v_H}^* \in \arg \max_{r \in [0, 1]} r \hat{U}_{1,v_H}(r_{v_H}^*, m_L^*, f_3^*(\mathbf{M})) + (1 - r) \hat{U}_{1,v_H}(r_{v_H}^*, m_H^*, f_3^*(\mathbf{M})). \quad (\text{D.3})$$

The main theorem in Simon and Zame (1990) implies the existence of such an $(r_{v_H}^*, f_3^*(\mathbf{M}))$. Set $r_{v_H}^*(\mathbf{M})(m_L^*) = r_{v_H}^*$.

²⁸Formally, if $n = 1$, $\gamma^{**}(\bar{\mu}_n)$ is the mechanism that sells the good with probability 1 at a price of v_L , that is, $\beta^{\mathbf{M}}(\bar{\mu}_n | v_L) = \beta^{\mathbf{M}}(\bar{\mu}_n | v_H) = 1$ and $(q^{\mathbf{M}}(\bar{\mu}_n), x^{\mathbf{M}}(\bar{\mu}_n)) = (1, v_L)$. For $n \geq 2$, $\gamma^{**}(\bar{\mu}_n)$ is the mechanism such that $\beta^{\mathbf{M}}(\bar{\mu}_{n-2} | v_L) = 1$, $\beta^{\mathbf{M}}(\bar{\mu}_{n-2} | v_H) = (\bar{\mu}_{n-2}/\bar{\mu}_n)(1 - \bar{\mu}_{n-2}/\bar{\mu}_n) = 1 - \beta^{\mathbf{M}}(1 | v_H)$ and sets $(q^{\mathbf{M}}(\bar{\mu}_{n-2}), x^{\mathbf{M}}(\bar{\mu}_{n-2})) = (0, 0)$ and $(q^{\mathbf{M}}(1), x^{\mathbf{M}}(1)) = (1, v_L + (1 - \delta^{n-1})\Delta v)$.

²⁹We cannot ensure that truth-telling holds for mechanisms in $\mathcal{M}_C^3$, which is why we need the extra piece of notation.

As we did before, whenever $\nu_3(\mu_t, \mu_{t+1}, r_{v_H}^*) = \bar{\mu}_n$ for $n \geq 1$, we can define $\phi_3(\mathbf{M}, \bar{\mu}_n)$ as the weight on $\gamma^*(\bar{\mu}_n)$ implied by $f_3^*(\mathbf{M}, \bar{\mu}_n)$.

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

$$\pi_{tv}^*(h_B^t, \mathbf{M}) = \begin{cases} 1 & \text{if } \mathbf{M} \in \mathcal{M}_C^0 \cup \mathcal{M}_C^3, \\ \pi_{v_H}^*(\mathbf{M}) & \text{if } \mathbf{M} \in \mathcal{M}_C^2 \text{ and } v = v_H, \\ 0 & \text{otherwise,} \end{cases} \quad (\text{D.4})$$

and conditional on participation,

$$r_{tv}^*(h_B^t, \mathbf{M}) = \begin{cases} \delta_v & \text{if } \mathbf{M} \in \mathcal{M}_C^0, \\ r_v^*(\mathbf{M}) & \text{otherwise} \end{cases}. \quad (\text{D.5})$$

D.2 Full specification of the PBE assessment

To complete the PBE assessment for $G(\mu_0)$, 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 (D.6)).

To simplify notation, we do not explicitly include the buyer's participation decision. Instead, we follow the convention that if the buyer rejects, then $\mu_{t+1} = \emptyset$ and $(q_t, x_t) = (0, 0)$. Given a mechanism $\mathbf{M}_t$, let $z_{(\mu_{t+1}, (q_t, x_t))}(\mathbf{M}_t)$ denote the tuple $(\mathbf{M}_t, \mu_{t+1}, q_t, x_t)$, which summarizes the period- t outcomes from offering $\mathbf{M}_t$. In particular, we let $z_\emptyset(\mathbf{M}) \equiv z_{(\emptyset, (0,0))}(\mathbf{M})$. With this notation, any public history at the beginning of period t can be written as $(h^{t-1}, z_{(\mu_t, q_{t-1}, x_{t-1})}(\mathbf{M}_t), \omega_t)$, with the convention that when $t = 0$, $h^0 = \{\omega_0\}$ for some $\omega_0 \in [0, 1]$. 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 value is v_H . Formally,

