DISCUSSION PAPER SERIES

DP13967

OPTIMAL MECHANISM FOR THE SALE OF A DURABLE GOOD

Laura Doval and Vasiliki Skreta

INDUSTRIAL ORGANIZATION

CE PR

OPTIMAL MECHANISM FOR THE SALE OF A DURABLE GOOD

Laura Doval and Vasiliki Skreta

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- Industrial Organization

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

OPTIMAL MECHANISM FOR THE SALE OF A DURABLE GOOD

Abstract

We show that posted prices are the optimal mechanism to sell a durable good to a privately informed buyer when the seller has limited commitment in an infinite horizon setting. We provide a methodology for mechanism design with limited commitment and transferable utility. Whereas in the case of commitment, subject to the buyer's truthtelling and participation constraints, the seller's problem is a decision problem, in the case of limited commitment, the seller's problem corresponds to an intrapersonal game, where different "incarnations" of the seller represent the different beliefs he may have about the buyer's valuation.

Laura Doval - lauradoval@gmail.com
California Institute of Technology

Vasiliki Skreta - vskreta@gmail.com
UT Austin and UCL and CEPR

Optimal mechanism for the sale of a durable good[*]

Laura Doval[†]   Vasiliki Skreta[‡]

August 19, 2019

Abstract

We show that posted prices are the optimal mechanism to sell a durable good to a privately informed buyer when the seller has limited commitment in an infinite horizon setting. We provide a methodology for mechanism design with limited commitment and transferable utility. Whereas in the case of commitment, subject to the buyer's truthtelling and participation constraints, the seller's problem is a decision problem, in the case of limited commitment, the seller's problem corresponds to an intrapersonal game, where different "incarnations" of the seller represent the different beliefs he may have about the buyer's valuation.

KEYWORDS: mechanism design, limited commitment, intrapersonal equilibrium, information design, self-generation, posted price

JEL CLASSIFICATION: D84, D86

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

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

1

1 Introduction

A classic problem in mechanism design is how to sell a good to a buyer so as to maximize the seller's revenue. The answer is surprisingly simple: out of the many ways in which a seller could achieve this feat, a posted price is optimal. This result does not depend on the length of the interaction between the buyer and the seller: even if they would interact for infinitely many periods, the revenue-maximizing mechanism is to post the same price in each period. An important assumption behind this result is that the seller can commit to the whole sequence of mechanisms that the buyer faces. The commitment assumption is important because the optimal mechanism is, in general, *time-inconsistent*. When the buyer does not buy the good at the posted price, the seller has to be able to not lower the price even if there is common knowledge of unrealized gains from trade.

However, much less is known about how to sell a durable good when the seller can commit to today's mechanism, but not to the mechanism he will offer if no sale occurs. In an infinite horizon game in which the seller may only post a price in each period, Ausubel and Deneckere (1989) show the seller can achieve monopoly profits provided that (i) the buyer's valuation is drawn from a continuum, with the lowest valuation being smaller than the seller's marginal cost (i.e., the *no gap* case), and (ii) we consider the limit case in which the seller is fully patient. In a finite horizon game in which the seller can choose mechanisms in each period, Skreta (2006) shows the seller's optimal mechanism under limited commitment is a sequence of posted prices when the buyer's valuation is drawn from a continuum. The rest of the literature on the sale of a durable good assumes the seller can only post a price each period, and analyzes the effects on monopoly profits of allowing the seller to revise the chosen prices (see, e.g., Stokey (1981); Bulow (1982); Gul et al. (1986)).

In this paper, we study an *infinite-horizon mechanism-selection* game between an uninformed seller, who owns one unit of a durable good, and a privately informed buyer. The buyer's valuation is binary and fully persistent. At the beginning of each period, as long as the good has not been sold, the seller may offer the buyer a *mechanism*, the rules of which determine the allocation for that period. While the seller is committed to the rules of the mechanism for that period, he cannot commit to the mechanism he will offer the buyer in the event that the allocation is no trade.

The main result of this paper is that, among all mechanisms, posted prices

are optimal when the seller has limited commitment. Theorem 1 characterizes the revenue-maximizing Perfect Bayesian equilibrium and a strategy profile that achieves it. The latter is such that, as long as a sale has not occurred, the seller will choose a mechanism that can be implemented as a posted price. Contrary to the case in which the seller has commitment to the entire sequence of mechanisms, the sequence of posted prices is not constant, reflecting the fact that the seller makes inferences about the buyer's valuation in the event that no trade happens.

The optimality of posted prices echoes the result for the case in which the seller has commitment, showing that the *shape* of the mechanism does not depend on whether there is commitment to short- or long-term mechanisms. Moreover, this result provides a microfoundation for the strategy space used by the literature that studies the sale of a durable good by a monopolist. Inasmuch as we are interested in understanding the best outcome for the monopolist, assuming he can post prices is without loss of generality.

The key conceptual innovation of the analysis is to translate the problem of finding the revenue-maximizing PBE in the mechanism selection game between the seller and the buyer into an *auxiliary* problem, which only involves the seller. This auxiliary problem is an *intrapersonal game*, where the seller is the only player, and each of his "incarnations" represents the different beliefs he may have about the buyer's valuation. This step establishes a formal connection between the literatures on mechanism design with limited commitment and that on time inconsistency. Time inconsistency of optimal mechanisms under full commitment lies at the heart of the difficulties in mechanism design with limited commitment (Laffont and Tirole (1990)). This paper is the first to show that these two problems, mechanism design with limited commitment and intrapersonal games, are formally related and to use the solution of one to solve the other.

To understand how we arrive to this auxiliary problem, a review of the four main steps involved in the proof of Theorem 1 is useful. First, we rely on the results in our previous work, Doval and Skreta (2018), to simplify the class of mechanisms the seller will offer in any equilibrium of the game and, more importantly, the buyer's equilibrium behavior. This step reduces the search of the optimal sequence of mechanisms to those that satisfy, loosely speaking, a sequence of *participation* and *incentive compatibility* constraints. Second, we derive necessary conditions that a revenue-maximizing PBE assessment must satisfy. These necessary conditions are the dynamic analogue of establishing which participation and incentive compatibility constraints are binding in the case of commitment. We use these conditions to obtain a recursive program which provides an upper bound

3

on the seller's payoff in the payoff-maximizing PBE, analogous to the maximization of the *virtual surplus* in the case of commitment. Indeed, as in the case of commitment, in this recursive program, the seller's payoff is written solely as a function of the allocation: the transfers have been removed by using the binding constraints.

The third step shows this recursive program has a unique solution. Whereas in the case of commitment the maximization of the virtual surplus is a decision problem, it corresponds to an *intrapersonal game* in the case of limited commitment. Indeed, the best response condition of the intrapersonal equilibrium captures the restrictions imposed by sequential rationality on the seller's behavior. Whenever the seller chooses an allocation that involves no trade, he internalizes how his beliefs will change and how that change will determine the choice of mechanisms in the continuation.

The output of this third step is a mapping that assigns a choice of a mechanism to each belief the seller may hold about the buyer's valuation. The final step is to return to our dynamic incomplete information game and show how to use this mapping to construct the revenue-maximizing PBE assessment. Crucial to this step is a result in our previous work that allows us to apply self-generation techniques as in Abreu et al. (1990); Athey and Bagwell (2008) to characterize the whole set of equilibrium payoffs.

Notwithstanding the conceptual contributions of the analysis, our work also makes a methodological contribution in that it provides a *recipe* for analyzing mechanism design problems with limited commitment with transferable utility. Indeed, the four steps we described above follow very closely the prototypical steps in classical mechanism design: (i) invoke the revelation principle to simplify the class of mechanisms and the agent's behavior, (ii) find the binding constraints, (iii) maximize the virtual surplus, and (iv) show the solution satisfies any constraints that may have been ignored. As will become clear from the analysis that follows, these four steps will also be necessary in other problems of mechanism design with limited commitment and transferable utility beyond the one we consider here.

Our goal is to highlight the conceptual and technical nuances that arise in an infinite-horizon *mechanism selection* game when the designer has limited commitment, and thus our focus on the case of binary valuations. This allows us to bring to the forefront the complexities that arise because of the designer's rich action space. The extension to the case in which the buyer's valuation is drawn from a

4

continuum is of interest, and we plan to address it in future work. When possible, we highlight in footnotes how the results presented extend to the continuum.

Related Literature: The paper contributes mainly to four strands of literature. The first strand, similar to this paper, derives optimal mechanisms when the designer has limited commitment, but unlike this paper, these studies consider finite horizon settings (see Laffont and Tirole (1988); Skreta (2006, 2015); Deb and Said (2015); Fiocco and Strausz (2015); Beccuti and Möller (2018)). The second strand, similar to this paper, studies infinite horizon problems in which the designer has limited commitment and faces an agent with one of finitely many types.[1] Given the difficulties with the revelation principle when the designer has limited commitment (Laffont and Tirole (1988); Bester and Strausz (2001)), the works on this second strand either impose restrictions on the class of contracts that can be offered (e.g., Strulovici (2017); Gerardi and Maestri (2018)), or on the solution concept (e.g., Acharya and Ortner (2017)). This paper contributes to the literature in two ways. First, by appealing to the results in Doval and Skreta (2018), we avoid imposing restrictions on the length of the interaction, the class of mechanisms, and the solution concept. Second, whereas we exploit the specifics of our problem to show equilibrium existence, the methodology suggested by the paper can be imported into other settings of interest and, as we discuss throughout the paper, it can be particularly powerful in the case of transferable utility.

The third strand is the literature that starts with 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 et al. (1986)) study price dynamics[2] and establish (under some conditions) Coase's conjecture.[3] Related to this literature is the problem of dynamic bargaining with one-sided incomplete information, where an uninformed proposer each period makes a take-it-or-leave-it offer to

---

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

[2]Even when the seller can commit not to revise prices, price dynamics can arise if there are demand or cost changes. See, for example, Stokey (1979), Conlisk et al. (1984), Sobel (1991), Garrett (2016), and Ortner (2017).

[3]There is also a literature that analyzes different variations on the sale of a durable good model to understand whether Coase's conjecture still obtains. See, for instance, McAfee and Wiseman (2008); Board and Pycia (2014); Nava and Schiraldi (2019).

a privately informed receiver. Sobel and Takahashi (1983), Fudenberg et al. (1985), and Ausubel and Deneckere (1989) characterize the equilibria of this game for the case of finite horizon, infinite horizon and gap, infinite horizon and no gap, respectively. In an analogous vein, Burguet and Sakovics (1996), McAfee and Vincent (1997), Caillaud and Mezzetti (2004), and Liu et al. (2019) study equilibrium reserve-price dynamics without commitment in different settings.[4] The common thread in all these papers is that the seller's inability to commit reduces monopoly profits.

The last strand is the literature on games with time-inconsistent preferences, which builds on the seminal works of Strotz (1955), Peleg and Yaari (1973), and Pollak (1968). Our contribution to this literature is to show that optimal mechanisms under limited commitment can be obtained as a solution to an intrapersonal game, even when both the seller and the buyer have time consistent preferences. In this way, we extend the range of applications of this literature, which has already seen applications ranging from growth and development (Bernheim et al. (2015)) to saving decisions (Harris and Laibson (2001)).

Organization   The rest of the paper is organized as follows. Section 2 describes the model. Section 3 formally states the main result and describes the main steps in the proof of Theorem 1. Section 3.1 summarizes the results in Doval and Skreta (2018) used to simplify the analysis that follows; Section 3.2 derives the recursive formulation that is the basis of the intrapersonal game; Section 3.3 characterizes the unique equilibrium of the intrapersonal game; Section 3.4 uses the intrapersonal equilibrium to build a PBE assessment that delivers the revenue-maximizing PBE. Section 4 concludes.

2   Model

Primitives: Two players, a seller and a buyer, interact over infinitely many periods. The seller owns one unit of a durable good to which he attaches value 0. The buyer has private information: before her interaction with the seller starts, she observes her valuation $v \in \{v_L, v_H\} \equiv V$, with $0 \leq v_L < v_H$. Let μ_0 denote

---

[4]Dilmé and Li (Forthcoming) study revenue management when the seller cannot commit to revise prices.

the probability that the buyer's valuation is v_H at the beginning of the game. In what follows, we denote by $\Delta(V)$ the set of distributions on V.

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

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 what a *mechanism* is; Remark 1 at the end of this section explains why this definition of a mechanism entails no loss of generality. A mechanism, $\mathbf{M} = \langle (V, \beta^{\mathbf{M}}, \Delta(V)), (q^{\mathbf{M}}, x^{\mathbf{M}}) \rangle$, consists of a communication device $\beta^{\mathbf{M}}$, which maps a report $v \in V$ to a distribution with finite support on $\Delta(V)$, and an allocation rule $(q^{\mathbf{M}}, x^{\mathbf{M}})$, which maps each element μ' in $\Delta(V)$ to a probability of trade, $q^{\mathbf{M}}(\mu')$, and a payment from the buyer to the seller, $x^{\mathbf{M}}(\mu')$.[5] Denote by $\mathcal{M}$ the set of all such mechanisms.

Timing: If in period t, the good is yet to be traded, then

$t.0$ Both players observe a draw from a correlating device $\omega \in [0, 1]$.[6]

$t.1$ The seller offers the buyer a mechanism $\mathbf{M} = (\langle V, \beta^{\mathbf{M}}, \Delta(V) \rangle, (q^{\mathbf{M}}, x^{\mathbf{M}}))$.

$t.2$ The buyer observes the mechanism and decides to participate ($p = 1$) or not ($p = 0$).

$t.2.1$ If the buyer does not participate in the mechanism, the good is not sold and no payments are made; that is, the allocation is $(q, x) = (0, 0)$.

$t.2.2$ If the buyer participates in the mechanism,

    i. She chooses a report $v' \in V$, which is unobserved by the seller,

    ii. The buyer's report v' determines the distribution, $\beta^{\mathbf{M}}(\cdot|v')$, from which a posterior, μ', is drawn. This, in turn, determines the probability of

---

[5]A priori, we could have allowed a mechanism to offer a randomization over allocations, i.e. a randomization over $\{0, 1\} \times \mathbb{R}$. However, because both players have quasilinear payoffs, the seller does not benefit from randomizing over transfers (see Lemma 1 in Doval and Skreta (2019) for a proof). He may, however, benefit from randomizing over whether the good is traded.

[6]Such correlating devices are standard in repeated games (see, e.g., Fudenberg and Maskin (1986)).

trade, $q^{\mathbf{M}}(\mu')$, and the payment $x^{\mathbf{M}}(\mu')$. Both the buyer and the seller observe the drawn posterior and the realized allocation. If the good is not traded, the game proceeds to period $t+1$.

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

$$h^t = (\omega_0, \mathbf{M}_0, p_0, \mu'_0, (q_0, x_0), \dots, \omega_{t-1}, \mathbf{M}_{t-1}, p_{t-1}, \mu'_{t-1}, (q_{t-1}, x_{t-1}), \omega_t),$$

where $p_r \in \{0,1\}$ denotes the buyer's participation decision, with the restriction that if $p_r = 0$, then $\mu'_r = \varnothing$, i.e., no signal is generated, and $(q_r, x_r) = (0,0)$. Let H^t denote the set of all period t public histories; they capture what the seller knows through period t. A strategy for the seller is then $\Gamma : \cup_{t=0}^\infty H^t \mapsto \Delta(\mathcal{M})$.

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

$$h^t_B = (\omega_0, \mathbf{M}_0, p_0, v_0, \mu'_0, (q_0, x_0), \dots, \omega_{t-1}, \mathbf{M}_{t-1}, p_{t-1}, v_{t-1}, \mu'_{t-1}, (q_{t-1}, x_{t-1}), \omega_t).$$

Let H^t_B denote the set of buyer histories through period t. The buyer also knows her valuation and hence, when her valuation is v, a history through period t is an element of $\{v\} \times H^t_B$. The buyer's participation strategy is $\pi_v : \cup_{t=0}^\infty (H^t_B \times \mathcal{M}) \mapsto [0,1]$. Conditional on participating in the mechanism $\mathbf{M}$, her reporting strategy is a distribution $r_v(h^t_B, \mathbf{M}, 1) \in \Delta(V)$ for each of her types v and each $h^t_B \in H^t_B$.

A belief for the seller at the beginning of time t, history h^t, is a distribution $\mu(h^t) \in \Delta(V \times H^t_B(h^t))$, where $H^t_B(h^t)$ is the set of buyer histories consistent with the public history, h^t.

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

1. Given $\mu(h^t)$, $\Gamma(h^t)$ is sequentially rational given $(\pi_v, r_v)_{v \in V}$,

2. Given $\Gamma(h^t)$, $\pi_v(h^t_B, \cdot), r_v(h^t_B, \cdot, 1)$ are sequentially rational for all $h^t_B \in H^t_B$ and $v \in V$,

3. $\mu(h^t)$ is derived via Bayes' rule where possible (see Definition 4 in Section A.1).[7]

---

[7]Section A.1 contains the formal statement of Bayes' rule where possible. To wit, we impose the

8

We denote by $\mathcal{E}^*(\mu_0) \subseteq \mathbb{R}^3$ the set of PBE payoffs of $G^\infty(\mu_0)$ and we use (u_S, u_H, u_L) to denote a generic element of $\mathcal{E}^*(\mu_0)$, where u_S is the seller's payoff and u_H, u_L denote the buyer's payoff when her type is v_H, v_L, respectively.

The seller's highest equilibrium payoff in $G^\infty(\mu_0)$ is of particular interest to us:

$$u_S^*(\mu_0) = \max\{u_s : (u_S, u_H, u_L) \in \mathcal{E}^*(\mu_0)\}. \tag{1}$$

Theorem 1 characterizes $u_S^*(\mu_0)$ and an assessment that achieves it for all $\mu_0 \in \Delta(V)$, which we denote by $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$.

In what follows, we abuse notation in the following way. Because valuations are binary, we can think of an element in $\Delta(V)$ (a distribution over v_L and v_H) as an element of the interval $[0, 1]$ (the probability assigned to v_H). We use the latter formulation in what follows. That is, whereas the mechanism outputs a distribution over v_L and v_H, we index this distribution by the probability of v_H.

We conclude Section 2 by commenting on our modeling assumptions in Remark 1. The reader interested in Theorem 1 may skip Remark 1 without loss of continuity:

Remark 1. The main result in Doval and Skreta (2018) implies the equilibrium payoffs of $G^\infty(\mu_0)$ coincide with the equilibrium payoffs of an *abstract* mechanism-selection game, where at the beginning of each period t, the designer may offer the agent any mechanism from a set $\widetilde{\mathcal{M}}$ that includes $\mathcal{M}$. To be precise, in Doval and Skreta (2018), we endow the designer with a set of input and output messages $(M_i, S_i)_{i \in \mathcal{I}}$ and allow him to choose any mechanism $\tilde{\mathbf{M}} = \langle (M^{\tilde{\mathbf{M}}}, \beta^{\tilde{\mathbf{M}}}, S^{\tilde{\mathbf{M}}}), \alpha^{\tilde{\mathbf{M}}} \rangle$, where

1. $(M^{\tilde{\mathbf{M}}}, S^{\tilde{\mathbf{M}}}) = (M_i, S_i)$ for some $i \in \mathcal{I}$,

2. $M^{\tilde{\mathbf{M}}}$ is finite, $|M^{\tilde{\mathbf{M}}}| \geq |V|$, and $S^{\tilde{\mathbf{M}}}$ includes an image of $\Delta(V)$,

3. $\beta^{\tilde{\mathbf{M}}} : M^{\tilde{\mathbf{M}}} \mapsto \Delta^*(S^{\tilde{\mathbf{M}}})$ is the mechanism's communication device, and[8]

4. $\alpha^{\tilde{\mathbf{M}}} : S^{\tilde{\mathbf{M}}} \mapsto \Delta^*(\mathcal{A})$ is the mechanism's allocation rule, where $\mathcal{A}$ is the set of allocations.

---

following requirement. Let h and h' be two consecutive information sets for the seller. The beliefs at h' are obtained via Bayes' rule from the beliefs at h if one of the following holds: (i) conditional on reaching h, h' is reached with positive probability under the equilibrium strategy profile, or (ii) conditional on reaching h, h' is reached from h through a deviation by the seller and the buyer playing the equilibrium strategy profile.

[8]For any set C, $\Delta^*(C)$ denotes the distributions with finite support on C.

In the case of the sale of a durable good, $\mathcal{A} = \{0,1\} \times \mathbb{R}$. The main result in Doval and Skreta (2018) implies that the set of payoffs of $G^\infty(\mu_0)$ coincides with the set of payoffs of the abstract mechanism selection game when $\mathcal{A} = \{0,1\} \times \mathbb{R}$. Thus, it is without loss of generality to study the game $G^\infty(\mu_0)$ in order to characterize the seller's highest equilibrium payoff when he can choose *any* mechanism. [9]

3 Main Result

Section 3 contains the main result of the paper: the optimal mechanism for a seller of a durable good is a sequence of posted prices. To state the result, we proceed as follows. First, we define what it means for the seller to attain a certain equilibrium payoff via a sequence of posted prices, when the seller's action space consists of mechanisms (Definition 1). This is enough to informally state our main result. Second, we describe a PBE assessment, $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$, that achieves $u_S^*(\mu_0)$ and argue that, indeed, $u_S^*(\mu_0)$ can be implemented via a sequence of posted prices. The rest of this section describes the main steps to prove $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$ indeed achieves $u_S^*(\mu_0)$.

Definition 1. Let u_S denote a seller's payoff in $\mathcal{E}^*(\mu_0)$. The payoff u_S can be *implemented via a sequence of posted prices* if a PBE assessment $\langle \Gamma, (\pi_v, r_v)_{v \in V}, \mu \rangle$ exists such that in each history h^t, the seller's mechanism $\mathbf{M}$ satisfies the following: there exist $\mu_1' \neq \mu_2'$ and a price $p(h^t)$ such that

1. $\beta^{\mathbf{M}}(\mu_1'|v_L) = \beta^{\mathbf{M}}(\mu_2'|v_H) = 1$,
2. $(q^{\mathbf{M}}(\mu_1'), x^{\mathbf{M}}(\mu_1')) = (0,0)$ and $(q^{\mathbf{M}}(\mu_2'), x^{\mathbf{M}}(\mu_2')) = (1, p(h^t))$.

Thus, a mechanism is a posted price whenever it can be implemented as if the buyer could choose one of two options: either she buys the good at price $p(h^t)$, or she does not buy the good, in which case she pays nothing. Note that when the seller posts a price in each period, he does not use *noisy* communication devices: by observing the posterior, the seller knows which report the buyer submitted into the mechanism. However, the buyer's report need not be truthful in this implementation.

Theorem 1. Let $u_S^*(\mu_0)$ denote the seller's maximum revenue in a PBE. Then, $u_S^*(\mu_0)$ can be implemented via a sequence of posted prices.

---

[9]For precisely this reason, we denote $G^\infty(\mu_0)$ the *canonical game* in Doval and Skreta (2018).

10

Theorem 1 shows that, amongst all trading protocols, posted prices are optimal. This provides a microfoundation for the seller's strategy space in the literature that studies the sale of a durable good.

Theorem 1 echoes the result for the case in which the seller has commitment and faces one buyer, where it is also well known that a posted price is optimal. However, in the latter case, the seller would post the same price at all histories, which would fail to be sequentially rational in $G^{\infty}(\mu_0)$. As we show in the analysis that follows, we find that the seller's mechanism determines a price path that is decreasing in the seller's belief that the buyer's valuation is v_H, as in the literature that analyzes the Coase conjecture.

PBE assessment that achieves $u_S^*(\mu_0)$:   The proof of Theorem 1 singles out a particular PBE assessment, $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$, that achieves $u_S^*(\mu_0)$, which we now define informally (the formal statement is in Section E.2). After defining it, we use it to argue that $u_s^*(\mu_0)$ can be implemented via a sequence of posted prices. In what follows, the ratio v_L/v_H is important and we denote it by $\overline{\mu}_1$. To see why $\overline{\mu}_1$ is important, recall the following fact from second-degree price discrimination: when the seller sells at a price of v_L, he leaves rents, $\Delta v \equiv v_H - v_L$, to v_H, so that when v_H has probability μ, the seller's revenue, v_L, can be written as

$$v_L = \mu(v_H - \Delta v) + (1 - \mu)v_L = \mu v_H + (1 - \mu)(v_L - \frac{\mu}{1 - \mu}\Delta v) = \mu v_H + (1 - \mu)\hat{v}_L(\mu).$$

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 the fact that when v_L is served, so is v_H, which leaves rents to v_H. $\overline{\mu}_1$ is the belief at which $\hat{v}_L(\mu) = 0$. At that belief, the seller is indifferent between serving the buyer for both of her valuations and excluding the low-valuation buyer.

We focus here and in the rest of this section on the case in which $v_L > 0$; we return to the (less interesting) case of $v_L = 0$ at the end of this section. When $v_L > 0$, we show a sequence $0 = \overline{\mu}_0 < \overline{\mu}_1 = v_L/v_H < \cdots < \overline{\mu}_n < \ldots$ exists such that the following holds:[10]

---

[10]Hart and Tirole (1988) compare price paths and the seller's revenue across different sale modes (sale vs. leasing) in a finite-horizon setting with limited commitment and binary valuations. As in their model with sale, the optimal price path in our model is determined by the beliefs at which the seller is indifferent between trading with the low-valuation buyer in n versus $n + 1$ periods.

11

1. Along the equilibrium path:

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

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

2. Off the equilibrium path, the seller's strategy coincides with the above, except that when $\mu^*(h^t) = \overline{\mu}_n$ for some $n \geq 1$, the seller may randomize between the mechanism he offers on the path of play when his belief is $\overline{\mu}_n$ and the one he offers on the path of play when his belief is $\overline{\mu}_{n-1}$.[11]

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

Implementation via posted prices: We now argue that if the assessment described above achieves $u_S^*(\mu_0)$, then $u_S^*(\mu_0)$ can be achieved via a sequence of posted prices. Clearly, when the seller's beliefs are below $\overline{\mu}_1$, the seller's mechanism corresponds to selling the good at a price of v_L. Consider then the case in which the seller's beliefs are in $[\overline{\mu}_1, \overline{\mu}_2)$. Note that when the buyer's valuation is v_H and the realized allocation is trade, then her payoff is $v_H - v_L - (1-\delta)\Delta v = \delta\Delta v$. On the other hand, when the buyer's valuation is v_H and the realized allocation is no trade, then the seller's beliefs next period are $\overline{\mu}_0 = 0$, so that the buyer's continuation payoff is $\delta\Delta v$. That is, the buyer with valuation v_H is indifferent between obtaining the good at price $v_L + (1-\delta)\Delta v$ and not obtaining the good,

---

However, because their model has a finite horizon and they define the buyer's value as a flow payoff from consuming the good, it is not immediate to compare the price paths in both models.

[11]The need for mixing arises for technical reasons: it ensures that the buyer's continuation payoffs when her valuation is v_H are upper-semicontinuous and, thus, guarantees that a best response exists after any deviation by the seller (see Section E.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.

and paying a price of v_L in the next period. Since the buyer with valuation is v_H is indifferent between these two options, she is willing to mix between buying at price $v_L + (1-\delta)\Delta v$ and not obtaining the good. She does so in a way that the seller's belief is $\overline{\mu}_0$ when the allocation is $(0,0)$. Since $\overline{\mu}_0 = 0$, it implies that when the seller's belief is in $[\overline{\mu}_1, \overline{\mu}_2)$, the buyer buys with probability 1 at a price of $v_L + (1-\delta)\Delta v$. Working recursively through the equations, one can show that when the seller's prior is in $[\overline{\mu}_n, \overline{\mu}_{n+1})$, the mechanism is equivalent to posting a price of $v_L + (1-\delta^n)\Delta v$. In this case, the low-valuation buyer chooses the $(0,0)$ allocation, whereas the high-valuation buyer mixes so that the seller's belief is $\overline{\mu}_{n-1}$ when the allocation is $(0,0)$.

The rest of this section describes the steps to show the assessment, $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$, achieves $u_S^*(\mu_0)$. Section 3.1 applies the results from Doval and Skreta (2018) to simplify the search for the revenue-maximizing PBE. Section 3.2 derives necessary conditions that the PBE assessment, $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$, must satisfy. They are the basis for the formulation of a recursive problem, the intrapersonal game, studied in Section 3.3. The main output of Section 3.3 is a mapping $\gamma^* : \Delta(V) \mapsto \mathcal{M}$, which describes for each prior μ_0' the seller may have about the buyer, his optimal choice of mechanism. Section 3.4 discusses then how we use γ^* to construct $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$. Finally, we appeal to self-generation techniques to show $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$ is indeed a PBE assessment.

3.1 Revelation Principle

In Doval and Skreta (2018), we showed that it is without loss of generality to focus on a particular class of assessments to characterize *all* equilibrium payoffs in $\mathcal{E}^*(\mu_0)$. Lemma 1 below summarizes the properties of these assessments for future reference:

Lemma 1 (Doval and Skreta (2018)). *Let $(u_S, u_H, u_L) \in \mathcal{E}^*(\mu_0)$. Then, there exists a PBE assessment $\langle \Gamma, (\pi_v, r_v)_{v \in V}, \mu \rangle$ of $G^\infty(\mu_0)$ that achieves payoff (u_S, u_H, u_L) and satisfies the following properties:*

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

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

13

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

Part 1 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 type and each public history. For completeness, these constraints are stated in Section A.2; see Equation PC_{v,h^t} and Equation IC_{v,h^t}. As in the case of commitment to long-term mechanisms, part 1 simplifies the analysis of the buyer's behavior, by reducing it to a series of constraints. Part 2 implies we can interpret the seller's choice of a communication device as a choice of a distribution over posteriors that satisfies a Bayes' plausibility constraint (see Equation $\text{BC}_{\mu(h^t)}$ in Section A.2). As a consequence, our analysis involves aspects of *information design*. Part 3 implies the set of PBE payoffs of $G^\infty(\mu_0)$ coincides with the set of *Public PBE* payoffs of $G^\infty(\mu_0)$ (Athey and Bagwell (2008)), allowing us to invoke self-generation techniques as in Abreu et al. (1990); Athey and Bagwell (2008) to argue the assessment we construct in Section 3.4 is indeed a PBE assessment.

In what follows, by a PBE assessment, we mean one that satisfies the conditions of Lemma 1.

3.2 Binding constraints and virtual surplus

To introduce the main result of Section 3.2, Proposition 1, we introduce the definition of incentive efficiency (see Bester and Strausz (2001)):[12]

Definition 2 (Incentive efficiency). Let $\langle \Gamma, (\pi_v, r_v)_{v \in V}, \mu \rangle$ be a PBE assessment and (u_S, u_H, u_L) its associated payoff. $\langle \Gamma, (\pi_v, r_v)_{v \in V}, \mu \rangle$ is incentive efficient if no $u'_S > u_S$ exists such that $(u'_S, u_H, u_L) \in \mathcal{E}^*(\mu_0)$.

In an incentive efficient PBE, the seller at the beginning of the game is earning his best payoff consistent with the buyer's equilibrium payoff.

Proposition 1. *If $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$ is the PBE assessment that delivers the seller's best payoff in $\mathcal{E}^*(\mu_0)$, $u_S^*(\mu_0)$, the following hold:*

---

[12]Incentive efficiency as defined in Bester and Strausz (2001) is related, but different from, incentive efficiency as defined in Holmström and Myerson (1983). While the definition in Holmström and Myerson (1983) would allow for Pareto improvements for both the buyer and the seller, the definition in Bester and Strausz (2001) only considers Pareto improvements for the seller.

1. *Without loss of generality, if the buyer rejects the seller's equilibrium choice of mechanism at history h^t, the seller assigns probability 1 to the buyer's valuation being v_H.*

2. *For every history on the path of play, h^t, $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle|_{h^t}$ is incentive efficient.*

3. *For all histories on the path of $(\Gamma^*, (\pi_v^*, r_v^*)_{v \in V})$, the buyer is indifferent between participating in the mechanism and not when her valuation is v_L.*

4. *For all histories on the path of $(\Gamma^*, (\pi_v^*, r_v^*)_{v \in V})$, the buyer is indifferent between reporting v_H and v_L when her valuation is v_H.*

Part 1 implies that it is without loss of generality to assume that the buyer's value of rejecting the seller's equilibrium choice of mechanism is 0: when the seller assigns probability 1 to v_H, the seller asks for a payment of v_H. Part 1 follows from two observations: (i) By Lemma 1, the buyer's rejection of the seller's equilibrium offer is an off the path event, so that the seller's beliefs are not pinned down by Bayes' rule, and (ii) by setting the beliefs of the seller to 1 on that event, we only relax the buyer's participation constraint without affecting the seller's incentives.[13]

Part 2 is a consequence of the result that it is without loss of generality to focus on Public PBE. While in general it is not true that $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle|_{h^t}$ is a PBE assessment of $G^\infty(\mu^*(h^t))$, this is indeed the case when $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$ is a Public PBE assessment (see Proposition 1 in Doval and Skreta (2019)). Part 2 then says that along the equilibrium path the seller's continuation payoff is the highest equilibrium payoff *consistent* with the buyer's equilibrium payoff. Thus, while the seller may not be earning $u_S^*(\mu^*(h^t))$ for a history h^t on the equilibrium path, the payoffs starting from h^t are the highest equilibrium payoffs consistent with the buyer's payoff at h^t.

Parts 3 and 4 imply that the standard observations from second-degree price discrimination hold along the path of play of $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$, namely, that all surplus is extracted from the buyer when her valuation is v_L and the buyer's payoff coincides with her information rents when her valuation is v_H. Clearly, they must hold at the initial history, because either of them failing implies the seller is "leaving money on the table." That they hold after every history on the

---

[13]Part 1 of Proposition 1 is not a consequence of using PBE as opposed to sequential equilibrium, since the buyer could always tremble and reject the mechanism with a higher probability when her valuation is v_H.

path of play is a result of discounting and the linearity of payoffs in transfers, which allows us to distribute the players' payoffs across time without affecting (and sometimes even increasing) payoffs at some history.[14]

It follows from parts 3 and 4 that the participation constraint for v_L and the incentive compatibility constraint for v_H bind along the equilibrium path.[15] Recall that when the seller has commitment to long-term mechanisms, these constraints are the ones that we use to replace the transfers out of the seller's payoffs, so that they are expressed solely in terms of the allocation. In what follows, we show that the same can be done when analyzing the revenue-maximizing PBE when the seller can only commit to short-term mechanisms.

Let $u_S^*(\mu_0)$ denote the seller's highest equilibrium payoff in $G^\infty(\mu_0)$ and let $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$ denote the PBE assessment that achieves it. Lemma 1 implies that we can write $u_S^*(\mu_0)$ as follows:

$$u_S^*(\mu_0) = \sum_{v \in V} \mu_0(v) \sum_{\mu' \in \Delta(V)} \beta^{\mathbf{M}_0^*}(\mu'|v)[x^{\mathbf{M}_0^*}(\mu') + \delta(1 - q^{\mathbf{M}_0^*}(\mu'))U_S^*(\mathbf{M}_0^*, 1, \mu', 0, x^{\mathbf{M}_0^*})], \quad (2)$$

where $U_S^*(h^1)$ denotes the seller's payoffs under $(\Gamma^*, (\pi_v^*, r_v^*)_{v \in V})$ from $t = 1$ onwards. By definition, $u_S^*(\mu_0)$ is the highest payoff the seller can achieve, given the continuation strategy, from among all mechanisms $\mathbf{M} \in \mathcal{M}$ that satisfy the participation and incentive compatibility constraints at the initial history. Proposition 1 implies $\mathbf{M}_0^*$ satisfies the following three properties. First, the buyer's payoff is 0 when her valuation is v_L, so that the following holds:

$$\sum_{\mu' \in \Delta(V)} \beta^{\mathbf{M}_0^*}(\mu'|v_L)v_L q^{\mathbf{M}_0^*}(\mu') = \sum_{\mu' \in \Delta(V)} \beta^{\mathbf{M}_0^*}(\mu'|v_L)x^{\mathbf{M}_0^*}(\mu'|v_L), \quad (3)$$

where we have already replaced the buyer's continuation value with 0. Second, when the buyer's valuation is v_H, she is indifferent between reporting v_H and v_L,

---

[14]Parts 2-4 may sound counterintuitive from a dynamic game perspective: there are games where, to sustain high payoffs today, low continuation payoffs are needed, even on the equilibrium path. Importantly, when we prove Proposition 1, we show that when we modify the strategy profile from history h^t onwards, we do not upset the equilibrium constraints at the histories that precede h^t, thereby showing that the new assessment is also a PBE assessment.

[15]Although with binary types, the incentive constraint for v_H binds as a result of revenue maximization, with a continuum of types the *adjacent downward-looking* incentive constraints obtain because incentive compatibility implies the envelope representation of payoffs. In the online appendix to Doval and Skreta (2018), we show how to obtain the envelope representation of payoffs in the abstract mechanism selection game we study there; it is immediate to show that it holds in this game as well. Thus, with a continuum of types, we obtain a representation of the seller's payoff similar to the one in this section.

so that the following equality holds:

$$\sum_{\mu'\in\Delta(V)} (\beta^{\mathbf{M}_0^*}(\mu'|v_H) - \beta^{\mathbf{M}_0^*}(\mu'|v_L))(v_H q^{\mathbf{M}_0^*}(\mu') + (1 - q^{\mathbf{M}_0^*}(\mu')) U_{H|L}^*(h^1)$$

$$= \sum_{\mu'\in\Delta(V)} (\beta^{\mathbf{M}_0^*}(\mu'|v_H) - \beta^{\mathbf{M}_0^*}(\mu'|v_L)) x^{\mathbf{M}_0^*}(\mu'), \tag{4}$$

where the notation $U_{H|L}^*(h^1)$ signifies that the continuation value of a buyer of valuation v_H is the utility that she would obtain from the allocation that corresponds to v_L – these are her information rents. Finally, the induced distribution over posteriors averages out to μ_0. That is, letting $\tau^{\mathbf{M}_0^*}(\mu_0, \mu') = \sum_{v\in V}\mu_0(v)\beta^{\mathbf{M}_0^*}(\mu'|v)$ denote the probability that posterior μ' is induced, we have

$$\sum_{\mu'\in\Delta(V)} \tau^{\mathbf{M}_0^*}(\mu_0, \mu')\mu' = \mu_0. \tag{5}$$

Replacing transfers out of seller's payoff: We can replace equations 3-5 in Equation 2 to obtain

$$u_S^*(\mu_0) = \frac{\sum_{\mu'\in\Delta(V)} \tau^{\mathbf{M}_0^*}(\mu_0, \mu') \left[q^{\mathbf{M}_0^*}(\mu')(\mu' v_H + (1 - \mu')\hat{v}_L(\mu_0)) + (1 - q^{\mathbf{M}_0^*}(\mu')) \times \right.}{\delta \left(U_S^*(h^1) + (\frac{\mu'}{1-\mu'} - \frac{\mu_0}{1-\mu_0})(1 - \mu')U_{H|L}^*(h^1) \right) \right]}, \tag{6}$$

where

$$\hat{v}_L(\mu_0) = v_L - \frac{\mu_0}{1 - \mu_0}\Delta v,$$

is the buyer's *virtual value* when her valuation is v_L and the seller assigns probability μ_0 to the buyer's valuation being v_H.

Equation 6 says that we can think of the seller's optimal mechanism at $t = 0$ as choosing a distribution over posteriors, $\tau^{\mathbf{M}_0^*}$, and, for each posterior he induces, a probability of trade, $q^{\mathbf{M}_0^*}$. If at a given posterior, μ', he sells the good (i.e., $q^{\mathbf{M}_0^*}(\mu') = 1$), he obtains the expected *virtual surplus*, where the expectation is taken with respect to μ', but the low-valuation buyer's virtual value is evaluated at μ_0. That the low-valuation buyer's virtual value is evaluated at μ_0 instead of at μ' is intuititive. After all, the seller assigns probability μ_0 to the buyer's valuation being v_H so that, whenever he sells to both buyer types (i.e., $\mu' > 0$), he leaves rents to v_H with probability μ_0. If at a given posterior, μ', he delays trade (i.e., $q^{\mathbf{M}_0^*}(\mu') = 0$), then he receives his continuation payoff. However, it is modified

17

by the buyer's continuation rents: the seller internalizes that, whenever trade is delayed, the buyer receives continuation rents when her valuation is v_H. Note the continuation rents enter with a term that depends on μ_0: the seller assigns probability μ_0 to v_H, and hence that is the rate at which he pays rents (this time in terms of continuation values) to the buyer.

It follows from Proposition 1 that an expression similar to that of Equation 6 holds for any history h^t on the path of $(\Gamma^*, (\pi^*_v, r^*_v)_{v \in V})$. That is,[16]

$$
U^*_S(h^t) = \frac{\sum_{\mu' \in \Delta(V)} \tau^{\mathbf{M}^*_t}(\mu^*(h^t), \mu') \left[q^{\mathbf{M}^*_t}(\mu')(\mu' v_H + (1-\mu')\hat{v}_L(\mu^*(h^t))) + (1 - q^{\mathbf{M}^*_t}(\mu')) \times \right.}{\delta \left(U^*_S(h^{t+1}) + (\frac{\mu'}{1-\mu'} - \frac{\mu^*(h^t)}{1-\mu^*(h^t)})(1-\mu')U^*_{H|L}(h^{t+1}) \right) \right]},
$$

(7)

where $\sum_{\mu' \in \Delta(V)} \tau^{\mathbf{M}^*_t}(\mu^*(h^t), \mu')\mu' = \mu^*(h^t)$.

Part 2 implies that this is payoff is the best one the seller can achieve in $\mathcal{E}^*(\mu^*(h^t))$, consistent with the buyer's payoff from h^t onwards. That is, from amongst all the mechanisms that satisfy the participation and incentive compatibility constraints given the buyer's equilibrium payoffs, $\mathbf{M}^*_t$ delivers the seller the highest payoff at h^t.

As in Equation 6, Equation 7 shows that the seller's mechanism at h^t can be thought of as a distribution over posteriors, $\tau^{\mathbf{M}^*_t}$, and a probability of trade, $q^{\mathbf{M}^*_t}$. $(\tau^{\mathbf{M}^*_t}, q^{\mathbf{M}^*_t})$ are chosen so as to maximize a version of the virtual surplus, evaluated at the seller's prior belief at h^t, $\mu^*(h^t)$. As in the discussion of Equation 6, the virtual value $\hat{v}_L$ is evaluated at $\mu^*(h^t)$ because $\mu^*(h^t)$ is the probability the seller assigns to the buyer's valuation being v_H at h^t.

Moreover, by expanding the above expressions and using the Bayes' consistency conditions, we can show that if $h^{t+1} = (\mathbf{M}^*_t, 1, \mu', (0, x^{\mathbf{M}^*_t}(\mu')))$ is a history on the path of $(\Gamma^*, (\pi^*_v, r^*_v)_{v \in V})$, then:

$$
U^*_S(h^{t+1}) + \left(\frac{\mu'}{1-\mu'} - \frac{\mu^*(h^t)}{1-\mu^*(h^t)} \right)(1-\mu')U^*_{H|L}(h^{t+1}) =
$$
$$
\sum_{\tilde{\mu} \in \Delta(V)} \tau^{\mathbf{M}^*_{t+1}}(\mu', \tilde{\mu}) \left[q^{\mathbf{M}^*_{t+1}}(\tilde{\mu})(\tilde{\mu}v_H + (1-\tilde{\mu})\hat{v}_L(\mu^*(h^t))) + (1 - q^{\mathbf{M}^*_{t+1}}(\tilde{\mu})) \times \right.
$$
$$
\left. \delta \left(U^*_S(h^{t+2}) + (\frac{\tilde{\mu}}{1-\tilde{\mu}} - \frac{\mu^*(h^t)}{1-\mu^*(h^t)})(1-\tilde{\mu})U^*_{H|L}(h^{t+2}) \right) \right],
$$

(8)

where we use that in history h^{t+1} the seller assigns probability μ' to the buyer's valuation being v_H.

---

[16]Equation 7 is obtained by replacing the corresponding versions of equations 3-5 at history h^t.

18

The above expression shows how the term $\left(\frac{\mu'}{1-\mu'} - \frac{\mu^*(h^t)}{1-\mu^*(h^t)}\right)(1-\mu')U^*_{H|L}(h^{t+1})$ adjusts the seller's continuation values. If at history h^t, the seller's mechanism does not sell the good to the buyer with some probability, he understands that the buyer receives continuation rents $U^*_{H|L}(h^{t+1})$, which have to be reflected in the payments of $\mathbf{M}^*_t$. Because the seller assigns probability $\mu^*(h^t)$ to the buyer's valuation being v_H, he pays those rents with probability $\mu^*(h^t)$. Hence, he adjusts the continuation values to reflect his perceived probability of paying these future rents, so that from his perspective, the virtual value $\hat{v}_L(\cdot)$ is computed at $\mu^*(h^t)$ at all periods after t.