$$T(\mu_0, z) = \begin{cases} \mu' & \text{if } z = z_{(\mu', \cdot)}(\mathbf{M}) \text{ and } \mathbf{M} \in \mathcal{M}_C^0, \\ 1 & \text{if } z = z_\emptyset(\mathbf{M}) \text{ and } \mathbf{M} \in \mathcal{M}_C^0 \cup \mathcal{M}_C^3 \\ & \text{or} \\ & z = z_{(\mu', \cdot)}(\mathbf{M}) \text{ and } \mathbf{M} \in \mathcal{M}_C^1 \cup \mathcal{M}_C^2, \\ \mu_0 & \text{if } z = z_\emptyset(\mathbf{M}) \text{ and } \mathbf{M} \in \mathcal{M}_C^1, \\ \nu_2(\mu_0, \pi_{v_H}^*(\mathbf{M})) & \text{if } z = z_\emptyset(\mathbf{M}) \text{ and } \mathbf{M} \in \mathcal{M}_C^2, \\ \nu_3(\mu_0, \mu', r_{v_H}^*(\mathbf{M}), r_{v_L}^*(\mathbf{M})) & \text{if } z = z_{(\mu', \cdot)}(\mathbf{M}) \text{ and } \mathbf{M} \in \mathcal{M}_C^3 \end{cases}. \quad (\text{D.6})$$

Let μ_0 denote the seller's prior in $G(\mu_0)$. Define $\mu^*(h^0) = \mu_0$ and $\Gamma^*(h^0) = \gamma^*(\mu_0)$, 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

$$\mu_t^*(h^t) = T(\mu_t^*(h^{t-1}), z(h^t)). \quad (\text{D.7})$$

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) $z = z_\emptyset(\mathbf{M})$, $\mu^*(h^t) \notin \{\bar{\mu}_n\}_{n \geq 1}$ and $\mathbf{M} \in \mathcal{M}_C^2$, (iii) $z = z_{(\mu', (0, \cdot))}(\mathbf{M})$, $\mu^*(h^t) \notin \{\bar{\mu}_n\}_{n \geq 1}$, 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

$$\Gamma_t^*(h^t)(\mathbf{M}) = \mathbb{1}[\mathbf{M} = \gamma^*(T(\mu_t^*(h^{t-1}), z(h^t)))]. \quad (\text{D.8})$$

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

$$\Gamma_t^*(h^t)(\mathbf{M}) = \begin{cases} \phi_2(\mathbf{M}, \bar{\mu}_n) & \text{if } \mathbf{M} = \gamma^*(\bar{\mu}_n), \\ 1 - \phi_2(\mathbf{M}, \bar{\mu}_n) & \text{if } \mathbf{M} = \gamma^{**}(\bar{\mu}_n), \\ 0 & \text{otherwise} \end{cases} \quad (\text{D.9})$$

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

$$\Gamma_t^*(h^t)(\mathbf{M}) = \begin{cases} \phi_3(\mathbf{M}, \bar{\mu}_n) & \text{if } \mathbf{M} = \gamma^*(\bar{\mu}_n), \\ 1 - \phi_3(\mathbf{M}, \bar{\mu}_n) & \text{if } \mathbf{M} = \gamma^{**}(\bar{\mu}_n), \\ 0 & \text{otherwise} \end{cases} \quad (\text{D.10})$$

Equations (D.4)–(D.5), together with Equations (D.7)–(D.10), define a PBE assessment where the system of beliefs is derived from the strategy profile where possible.

D.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 Appendix C.1 and Appendix D.1. That this is true for the buyer given the seller's strategy follows from the construction in Appendix D.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 $\mathcal{M}_C^0$, $\mathbf{M}'_t \in \mathcal{M}_C^0$ exists 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 $u_S^*(\mu_t^*(h^t))$.