By contrast, Equation 7 evaluated at period $t+1$ shows that when the seller's prior is μ', he chooses his mechanism taking into account that he leaves rents with probability μ' to the buyer. Whenever $\mu' \neq \mu^*(h^t)$, the seller in period $t+1$ does not internalize the cost he may impose on his "predecessor" because he evaluates leaving rents to the high-valuation buyer differently than how the seller in period t does.

One last implication of incentive efficiency is that, along the equilibrium path, the seller's beliefs together with the buyer's rents pin down the strategy profile:

Corollary 1. Let $\langle \Gamma^*, (\pi^*_v, r^*_v)_{v\in V}, \mu^* \rangle$ be the PBE assessment that delivers payoff $u^*_S(\mu_0)$. If there are two histories, h^τ, h^t, on the path of play such that $\mu^*(h^t) = \mu^*(h^\tau)$ and $U^*_{H|L}(h^t) = U^*_{H|L}(h^\tau)$, then $\Gamma^*(h^\tau)$ and $\Gamma^*(h^t)$ are payoff equivalent for the seller.

Corollary 1 implies that along the equilibrium path, we can think of $U^*_S(h^t)$ and $U^*_{H|L}(h^t)$ as functions of the seller's belief $\mu^*(h^t)$. In a slight abuse of notation, we denote them by $U^*_S(\mu^*(h^t))$ and $U^*_{H|L}(\mu^*(h^t))$.

Corollary 1, Equation 7, and Equation 8 are the basis for the recursive formulation we use to characterize the seller's highest equilibrium payoff, $u^*_S(\mu_0)$. Indeed, let h^t be a history such that the seller assigns probability μ^* to v_H. Define

$$R^*(\mu', \mu^*) = U^*_S(\mu') + \left(\frac{\mu'}{1-\mu'} - \frac{\mu^*}{1-\mu^*}\right)(1-\mu')U^*_{H|L}(\mu'),$$

and let

$$R^*(\mu^*, \mu^*) = \tag{9}$$
$$\sum_{\mu' \in \Delta(V)} \tau^{\mathbf{M}^*_t}(\mu^*, \mu')(q^{\mathbf{M}^*_t}(\mu')(\mu' v_H + (1-\mu')\hat{v}_L(\mu^*)) + (1 - q^{\mathbf{M}^*_t}(\mu'))\delta R^*(\mu', \mu^*)).$$

19

Consider now the following problem:

Problem 1. Find a policy $(\tau, q) : \Delta(V) \mapsto \Delta(\Delta(V)) \times [0,1]^{\Delta(V)}$ and a value function $R^{(\tau,q)}$ such that

1. For all $\mu_0 \in \Delta(V)$, $\int \mu' \tau(\mu_0, d\mu') = \mu_0$,

2. For all $\mu_0, \mu' \in \Delta(V)$ such that $\mu_0 < 1$,

$$R^{(\tau,q)}(\mu', \mu_0) = \int \left[q(\mu', \tilde{\mu})(\tilde{\mu} v_H + (1 - \tilde{\mu})\hat{v}_L(\mu_0)) + \delta(1 - q(\mu', \tilde{\mu}))R^{(\tau,q)}(\tilde{\mu}, \mu_0) \right] \tau(\mu', d\tilde{\mu}),$$

3. For all $0 \le \mu_0 < 1$, $(\tau(\mu_0, \cdot), q(\mu_0, \cdot))$ solves

$$\max_{\tau', q'} \int \left[q'(\mu')(\mu' v_H + (1 - \mu')\hat{v}_L(\mu_0)) + \delta(1 - q'(\mu'))R^{(\tau,q)}(\mu', \mu_0) \right] \tau'(d\mu')$$

where $\int \mu' \tau'(d\mu') = \mu_0$ and $q'(\cdot) \in [0,1]$. For $\mu_0 = 1$, $q(1,1) = 1$.

A solution to Problem 1 specifies (i) for each prior, μ_0, a Bayes' plausible distribution over posteriors (part 1) and a probability of trade for each induced posterior, and (ii) continuation values that are consistent with the policy (part 2). Moreover, $\langle(\tau, q), R^{(\tau,q)}\rangle$ is a solution if for each prior belief the seller may hold, he does not have a one-shot deviation from the policy, (τ, q), given the continuation values (part 3).

Standard dynamic programming arguments imply that if we find a solution, $\langle(\tau^*, q^*), R^{(\tau^*,q^*)}\rangle$, to Problem 1, we will have found a solution to the system defined by Equation 7 for every h^t on the path of play. Because this system of equations satisfies the necessary conditions in Proposition 1, $R^{(\tau^*,q^*)}(\mu_0, \mu_0)$ is an upper bound on $u_S^*(\mu_0)$. If we then show that there is a PBE that attains $R^{(\tau^*,q^*)}(\mu_0, \mu_0)$, we have characterized the seller's best PBE payoff using the solution to the recursive program.

3.3 Intrapersonal game

Problem 1 is not a decision problem, but a *game*. The players in this game represent the seller holding different beliefs about the buyer's valuation being v_H. In

this way, Problem 1 is a formal representation of an idea already present in the literature on the sale of a durable good: when the seller cannot commit not to lower the prices in the future, the seller today competes against his future selves.[17]

Contrast the analysis so far with the solution if we did not require that the seller's strategy be sequentially rational. An analogue of parts 3 and 4 of Proposition 1 would also hold in that case. Moreover, we would use them to replace the transfers out of the seller's problem to obtain a program in which the seller maximizes the virtual surplus. However, under commitment, the maximization of the virtual surplus is now a *decision problem*: in the event that he does not sell the good, the seller does not need to consider that his mechanism may not be optimal given the information he has now learned.

Notwithstanding the aforementioned difference, Problem 1 shares a useful similarity with the usual approach when the designer has commitment: in both cases, the buyer's behavior has been reduced to a system of equations and is no longer a player.[18] Moreover, under the assumption of transferable utility, these equations allow us to express the seller's problem only in terms of choosing the allocation and the communication device. This result is a consequence of the application of the revelation principle in Doval and Skreta (2018), and thus, we expect that a similar simplification is feasible in settings other than the one we study here.

Problems like Problem 1 have a long tradition in economics. Indeed, it is an example of an intrapersonal equilibrium (see Pollak (1968); Peleg and Yaari (1973); Harris and Laibson (2001); Bernheim et al. (2015)).[19] For future reference, we

---

[17]McAfee and Wiseman (2008) put forward this intuition in the introduction to their paper. In a sense, it already appears in Coase (1972) where Coase points out

> [...] In these circumstances, why should the landowner continue to hold MQ off the market? The original landowner could obviously improve his position by selling more land since he could by these means acquire more money. It is true that this would reduce the value of the land OM owned by those who had previously bought land from him loss would fall on them, not on him.

[18]Contrast this to the approach in Freixas et al. (1985); Laffont and Tirole (1987, 1988); Bester and Strausz (2001); Gerardi and Maestri (2018) among others, where the designer can offer menus of contracts and observes the agent's choice out of the menu. In these papers, the offered menu together with the agent's choice out of the menu determine the designer's beliefs about the agent's type. This makes it difficult to separate the design of the allocation from the "design" of the information that is revealed by the choice out of the menu.

[19]Indeed, the system defined by Equation 7 and Equation 8 has analogues in the work of Harris and Laibson (2001) and Bernheim et al. (2015) on hyperbolic discounting.

record this observation in Definition 3 below.

Definition 3. A policy $(\tau, q) : \Delta(V) \mapsto \Delta(\Delta(V)) \times [0,1]^{\Delta(V)}$ and a value function $R^{(\tau,q)}$ constitute an *intrapersonal equilibrium* if they solve Problem 1.

The main result of this section shows that an intrapersonal equilibrium exists in the game that we study. This is stated formally in Theorem 2 below (note the sequence in the statement is the one described after the statement of Theorem 1).

Theorem 2. *There exists a unique intrapersonal equilibrium,* $\langle (\tau^*, q^*), R^{(\tau^*, q^*)} \rangle$. *Indeed, there exists a sequence* $0 = \bar{\mu}_0 < \bar{\mu}_1 = v_L / v_H < \cdots < \bar{\mu}_n < \dots$ *such that if* $\mu_0 \in [\bar{\mu}_i, \bar{\mu}_{i+1})$,

1. *and* $i = 0$, *then* $\tau^*(\mu_0, \mu_0) = 1, q^*(\mu_0, \mu_0) = 1$, *while*

2. *if* $i \geq 1$, *then* $\tau^*(\mu_0, 1) = 1 - \tau^*(\mu_0, \bar{\mu}_{i-1}) = (\mu_0 - \bar{\mu}_{i-1}) / (1 - \bar{\mu}_{i-1})$ *and* $q^*(\mu_0, 1) = 1 = 1 - q^*(\mu_0, \bar{\mu}_{i-1})$.

Theorem 2 says that when the seller's prior is low (i.e., $\mu_0 < \bar{\mu}_1$), he sells with probability 1 and transmits no information, whereas when the seller's prior is high (i.e., $\bar{\mu}_1 \leq \mu_0$), he induces two posteriors: one in which trade happens ($\mu' = 1$) and one in which trade is delayed ($\mu' = \bar{\mu}_{i-1}$).

Theorem 2 is the basis for the equilibrium constructed in Theorem 1. Indeed, the communication device and the probability of trade used by the seller to achieve his maximum PBE payoff are those from the intrapersonal equilibrium.[20] Theorem 2, however, says nothing about the transfers or the buyer's behavior in the PBE assessment. After all, in the intrapersonal game, the seller is the only player.

The proof of Theorem 2 is constructive. The rest of this section sketches out the main steps and provides intuition for Theorem 2. The reader interested in understanding how we move from the intrapersonal equilibrium to a PBE assessment that delivers payoff $u_S^*(\mu_0)$ may skip straight to Section 3.4, where we tackle the last step of our construction.

Necessary conditions: To show that an intrapersonal equilibrium exists and is unique, we proceed as follows. Proposition 2 and Proposition 3 below describe necessary conditions that an intrapersonal equilibrium must satisfy. We use

---

[20]More precisely, the communication device used by the seller is derived from the distribution over posteriors in the intrapersonal equilibrium.

these properties to build a candidate policy, (τ^*, q^*), and its continuation values, $R^{(\tau^*, q^*)}$, and show they constitute an equilibrium.

Proposition 2. *Let* $\langle (\tau^*, q^*), R^{(\tau^*, q^*)} \rangle$ *denote an intrapersonal equilibrium. Then, the following hold:*

1. *If* $\mu_0 < \bar{\mu}_1$, *then* $q^*(\mu_0, \mu') = 1$ *for all* $\mu' \in \Delta(V)$,

2. *If* $\mu_0 > \bar{\mu}_1$, *then the seller places positive probability on two beliefs,* $\{\mu_D(\mu_0), 1\}$, *and sets* $q^*(\mu_0, \mu_D(\mu_0)) = 0$ *and* $q^*(\mu_0, 1) = 1$.

The intuition for Proposition 2 is as follows. To understand part 1, recall that the solution when the seller has commitment and his prior is below $\bar{\mu}_1$ is to sell to both types of the buyer with probability 1. This is still a solution under limited commitment and is clearly the best that the seller can do. To understand part 2, note that in the intrapersonal equilibrium, the seller's problem is as follows:

$$R^{(\tau^*, q^*)}(\mu_0, \mu_0) = \max_{\tau, q} \int_0^1 \Big(q(\mu')(\mu' v_H + (1 - \mu')\hat{v}_L(\mu_0)) + (1 - q(\mu'))\delta R^{(\tau^*, q^*)}(\mu', \mu_0) \Big) \tau(d\mu')$$

$$= \max_\tau \int_0^1 \max\{\mu' v_H + (1 - \mu')\hat{v}_L(\mu_0), \delta R^{(\tau^*, q^*)}(\mu', \mu_0)\} \tau(d\mu'). \tag{10}$$

Indeed, given the continuation values, the seller can obtain whatever is best between selling today (with value $\mu' v_H + (1 - \mu')\hat{v}_L(\mu_0)$) and delaying trade (with value $\delta R^{(\tau^*, q^*)}(\mu', \mu_0)$), by choosing $q(\cdot)$ appropriately.

Equation 10 shows the seller's problem is like an *information design* problem.[21] Starting from an equilibrium $\langle (\tau^*, q^*), R^{(\tau^*, q^*)} \rangle$, it follows that the seller always has a best response in which he uses at most two posteriors. We show in Section C.1 that (i) the seller must use two posteriors, one at which he sets $q = 0$ and another at which he sets $q = 1$, and (ii) that this solution is the unique best response. That the seller induces a posterior of 1 when he sets $q = 1$ is intuitive: the largest payoff he can get when he sells is v_H. By setting $q = 0$ for at most one posterior, the seller maximizes the probability with which he induces a posterior of 1.[22]

---

[21] However, we cannot just invoke concavification-style arguments because the continuation values, $R^{(\tau^*, q^*)}$, may fail to be upper semi-continuous (see the discussion at the end of this section).

[22] The linearity of the seller's payoff in the allocation and in the posterior beliefs in the event of trade matters for this argument. The same linearity obtains when the buyer's valuations are drawn from a continuum.

23

The result in Proposition 2 implies we can reduce the construction of the intrapersonal equilibrium to the construction of the belief $\mu_D(\mu_0)$ at which the seller sets $q = 0$ when $\mu_0 > \overline{\mu}_1$. Suppose $\langle \mu_{D*}, R^{\mu_{D*}} \rangle$ is an intrapersonal equilibrium, where $R^{\mu_{D*}}$ are the continuation values implied by the policy that sets $q^*(\mu_0, \mu_{D*}(\mu_0)) = 0$. The martingale property of beliefs implies

$$\tau^*(\mu_0, \mu_{D*}(\mu_0)) = \frac{1 - \mu_0}{1 - \mu_{D*}(\mu_0)},$$

so that part 2 of Problem 1 implies

$$R^{\mu_{D*}}(\mu_0, \mu_0) = \frac{\mu_0 - \mu_{D*}(\mu_0)}{1 - \mu_{D*}(\mu_0)} v_H + \frac{1 - \mu_0}{1 - \mu_{D*}(\mu_0)} R^{\mu_{D*}}(\mu_{D*}(\mu_0), \mu_0). \quad (11)$$

Proposition 3 characterizes the properties of μ_{D*}, under the assumption that such an equilibrium exists:

Proposition 3. *Suppose* $\langle \mu_{D*}, R^{\mu_{D*}} \rangle$ *is an intrapersonal equilibrium. Then, the following hold:*

1. *If* $\overline{\mu}_1 < \mu_0$ *is such that* $\mu_{D*}(\mu_0) < \overline{\mu}_1$, *then* $\mu_{D*}(\mu_0) = 0$.

2. *If* $\overline{\mu}_1 < \mu_0$, *then there exists* $N_{\mu_0} < \infty$ *such that the seller's belief drops below* $\overline{\mu}_1$ *after* N_{μ_0} *periods, i.e.,* $\mu_{D*}^{(N_{\mu_0})}(\mu_0) \leq \overline{\mu}_1$.

3. *If* $\mu_0 < \mu_0'$, *then* $\mu_{D*}(\mu_0) \leq \mu_{D*}(\mu_0')$.

4. *If* $N_{\mu_0} = N_{\mu_0'}$, *then* $\mu_{D*}(\mu_0) = \mu_{D*}(\mu_0')$.

The proof is in Appendix C.2. In what follows, we provide intuition for the result, which in turn also provides intuition for the structure of the intrapersonal equilibrium in Theorem 2.

Part 1 is immediate. Recall from Proposition 2 that a seller with prior $\mu_0' < \overline{\mu}_1$ trades immediately. Thus, if a seller with prior $\overline{\mu}_1 < \mu_0$ sets $\mu_{D*}(\mu_0) = \tilde{\mu} \in (0, \overline{\mu}_1)$, he knows that he trades with v_L in the next period. That is, $R^{\mu_{D*}}(\tilde{\mu}, \mu_0) = \tilde{\mu} v_H + (1 - \tilde{\mu}) \hat{v}_L(\mu_0)$. By setting $\mu_{D*}(\mu_0) = 0$, the seller with prior μ_0 maximizes the probability of trading with the high-valuation buyer today.

Part 2 says that trade with the buyer when her valuation is v_H happens in finitely many periods. It does not say, however, that the same is true when the buyer's valuation is v_L, unless $0 < \overline{\mu}_1$. We now explain why the seller taking infinitely many periods for his belief to update below $\overline{\mu}_1$, conditional on the event of not

allocating the good, cannot be part of an intrapersonal equilibrium. Starting from any prior, $\overline{\mu}_1 < \mu_0$, the repeated application of the function μ_{D^*} induces a decreasing sequence of beliefs, $\mu_m = \mu_{D^*}^{(m)}(\mu_0)$, all of them above $\overline{\mu}_1$. But If μ_m remains above $\overline{\mu}_1$ for every finite m, the probability $\tau^*(\mu_m, 1)$ becomes small. Furthermore, a seller with a prior μ_m always has the possibility of placing weight μ_m on 1 and the remaining weight on 0. We show that for m large enough, this deviation is profitable. Thus, whereas the seller with prior belief, μ_0, would benefit from his successors never splitting their beliefs below $\overline{\mu}_1$, taking infinitely many periods to update below $\overline{\mu}_1$ cannot be part of an intrapersonal equilibrium.

Part 2 echoes the results in Fudenberg et al. (1985), Gul et al. (1986), and Ausubel and Deneckere (1989) that in the "gap" case, trade happens with both types of the buyer in finitely many periods. We do not phrase it in these terms because when $v_L = 0$, which corresponds to the no-gap case in our model, part 2 holds; however, in the unique intrapersonal equilibrium, the seller with prior $\mu_0 = 0$ does not allocate the good to the buyer whose valuation is v_L. Because once we reach a posterior of 0 beliefs do not change, trade never happens with the low-valuation buyer.

Parts 3-4 say the policy is monotone in that if $\mu_0 < \mu_0'$, the seller with belief μ_0' chooses a higher posterior at which to delay trade than the seller with belief μ_0. This is intuitive. When $0 < \overline{\mu}_1$, the rents for the high-valuation buyer are determined by how long it takes to trade with v_L: the longer it takes, the lower the rents. The seller's prior is also the probability with which he pays rents to the buyer. Thus, the higher the prior, the higher the incentive for the seller to delay trade with v_L so as to make the rents for v_H smaller.[23]

Part 4 is important in the construction of the intrapersonal equilibrium when $0 < \overline{\mu}_1$. We can classify the priors, μ_0, according to how many periods it takes for μ_0 to drop below $\overline{\mu}_1$. That is, define

$$D_n = \{\mu_0 \in \Delta(V) : \mu_{D^*}^{(n)}(\mu_0) = 0\}, \tag{12}$$

where we appeal to Proposition 3 to say that the "final" belief at which the seller updates is 0 and we define $D_0 = \{0\}$. Part 4 says that if $\mu_0, \mu_0' \in D_n$, they delay trade at the same posterior. Otherwise, by setting $\mu_{D^*}(\cdot)$ to be $\min\{\mu_{D^*}(\mu_0), \mu_{D^*}(\mu_0')\}$,

---

[23]Note that when $\overline{\mu}_1 = 0$, the seller sets $\mu_D(\mu_0) = 0$ for all $0 < \mu_0$: in this case, the seller never pays rents to the buyer, so that for each prior, the seller sells to the high-valuation buyer with the maximum possible probability, μ_0.

the seller (weakly) increases the probability of immediately trading with v_H without affecting how long it takes to trade with v_L (both $\mu_{D*}(\mu_0), \mu_{D*}(\mu'_0)$ correspond to beliefs for which the seller trades in $n-1$ periods with v_L.)

Part 4 suggests the "smallest" element of D_n[24] is of relevance: this prior, denoted by $\overline{\mu}_n$ in what follows, is the "largest" prior in D_{n-1} and the smallest in D_n. The equilibrium policy, (τ^*, q^*), in Theorem 2 selects the policy of $\overline{\mu}_n$ so that it is an element of D_n. We break the indifference between n and $n-1$ periods until updating to 0 in favor of n.

The sequence $\{\overline{\mu}_n\}_{n\geq 0}$: We now construct the cutoffs $\overline{\mu}_n$. We set $\overline{\mu}_0 = 0$ and recall that $\overline{\mu}_1$ denotes the ratio v_L/v_H. Note that it satisfies that

$$\overline{\mu}_1 v_H + (1 - \overline{\mu}_1)\hat{v}_L(\overline{\mu}_1) = \overline{\mu}_1 v_H + (1 - \overline{\mu}_1)\delta \underbrace{\hat{v}_L(\overline{\mu}_1)}_{R^{\mu_{D*}}(0,\overline{\mu}_1)}, \tag{13}$$

because $\hat{v}_L(\overline{\mu}_1) = 0$. Equation 13 shows that, when his prior is $\overline{\mu}_1$, the seller is indifferent between trading today with both types of the buyer, or trading today with the buyer whose type is v_H and waiting until tomorrow to trade with the buyer whose type is v_L. Set $\mu_{D*}(\overline{\mu}_1) = 0 = \overline{\mu}_0$ and $\tau(\overline{\mu}_1, 0) = 1 - \overline{\mu}_1$. That is, we break the indifference of the seller with belief $\overline{\mu}_1$ between trading immediately with v_L and delaying trade with v_L for one period in favor of delaying trade.

For $n \geq 1$, define inductively $\overline{\mu}_{n+1}$ to be the prior of the seller such that

$$\frac{\overline{\mu}_{n+1} - \overline{\mu}_n}{1 - \overline{\mu}_n} v_H + \frac{1 - \overline{\mu}_{n+1}}{1 - \overline{\mu}_n}\delta R^{\mu_{D*}}(\overline{\mu}_n, \overline{\mu}_{n+1}) = \frac{\overline{\mu}_{n+1} - \overline{\mu}_{n-1}}{1 - \overline{\mu}_{n-1}} v_H + \frac{1 - \overline{\mu}_{n+1}}{1 - \overline{\mu}_{n-1}}\delta R^{\mu_{D*}}(\overline{\mu}_{n-1}, \overline{\mu}_{n+1}). \tag{14}$$

That is, when the seller's prior is $\overline{\mu}_{n+1}$, he is indifferent between taking $n+1$ periods to trade with v_L or taking n periods to trade with v_L. Note that Equation 14 only depends on the policy for $\overline{\mu}_m, m \leq n$, which has already been defined.