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 $\mu_t^*(h^t)$ offers a mechanism $\mathbf{M}_t$ in $\mathcal{M}_C^2$. Let $\pi_{v_H}^*(h_B^t, \mathbf{M}_t)$, $r_{v_H}^*(h_B^t, \mathbf{M}_t)$ denote the buyer's best response as constructed in Appendix D.1. Denote by $v_2 \equiv v_2(\mu_t(h^t), \pi_{v_H}^*(h_B^t, \mathbf{M}_t))$ the seller's belief that the buyer's value is v_H when he observes nonparticipation. The seller's payoff is then

$$\begin{aligned} & \mu_t^*(h^t) \pi_{v_H}^*(h_B^t, \mathbf{M}_t) \int_{\Delta(V) \times A} [x + \delta(1 - q)v_H] \left(\sum_{v \in V} \varphi^{\mathbf{M}_t}(d(\mu, q, x)|v) r_{v_H}^*(h_B^t, \mathbf{M}_t)(v) \right) \\ & + (1 - \mu_t^*(h^t) \pi_{v_H}^*(h_B^t, \mathbf{M}_t)) \delta u_S^*(v_2), \end{aligned} \quad (\text{D.11})$$

where the continuation values after rejection are constructed using the equilibrium strategy when the seller has posterior ν_2 .³⁰ Consider an alternative mechanism, $\mathbf{M}'_t$:

$$\begin{aligned}\beta^{\mathbf{M}'_t}(1|v_H) &= \pi_{tv_H}^*(h_B^t, \mathbf{M}_t), \beta^{\mathbf{M}'_t}(\nu_2|v_H) = (1 - \pi_{tv_H}^*(h_B^t, \mathbf{M}_t)), \beta^{\mathbf{M}'_t}(\nu_2|v_L) = 1, \\ q^{\mathbf{M}'_t}(\nu_2) &= x^{\mathbf{M}'_t}(\nu_2) = 0, \\ x^{\mathbf{M}'_t}(1) &= \int_{\Delta(V) \times \mathcal{A}} x \left(\sum_{v \in V} \varphi^{\mathbf{M}_t}(d(\mu, q, x)|v) r_{tv_H}^*(h_B^t, \mathbf{M}_t)(v) \right), \\ q^{\mathbf{M}'_t}(1) &= \int_{\Delta(V) \times \mathcal{A}} q \left(\sum_{v \in V} \varphi^{\mathbf{M}_t}(d(\mu, q, x)|v) r_{tv_H}^*(h_B^t, \mathbf{M}_t)(v) \right).\end{aligned}$$

If the buyer participates and truthfully reports her type, $\mathbf{M}'_t$ gives the seller a payoff equal to the expression in Equation (D.11). 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 $u_S^*(\mu_t^*(h^t))$. Indeed, we have

$$\begin{aligned}U_S(h^t, \mathbf{M}_t) &= \int_{\Delta(V)} [x^{\mathbf{M}'_t}(\mu_{t+1}) + (1 - q^{\mathbf{M}'_t}(\mu_{t+1})) \delta u_S^*(\mu_{t+1})] \tau^{\mathbf{M}'_t}(d\mu_{t+1}) \\ &\leq VS((\tau^{\mathbf{M}'_t}, q^{\mathbf{M}'_t}), (0, u_H^{\mathbf{M}}, u_S^*), \mu_t^*(h^t)) \\ &\leq VS((\tau^{\mathbf{M}'_t}, q^{\mathbf{M}'_t}), (0, u_H^{\mathbf{M}}, u_S^*)_{\mu_{t+1} < \mu_t^*(h^t)}, (0, u_H^{\mathbf{M}}, u_S^*)_{\mu_{t+1} \geq \mu_t^*(h^t)}, \mu_t^*(h^t)) \\ &\leq \max_{(\tau, q), u_{\mu_{t+1} > \mu_t^*(h^t)}} VS((\tau^{\mathbf{M}'_t}, q^{\mathbf{M}'_t}), (0, u_H^*, u_S^*)_{\mu_{t+1} < \mu_t^*(h^t)}, (u_L, u_H, u_S)_{\mu_{t+1} \geq \mu_t^*(h^t)}, \mu_t^*(h^t)) \\ &= u_S^*(\mu_t^*(h^t)).\end{aligned}$$

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, whereas 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 $\mu_{t+1} < \mu_t^*(h^t)$, the term multiplying u_H in the virtual surplus is negative. Thus, the seller with belief $\mu_t^*(h^t)$ prefers that whenever his posterior belief μ_{t+1} is below $\mu_t^*(h^t)$, the high-valuation buyer's continuation payoff is that which minimizes her rents, that

³⁰Recall that the seller's continuation payoff only depends on his belief that the buyer's value is v_H ; it is only the high-valuation buyer's payoff, which may be different from u_H^* off the path of play.

is, $u_H^*(\mu_{t+1})$. The last inequality follows by definition. Proposition 1 and the results in Appendix C.2 imply the last equality. We conclude that the seller has no one-shot deviations.