Lemma 4 in the Appendix shows $\{\overline{\mu}_n\}_{n\geq 0}$ is an increasing sequence. These cutoffs are precisely those in the statement of Theorem 2. Indeed, we set $D_n = [\overline{\mu}_n, \overline{\mu}_{n+1})$ for $n \geq 0$ and $\mu_{D*}(\mu_0) = \overline{\mu}_{n-1}$ for $\mu_0 \in D_n$. The policy (τ^*, q^*) satisfies all the properties of Proposition 3. We verify in Section C.3 that it constitutes an intrapersonal equilibrium. Uniqueness follows from Proposition 2 and Proposition 3 in Doval and Skreta (2019), where we show that although seller incarnations

---

[24]Technically, at this point, an infimum rather than a minimum.

with beliefs $\{\overline{\mu}_n\}_{n\geq 1}$ have multiple best responses, the one specified above is the only one that can be part of an intrapersonal equilibrium.

Although the intrapersonal game has a unique equilibrium, the game between the seller and the buyer does not (see Section IV in Doval and Skreta (2019)). This finding is a reflection of a deeper observation: we derived the intrapersonal game from necessary conditions for revenue maximization, whereas not all equilibria in the game between the seller and the buyer give the seller his best equilibrium payoff.

Tie-breaking and upper semicontinuity: Before proceding with the construction of the PBE assessment, $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v\in V}, \mu^* \rangle$, we make a final observation about the properties of the intrapersonal equilibrium. Figure 1 below illustrates that when $\mu_0 \in [\overline{\mu}_n, \overline{\mu}_{n+1})$, the continuation values $R^{(\tau^*, q^*)}(\cdot, \mu_0)$ fail to be upper semi-continuous at beliefs $\{\overline{\mu}_k : k \geq n+1\}$:[25]

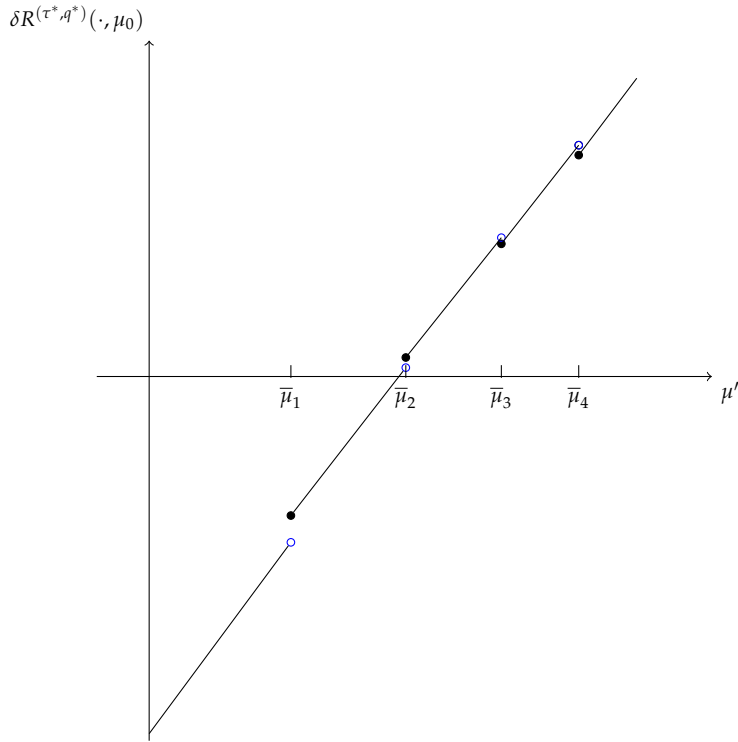[25]Lemma 3 in Doval and Skreta (2019) shows formally that upper semicontinuity fails at these beliefs.

Figure 1: Illustration of the failure of upper-semicontinuity for $\mu_0 \in [\overline{\mu}_2, \overline{\mu}_3)$: blue empty circles indicate the discontinuity

Conceptually, the failure of upper semicontinuity illustrates that sellers with different prior beliefs *disagree* about how ties should be broken. If $\mu_0 \in [\overline{\mu}_n, \overline{\mu}_{n+1})$, the seller with belief μ_0 finds that sellers with beliefs $\overline{\mu}_k, k \geq n + 1$ take too long to trade with the buyer when her valuation is v_L. However, inducing posteriors greater than $\overline{\mu}_{n+1}$ when $\mu_0 < \overline{\mu}_{n+1}$ is never optimal: any weight on such posteriors can be split between $\overline{\mu}_{n-1}$ and 1, which are the optimal choices for the seller when his prior $\mu_0 \in [\overline{\mu}_n, \overline{\mu}_{n+1})$. This observation, together with the property that seller incarnations with priors lower than μ_0 break ties in favor of the seller with prior μ_0, guarantees an intrapersonal equilibrium exists.

Note that conflicts in tie-breaking are usually an issue for equilibrium existence in the literature on intrapersonal games. A similar issue arises in sequential voting games (see Duggan (2006) and the references therein). Two properties of our game seem important in managing to sidestep these issues: (i) the agreement between a seller with belief μ_0 with how sellers with lower priors break ties, and (ii) the

seller with belief $\mu_0 \in [\overline{\mu}_n, \overline{\mu}_{n+1})$ can always choose not to put a seller with belief $\mu' \geq \overline{\mu}_{n+1}$ on the move.

3.4 From the intrapersonal equilibrium to the PBE assessment

Section 3.4 describes how we use the intrapersonal equilibrium policy, (τ^*, q^*), to construct an assessment, $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$, that gives the seller his highest equilibrium payoff in $G^\infty(\mu_0)$.

Theorem 2 delivers an allocation rule and a distribution over posteriors for each prior belief the seller may hold. We now use the seller's policy in the intrapersonal equilibrium to construct a mapping, $\gamma^* : \Delta(V) \mapsto \mathcal{M}$, that assigns a mechanism to each belief the seller may have about the buyer (omitted details of this construction are in Section D.1). Given the seller's prior, μ_0, we can use the distribution over posteriors, $\tau^*(\mu_0, \cdot)$, to construct a communication device $\beta_{\mu_0}^*$, which under the assumption of truthtelling, implements the same distribution over posteriors as τ^*. Indeed, if $\mu_0 \in D_n$,

$$\beta_{\mu_0}^*(\overline{\mu}_{n-1}|v_L) = 1 \text{ and } \beta_{\mu_0}^*(\overline{\mu}_{n-1}|v_H) = \frac{\overline{\mu}_{n-1}\tau^*(\mu_0, \overline{\mu}_{n-1})}{\mu_0}.$$

Moreover, we can use the allocation rule $q^*(\mu_0, \cdot)$ to construct the probability of trade:

$$q_{\mu_0}^*(\mu') = q^*(\mu_0, \mu') \text{ for } \mu' \in \{\overline{\mu}_{n-1}, 1\}.$$

Finally, using Equations 3 and 4, we can use the allocation rule, $q^*(\mu_0, \cdot)$, and the communication device, $\beta_{\mu_0}^*$, to construct the transfers. Indeed, if $\mu_0 \in D_n$,

$$x_{\mu_0}^*(\overline{\mu}_{n-1}) = 0 \text{ and } x_{\mu_0}^*(1) = v_L + (1 - \delta^n)\Delta v. \tag{15}$$

The mapping γ^* then maps each prior, μ_0, to $\langle (V, \beta_{\mu_0}^*, \Delta(V)), (q_{\mu_0}^*(\cdot), x_{\mu_0}^*(\cdot)) \rangle$. Moreover, we construct the buyer's rents under this mechanism as a function of the seller's prior. Indeed, if $\mu_0 \in D_n$, the buyer's continuation payoff when her valuation is v_H is given by:

$$u_H^*(\mu_0) = \delta^n \Delta v. \tag{16}$$

To construct a PBE assessment, we need to specify the buyer's and the seller's strategy after every history. We do so in Sections E.1 and E.2 in Appendix E. In

particular, Section E.2 shows how to jointly construct a strategy for the seller, Γ^*, and a system of beliefs, μ^*, such that after every history, the seller's choice of mechanism is determined by $\Gamma^*(h^t) = \gamma^*(\mu^*(h^t))$. Moreover, $\mu^*(h^t)$ is consistent with the buyer's strategy constructed in Section E.1.

The PBE assessment we construct has the following property at all histories h^t. If at history h^t, the seller has beliefs $\mu^*(h^t)$, the seller's payoff is $R^{(\tau^*,q^*)}(\mu^*(h^t), \mu^*(h^t))$ and the buyer's payoff is 0 if her valuation is v_L and is $u_H^*(\mu^*(h^t))$ if her valuation is v_H (recall equation (16)). Section E.3 verifies that given these continuation payoffs, at history h^{t-1}, neither the buyer nor the seller have a one-shot deviation from the equilibrium strategy profile. In the language of Abreu et al. (1990), the strategy at h^{t-1}, together with the continuation payoffs at h^t, decompose the payoffs at h^{t-1}. However, without knowing whether $(R^{(\tau^*,q^*)}(\mu^*(h^t), \mu^*(h^t)), 0, u_H^*(\mu^*(h^t)))$ is itself an equilibrium payoff in $G^\infty(\mu^*(h^t))$, this is not enough to conclude $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$ is a PBE assessment.

In Doval and Skreta (2019), we lay out the definitions and statements needed to show that self-generation techniques apply to our setting, so that the above steps verify $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$ is a PBE assessment. Ours is a dynamic incomplete information game with persistent types, like the one studied in Athey and Bagwell (2008).[26] The authors show that under a restriction on the solution concept, which they term Public PBE, equilibrium payoffs can be characterized using techniques like in Abreu et al. (1990).[27] More precisely, a Public PBE is a PBE assessment where each player's strategy depends on the public history of the game and their current private information.

Although in general the set of Public PBE payoffs is a strict subset of the PBE payoffs, it is not in the game we study.[28] As we discussed in Section 3.2, it follows from Doval and Skreta (2018) that it is without loss of generality to focus on PBE assessments where the buyer's strategy depends only on her valuation and the public history. Thus, self-generation techniques can be applied to our game to characterize not only the seller's best equilibrium payoff, but all equilibrium payoffs in our game.

---

[26]To be sure, in Athey and Bagwell (2008), types are persistent, but not fully persistent as in our game.

[27]Cole and Kocherlakota (2001) introduce similar techniques, but in dynamic games with *full support*. That is, the information structure in the game is such that no player can infer from his signal realizations that another player has deviated.

[28]Notwithstanding this difficulty, Athey and Bagwell (2008) show that, under some conditions, the most collusive outcome can be achieved by using Public PBE strategy profiles.

We close Section 3 by illustrating some qualitative features of the model by means of pictures and discussing the case $v_L = 0$:

Illustrations: Although the recursive nature of the equations defining the cut-offs $\{\bar{\mu}_n\}_{n \geq 0}$ make performing comparative statics somewhat difficult, in what follows, we reproduce the price path and the seller's payoff as a function of the seller's prior and the discount factor for a specific parameterization of the model. In particular, the figures are plotted using $v_L = 1$ and $v_H = 5$ (which implies $\bar{\mu}_1 = 0.2$).
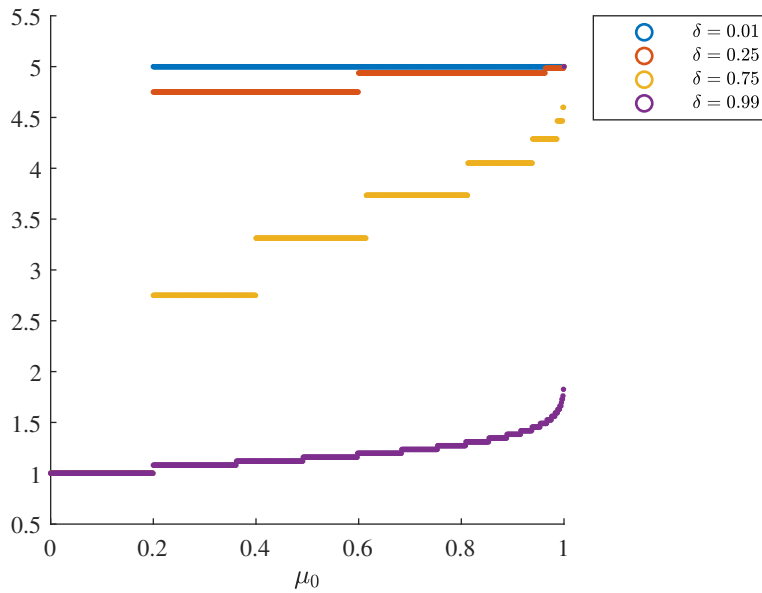

Figure 2: Price dynamics as a function of seller's prior – different discount factors.

Figure 2 plots the price dynamics as a function of the discount factor. Note that for low values of the discount factor ($\delta = 0.01$ in blue), the solution is almost like the commitment solution. In a sense, the seller's impatience affords him commitment power, so that he can keep the prices close to v_H as long as his prior is above $\bar{\mu}_1$. At the other extreme, when the seller's discount factor is high ($\delta = 0.99$ in purple), the logic behind the Coase conjecture applies and the seller quickly lowers the prices. As a consequence, the seller's revenue is lower for higher discount factors. We illustrate this in Figure 3:
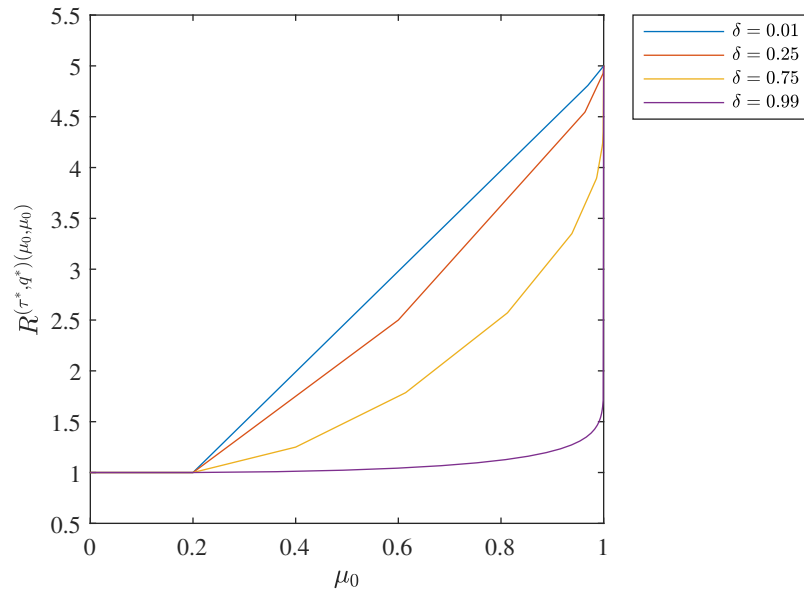
31

Figure 3: Seller's payoff as a function of seller's prior – different discount factors.

Figure 3 may give the impression that the seller's revenue is monotonically decreasing in the discount factor for a given prior. It is not, however, as Figure 4 shows:
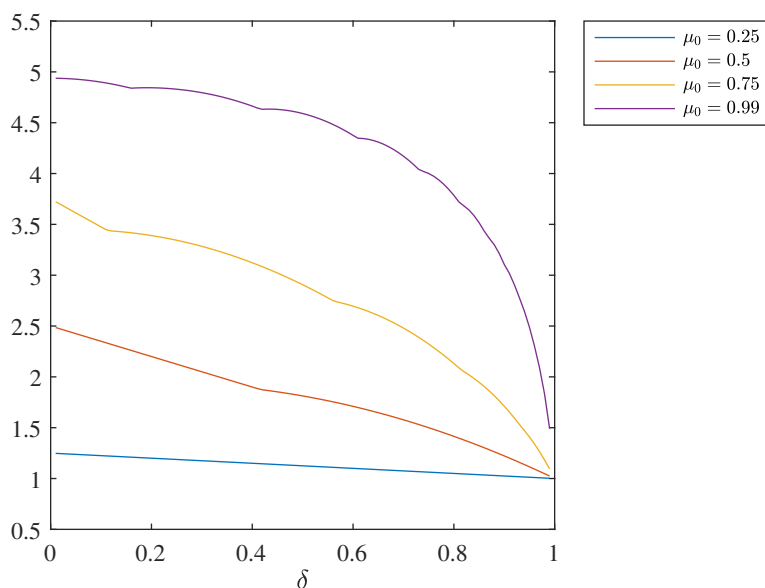
Figure 4: Seller's payoff as a function of the discount factor – different priors.

The case $v_L = 0$ The case of $v_L = 0$ shares features with the *no gap case* in Ausubel and Deneckere (1989). The 'no gap' case refers to the possibility that the seller's value is above the lowest possible buyer's value (in our case this means $v_L \leq 0$). As we argue next, when $v_L = 0$, the seller can sustain the full commitment profit even with limited commitment.

In particular, at the seller optimal equilibrium revenue is $\mu_0 v_H$, regardless of the discount factor (and not just in the limit as in Ausubel and Deneckere (1989)). From the perspective of the intrapersonal equilibrium, regardless of his belief, μ_0, as long as $\mu_0 > 0$, the seller prefers that the seller with belief 0 breaks ties in favor of no trade. In that case, each seller with belief $\mu_0 > 0$ can trade with maximal probability with v_H because the "threat" of no further trade is credible.

The reader may then be tempted to draw these conclusions more broadly. However, that the seller can achieve the commitment profits when $v_L = 0$ for all discount factors is an artifact of the binary case. With finitely many types, the lowest of which is below 0, the seller would eventually trade with all positive ones in a price-posting game. Thus, with finitely many types, the seller can always avoid dropping the price below the smallest, yet positive, valuation.

33

4 Conclusions

This paper makes two contributions to the literature on mechanism design with limited commitment. First, we characterize the optimal mechanism for the sale of a durable good when the seller has limited commitment and interacts with the buyer for infinitely many periods. Second, we provide a *recipe* for solving problems of mechanism design with limited commitment with transferable utility, which applies to settings other than the ones we study here.

Indeed, the steps followed in the proof of our main result closely follow the standard steps in classical mechanism design. To wit, the prototypical procedure to analyze problems of mechanism design with commitment and transferable utility follows these steps: (i) Apply the revelation principle to simplify the space of mechanisms and the features of the agent's strategy profile (participation and truthtelling), (ii) appeal to quasilinearity and single-crossing to derive an envelope representation of payoffs that replaces the transfers out of the designer's payoff, (iii) solve the designer's *decision problem*, and (iv) verify that the solution satisfies all constraints that (may) have been ignored.

Note that these steps are mainly the ones we followed in Sections 3.1-3.4. Indeed, thanks to the result in our previous work, Doval and Skreta (2018), we are able to reduce the buyer's behavior to a series of participation and truthtelling constraints, just as we do in the case of full commitment. The main difference is step (iii): with limited commitment, the designer's problem is an intrapersonal game. This is a reflection of a deeper observation: optimal mechanisms under full commitment often fail to be sequentially rational, whereas the best response condition of the intrapersonal equilibrium captures the restrictions imposed by sequential rationality on the designer's behavior.

The analysis so far has focused on the case of binary valuations, which is consistent with recent papers in the literature that studies infinite-horizon principal-agent problems with limited commitment, like Strulovici (2017) and Gerardi and Maestri (2018). The case of binary valuations allows us to bring to the forefront the conceptual innovations that arise in an infinite horizon game between two long-run players, where one of them, the seller, has a rich action space. Nevertheless, the extension to the case in which buyer valuations are drawn from a continuum is of interest, and we plan to address it in future work.

References

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

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

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

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

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

BERNHEIM, B. D., D. RAY, AND Ş. YELTEKIN (2015): "Poverty and self-control," *Econometrica*, 83, 1877–1911.

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

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

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

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

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

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

COLE, H. L. AND N. KOCHERLAKOTA (2001): "Dynamic games with hidden actions and hidden states," *Journal of Economic Theory*, 98, 114–126.

CONLISK, J., E. GERSTNER, AND J. SOBEL (1984): "Cyclic pricing by a durable goods monopolist," *The Quarterly Journal of Economics*, 99, 489–505.

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

DILMÉ, F. AND F. LI (Forthcoming): "Revenue Management without Commitment: Dynamic Pricing and Periodic Flash Sales," *The Review of Economic Studies*.

DOVAL, L. AND V. SKRETA (2018): "Mechanism Design with Limited Commitment," *arXiv preprint arXiv:1811.03579*.

——— (2019): "Supplement to "Optimal mechanism for the sale of a durable good"," .

DUGGAN, J. (2006): "Endogenous voting agendas," *Social Choice and Welfare*, 27, 495–530.

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

FREIXAS, X., R. GUESNERIE, AND J. TIROLE (1985): "Planning under incomplete information and the ratchet effect," *The Review of Economic Studies*, 52, 173–191.

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

FUDENBERG, D. AND E. MASKIN (1986): "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information," *Econometrica*, 54, 533–554.

GARRETT, D. F. (2016): "Intertemporal price discrimination: Dynamic arrivals and changing values," *American Economic Review*, 106, 3275–99.

GERARDI, D. AND L. MAESTRI (2018): "Dynamic Contracting with Limited Commitment and the Ratchet Effect," Tech. rep.

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

HARRIS, C. AND D. LAIBSON (2001): "Dynamic choices of hyperbolic consumers," *Econometrica*, 69, 935–957.

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

HOLMSTRÖM, B. AND R. B. MYERSON (1983): "Efficient and durable decision rules with incomplete information," *Econometrica*, 1799–1819.

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

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

——— (1990): "Adverse selection and renegotiation in procurement," *The Review of Economic Studies*, 57, 597–625.

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

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

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

NAVA, F. AND P. SCHIRALDI (2019): "Differentiated durable goods monopoly: a robust Coase conjecture," *American Economic Review*, 109, 1930–68.

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

PELEG, B. AND M. E. YAARI (1973): "On the existence of a consistent course of action when tastes are changing," *The Review of Economic Studies*, 40, 391–401.

POLLAK, R. A. (1968): "Consistent planning," *The Review of Economic Studies*, 35, 201–208.

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

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

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

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

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

STOKEY, N. L. (1979): "Intertemporal price discrimination," *The Quarterly Journal of Economics*, 355–371.

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

STROTZ, R. H. (1955): "Myopia and inconsistency in dynamic utility maximization," *The Review of Economic Studies*, 23, 165–180.

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

A   Omitted formal statements and equations

A.1   Bayes' rule where possible

In this section, we define formally what we mean by Bayes' rule where possible in Section 2. Note that in the game under consideration, the public history h^t represents an information set for the seller with nodes $h_B^t \in H_B^t(h^t)$. In what follows, we use the notation v, v' to denote nodes in the game.

Fix a strategy profile $(\Gamma, (\pi_v, r_v)_{v \in V})$, and two nodes v and v' such that v precedes v'. We can use the strategy profile to define a probability, $P^{(\Gamma, (\pi_v, r_v)_{v \in V})}(v'|v)$, of reaching node v' conditional on being at node v. Extend this probability to all nodes by making it 0 for nodes v' that do not succeed v.

Say that information set h^t precedes information set h^{t+1}, or that h^t, h^{t+1} are consecutive information sets if there exists a mechanism, **M**, such that either of the following hold:

(a) there is a posterior, μ', such that $\sum_{v \in V} \beta^{\mathbf{M}}(\mu'|v) > 0$[29] and
$h^{t+1} = (h^t, \mathbf{M}, 1, \mu', (0, x^{\mathbf{M}}(\mu')), \omega_{t+1})$, or

(b) $h^{t+1} = (h^t, \mathbf{M}, 0, \varnothing, (0,0), \omega_{t+1})$.

Fix an assessment, $\langle \Gamma, (\pi_v, r_v)_{v \in V}, \mu \rangle$, and two consecutive information sets h^t, h^{t+1}. Say that h^{t+1} is reached with positive probability from h^t under $\langle \Gamma, (\pi_v, r_v)_{v \in V}, \mu \rangle$, if

$$P^{\langle \Gamma, (\pi_v, r_v)_{v \in V}, \mu \rangle}(h^{t+1}|h^t) \equiv \sum_{v \in h^t, v' \in h^{t+1}} \mu^*(v|h^t) P^{(\Gamma, (\pi_v, r_v)_{v \in V})}(v'|v) > 0$$

Say that h^{t+1} can be reached from h^t through a deviation by the seller if there exists Γ' such that $P^{\langle \Gamma', (\pi_v, r_v)_{v \in V}, \mu \rangle}(h^{t+1}|h^t) > 0$.

Definition 4. An assessment $\langle \Gamma, (\pi_v, r_v)_{v \in V}, \mu \rangle$ satisfies Bayes' rule where possible if for all $t \geq 0$ and for all consecutive h^t, h^{t+1}, $\mu(v'|h^{t+1})$ is obtained via Bayes' rule from $\mu(\cdot|h^t)$ if either

1. $P^{\langle \Gamma, (\pi_v, r_v)_{v \in V}, \mu \rangle}(h^{t+1}|h^t) > 0$, or

2. h^{t+1} can be reached from h^t through a deviation by the seller.

A.2 Participation, truthtelling, and Bayes' plausibility constraints

Lemma 1 implies that for any payoff $(u_S, u_H, u_L) \in \mathcal{E}^*(\mu_0)$, we can find a PBE assessment, $\langle \Gamma, (\pi_v, r_v)_{v \in V}, \mu \rangle$, such that the following constraints are satisfied. For each period t and each public history h^t, the buyer must find it optimal to participate in the mechanism chosen by the seller:

$$\sum_{\mu' \in \Delta(V)} \beta^{\mathbf{M}_t^*}(\mu'|v)(vq^{\mathbf{M}_t^*}(\mu') - x^{\mathbf{M}_t^*}(\mu') + (1 - q^{\mathbf{M}_t^*}(\mu'))\delta U_v(h^t, \mathbf{M}_t^*, 1, \mu', (0, x^{\mathbf{M}_t^*})))$$
$$\geq \delta U_v((h^t, \mathbf{M}_t^*, \varnothing, (0,0))), \tag{PC_{v,h^t}}$$

where we use $U_v(h^{t+1})$ to denote the buyer's continuation payoffs at history h^{t+1} when her valuation is v under strategy profile $(\Gamma^*, (\pi_v^*, r_v^*)_{v \in V})$. Second, the buyer

---

[29] In Doval and Skreta (2018), we argue that it is without loss of generality to prune from the tree all histories that correspond to posteriors that cannot be generated with positive probability by the mechanism.

must find it optimal to report her type truthfully. That is, for all $v \in V$ and $v' \neq v$, we have that

$$\sum_{\mu' \in \Delta(V)} \beta^{\mathbf{M}_t^*}(\mu'|v)(vq^{\mathbf{M}_t^*}(\mu') - x^{\mathbf{M}_t^*}(\mu') + (1 - q^{\mathbf{M}_t^*}(\mu'))\delta U_v(h^t, \mathbf{M}_t^*, 1, \mu', (0, x^{\mathbf{M}_t^*})))$$

$$\geq \sum_{\mu' \in \Delta(V)} \beta^{\mathbf{M}_t^*}(\mu'|v')(vq^{\mathbf{M}_t^*}(\mu') - x^{\mathbf{M}_t^*}(\mu') + (1 - q^{\mathbf{M}_t^*}(\mu'))\delta U_v(h^t, \mathbf{M}_t^*, 1, \mu', (0, x^{\mathbf{M}_t^*}))).$$

$$(\text{IC}_{v,h^t})$$

Finally, if the mechanism outputs posterior μ', it must be that

$$\mu'[\sum_{v \in V} \beta^{\mathbf{M}_t^*}(\mu'|v)\mu(h^t)(v)] = \mu(h^t)\beta^{\mathbf{M}_t^*}(\mu'|v_H). \qquad (\text{BC}_{\mu(h^t)})$$

That is, from the communication device, we can infer a distribution over posteriors $\tau^{\mathbf{M}_t^*}(\mu(h^t), \cdot) \in \Delta(\Delta(V))$ such that $\tau^{\mathbf{M}_t^*}(\mu(h^t), \mu') = \sum_{v \in V} \mu(h^t)(v)\beta^{\mathbf{M}_t^*}(\mu'|v)$.

B Proof of Proposition 1

Part 1: The proof of this part is immediate, and hence we omit it.

Part 2: For this proof, we rely on Proposition 1 in Doval and Skreta (2019), which shows that if $\langle \Gamma, (\pi_v, r_v)_{v \in V}, \mu \rangle$ is a PBE assessment of $G^\infty(\mu_0)$, then the continuation payoffs at history h^t are equilibrium payoffs of $G^\infty(\mu(h^t))$. Let $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$ denote the PBE assessment that gives the seller payoff $u_S^*(\mu_0)$. Toward a contradiction, suppose that h^t is a history on the path of play such that $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle|_{h^t}$ is not incentive-efficient. Thus, $U_S^*(h^t) < \max\{u_S : (u_S, U_H^*(h^t), U_L^*(h^t)) \in \mathcal{E}^*(\mu^*(h^t))\}$. Then, there is a strategy profile in $G^\infty(\mu^*(h^t))$ that gives the seller a higher payoff and the buyer the same payoff at h^t. We can modify $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$ from h^t onwards and increase the seller's payoff, while still being a PBE assessment, contradicting the definition of $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$.

Part 3: We first show that in the first period, v_L has to be indifferent between participating or not in the mechanism. Let $U_L^*(h^0, \mathbf{M}_0^*)$ denote the buyer's equilibrium payoff at the initial history in the PBE when the seller offers mechanism $\mathbf{M}_0^*$ according to $\Gamma^*(h^0)$. If $U_L^*(h^0, \mathbf{M}_0^*) > 0$, consider the following modification to the seller's strategy profile. At the initial history, instead of offer mechanism $\mathbf{M}_0^*$, the seller offers mechanism $\mathbf{M}_0'$ that coincides with mechanism $\mathbf{M}_0^* \equiv \Gamma^*(\emptyset)$,

except that all transfers $x^{\mathbf{M}_0^*}(\mu')$ such that $\sum_{v \in V} \beta^{\mathbf{M}_0^*}(\mu'|v) > 0$, they are raised by $U_L^*(h^0, \mathbf{M}_0^*)$. Modify the buyer's strategy at h^0 so that $\pi_v^*(h^0, \mathbf{M}_0') = \pi_v^*(h^0, \mathbf{M}_0^*)$ and $r_v^*(h^0, \mathbf{M}_0') = r_v^*(h^0, \mathbf{M}_0^*)$; leave the buyer's strategy everywhere else unchanged. Finally, modify the seller's beliefs so that when the buyer rejects $\mathbf{M}_0'$, he assigns probability 1 to the buyer's valuation being v_H. It is standard to check that the seller's and buyer's sequential rationality constraints hold[30] and the seller's payoff increases by $U_L^*(h^0) > 0$.

To see that part 3 holds after every history on the path of play, let h^t denote the shortest-length public history on the path of play such that a mechanism $\mathbf{M}_t^*$ is offered by the seller at h^t such that $U_L^*(h^t, \mathbf{M}_t^*) > 0$. Following similar steps as above, we can modify the mechanism at h^t, so that the buyer's payoff starting from h^t is 0 when her valuation is v_L and the seller's payoff at h^t goes up by $U_L^*(h^t, \mathbf{M}_t^*)$; adjusting the strategies of the buyer so that it is a best response to participate and truthfully report her type in the modified mechanism. Letting $h^t = (h^{t-1}, \mathbf{M}_{t-1}^*, \mu', 0, x^{\mathbf{M}_{t-1}^*}(\mu'))$, modify $\mathbf{M}_{t-1}^*$ so that transfers are lowered by $\delta(1 - q^{\mathbf{M}_{t-1}^*}(\mu'))U_L^*(h^t, \cdot)$. Clearly, the change is revenue neutral for the seller, it satisfies that $U_L^*(h^t, \cdot) = U_L^*(h^{t-1}, \cdot) = 0$, and it is standard to check that the buyer's incentives and payoffs remain unchanged.

Part 4: Let $\langle \Gamma^*, (\pi_v^*, r_v^*)_{v \in V}, \mu^* \rangle$ denote the PBE assessment that delivers the highest payoff for the seller. Clearly, if the mechanism used in the initial history satisfies that

$$|\text{supp } \mu_0(v)\beta^{\mathbf{M}_0^*}(\cdot|v)| = 1,$$

then the buyer is indifferent between reporting v_L and v_H at h^0.

Consider then the case in which $|\text{supp } \mu_0(v)\beta^{\mathbf{M}_0^*}(\cdot|v)| > 1$. Then, Bayes' consistency implies that two beliefs, μ_H', μ_L', exist such that $\beta^{\mathbf{M}_0^*}(\mu_H'|v_H) > \beta^{\mathbf{M}_0^*}(\mu_H'|v_L)$ and $\beta^{\mathbf{M}_0^*}(\mu_L'|v_H) < \beta^{\mathbf{M}_0^*}(\mu_L'|v_L)$. Let $\epsilon_H = U_H^*(h^0, \mathbf{M}_0^*) - U_{H|L}^*(h^0, \mathbf{M}_0^*)$, where $U_H^*(h^0, \mathbf{M}_0^*)$ denotes the equilibrium payoff of the high type at history h^0 and $U_{H|L}^*$ denotes the buyer's payoff when her type is v_H, she reports v_L into the

---

[30] Because the original assessment is a PBE, sequential rationality of the seller's strategy profile implies that he does not have an incentive to deviate to mechanisms different from $\mathbf{M}_0^*$ and $\mathbf{M}_0'$. Clearly, given the buyer's strategy profile, offering $\mathbf{M}_0'$ dominates offering $\mathbf{M}_0^*$. As for the buyer, note that conditional on participating in $\mathbf{M}_0'$, reporting her valuation truthfully is a best response because it was a best response in the original PBE. Because in the new assessment, participating in $\mathbf{M}_0'$ yields a non-negative payoff for the buyer, whereas rejection guarantees a payoff of 0, participating is a best response for the buyer.

mechanism $\mathbf{M}_0^*$, and then plays according to her equilibrium strategy. By assumption, $\epsilon_H > 0$.

Then, define $\Delta_{\mu'_H}, \Delta_{\mu'_L}$ to be such that

$$\beta^{\mathbf{M}_0^*}(\mu'_H | v_L) \Delta_{\mu'_H} = \beta^{\mathbf{M}_0^*}(\mu'_L | v_L) \Delta_{\mu'_L},$$
$$\epsilon_H = \beta^{\mathbf{M}_0^*}(\mu'_H | v_H) \Delta_{\mu'_H} - \beta^{\mathbf{M}_0^*}(\mu'_L | v_H) \Delta_{\mu'_L}.$$

Modify the transfers of mechanism $\mathbf{M}_0^*$ so that $x^{\mathbf{M}_0'}(\mu'_L) = x^{\mathbf{M}_0^*}(\mu'_L) - \Delta_{\mu'_L}, x^{\mathbf{M}_0'}(\mu'_H) = x^{\mathbf{M}_0^*}(\mu'_H) + \Delta_{\mu'_H}$. Modify the buyer's strategy at h^0 so that $\pi_v^*(h^0, \mathbf{M}_0') = \pi_v^*(h^0, \mathbf{M}_0^*)$ and $r_v^*(h^0, \mathbf{M}_0') = r_v^*(h^0, \mathbf{M}_0^*)$; leave the buyer's strategy everywhere else unchanged. Finally, modify the seller's beliefs so that when the buyer rejects $\mathbf{M}_0'$, he assigns probability 1 to the buyer's valuation being v_H. The seller's revenue goes up by $\mu_0 \epsilon_H$ and the buyer's best response is still to participate in the mechanism and truthfully report her type.[31]

To see that along the path of play, v_H must be indifferent between truthfully reporting and reporting v_L, let h^t denote the shortest-length history on the path of play such that the seller offers a mechanism $\mathbf{M}_t^*$ such that the buyer strictly prefers to report v_H than to report v_L when her type is v_H.[32] As before, at least two posteriors must be generated in the mechanism. Let $\epsilon(h^t) > 0$ denote the difference in payoffs between reporting v_H and reporting v_L. Following the construction in the previous paragraph, modify the strategy profile at h^t so that the seller, instead of offering $\mathbf{M}_t^*$, offers the modified version $\mathbf{M}_t'$. Mechanism $\mathbf{M}_t'$ is as $\mathbf{M}_t^*$, except that the transfers have been modified as we did for the case $t = 0$, so that the buyer is now indifferent between reporting v_H and v_L, when her valuation is v_H. Modify the buyer's strategy at h^t so that $\pi_v^*(h^t, \mathbf{M}_t') = \pi_v^*(h^t, \mathbf{M}_t^*)$ and $r_v^*(h^t, \mathbf{M}_t') = r_v^*(h^t, \mathbf{M}_t^*)$; for now, leave the buyer's strategy everywhere else unchanged. Finally, modify the seller's beliefs so that when the buyer rejects $\mathbf{M}_t'$, he assigns probability 1 to the buyer's valuation being v_H. To keep the incentives at the histories preceding h^t the same, we proceed as follows. Let $h^t = (h^{t-1}, \mathbf{M}_{t-1}^*, 1, \mu', (0, x^{\mathbf{M}_{t-1}^*}(\mu')))$. Modify the strategy at h^{t-1} so that instead

<hr>

[31]To see that participating with probability one and truthtelling are best responses for the buyer, note that when the buyer announces v_L, her expected payment (averaging out over posteriors) is unchanged, whereas when she announces v_H her expected payment goes up by ϵ_H. This means that when her valuation is v_H the buyer is now indifferent between reporting v_L and v_H, whereas the buyer now has a strict incentive to report v_L when her valuation is v_L.

[32]Because without loss of generality we can focus on assessments where the buyer is truthful, if she is not indifferent, she must strictly prefer to tell the truth.

of offering $\mathbf{M}^*_{t-1}$, the seller offers $\mathbf{M}'_{t-1}$. This mechanism coincides with $\mathbf{M}^*_{t-1}$ except that $x^{\mathbf{M}^*_{t-1}}(\mu')$ is lowered by $\Delta_{\mu'} = \delta(1 - q^{\mathbf{M}^*_{t-1}}(\mu'))\epsilon_H(h^t)$. Note that this is payoff neutral for the seller: his increase in payoff at $(h^t, \mathbf{M}'_t)$ is exactly offset by his payoff decrease at h^{t-1}, $\Delta_{\mu'}$. However, by lowering $x^{\mathbf{M}^*_{t-1}}(\mu')$, we may have relaxed the participation constraint of v_L if $1 - \mu^*(h^{t-1}) > 0$, in which case, we increase all transfers by $\epsilon_L = \beta^{\mathbf{M}^*_t}(\mu'|v_L)\Delta_{\mu'} > 0$. Again, modify the buyer's strategy so that her best response at $(h^{t-1}, \mathbf{M}'_{t-1})$ is the same as at $(h^{t-1}, \mathbf{M}^*_{t-1})$ and modify the seller's beliefs so that when the buyer rejects $\mathbf{M}_{t-1}$, he assigns probability 1 to the buyer's valuation being v_H. The best response conditions from the original assessment imply the new assessment is also a PBE assessment. In this new assessment, the seller's payoff (weakly) increases, which contradicts incentive efficiency if $1 - \mu^*(h^{t-1}) > 0$.

C   Proofs of Section 3.3

C.1   Proof of Proposition 2

The proof follows from the following results:

Lemma 2 (Properties of the continuation values, $R^{(\tau,q)}$). *Let $R^{(\tau,q)}$ denote the continuation values induced by a policy (τ, q). Then,*

1. *If $\mu_0 < \bar{\mu}_1$, then $R^{(\tau,q)}(\mu', \mu_0) \leq \mu' v_H + (1 - \mu')\hat{v}_L(\mu_0)$,*

2. *If $\mu_0 > \bar{\mu}_1$, then $R^{(\tau,q)}(\mu', \mu_0) \leq \mu' v_H$.*

Proof. Note the contraction mapping theorem implies $R^{(\tau,q)}$ is well defined for any policy, (τ, q). Fix a prior μ_0, and consider the operator T_{μ_0} that takes bounded function $w : \Delta(V) \mapsto \mathbb{R}$ into

$$T_{\mu_0}(w)(\mu') = \int \left(q(\mu', \tilde{\mu})(\tilde{\mu} v_H + (1 - \tilde{\mu})\hat{v}_L(\mu_0)) + (1 - q(\mu', \tilde{\mu}))\delta w(\tilde{\mu}) \right) \tau(\mu', d\tilde{\mu})$$

To show part 1 holds, let w be a bounded function such that $w(\tilde{\mu}) \leq \tilde{\mu} v_H + (1 - \tilde{\mu})\hat{v}_L(\mu_0)$. Then,

$$T_{\mu_0}(w(\mu')) \leq \int \left(q(\mu', \tilde{\mu})(\tilde{\mu} v_H + (1 - \tilde{\mu})\hat{v}_L(\mu_0)) + (1 - q(\mu', \tilde{\mu}))\delta(\tilde{\mu} v_H + (1 - \tilde{\mu})\hat{v}_L(\mu_0)) \right) \tau(\mu', d\tilde{\mu})$$
$$\leq \mu' v_H + (1 - \mu')\hat{v}_L(\mu_0),$$

43

where the inequalities follow from (i) the bound on w, (ii) $\hat{v}_L(\mu_0) > 0$, and (iii) Bayes' plausibility. This implies that for all $\mu' \in \Delta(V)$, $R^{(\tau,q)}(\mu',\mu_0) \leq \mu'v_H + (1-\mu')\hat{v}_L(\mu_0)$.

Similarly, to show part 2 holds, let w be a bounded function such that $w(\tilde{\mu}) \leq \tilde{\mu}v_H$. Then,

$$T_{\mu_0}(w(\mu')) \leq \int \left(q(\mu',\tilde{\mu})(\tilde{\mu}v_H + (1-\tilde{\mu})\hat{v}_L(\mu_0)) + (1 - q(\mu',\tilde{\mu}))\delta\tilde{\mu}v_H \right) \tau(\mu',d\tilde{\mu})$$
$$\leq \mu'v_H,$$

where the inequalities follow from (i) the bound on w, (ii) $\hat{v}_L(\mu_0) < 0$, and (iii) Bayes' plausibility. This implies that for all $\mu' \in \Delta(V)$, $R^{(\tau,q)}(\mu',\mu_0) \leq \mu'v_H$. □

Proposition 4. *Let $\langle (\tau^*, q^*), R^{(\tau^*,q^*)} \rangle$ be an intrapersonal equilibrium. Then,*

1. *For all $\mu_0 < \overline{\mu}_1$, $q^*(\mu_0,\cdot) = 1$,*

2. *For all $\mu_0 > \overline{\mu}_1$, if $q^*(\mu_0,\mu') = 1$ and the seller induces μ' with positive probability, then $\mu' = 1$.*

3. *For all $\mu_0 > \overline{\mu}_1$, $\int_{\mu_0}^1 \tau^*(\mu_0,d\mu') = \tau^*(\mu_0,1)$, and*

4. *For all $\mu_0 > \overline{\mu}_1$, there exists $\mu^D < \mu_0$ such that $\tau^*(\mu_0,\mu_D) = 1 - \tau^*(\mu_0,1) > 0$ and $q^*(\mu_0,\mu_D) = 0 = 1 - q^*(\mu_0,1)$.*

Proof. To prove part 1, note Lemma 2 implies that

$$R^{(\tau^*,q^*)}(\mu_0,\mu_0) \leq \mu_0 v_H + (1-\mu_0)\hat{v}_L(\mu_0),$$

and this bound is achieved by setting $q^*(\mu_0,\cdot) = 1$.

To prove part 2, note that whenever $\mu_0 > \overline{\mu}_1$, for any posterior μ' it follows that

$$\mu'v_H + (1-\mu')\hat{v}_L(\mu_0) < \mu'v_H + (1-\mu')\delta\hat{v}_L(\mu_0), \tag{17}$$

because $\hat{v}_L(\mu_0) < 0$ and $\delta < 1$. Thus, instead of setting $q^*(\mu_0,\mu') = 1$, which yields the payoff on the left-hand side of Equation 17, the seller can split μ' between 0 and 1, setting $q^*(\mu_0,1) = 1 = 1 - q^*(\mu_0,0)$, which yields the payoff on the right-hand side of Equation 17.

To prove part 3, note Bayes' plausibility implies that whenever $\tau^*(\mu_0, \mu_0) \neq 1$, $\int_{\mu_0}^1 \tau^*(\mu_0, \mu') \neq 0$. Then,

$$
R^{(\tau^*,q^*)}(\mu_0, \mu_0) = \int \left(q^*(\mu_0, \mu')(\mu' v_H + (1 - \mu')\hat{v}_L(\mu_0)) + (1 - q^*(\mu_0, \mu'))\delta R^{(\tau,q)}(\mu', \mu_0) \right) \tau^*(\mu_0, d\mu')
$$

$$
\leq \int_0^{\mu_0} \left(q^*(\mu_0, \mu')(\mu' v_H + (1 - \mu')\hat{v}_L(\mu_0)) + (1 - q^*(\mu_0, \mu'))\delta R^{(\tau,q)}(\mu', \mu_0) \right) \tau^*(\mu_0, d\mu')
$$

$$
+ \delta \int_{\mu_0}^1 \mu' v_H \tau^*(\mu_0, d\mu')
$$

$$
< \int_0^{\mu_0} \left(q^*(\mu_0, \mu')(\mu' v_H + (1 - \mu')\hat{v}_L(\mu_0)) + (1 - q^*(\mu_0, \mu'))\delta R^{(\tau,q)}(\mu', \mu_0) \right) \tau^*(\mu_0, d\mu')
$$

$$
+ \int_{\mu_0}^1 \mu' v_H \tau^*(\mu_0, d\mu'),
$$

where the first inequality follows from Lemma 2 and the second one from $\delta < 1$. However, the upper bound can be achieved by moving all the mass above μ_0 to 1.