APPENDIX E: 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 $\mathcal{M}_{\mathcal{I}}$ is such that M_i is finite for all $i \in \mathcal{I}$. This is without loss of generality when V is finite (cf. Doval and Skreta (2022)). It affords two important simplifications. First, $\mathcal{M}_{\mathcal{I}}$ 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 (cf. Aumann (1964)). Second, given a public history h^t , the set of buyer histories consistent with h^t , $H_B^t(h^t)$, is finite and, therefore, the support of $\mu_t(h^t) \in \Delta(V \times H_B^t(h^t))$ is finite.

Given the buyer's participation and reporting strategy and a mechanism $\mathbf{M}_t$ define a distribution over $\Delta(S^{\mathbf{M}_t} \times \{0, 1\} \times \mathbb{R})$ such that, for all measurable subsets $S' \times A' \subseteq S^{\mathbf{M}_t} \times \{0, 1\} \times \mathbb{R}$,

$$\rho^{(\pi, r)}(S' \times A' | v, h_B^t, \mathbf{M}_t) = \pi_{tv}(h_B^t, \mathbf{M}_t) \sum_{m \in M^{\mathbf{M}_t}} \varphi^{\mathbf{M}_t}(S' \times A' | m) r_{tv}(h_B^t, \mathbf{M}_t)(m).$$

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 payoff is given by

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

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

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

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:

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

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

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

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.

REFERENCES

- Abreu, Dilip, David Pearce, and Ennio Stacchetti (1990), "Toward a theory of discounted repeated games with imperfect monitoring." *Econometrica*, 58, 1041–1063. [872]
- Acharya, Avidit and Juan Ortner (2017), "Progressive learning." *Econometrica*, 85, 1965–1990. [868]
- Aliprantis, Charalambos and Kim Border (2013), *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer-Verlag, Berlin. [869]
- Athey, Susan and Kyle Bagwell (2008), "Collusion with persistent cost shocks." *Econometrica*, 76, 493–540. [872, 887, 905]
- Aumann, Robert J. (1964), "Mixed and behavior strategies in infinite extensive games." Tech. rep., Princeton University. [870, 911]
- Aumann, Robert J., Michael Maschler, and Richard E. Stearns (1995), "Repeated games with incomplete information." MIT press. [868]
- Ausubel, Lawrence M. and Raymond J. Deneckere (1989), "Reputation in bargaining and durable goods monopoly." *Econometrica*, 511–531. [868]
- Beccuti, Juan and Marc Möller (2018), "Dynamic adverse selection with a patient seller." *Journal of Economic Theory*, 173, 95–117. [867, 868]

Bester, Helmut and Roland Strausz (2001), “Contracting with imperfect commitment and the revelation principle: The single agent case.” *Econometrica*, 69, 1077–1098. [867, 868]

Bester, Helmut and Roland Strausz (2007), “Contracting with imperfect commitment and noisy communication.” *Journal of Economic Theory*, 136, 236–259. [867]

Board, Simon and Marek Pycia (2014), “Outside options and the failure of the Coase conjecture.” *American Economic Review*, 104, 656–671. [868]

Breig, Zachary (2022), “Repeated contracting without commitment.” *Journal of Economic Theory*, 204, 105514. [867]

Brzustowski, Thomas, Alkis Georgiadis-Harris, and Balázs Szentes (2023), “Smart contracts and the Coase conjecture.” *American Economic Review*, 113, 1334–1359. [868]

Bulow, Jeremy I. (1982), “Durable-goods monopolists.” *Journal of Political Economy*, 90, 314–332. [868]

Carrillo, Juan D. and Thomas Mariotti (2000), “Strategic ignorance as a self-disciplining device.” *The Review of Economic Studies*, 67, 529–544. [868]

Coase, Ronald H. (1972), “Durability and monopoly.” *The Journal of Law and Economics*, 15, 143–149. [868]

Deb, Rahul and Maher Said (2015), “Dynamic screening with limited commitment.” *Journal of Economic Theory*, 159, 891–928. [867]

Denicolo, Vincenzo and Paolo G. Garella (1999), “Rationing in a durable goods monopoly.” *The RAND Journal of Economics*, 44–55. [875, 876]