To prove part 4, we first make the following observation. Suppose the optimal policy for the seller μ_0 placed positive mass on three posteriors, $\mu_1', \mu_2', 1$. Part 3 implies $\mu_2' \leq \mu_0$. Then, the seller's payoff at prior μ_0 is given by

$$
R^{(\tau^*,q^*)}(\mu_0, \mu_0) = \tau_1 \delta R^{(\tau^*,q^*)}(\mu_1', \mu_0) + \tau_2 \delta R^{(\tau^*,q^*)}(\mu_2', \mu_0) + (1 - \tau_1 - \tau_3)v_H,
$$

where $\tau_1 \mu_1' + \tau_2 \mu_2' + (1 - \tau_1 - \tau_2)1 = \mu_0$, $0 < \tau_1, \tau_2$ and $\tau_1 + \tau_2 < 1$. It follows from part 2 of Lemma 2 that the above is dominated by a policy that sets weight τ_1 on μ_1' and $\tau_2 \mu_2' + (1 - \tau_1 - \tau_2)$ on posterior 1.

Then, fix $\mu_0 > \bar{\mu}_1$ and assume the seller delays trade with positive probability at more than one posterior. Part 3 implies $R^{(\tau^*,q^*)}(\mu_0, \mu_0)$ can be written as:

$$
R^{(\tau^*,q^*)}(\mu_0, \mu_0) = (1 - \tau^*(\mu_0, 1)) \int_0^{\mu_0} \delta R^{(\tau^*,q^*)}(\mu', \mu_0) \frac{\tau^*(\mu_0, d\mu')}{1 - \tau^*(\mu_0, 1)} + \tau^*(\mu_0, 1)v_H.
$$

Let $\mu_D = \int_0^{\mu_0} (\mu'/(1 - \tau^*(\mu_0, 1)))\tau^*(\mu_0, d\mu')$ and define:

$$
R^D = \int_0^{\mu_0} \delta R^{(\tau^*,q^*)}(\mu', \mu_0) \frac{\tau^*(\mu_0, d\mu')}{1 - \tau^*(\mu_0, 1)}.
$$

Let $D = \{(\mu', R^{(\tau^*,q^*)}(\mu', \mu_0)) : \mu' \in \operatorname{supp} \tau^*(\mu_0, \cdot) \wedge \mu' \leq \mu_0\}$. Let co D denote the convex hull of D. Note $(\mu_D, R^D) \in cl(\operatorname{co} D) = \operatorname{co}D$, where the last equality follows from $D \subseteq \mathbb{R}^2$.

Consider $R^* = \max\{r : (\mu_D, r) \in \text{co } D\}$. By Caratheodory's theorem, (μ_D, R^*) can be written as the convex combination of $m \leq 2$ points in D. If $m = 1$, the policy that splits between μ_D and 1 is a strict improvement, a contradiction. (It is a strict improvement because it implies (μ_D, R^*) is an extreme point of D and so cannot be achieved by convex combinations of points in D.)

Suppose now that $(\mu_D, R^*) = \sum_{i=1}^{2} \lambda_i (\mu_i, r_i)$, with $(\mu_i, r_i) \in D$. Then, the previous observation implies that for $i, j \in \{1, 2\}, i \neq j$, the policy that puts weight $(1 - \tau^*(\mu_0, 1))\lambda_i$ on μ_i and $\tau^*(\mu_0, 1) + (1 - \tau^*(\mu_0, 1))(1 - \lambda_i)\mu_j$ on 1 strictly improves the seller's payoff when his prior is μ_0. Thus, it has to be the case that

$$(1 - \tau^*(\mu_0, 1))\lambda_1 \delta R^{(\tau^*, q^*)}(\mu_1, \mu_0) + (\tau^*(\mu_0, 1) + (1 - \tau^*(\mu_0, 1))(1 - \lambda_2))v_H$$
$$= (1 - \tau^*(\mu_0, 1))(1 - \lambda_1)\delta R^{(\tau^*, q^*)}(\mu_2, \mu_0) + (\tau^*(\mu_0, 1) + (1 - \tau^*(\mu_0, 1))\lambda_1)v_H.$$

so that all points lie in the same line. Note that, by definition,

$$(1 - \tau^*(\mu_0, 1))R^D + \tau^*(\mu_0, 1)v_H \leq (1 - \tau^*(\mu_0, 1))R^* + \tau^*(\mu_0, 1)v_H$$
$$\leq (1 - \tau^*(\mu_0, 1))\lambda_1 \delta R^{(\tau^*, q^*)}(\mu_1, \mu_0) + (\tau^*(\mu_0, 1) + (1 - \tau^*(\mu_0, 1))(1 - \lambda_2))v_H,$$

so the seller has a (strict) deviation when his prior is μ_0. This completes the proof of part 4. □

C.2   Proof of Proposition 3

Part 1: Note that if $\mu_{D^*} \equiv \mu_{D^*}(\mu_0) < \bar{\mu}_1$, then $R^{\mu_{D^*}}(\mu_{D^*}, \mu_0) = (1 - \mu_{D^*})\hat{v}_L(\mu_0) + \mu_{D^*}v_H$. Consider the following alternative policy for μ_0, where $\mu_{D'}(\mu_0) = 0$. This is achieved by $\tau'(\mu_0, 1) = \tau^*(\mu_0, 1) + \mu_{D^*}$ and $\tau'(\mu_0, 0) = 1 - \tau^*(\mu_0, 1) - \mu_{D^*}$, where $\tau^*(\mu_0, \cdot)$ are the weights on μ_{D^*} and 1 in the original best response. The difference in payoffs is given by:

$$\mu_{D^*}v_H + (1 - \tau^*(\mu_0, 1))\delta[\hat{v}_L(\mu) - \mu_{D^*}v_H - (1 - \mu_{D^*})\hat{v}_L(\mu_0)] - \mu_{D^*}\delta\hat{v}_L(\mu_0) =$$
$$= \mu_{D^*}v_H + (1 - \tau^*(\mu_0, 1)\delta\mu_{D^*}[\hat{v}_L(\mu_0) - v_H] - \mu_{D^*}\delta\hat{v}_L(\mu)$$
$$= \mu_{D^*}(v_H - \delta\hat{v}_L(\mu_0)) - (1 - \tau^*(\mu_0, 1))\mu_{D^*}\delta(v_H - \hat{v}_L(\mu_0))$$
$$= \mu_{D^*}v_H(1 - \tau^*(\mu_0, 1)\delta) - \mu_{D^*}\delta\hat{v}_L(\mu_0)\tau^*(\mu_0, 1) > 0,$$

where the last inequality follows from noting $\hat{v}_L(\mu_0) < 0$. This contradicts that $\mu_{D^*}(\mu_0)$ is a best response. Hence, part 1 holds.

46

Part 2: Let $\overline{\mu}_1 \leq \mu_0$ denote the seller's prior. Toward a contradiction, assume that for all finite n, $\mu_n = \mu_{D^*}^{(n)}(\mu_0) > \overline{\mu}_1$. Fix $\epsilon > 0$. Note that there exists M_ϵ such that for all $M_\epsilon \leq m$, we have

$$\mu_m - \mu_{m+1} < \epsilon(1 - \mu_{m+1}), \tag{18}$$

so that as m goes to infinity, the seller puts smaller weight on posterior 1. To see that Equation 18 holds, assume to the contrary that for all m, we have

$$\epsilon(1 - \mu_{m+1}) \leq \mu_m - \mu_{m+1}.$$

Adding up from $m = 0$ to $m = N - 1$, we obtain

$$\epsilon \sum_{i=1}^{N} (1 - \mu_i) \leq \mu_0 - \mu_N. \tag{19}$$

Now, the right-hand side of Equation 19 is bounded above by $\mu_0 - \overline{\mu}_1$ because $\overline{\mu}_1 < \mu_N$. The left-hand side of Equation 19 is bounded below by $N\epsilon(1 - \mu_0)$ because $\mu_0 \geq \mu_i$ for all $i \geq 0$. We then obtain the following chain of inequalities:

$$N\epsilon(1 - \mu_0) < \epsilon \sum_{i=1}^{N} (1 - \mu_i) \leq \mu_0 - \mu_N < \mu_0 - \overline{\mu}_1,$$

and this cannot possibly hold for all N. Thus, Equation 18 holds.

Now fix $\epsilon > 0$ such that

$$\frac{\epsilon}{1 - \delta(1 - \epsilon)} < \overline{\mu}_1(1 - \delta). \tag{20}$$

Note the right-hand side of Equation 20 is bounded above by $\mu_m(1 - \delta)$ for all $m \geq 0$. Take $M_\epsilon \leq m$ and note

$$R^{\mu_{D^*}}(\mu_m, \mu_m) < \epsilon v_H + (1 - \epsilon)\delta R^{\mu_{D^*}}(\mu_{m+1}, \mu_m). \tag{21}$$

To see that Equation 21 holds, note

$$
\begin{aligned}
R^{\mu_{D^*}}(\mu_m, \mu_m) &= \frac{\mu_m - \mu_{m+1}}{1 - \mu_{m+1}} v_H + \frac{1 - \mu_m}{1 - \mu_{m+1}} \delta R^{\mu_{D^*}}(\mu_{m+1}, \mu_m) \\
&= \frac{\mu_m - \mu_{m+1}}{1 - \mu_{m+1}} (v_H - \delta R^{\mu_{D^*}}(\mu_{m+1}, \mu_m)) + \delta R^{\mu_{D^*}}(\mu_{m+1}, \mu_m) \\
&< \epsilon v_H + (1 - \epsilon)\delta R^{\mu_{D^*}}(\mu_{m+1}, \mu_m),
\end{aligned}
$$

where the inequality follows from $v_H > \delta R^{\mu_{D*}}(\mu_{m+1}, \mu_m)$ and $M_\epsilon \leq m$. Because Equation 21 holds for all $M_\epsilon \leq m$, applying it recursively, we obtain

$$R^{\mu_{D*}}(\mu_m, \mu_m) < \frac{v_H \epsilon}{1 - \delta(1 - \epsilon)} + \lim_{N \to \infty} (\delta(1 - \epsilon))^N R^{\mu_{D*}}(\mu_{m+N}, \mu_m) = \frac{v_H \epsilon}{1 - \delta(1 - \epsilon)},$$

because $R^{\mu_{D*}}(\cdot, \mu_m)$ is bounded and $\delta(1 - \epsilon) < 1$. Now, the following policy is always feasible for μ_m: place probability μ_m on 1 and $1 - \mu_m$ on 0. This yields

$$\mu_m(1 - \delta)v_H + v_L \delta,$$

when $\overline{\mu}_1 > 0$ and

$$\mu_m v_H,$$

when $\overline{\mu}_1 = 0$. Suppose $0 < \overline{\mu}_1$ (the other case is similar). Because splitting between 0 and 1 is not optimal, we must have

$$\mu_m(1 - \delta)v_H + v_L \delta \leq R^{\mu_{D*}}(\mu_m, \mu_m) < \frac{\epsilon}{1 - \delta(1 - \epsilon)}v_H < v_H(1 - \delta)\overline{\mu}_1,$$

where the last inequality follows from Equation 20. This is a contradiction to $\langle \mu_{D*}, R^{\mu_{D*}} \rangle$ being an intrapersonal equilibrium. Thus, starting from μ_0, a finite N exists such that $\mu_{D*}^{(N)}(\mu_0) \leq \overline{\mu}_1$.

Lemma 3 below is used in the proof of Proposition 3, part 3, and it is the first step in the induction in the proof of Proposition 3, part 4:

Lemma 3. *Let* $\overline{\mu}_1 < \mu_0 < \mu_0''$ *be such that* $\mu_{D*}(\mu_0) = \mu_{D*}(\mu_0'') = 0$. *Then, for all* $\mu_0' \in (\mu_0, \mu_0'')$, $\mu_{D*}(\mu_0') = 0$.

Proof. Suppose $\mu_0 < \mu_0' < \mu_0''$ are such that $\mu_{D*}(\mu_0) = \mu_{D*}(\mu_0'') = 0$, $\mu_D \equiv \mu_{D*}(\mu_0') \geq \overline{\mu}_1$.[33] Then, the following must be true:

$$\frac{\mu_0' - \mu_D}{1 - \mu_D}v_H + \frac{1 - \mu_0'}{1 - \mu_D}\delta R^{\mu_{D*}}(\mu_D, \mu_0') \geq \mu v_H + (1 - \mu_0')\delta \hat{v}_L(\mu_0').$$

Rewriting the above expression, we obtain:

$$g(\mu_0') \equiv \left[\frac{\mu_0' - \mu_D}{1 - \mu_D} - \mu_0'\right]v_H + \frac{1 - \mu_0'}{1 - \mu_D}\delta R^{\mu_{D*}}(\mu_D, \mu_0') + (1 - \mu_0')\delta(-\hat{v}_L(\mu_0')) \geq 0. \tag{22}$$

---

[33] By Proposition 3, part 1, if $\mu_D(\mu_0') \neq 0$, then $\mu_D(\mu_0') \geq \overline{\mu}_1$.

In Section V.1 in Doval and Skreta (2019), we show the continuation values $R^{\mu_{D^*}}$ can be written as:

$$R^{\mu_{D^*}}(\mu_D, \mu_0') = v_H \underbrace{\left(\frac{\mu_D - \mu_1}{1 - \mu_1} + (1 - \mu_D) \sum_{i=1}^{N_D-1} \frac{\delta^i}{1 - \mu_i} \frac{\mu_i - \mu_{i+1}}{1 - \mu_{i+1}} \right)}_{\alpha(\mu_D)} + \underbrace{(1 - \mu_D)\delta^{N_D}}_{\gamma(\mu_D)} \hat{v}_L(\mu_0'),$$

where $\mu_1 > \mu_2 > \cdots > \mu_{N_D-1} \geq \overline{\mu}_1 > 0 = \mu_{N_D}$ are the posteriors generated by the policy at which trade is delayed. Note $R^{\mu_{D^*}}(\mu_D, \mu_0')$ is differentiable in μ_0' and

$$\frac{\partial}{\partial \mu} R^{\mu_{D^*}}(\mu_D, \mu)|_{\mu=\mu_0'} = -\gamma(\mu_D)\frac{\Delta v}{(1 - \mu_0')^2}.$$

Therefore, g is differentiable with respect to μ_0', and we obtain the following expression for $g'(\mu_0')$:

$$
\begin{aligned}
g'(\mu_0') &= \frac{\mu_D}{1 - \mu_D} v_H - \frac{\delta R^{\mu_{D^*}}(\mu_D, \mu_0')}{1 - \mu_D} + \delta \hat{v}_L(\mu_0') + \delta(1 - \mu_0')\left[1 - \frac{\gamma(\mu_D)}{1 - \mu_D} \right] \frac{\Delta v}{(1 - \mu_0')^2} \\
&= v_H \frac{\mu_D}{1 - \mu_D} - \frac{\delta}{1 - \mu_D}\left(\alpha(\mu_D)v_H + \gamma(\mu_D)\hat{v}_L(\mu_0') \right) + \delta \hat{v}_L(\mu_0') + \delta \frac{\Delta v}{1 - \mu_0'}\left[1 - \frac{\gamma(\mu_D)}{1 - \mu_D} \right] \\
&= v_H \left[\frac{\mu_D - \delta\alpha(\mu_D)}{1 - \mu_D} \right] + \delta \left(\hat{v}_L(\mu_0') + \frac{\Delta v}{1 - \mu_0'} \right)\left[1 - \frac{\gamma(\mu_D)}{1 - \mu_D} \right] \\
&= v_H \left[\frac{\mu_D - \delta\alpha(\mu_D)}{1 - \mu_D} \right] + \delta v_H \left[1 - \frac{\gamma(\mu_D)}{1 - \mu_D} \right].
\end{aligned}
$$

Now, note $\alpha(\mu_D) \leq \mu_D$. Intuitively, the largest probability that μ_D can trade with v_H is μ_D. To see this formally, we use that[34]

$$
\begin{aligned}
\alpha(\mu_D) &= \frac{\mu_D - \mu_1}{1 - \mu_1} + (1 - \mu_D) \sum_{i=1}^{N_D-1} \frac{\delta^i}{1 - \mu_i} \frac{\mu_i - \mu_{i+1}}{1 - \mu_{i+1}} \\
&\leq \frac{\mu_D - \mu^1}{1 - \mu^1} + (1 - \mu_D)\frac{\mu_1}{1 - \mu_1} \sum_{i=1}^{N_D-1} \delta_i,
\end{aligned}
$$

where the inequality follows from noting that the expression inside the sum is increasing in μ_i and decreasing in μ_{i+1}, and that $\mu_D - \mu_1 + (1 - \mu_D)\mu_1 \sum_{i=1}^{N_D-1} \delta_i \leq \mu_D(1 - \mu_1)$.

---

[34]See Section V.1 in Doval and Skreta (2019).

It then follows that $g'(\mu_0') \geq 0$, which contradicts that μ_0'' finds it optimal to set $\mu_{D^*}(\mu'') = 0$. $\qquad\square$

Part 3: We use the following property of a policy where $\mu_{D^*}(\cdot)$ is weakly increasing. If we let N_{μ_0} denote the smallest value of n such that $\mu_{D^*}^{(n)}(\mu_0) \leq \overline{\mu}_1$, then N_{μ_0} is weakly increasing in μ_0.[35] To see this, let $\mu_0 < \mu_0'$, and inductively define $\mu_{i+1} = \mu_{D^*}(\mu_i)$ and similarly, $\mu_{i+1}' = \mu_{D^*}(\mu_i')$. Clearly, we have that $\mu_0 > \mu_1 > \cdots > \mu_N > \ldots$ and similarly, $\mu_0' > \mu_1' > \cdots > \mu_N' > \ldots$. Moreover, monotonicity implies $\mu_1 = \mu_{D^*}(\mu_0) \leq \mu_1' = \mu_{D^*}(\mu_0')$ and inductively $\mu_N < \mu_N'$. Lemma 3 implies that if $\mu_{D^*}(\mu_N') = 0$, then letting n be the smallest m such that $\mu_m \leq \mu_N'$, we have $\mu_{D^*}(\mu_m) = 0$. Hence, $N_{\mu_0} \leq N_{\mu_0'}$.

Let μ_0 denote the smallest prior above $\overline{\mu}_1$ such that there exists $\mu_0' > \mu_0$ with $\mu_{D^*}(\mu_0) > \mu_{D^*}(\mu_0')$.[36]

Now consider a policy for μ_0 that splits the weight $\tau(\mu_0, \mu_{D^*}(\mu_0))$ between $\mu_{D^*}(\mu_0')$ and 1. The payoff from this policy is:

$$v_H \left[\frac{\mu_0 - \mu_{D^*}(\mu_0)}{1 - \mu_{D^*}} + \frac{1 - \mu_0}{1 - \mu_{D^*}} \frac{\mu_{D^*}(\mu_0) - \mu_{D^*}(\mu_0')}{1 - \mu_{D^*}(\mu_0')} \right] + \delta \frac{1 - \mu}{1 - \mu_{D^*}(\mu_0')} R^{\mu_{D^*}}(\mu_{D^*}(\mu_0'), \mu_0).$$

Consider the difference between the payoff of the above policy and that of the policy that the seller employs when his prior is μ_0:

$$\frac{1 - \mu_0}{1 - \mu_{D^*}(\mu_0)} \frac{\mu_{D^*}(\mu_0) - \mu_{D^*}(\mu_0')}{1 - \mu_{D^*}(\mu_0')} v_H + \delta \frac{1 - \mu_0}{1 - \mu_{D^*}(\mu_0')} R^{\mu_{D^*}}(\mu_{D^*}(\mu_0'), \mu_0)$$
$$- \delta \frac{1 - \mu_0}{1 - \mu_{D^*}(\mu_0)} R^{\mu_{D^*}}(\mu_{D^*}(\mu_0), \mu_0). \qquad (23)$$

Optimality of μ_0''s policy implies

$$\frac{\mu_0' - \mu_{D^*}(\mu_0')}{1 - \mu_{D^*}(\mu_0')} v_H + \delta \frac{1 - \mu_0'}{1 - \mu_{D^*}(\mu_0')} R^{\mu_{D^*}}(\mu_{D^*}(\mu_0'), \mu_0')$$
$$\geq \frac{\mu_0' - \mu_{D^*}(\mu_0)}{1 - \mu_{D^*}(\mu_0)} v_H + \delta \frac{1 - \mu_0'}{1 - \mu_{D^*}(\mu_0)} R^{\mu_{D^*}}(\mu_{D^*}(\mu_0), \mu_0'),$$

or, equivalently,

$$\delta \left[\frac{R^{\mu_{D^*}}(\mu_{D^*}(\mu_0'), \mu_0')}{1 - \mu_{D^*}(\mu_0')} - \frac{R^{\mu_{D^*}}(\mu_{D^*}(\mu_0), \mu_0)}{1 - \mu_{D^*}(\mu_0)} \right] \geq - \frac{(\mu_{D^*}(\mu_0) - \mu_{D^*}(\mu_0'))}{(1 - \mu_{D^*}(\mu_0))(1 - \mu_{D^*}(\mu_0'))} v_H.$$

---

[35] Proposition 3 implies $\mu_{D^*}(\mu_{D^*}^{(N_{\mu_0} - 1)}) = 0$.

[36] Note $\mu_0 > \overline{\mu}_1$ because we know $\mu_{D^*}(\overline{\mu}_1) = 0$.