Doval, Laura and Vasiliki Skreta (2022), “Mechanism design with limited commitment.” *Econometrica*, 90, 1463–1500. [865, 866, 867, 868, 869, 871, 876, 888, 891, 911]

Dworczak, Piotr and Giorgio Martini (2019), “The simple economics of optimal persuasion.” *Journal of Political Economy*, 127, 1993–2048. [889]

Fiocco, Raffaele and Roland Strausz (2015), “Consumer standards as a strategic device to mitigate ratchet effects in dynamic regulation.” *Journal of Economics & Management Strategy*, 24, 550–569. [867]

Fudenberg, Drew, David Levine, and Jean Tirole (1985), “Infinite horizon models of bargaining with one-sided uncertainty.” In *Game Theoretic Models of Bargaining*, volume 73, 79. Cambridge University Press. [868, 873, 881]

Fudenberg, Drew and Jean Tirole (1991a), *Game Theory*. MIT Press. [901]

Fudenberg, Drew and Jean Tirole (1991b), “Perfect Bayesian equilibrium and sequential equilibrium.” *Journal of Economic Theory*, 53, 236–260. [870]

Gentzkow, Matthew and Emir Kamenica (2016), “A Rothschild–Stiglitz approach to Bayesian persuasion.” *American Economic Review*, 106, 597–601. [889]

Gerardi, Dino and Lucas Maestri (2020), “Dynamic contracting with limited commitment and the ratchet effect.” *Theoretical Economics*, 15, 583–623. [867, 906]

Gul, Faruk, Hugo Sonnenschein, and Robert Wilson (1986), “Foundations of dynamic monopoly and the Coase conjecture.” *Journal of Economic Theory*, 39, 155–190. [866, 867, 868, 884, 886]

Habibi, Amir (2020), “Motivation and information design.” *Journal of Economic Behavior & Organization*, 169, 1–18. [868]

Hart, Oliver D. and Jean Tirole (1988), “Contract renegotiation and Coasian dynamics.” *The Review of Economic Studies*, 55, 509–540. [866, 873]

Kamenica, Emir and Matthew Gentzkow (2011), “Bayesian persuasion.” *American Economic Review*, 101, 2590–2615. [868]

Kolotilin, Anton (2018), “Optimal information disclosure: A linear programming approach.” *Theoretical Economics*, 13, 607–635. [889]

Laffont, Jean-Jacques and Jean Tirole (1988), “The dynamics of incentive contracts.” *Econometrica*, 1153–1175. [867]

Laffont, Jean-Jacques and Jean Tirole (1993), *A Theory of Incentives in Procurement and Regulation*. MIT press. [889]

McAfee, R. Preston and Thomas Wiseman (2008), “Capacity choice counters the Coase conjecture.” *The Review of Economic Studies*, 75, 317–331. [868]

Ortner, Juan (2017), “Durable goods monopoly with stochastic costs.” *Theoretical Economics*, 12, 817–861. [868]

Peski, Marcin (2022), “Bargaining with mechanisms.” *American Economic Review*, 112, 2044–2082. [868]

Rockafellar, Ralph Tyrell (1970), *Convex Analysis*. Princeton University Press. [889]

Simon, Leo K. and William R. Zame (1990), “Discontinuous games and endogenous sharing rules.” *Econometrica*, 861–872. [875, 887, 907]

Skreta, Vasiliki (2006), “Sequentially optimal mechanisms.” *The Review of Economic Studies*, 73, 1085–1111. [867, 868]

Skreta, Vasiliki (2015), “Optimal auction design under non-commitment.” *Journal of Economic Theory*, 159, 854–890. [867]

Sobel, Joel (1991), “Durable goods monopoly with entry of new consumers.” *Econometrica*, 59, 1455–1485. [868]

Sobel, Joel and Ichiro Takahashi (1983), “A multistage model of bargaining.” *The Review of Economic Studies*, 50, 411–426. [868]

Stokey, Nancy L. (1981), “Rational expectations and durable goods pricing.” *The Bell Journal of Economics*, 112–128. [868]

Wolinsky, Asher (1991), “Durable-good monopoly with inventories.” *Economics Letters*, 37, 339–343. [868]

Co-editor Simon Board handled this manuscript.

Manuscript received 25 August, 2020; final version accepted 26 July, 2023; available online 1 August, 2023.