To show μ_0 can improve by using the new policy, we only need to show

$$\frac{R^{\mu_{D^*}}(\mu_{D^*}(\mu_0'),\mu_0)}{1-\mu_{D^*}(\mu_0')} - \frac{R^{\mu_{D^*}}(\mu_{D^*}(\mu_0),\mu_0)}{1-\mu_{D^*}(\mu_0)} \geq \frac{R^{\mu_{D^*}}(\mu_{D^*}(\mu_0'),\mu_0')}{1-\mu_{D^*}(\mu_0')} - \frac{R^{\mu_{D^*}}(\mu_{D^*}(\mu_0),\mu_0')}{1-\mu_{D^*}(\mu_0)}$$

$$\Leftrightarrow$$

$$(\delta^{N_{\mu_{D^*}(\mu_0')}} - \delta^{N_{\mu_{D^*}(\mu_0)}})(-\hat{v}_L(\mu_0)) \leq (\delta^{N_{\mu_{D^*}(\mu_0')}} - \delta^{N_{\mu_{D^*}(\mu_0)}})(-\hat{v}_L(\mu_0')),$$

where recall that $N.$ denotes the first time at which updating leads to a posterior below $\overline{\mu}_1$ for a given prior. Note that a sufficient condition for the above inequality to hold is that $N_{\mu_{D^*}(\mu_0')} < N_{\mu_{D^*}(\mu_0)}$. That this is indeed the case follows from the observation before the proof because the policy below μ_0 (and hence for $\mu_{D^*}(\mu_0), \mu_{D^*}(\mu_0')$) satisfies monotonicity.

Because $N_{\mu_{D^*}(\mu_0')} < N_{\mu_{D^*}(\mu_0)}$ implies the expression in Equation 23 is non-negative, we see that without loss of generality, we can have $\mu_{D^*}(\mu_0) \leq \mu_{D^*}(\mu_0')$ and strictly so when $N_{\mu_{D^*}(\mu_0)} \neq N_{\mu_{D^*}(\mu_0')}$.

Part 4: The proof is by induction on $n \geq 1$. Define $D_0 = \{0\}$, and for $n \geq 1$,

$$D_n = \{\mu_0 : \mu_{D^*}^{(n)} = 0\},$$

to be the set of priors, μ_0, such that the seller updates his prior to 0 in n periods when his prior is μ_0. Let $\mathbf{P(n)}$ denote the following inductive statement:

$\mathbf{P}(n)$: If $\mu_0, \mu_0' \in D_n$, then $\mu_{D^*}(\mu_0) = \mu_{D^*}(\mu_0') \in D_{n-1}$.

Lemma 3 implies the inductive statement is true for $n = 1$, i.e., $\mathbf{P}(1) = 1$.

Suppose $\mathbf{P}(n) = 1$, and we show that $\mathbf{P}(n+1) = 1$. Let $\mu_0, \mu_0' \in D_{n+1}$ and assume without loss of generality that $\mu_0 < \mu_0'$. Toward a contradiction, suppose $\mu_{D^*}(\mu_0) \neq \mu_{D^*}(\mu_0')$; by part 3, it follows that $\mu_{D^*}(\mu_0) < \mu_{D^*}(\mu_0')$. The payoff for the seller when his prior is μ_0' from following his policy is

$$\frac{\mu_0' - \mu_{D^*}(\mu_0')}{1-\mu_{D^*}(\mu_0')}v_H + \delta\frac{1-\mu_0'}{1-\mu_{D^*}(\mu_0')}\left[\frac{\mu_{D^*}(\mu_0') - \overline{\mu}_{n-1}}{1-\overline{\mu}_{n-1}}v_H + \delta\frac{1-\mu_{D^*}(\mu_0')}{1-\overline{\mu}_{n-1}}R^{\mu_{D^*}}(\overline{\mu}_{n-1},\mu_0')\right],$$

where $\overline{\mu}_{n-1} = \mu_{D^*}(\mu_{D^*}(\mu_0'))$ updates. By the inductive hypothesis, $\mu_{D^*}(\mu_{D^*}(\mu_0')) = \mu_{D^*}(\mu_{D^*}(\mu_0)) = \overline{\mu}_{n-1}$.

Instead, if the seller follows μ_0's policy when his posterior is μ_0', his payoff would be

$$\frac{\mu_0' - \mu_{D^*}(\mu_0)}{1-\mu_{D^*}(\mu_0)}v_H + \delta\frac{1-\mu'}{1-\mu_{D^*}(\mu_0)}\left[\frac{\mu_{D^*}(\mu_0) - \overline{\mu}_{n-1}}{1-\overline{\mu}_{n-1}}v_H + \delta\frac{1-\mu_{D^*}(\mu_0)}{1-\overline{\mu}_{n-1}}R^{\mu_{D^*}}(\overline{\mu}_{n-1},\mu_0')\right],$$

where we use the inductive statement that $\mu_{D^*}(\mu_0), \mu_{D^*}(\mu'_0) \in D_{n-1}$ and hence $\mu_{D^*}(\mu_{D^*}(\mu_0)) = \mu_{D^*}(\mu_{D^*}(\mu'_0)) = \overline{\mu}_{n-1}$.

Taking differences, we have:

$$v_H \left[\frac{\mu'_0 - \mu_{D^*}(\mu'_0)}{1 - \mu_{D^*}(\mu'_0)} + \delta \frac{1 - \mu'_0}{1 - \mu_{D^*}(\mu'_0)} \frac{\mu_{D^*}(\mu'_0) - \overline{\mu}_{n-1}}{1 - \overline{\mu}_{n-1}} - \frac{\mu'_0 - \mu_{D^*}(\mu_0)}{1 - \mu_{D^*}(\mu_0)} - \delta \frac{1 - \mu'_0}{1 - \mu_{D^*}(\mu_0)} \frac{\mu_{D^*}(\mu_0) - \overline{\mu}_{n-1}}{1 - \overline{\mu}_{n-1}} \right]$$

$$= \frac{v_H(1 - \mu'_0)}{(1 - \mu_{D^*}(\mu_0))(1 - \mu_{D^*}(\mu'_0))} \left((\delta - 1)(\mu_{D^*}(\mu'_0) - \mu_{D^*}(\mu_0)) \right) < 0,$$

which contradicts the optimality of μ'_0's policy. Thus, $\mu_{D^*}(\mu_0) = \mu_{D^*}(\mu'_0)$ (by part 3 we cannot have $\mu_{D^*}(\mu'_0) < \mu_{D^*}(\mu_0)$). Hence, $\mathbf{P(n+1)} = 1$.

C.3   Proof of Theorem 2

As in the proof of Proposition 3, define

$$D_n = \{\mu_0 : \mu_{D^*}^{(n)}(\mu_0) = 0\},$$

to be the set of seller priors such that trade with v_L happens in n periods.

The proof of Theorem 2 is structured as follows. First, Lemma 4 shows the cutoffs $\{\overline{\mu}_n\}_{n \geq 0}$ form a strictly increasing sequence. Second, we construct the payoffs from (τ^*, q^*). Finally, we show $\langle (\tau^*, q^*), R^{(\tau^*, q^*)} \rangle$ is an intrapersonal equilibrium.

Lemma 4. *The sequence $\{\overline{\mu}_n\}_{n \geq 0}$ is strictly increasing in n for $n \geq 1$.*

Proof. Recall that we are defining $\overline{\mu}_0 = 0$ and $\overline{\mu}_1 = v_L/v_H$. Now let $n \geq 1$ and let $\mu_n \in D_n$ and $\mu_{n-1} \in D_{n-1}$. Fix a prior, $\mu_0 \geq \mu_n$. We claim that the difference

$$\Delta_n(\mu_0; \mu_n, \mu_{n-1}) = \left[\frac{\mu_0 - \mu_n}{1 - \mu_n} v_H + \frac{1 - \mu_0}{1 - \mu_n} \delta R^{\mu_{D^*}}(\mu_n, \mu_0) \right] - \left[\frac{\mu_0 - \mu_{n-1}}{1 - \mu_{n-1}} v_H + \frac{1 - \mu_0}{1 - \mu_{n-1}} \delta R^{\mu_{D^*}}(\mu_{n-1}, \mu_0) \right],$$

is increasing in μ_0 for any monotone simple policy μ_{D^*}.

The proof follows similar steps to those in the proof of Lemma 3. The arguments

52

in that proof imply Δ_n is differentiable in μ_0. Then,

$$\frac{\partial}{\partial \mu_0} \Delta_n(\mu_0; \mu_n, \mu_{n-1}) =$$

$$= \frac{v_H(\mu_n - \mu_{n-1})}{(1-\mu_n)(1-\mu_{n-1})} - \delta\left[\frac{R^{\mu_{D^*}}(\mu_n, \mu_0)}{1-\mu_n} - \frac{R^{\mu_{D^*}}(\mu_{n-1}, \mu_0)}{1-\mu_{n-1}}\right] + \delta(1-\mu_0)[\delta^{n-1} - \delta^n]\frac{\Delta v}{(1-\mu_0)^2}$$

$$= \frac{v_H(\mu_n - \mu_{n-1})}{(1-\mu_n)(1-\mu_{n-1})} - \delta\left[\frac{v_H \alpha(\mu_n)}{1-\mu_n} - \frac{v_H \alpha(\mu_n - 1)}{1-\mu_{n-1}}\right] + \delta \hat{v}_L(\mu_0)(\delta^{n-1} - \delta^n) + \delta[\delta^{n-1} - \delta^n]\frac{\Delta v}{(1-\mu_0)}$$

$$= \frac{v_H(\mu_n - \mu_{n-1})}{(1-\mu_n)(1-\mu_{n-1})} - \delta v_H\left[\frac{\mu_n - \mu_{n-1}}{(1-\mu_n)(1-\mu_{n-1})}\right] + \delta(\delta^{n-1} - \delta^n)v_H > 0,$$

where the last equality follows from using the definition of $\alpha(\cdot)$.

Now recall Equation 14, which we reproduce below for easy reference:

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

Note that the cutoff $\overline{\mu}_{n+1}$ is defined by $\Delta_{n+1}(\overline{\mu}_{n+1}; \overline{\mu}_n, \overline{\mu}_{n-1}) = 0$.

Note that $\overline{\mu}_{n+1} \neq \overline{\mu}_n$ if $n \geq 1$. If $\overline{\mu}_n = \overline{\mu}_{n+1}$, then

$$0 = \Delta_{n+1}(\overline{\mu}_{n+1}; \overline{\mu}_n, \overline{\mu}_{n-1}) = \delta R^{(\tau^*, q^*)}(\overline{\mu}_n, \overline{\mu}_n) - R^{(\tau^*, q^*)}(\overline{\mu}_n, \overline{\mu}_n) \leq 0,$$

a contradiction, unless $\delta = 1$. Because $\Delta_{n+1}(\cdot; \overline{\mu}_n, \overline{\mu}_{n-1})$ is increasing, $\overline{\mu}_{n+1} > \overline{\mu}_n$. This completes the proof. $\square$

An implication of Lemma 4 is that $\overline{\mu}_{n+1} = \inf\{\mu_0 : \mu_0 \in D_{n+1}\} = \sup\{\mu_0 : \mu_0 \in D_n\}$. The equilibrium we construct makes it so that $\overline{\mu}_{n+1} = \min\{\mu_0 : \mu_0 \in D_{n+1}\}$.

Given (τ^*, q^*) as in Theorem 2, we construct the continuation payoffs, $R^{(\tau^*, q^*)}$, for the seller for each of his priors, μ_0, and posteriors he may induce, μ'. If $\mu' < \overline{\mu}_1$,

$$R^{(\tau^*, q^*)}(\mu', \mu_0) = \mu' v_H + (1 - \mu')\hat{v}_L(\mu_0), \tag{24}$$

whereas if $\mu' \in [\overline{\mu}_n, \overline{\mu}_{n+1})$ for $n \geq 1$, we have

$$R^{(\tau^*, q^*)}(\mu', \mu_0) = \frac{\mu' - \overline{\mu}_{n-1}}{1 - \overline{\mu}_{n-1}} v_H + \delta \frac{1 - \mu'}{1 - \overline{\mu}_{n-1}} R^{(\tau^*, q^*)}(\overline{\mu}_{n-1}, \mu_0) \tag{25}$$

$$R^{(\tau^*, q^*)}(\overline{\mu}_n, \mu_0) = v_H \left[\frac{\overline{\mu}_n - \overline{\mu}_{n-1}}{1 - \overline{\mu}_{n-1}} + (1 - \overline{\mu}_n)\sum_{i=1}^{n-1} \frac{\delta^{n-i}}{1 - \overline{\mu}_i} \frac{\overline{\mu}_i - \overline{\mu}_{i-1}}{1 - \overline{\mu}_{i-1}}\right] + \delta^n(1 - \overline{\mu}_n)\hat{v}_L(\mu_0)$$

53

We now verify that for each $\mu_0 \geq \overline{\mu}_1$, $\langle \tau^*(\mu_0, \cdot) q^*(\mu_0, \cdot) \rangle$, is optimal given $R^{(\tau^*, q^*)}(\cdot, \mu_0)$. Recall that given $R^{(\tau^*, q^*)}(\cdot, \mu_0)$, the seller solves:

$$\max_{\tau, q} \int_0^1 [q(\mu_0, \mu')(\mu' v_H + (1 - \mu')\hat{v}_L(\mu_0)) + \delta(1 - q(\mu_0, \mu')) R^{(\tau^*, q^*)}(\mu', \mu_0)] d\tau(\mu_0, \mu')$$

(26)

subject to the constraints that $q(\mu_0, \mu') \in [0, 1]$ and $\int_0^1 \mu' d\tau(\mu_0, \mu') = \mu_0$.

Consider first the case in which $\mu_0 = \overline{\mu}_1$. Note that $\mu' v_H + (1 - \mu')\hat{v}_L(\overline{\mu}_1) = \mu' v_H \geq \delta R^{(\tau^*, q^*)}(\mu', \overline{\mu}_1) = \delta \alpha(\mu') v_H$ strictly so when $\mu' > 0$. Thus, $q^*(\overline{\mu}_1, \mu') = 1$, except possibly at $\mu = 0$. Moreover, the payoff of setting $q^*(\overline{\mu}_1, 0) = 0$ and $q^*(\overline{\mu}_1, 1) = 1$ (and splitting beliefs in this way) equals $\overline{\mu}_1 v_H$ and is exactly the same as the payoff of setting $q^*(\overline{\mu}_1, \cdot) = 1$ everywhere. Thus, $\tau^*(\overline{\mu}_1, 0) = 1 - \overline{\mu}_1 = 1 - \tau^*(\overline{\mu}_1, 1)$ is optimal.

Consider next $\mu_0 \in D_n$ for $n \geq 1$, if $n = 1$, then $\mu_0 > \overline{\mu}_1$. Proposition 4 implies that the seller must choose two posteriors, 1 and set $q = 1$, and $\mu_D < \mu_0$ at which he sets $q = 0$. Thus, we only need to rule out policies in which the seller delays trade at posteriors $\mu' < \overline{\mu}_{n-1}$. Toward a contradiction, let $\mu' \in (\overline{\mu}_m, \overline{\mu}_{m+1})$ for $m \leq n - 2$. The payoff to the seller in this case is $\delta R^{(\tau^*, q^*)}(\mu', \mu_0)$. Note that there exists $\alpha \in (0, 1)$ such that $\mu' = \alpha \overline{\mu}_m + (1 - \alpha)\overline{\mu}_{m+1}$. Moreover, because $R^{(\tau^*, q^*)}(\mu', \mu_0)$ is linear in μ' on $(\overline{\mu}_m, \overline{\mu}_{m+1})$, it follows that

$$R^{(\tau^*, q^*)}(\mu', \mu_0) = \alpha R^{(\tau^*, q^*)}(\overline{\mu}_m, \mu_0) + (1 - \alpha) \lim_{\tilde{\mu} \to \mu_{m+1}^{\star-}} R^{(\tau^*, q^*)}(\tilde{\mu}, \mu_0)$$

$$< \alpha R^{(\tau^*, q^*)}(\overline{\mu}_m, \mu_0) + (1 - \alpha) R^{(\tau^*, q^*)}(\overline{\mu}_{m+1}, \mu_0),$$

where the inequality follows because $R^{(\tau^*, q^*)}(\cdot, \mu_0)$ jumps up at $\mu' = \overline{\mu}_{m+1}$ (see Lemma 3 in Doval and Skreta (2019)).

Consider now the case $\mu' = \overline{\mu}_m$. Modify the policy so that the weight τ placed on $\overline{\mu}_m$ is placed instead on $\{\overline{\mu}_{m-1}, 1\}$, which is feasible because

$$\overline{\mu}_m = \frac{\overline{\mu}_m - \overline{\mu}_{m-1}}{1 - \overline{\mu}_{m-1}} \times 1 + \frac{1 - \overline{\mu}_m}{1 - \overline{\mu}_{m-1}} \times \overline{\mu}_{m-1}.$$

The change in the seller's payoff is proportional to

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

54

where the inequality follows from Lemma 2 in Doval and Skreta (2019). This inequality contradicts the optimality of the policy and, hence, it is not possible that the seller puts weight on $\overline{\mu}_m, m \leq n - 2$ when $\mu_0 \in [\overline{\mu}_n, \overline{\mu}_{n+1})$.

D Collected objects from intrapersonal equilibrium

In this section, we construct the following objects using the ingredients from the intrapersonal equilibrium:

1. The mapping $\gamma^* : \Delta(V) \mapsto \mathcal{M}$ that associates to each belief the seller may hold, μ_0, a mechanism $\gamma^*(\mu_0) \in \mathcal{M}$,

2. The mapping $u_H^* : \Delta(V) \mapsto \mathbb{R}$ that associates to each belief the seller may hold, μ_0, the buyer's payoff when her valuation is v_H and the seller employs the mechanism derived from the intrapersonal equilibrium,

3. The correspondence $\mathcal{U}_H^* : \Delta(V) \rightrightarrows \mathbb{R}$ that associates to each belief the seller may hold, μ_0, the set of feasible buyer's payoffs when her valuation is v_H, the seller employs the mechanism derived from the intrapersonal equilibrium, but does not necessarily break ties as in the intrapersonal equilibrium. That is, $\mathcal{U}_H^*(\mu_0) = u_H^*(\mu_0)$ whenever $\mu_0 \neq \overline{\mu}_i, i = 0, 1, \ldots$.

D.1 Mechanism and buyer's payoffs in the intrapersonal equilibrium

In this section, we construct the transfers and the continuation rents for the buyer, when her type is v_H, implied by the intrapersonal equilibrium. (The communication device and the probability of trade are the ones specified in Section 3.4.) We proceed "backwards," starting from $\mu_0 < \overline{\mu}_1$ and then $\mu_0 \in D_n$ for $n \geq 1$.

Let $\mu_0 < \overline{\mu}_1$. Using the binding participation constraint for the buyer when her type is v_L, we have

$$x_{\mu_0}^*(\mu_0) = v_L.$$

Note that the utility of the buyer of type v_H when the seller has belief $\mu_0 < \overline{\mu}_1$ is $u_H^*(\mu_0) = \Delta v$.

Fix $n \geq 1$. Consider now $\mu_0 \in [\overline{\mu}_n, \overline{\mu}_{n+1})$. Using the binding participation constraint for the buyer when her valuation is v_L, we have

$$x^*_{\mu_0}(\overline{\mu}_{n-1}) = 0.$$

To determine $x^*_{\mu_0}(1)$, we use the incentive compatibility constraint for the high type:

$$\beta^*_{\mu_0}(1|v_H)(v_H - x^*_{\mu_0}(1)) + \beta^*_{\mu_0}(\overline{\mu}_{n-1}|v_H)\delta u^*_H(\overline{\mu}_{n-1}) = \delta u^*_H(\overline{\mu}_{n-1})$$
$$x^*_{\mu_0}(1) = v_H - \delta u^*_H(\overline{\mu}_{n-1}).$$

We then construct $x^*_{\mu_0}(1)$ recursively. Set $n = 1$. Then, for $\mu \in [\overline{\mu}_1, \overline{\mu}_2)$,

$$x^*_{\mu_0}(1) = v_H - \delta\Delta v = v_L + (1-\delta)\Delta v$$
$$u^*_H(\mu_0) = \delta\Delta v,$$

and for $n \geq 2$, we obtain that for $\mu_0 \in [\overline{\mu}_n, \overline{\mu}_{n+1})$,

$$x^*_{\mu_0}(1) = v_H - \delta^n\Delta v = v_L + (1-\delta^n)\Delta v$$
$$u^*_H(\mu_0) = \delta^n\Delta v.$$

Note the construction of the transfers highlights that if the seller posts a price of $x^*_{\mu_0}(1)$ when his belief is μ_0, the buyer, when her type is v_H, is indifferent between accepting this price and rejecting it; whereas the low type is always better off rejecting this price because it lies above v_L.

Summing up, the intrapersonal equilibrium determines a map $\gamma^* : \Delta(V) \mapsto \mathcal{M}$ such that

$$\gamma^*(\mu_0) = \langle(V, \beta^*_{\mu_0}, \Delta(V)), (q^*_{\mu_0}, x^*_{\mu_0})\rangle. \tag{27}$$

The buyer's payoffs when her valuation is v_H are given by:

$$u^*_H(\mu_0) = \delta^i\Delta v, \mu_0 \in D_i, i = 0, \ldots.$$

Let $\mathcal{U}^*_H$ denote the following correspondence:

$$\mathcal{U}^*_H(\mu_0) = \begin{cases} u^*_H(\mu_0) & \text{if } \mu_0 \neq \overline{\mu}_i, i \geq 1 \\ [\delta^i\Delta v, \delta^{i-1}\Delta v] & \text{if } \mu_0 = \overline{\mu}_i \end{cases}. \tag{28}$$

For future use, note $\mathcal{U}^*_H$ is upper-hemicontinuous, convex-valued, and compact-valued.

E Proofs of Section 3.4

Appendix E is organized in three parts. Section E.1 constructs the buyer's strategy profile so as to specify her best responses after every history in the game. Section E.2 performs the same exercise for the seller, using the strategy from the intrapersonal equilibrium (recall Equation 27). Finally, in Section E.3, we show that if continuation values are specified as in the intrapersonal equilibrium, and the buyer's strategy is completed as in the first step, the seller has no one-shot deviations from the equilibrium strategy. We also argue the buyer has no one-shot deviations. Section I in Doval and Skreta (2019) lays out the formalisms that justify the use of self-generation arguments in our setting; this, in turn, implies that guaranteeing the absence of one-shot deviations given the continuation values is enough to conclude we have indeed constructed a PBE of $G^\infty(\mu_0)$, which gives the seller the same payoff as in the intrapersonal equilibrium.

E.1 Completing the buyer's strategy

Fix a public history h^t and let μ_0 denote the seller's beliefs at that public history.[37] To complete the buyer's strategy, we classify mechanisms, $\mathbf{M}$, in four categories:[38]

0. Mechanisms that given the intrapersonal equilibrium continuation values satisfy participation and truthtelling, that is, $\mathbf{M}$ such that the following hold:

$$\sum_{\mu' \in \Delta(V)} \beta^{\mathbf{M}}(\mu'|v_H)[v_H q^{\mathbf{M}}(\mu') - x^{\mathbf{M}}(\mu') + \delta(1 - q^{\mathbf{M}}(\mu')u_H^*(\mu')]$$

$$\geq \max\{0, \sum_{\mu' \in \Delta(V)} \beta^{\mathbf{M}}(\mu'|v_L)[v_H q^{\mathbf{M}}(\mu') - x^{\mathbf{M}}(\mu') + \delta(1 - q^{\mathbf{M}}(\mu')u_H^*(\mu')]\},$$

and

$$\sum_{\mu' \in \Delta(V)} \beta^{\mathbf{M}}(\mu'|v_L)[v_L q^{\mathbf{M}}(\mu') - x^{\mathbf{M}}(\mu')] \geq \max\{0, \sum_{\mu' \in \Delta(V)} \beta^{\mathbf{M}}(\mu'|v_H)[v_L q^{\mathbf{M}}(\mu') - x^{\mathbf{M}}(\mu')]\}.$$

Denote the set of these mechanisms $\mathcal{M}^0$.

---

[37]This prior will, of course, be an equilibrium object, but we suppress this from the notation to keep things simple.

[38]Gerardi and Maestri (2018) use a similar trick to complete the worker's strategy in their paper.

1. Mechanisms not in $\mathcal{M}^0$ such that the buyer, when her type is v_H, has no re-porting strategy, $\rho \in \Delta(V)$, such that $\sum_{v'} \beta^{\mathbf{M}}(\mu'|v')\rho(v')[v_H q^{\mathbf{M}}(\mu') - x^{\mathbf{M}}(\mu')] \geq 0$; let $\mathcal{M}^1$ denote the set of these mechanisms.

2. Mechanisms not in $\mathcal{M}^0$ such that the buyer, when her type is v_H, has a re-porting strategy, $\rho \in \Delta(V)$, such that $\sum_{v'} \beta^{\mathbf{M}}(\mu'|v')\rho(v')[v_H q^{\mathbf{M}}(\mu') - x^{\mathbf{M}}(\mu')] \geq 0$, but not when her type is v_L; let $\mathcal{M}^2$ denote the set of these mechanisms.

3. Mechanisms not in $\mathcal{M}^0$ such that both types have such a reporting strategy; let $\mathcal{M}^3$ denote the set of these mechanisms.

If $\mathbf{M} \in \mathcal{M}^1$, specify that the buyer rejects the mechanism for both types. 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 type. For each type v, let $r_v^*(\mathbf{M}, 1)$ denote an element of[39]

$$\arg\max_{\rho \in \Delta(V)} \sum_{v' \in V} \rho(v') \sum_{\mu' \in \Delta(V)} \beta^{\mathbf{M}}(\mu'|v')(v q^{\mathbf{M}}(\mu') - x^{\mathbf{M}}(\mu')).$$

If $\mathbf{M} \in \mathcal{M}^2$, specify that the buyer rejects when her type is v_L. Hence, without loss of generality, we can specify that if the seller observes that the buyer accepts the mechanism, the buyer's type is v_H. For type v_H, let $r_{v_H}^*(\mathbf{M}, 1)$ satisfy that if $r_{v_H}^*(\mathbf{M}, 1)(v) > 0$, then

$$\sum_{\mu' \in \Delta(V)} \beta^{\mathbf{M}}(\mu'|v)[v_H q^{\mathbf{M}}(\mu') - x^{\mathbf{M}}(\mu')] \geq \sum_{\mu' \in \Delta(V)} \beta^{\mathbf{M}}(\mu'|v')[v_H q^{\mathbf{M}}(\mu') - x^{\mathbf{M}}(\mu')]$$

for all $v' \in V$. Note we are using that the buyer's continuation payoffs are 0 conditional on her accepting the mechanism.

Then, v_H's payoff from participating in the mechanism $\mathbf{M}$ is given by

$$U_{1,v_H}(\mathbf{M}) = \sum_{v' \in V} r_{v_H}^*(\mathbf{M}, 1)(v') \sum_{\mu' \in \Delta(V)} \beta^{\mathbf{M}}(\mu'|v')[v_H q^{\mathbf{M}}(\mu') - x^{\mathbf{M}}(\mu') + \delta(1 - q^{\mathbf{M}}(\mu')) \times 0],$$

whereas the payoff from rejecting is

$$U_{0,v_H}(\pi_{v_H}, f) = \delta f(v_2(\mu_0, \pi_{v_H})).$$

---

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

58

where

$$v_2(\mu_0, \pi_{v_H}) = \frac{\mu_0(1 - \pi_{v_H})}{\mu_0(1 - \pi_{v_H}) + 1 - \mu_0}.$$

is the seller's belief that the buyer is of type v_H when observing a rejection, according to Bayes' rule, and $f(v_2(\pi_{v_H}))$ is a measurable selection from $\mathcal{U}_H^*(v_2(\pi_{v_H}))$. Note the payoff from rejecting is specified under the assumption that in the continuation, the equilibrium path coincides with that of the intrapersonal equilibrium when beliefs are $v_2(\pi_{v_H})$. We use, however, a selection from $\mathcal{U}_H^*$ to ensure that if needed, the seller randomizes between the posted prices when indifferent to help make the buyer's continuation problem well-behaved.

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

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

The main result in Simon and Zame (1990) implies a solution to the above problem exists, given the properties of $\mathcal{U}_H^*$ and the linearity in p of the objective. Let $\pi_{v_H}^*(\mathbf{M})$ denote this fixed point. Now, if $\pi_{v_H}^*(\mathbf{M}) < 1$ and $v_2(\mu_0, \pi_{v_H}^*(\mathbf{M})) = \overline{\mu}_i$ for some $i \geq 1$, then there exists a weight $\phi_2(\overline{\mu}_i, \mathbf{M}) \in [0, 1]$ that solves the following:

$$f_2^*(\mathbf{M}, \overline{\mu}_i) = (1 - \phi)\delta^{i-1}\Delta v + \phi\delta^i\Delta v. \tag{29}$$

This weight captures the probability with which the seller, when his prior is $\overline{\mu}_i$, mixes between $\gamma^*(\overline{\mu}_i)$ and $\gamma^*(\overline{\mu}_{i-1})$.

For a mechanism in $\mathcal{M}^2$, let $r_{v_L}^*(\mathbf{M}, 1)$ denote an element in

$$\arg\max_{\rho \in \Delta(V)} \sum_{v' \in V} \rho(v') \sum_{\mu' \in \Delta(V)} \beta^{\mathbf{M}}(\mu'|v')(v_L q^{\mathbf{M}}(\mu') - x^{\mathbf{M}}(\mu')).$$

Finally, if $\mathbf{M} \in \mathcal{M}^3$, specify that the buyer accepts for both types. If the seller observes that the buyer rejects the mechanism, he assigns probability 1 to the buyer's valuation being v_H. Thus, upon rejection, continuation payoffs are 0 for both types. Note that mechanisms in $\mathcal{M}^3$ satisfy the participation constraint given the continuation values of the intrapersonal equilibrium. Now, let m_L^* satisfy[40]

$$\sum_{\mu' \in \Delta(V)} \beta^{\mathbf{M}}(\mu'|m_L^*)(v_L q^{\mathbf{M}}(\mu') - x^{\mathbf{M}}(\mu')) \geq \sum_{\mu' \in \Delta(V)} \beta^{\mathbf{M}}(\mu'|v)(v_L q^{\mathbf{M}}(\mu') - x^{\mathbf{M}}(\mu')),$$

---

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

for all $v \in V$. Set $r_{v_L}^*(\mathbf{M}, 1)(m_L^*) = 1$. Let $\{m_H^*\} = V \setminus \{m_L^*\}$ and define

$$\Delta(V)^H = \{\mu' : \beta^{\mathbf{M}}(\mu'|m_H^*) > 0\}$$
$$\Delta(V)^L = \{\mu' : \beta^{\mathbf{M}}(\mu'|m_L^*) > 0\}.$$

Let μ_0 denote the seller's belief about the buyer's valuation being v_H. Let $r \in [0, 1]$ denote the weight the buyer assigns to m_L^* when her valuation is v_H:

$$\nu_3(\mu_0, \mu', r, \delta_{m_L^*}) = \begin{cases} 1 & \text{if } \mu' \in \Delta(V)^H \setminus \Delta(V)^L \\ \frac{\mu_0(r\beta^{\mathbf{M}}(\mu'|m_L^*)+(1-r)\beta^{\mathbf{M}}(\mu'|m_H^*))}{\mu_0(r\beta^{\mathbf{M}}(\mu'|m_L^*)+(1-r)\beta^{\mathbf{M}}(\mu'|m_H^*))+(1-\mu_0)\beta^{\mathbf{M}}(\mu'|m_L^*)} & \text{if } \mu' \in \Delta(V)^H \cap \Delta(V)^L \\ \frac{\mu_0 r}{\mu_0 r+(1-\mu_0)} & \text{if } \mu' \in \Delta(V)^L \setminus \Delta(V)^H \end{cases},$$

where we assume that if the seller observes an output message that is not consistent with m_L^*, he believes it was generated by the high-valuation buyer. This specification of beliefs does not conflict with Bayes' rule where possible: either m_H^* has positive probability in the optimal reporting strategy of the buyer when her valuation is v_H, in which case, ν_3 would be consistent with Bayes' rule, or it does not, in which case, Bayes' rule where possible places no restrictions on $\nu_3(\mu_0, \mu', \cdot)$ for $\mu' \in \Delta(V)^H \setminus \Delta(V)^L$.

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

$$U_{1,v_H}(r, m, f) = \sum_{\mu' \in \Delta(V)} \beta^{\mathbf{M}}(\mu'|m)(v_H q^{\mathbf{M}}(\mu') - x^{\mathbf{M}}(\mu') + \delta(1 - q^{\mathbf{M}}(\mu'))f(\nu_3(\mu_0, \mu', r, \delta_{m_L^*}))),$$

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

$$r_{v_H}^* \in \arg\max_{r \in [0,1]} rU_{1,v_H}(r_{v_H}^*, m_L^*, f_3^*(\mathbf{M})) + (1 - r)U_{1,v_H}(r_{v_H}^*, m_H^*, f_3^*(\mathbf{M})). \tag{30}$$

Again, we appeal to the result in Simon and Zame (1990) to argue that the properties of $\mathcal{U}_H^*$ and the objective function above imply the existence of such an $(r^*, f_3^*(\mathbf{M}))$. Set $r_{v_H}^*(\mathbf{M}, 1)(m_L^*) = r_{v_H}^*$.

As we did before, whenever $\nu_3(\mu_0, \mu', r_{v_H}^*, \delta_{m_L^*}) = \bar{\mu}_i$ for $i \geq 1$, we can define $\phi_3(\mathbf{M}, \bar{\mu}_i)$ as the weight on $\gamma^*(\bar{\mu}_i)$ implied by $f_3^*(\mathbf{M}, \bar{\mu}_i)$.

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

$$\pi_v^*(h_B^t, \mathbf{M}) = \begin{cases} 1 & \text{if } \mathbf{M} \in \mathcal{M}^0 \cup \mathcal{M}^3 \\ \pi_{v_H}^*(\mathbf{M}) & \text{if } \mathbf{M} \in \mathcal{M}^2 \text{ and } v = v_H \\ 0 & \text{otherwise} \end{cases} \tag{31}$$

and

$$r_v^*(h_B^t, \mathbf{M}, 1) = \begin{cases} \delta_v & \text{if } \mathbf{M} \in \mathcal{M}^0 \\ r_v^*(\mathbf{M}, 1) & \text{otherwise} \end{cases}. \tag{32}$$

E.2 Full specification of the PBE assessment

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

From the beginning of period t of the game until the end of that period, the following things are determined: (i) the seller's choice of mechanism, $\mathbf{M}$; (ii) the buyer's participation decision, $p \in (0,1)$; and (iii) if the buyer participates of the mechanism, the allocation and the output message. We let $z_\emptyset(\mathbf{M}) = (\mathbf{M}, 0, \emptyset, 0, 0)$ to denote the outcome when mechanism $\mathbf{M}$ is chosen, the buyer does not participate, no output message is produced, and no trade and no transfers occur. We let $z_q(\mathbf{M}, \mu') = (\mathbf{M}, 1, \mu', q, t)$ denote the outcome at the end of period t when mechanism $\mathbf{M}$ is chosen, the buyer participates (1), the output message is μ', and the allocation is (q, t). Of course, if $q = 1$, the game ends. Note that any public history h^t can be written as (h^{t-1}, z) for some z as defined above. Given h^t, let $z(h^t)$ denote the corresponding outcome z. Given any prior μ_0, define

$$T(\mu_0, z) = \begin{cases} \mu' & \text{if } z = z_q(\mathbf{M}, \mu') \text{ and } \mathbf{M} \in \mathcal{M}^0 \\ & \quad \text{if } z = z_\emptyset(\mathbf{M}) \text{ and } \mathbf{M} \in \mathcal{M}^0 \cup \mathcal{M}^3 \\ 1 & \qquad\qquad \text{or} \\ & \quad z = z_q(\mathbf{M}, \mu') \text{ and } \mathbf{M} \in \mathcal{M}^1 \cup \mathcal{M}^2 \\ \mu_0 & \text{if } z = z_\emptyset(\mathbf{M}) \text{ and } \mathbf{M} \in \mathcal{M}^1 \\ \nu_2(\mu_0, \pi_{v_H}^*(\mathbf{M})) & \text{if } z = z_q(\mathbf{M}, \mu') \text{ and } \mathbf{M} \in \mathcal{M}^2 \\ \nu_3(\mu_0, \mu', r_{v_H}^*(\mathbf{M}, 1), r_{v_L}^*(\mathbf{M}, 1)) & \text{if } z = z_q(\mathbf{M}, \mu') \text{ and } \mathbf{M} \in \mathcal{M}^3 \end{cases} \tag{33}$$

Let μ_0 denote the seller's prior in $G^\infty(\mu_0)$. Define $\mu^*(\emptyset) = \mu_0$ and $\Gamma^*(\emptyset) = \gamma^*(\mu_0)$, where γ^* is the strategy in the intrapersonal equilibrium as defined in Equation 27. For any public history h^t, define

$$\mu^*(h^t) = T(\mu^*(h^{t-1}), z(h^t)). \tag{34}$$

61

If either (i) $z = z_q(\mathbf{M}, \cdot)$ and $\mathbf{M} \in \mathcal{M}^0 \cup \mathcal{M}^1 \cup \mathcal{M}_2$, (ii) $z = z_\emptyset(\mathbf{M})$, $\mu^*(h^t) \notin \{\overline{\mu}_i\}_{i \geq 1}$ and $\mathbf{M} \in \mathcal{M}^2$, (iii) $z = z_q(\mathbf{M}, \cdot)$, $\mu^*(h^t) \notin \{\overline{\mu}_i\}_{i \geq 1}$ and $\mathbf{M} \in \mathcal{M}^3$, or (iv) $z = z_\emptyset(\mathbf{M})$ and $\mathbf{M} \in \mathcal{M}^0 \cup \mathcal{M}^1 \cup \mathcal{M}_3$, define

$$\Gamma^*(h^t)(\mathbf{M}) = \mathbb{1}[\mathbf{M} = \gamma^*(T(\mu^*(h^{t-1}), z(h^t)))]. \tag{35}$$

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

$$\Gamma^*(h^t)(\mathbf{M}) = \begin{cases} \phi_2(\mathbf{M}, \overline{\mu}_i) & \mathbf{M} = \gamma^*(\overline{\mu}_i) \\ 1 - \phi_2(\mathbf{M}, \overline{\mu}_i) & \mathbf{M} = \gamma^*(\overline{\mu}_{i-1}) \\ 0 & \text{otherwise} \end{cases}. \tag{36}$$

Finally, if $z(h^t) = z_q(\mathbf{M}, \cdot)$ for $\mathbf{M} \in \mathcal{M}^3$ and $\mu^*(h^t) = \overline{\mu}_i, i \geq 1$, define

$$\Gamma^*(h^t)(\mathbf{M}) = \begin{cases} \phi_3(\mathbf{M}, \overline{\mu}_i) & \mathbf{M} = \gamma^*(\overline{\mu}_i) \\ 1 - \phi_3(\mathbf{M}, \overline{\mu}_i) & \mathbf{M} = \gamma^*(\overline{\mu}_{i-1}) \\ 0 & \text{otherwise} \end{cases}. \tag{37}$$

Equations 31-32, together with equations 34- 37, define a PBE assessment where the system of beliefs is derived from the strategy profile where possible.

E.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 for the seller and the buyer constructed using the intrapersonal equilibrium. This implies that the payoffs in the intrapersonal equilibrium together with the belief operator defined in Equation 33 belong in the self-generating set. Section I in Doval and Skreta (2019) lays out the formalisms that justify the use of techniques á la Abreu et al. (1990) for the game under consideration.

Let $\mu^*(h^t)$ denote the seller's belief at history h^t that the buyer's type is v_H. We now show the seller cannot achieve a payoff higher than $R^{(\tau^*, q^*)}(\mu^*(h^t), \mu^*(h^t))$. Note that among mechanisms of type 0, the one that corresponds to the intrapersonal equilibrium is, by definition, the best that the seller can do. Hence, to show that there are no deviations, we need to show the seller cannot benefit from offering mechanisms of types 1-3.

Given the buyer's strategy, the payoff from offering a mechanism in $\mathcal{M}^1$ is $\delta R^{(\tau^*,q^*)}(\mu^*(h^t), \mu^*(h^t)) \leq R^{(\tau^*,q^*)}(\mu^*(h^t), \mu^*(h^t))$, and hence this deviation is not profitable. Let $\mathbf{M}$ denote a mechanism in $\mathcal{M}^2$ and let $\pi^*_{v_H}(h^t_B, \mathbf{M}), r^*_{v_H}(h^t_B, \mathbf{M}, 1)$ denote the buyer's best responses as constructed in Section E.1. Denote by $v_2 \equiv v_2(\mu^*(h^t), \pi^*_{v_H}(h^t_B, \mathbf{M}))$ the seller's belief that he is facing a buyer with valuation v_H when he observes non-participation. The seller's payoff is then

$$\mu^*(h^t)\pi^*_{v_H}(h^t_B, \mathbf{M}) \sum_{\mu' \in \Delta(V)} \left(\sum_{v \in V} \beta^{\mathbf{M}}(\mu'|v) r^*_{v_H}(h^t_B, \mathbf{M}, 1)(v) \right) [x^{\mathbf{M}}(\mu') + \delta(1 - q^{\mathbf{M}}(\mu'))v_H]$$

$$+ (1 - \mu^*(h^t)\pi^*_{v_H}(h^t_B, \mathbf{M}))R^{(\tau^*,q^*)}(v_2, v_2)), \tag{38}$$

where the continuation values after rejection are constructed using the policy from the intrapersonal equilibrium once the seller has posterior $v_2(\mu^*(h^t), \pi^*_{v_H}(h^t_B, \mathbf{M}))$.[41] Now, consider the following alternative mechanism, $\mathbf{M}'$:

$$\beta^{\mathbf{M}'}(1|v_H) = \pi^*_{v_H}(\mathbf{M}) \sum_{v \in V, \mu' \in \Delta(V)} \beta^{\mathbf{M}}(\mu'|v) r^*_{v_H}(\mathbf{M})(v)$$

$$\beta^{\mathbf{M}'}(v_2|v_H) = (1 - \pi^*_{v_H}(\mathbf{M}))$$

$$\beta^{\mathbf{M}'}(v_2|v_L) = 1$$

$$q^{\mathbf{M}'}(v_2) = x^{\mathbf{M}'}(v_2) = 0$$

$$x^{\mathbf{M}'}(1) = \sum_{v \in V} \sum_{\mu' \in \Delta(V)} \beta^{\mathbf{M}}(\mu'|v) r^*_{v_H}(\mathbf{M})(v) x^{\mathbf{M}}(\mu')$$

$$q^{\mathbf{M}'}(1) = \sum_{v \in V} \sum_{\mu' \in \Delta(V)} \beta^{\mathbf{M}}(\mu'|v) r^*_{v_H}(\mathbf{M})(v) q^{\mathbf{M}}(\mu')$$

Note that if the buyer participates and truthfully reports her type, $\mathbf{M}'$ gives the seller a payoff equal to the expression in Equation 38. Moreover, participating and truthfully reporting her type is a best response for the buyer. Then, we can write the payoff to mechanism, $\mathbf{M}'$, as

$$\sum_{\mu' \in \{v_2, 1\}} \tau^{\mathbf{M}'}(\mu^*(h^t), \mu')[x^{\mathbf{M}'}(\mu') + \delta(1 - q^{\mathbf{M}'}(\mu'))R^{(\tau^*,q^*)}(\mu', \mu')],$$

where, as in the main text,

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

---

[41]To keep notation simple, we do not use that the seller in the continuation may randomize in his choice.

Now, truthtelling implies $x^{\mathbf{M}'}(1) \leq v_H q^{\mathbf{M}'}(1) - \delta u_H^*(\mu'(\pi_{v_H}^*(\mathbf{M})))$, so that

$$\sum_{\mu' \in \{v_2, 1\}} \tau^{\mathbf{M}'}(\mu^*(h^t), \mu')[x^{\mathbf{M}'}(\mu') + \delta(1 - q^{\mathbf{M}'}(\mu'))R^{(\tau^*, q^*)}(\mu', \mu')]$$

$$\leq \sum_{\mu' \in \{v_2, 1\}} \tau^{\mathbf{M}'}(\mu^*(h^t), \mu') \left\{ q^{\mathbf{M}'}(\mu')(\mu' v_H + (1 - \mu')\hat{v}_L(\mu^*(h^t)) \right.$$

$$\left. + \delta(1 - q^{\mathbf{M}'}(\mu')) \left[R^{(\tau^*, q^*)}(\mu', \mu') + (1 - \mu')(\frac{\mu'}{1 - \mu'} - \frac{\mu^*(h^t)}{1 - \mu^*(h^t)}) u_H^*(\mu')\right] \right\},$$

which, by definition, is weakly lower than $R^{(\tau^*, q^*)}(\mu^*(h^t), \mu^*(h^t))$.

Similar steps imply the seller does not have a one-shot deviation to a mechanism of type 3.

Finally, note the construction of π_v^*, r_v^* in Section E.1 (Equations 31 and 32) ensures the buyer is best responding after every history given the continuation values.