Regret in durable-good monopoly

Rumen Kostadinov*
Department of Economics, McMaster University

November, 2025

Abstract

I study a dynamic model of durable-good monopoly where the seller minimises lifetime regret against the worst-case type of buyer. The optimal mechanism is time-consistent: at no point can the seller benefit from replacing it with another mechanism. Despite this, the optimal mechanism cannot be supported in an equilibrium without commitment. This is because the seller's regret is endogenously determined by the best counterfactual payoffs he can obtain against every type, and these payoffs vary with his commitment power. When the seller lacks commitment the good may not be sold to all types. However, in the limit as offers become frequent the good is sold immediately at a price equal to the lowest buyer value.

1 Introduction

How should a monopolist sell a single unit of a good when faced with uncertainty about the buyer's value? The classic answer is that the seller maximises expected profits given a prior distribution over the buyer's values. However, this probabilistic information may be unreliable, or missing altogether. The literature has responded to these concerns by developing robust approaches such as profit maximisation given

*I thank Sam Kapon for persistent encouragement and helpful discussions. I also received valuable feedback from Laura Doval and audiences at Waterloo, CEA, Stony Brook, and Midwest Theory. The paper benefited greatly from comments by three anonymous referees. I gratefully acknowledge financial support from the SSHRC Insight Grant 435-2024-0906.

a worst-case distribution of values (Carrasco et al., 2018, 2019) and regret minimisation given a worst-case distribution of values (Bergemann and Schlag, 2008, 2011).[1] Extensions of these approaches to dynamic models are scarce and tend to favour worst-case profit maximisation (Libgober and Mu, 2021; **?**). This paper, instead, explores regret minimisation in a dynamic setting where the seller offers prices to a buyer over an infinite horizon knowing only that the buyer's value lies in an interval $[\underline{v}, \bar{v}]$.

My approach can be motivated by the following reasoning. Suppose the seller learns the buyer's value v after selling the good. He then compares the profits he has made to the best profits he could have made if he had known the buyer's value was v at the start of the game. The resulting (hypothetical) improvement in profits is the seller's (lifetime) regret against type v of the buyer. Prior to the sale the seller anticipates the regret he will have against each possible type. Lacking probabilistic beliefs, the seller adopts the minmax regret decision rule introduced by Savage (1951): he minimises regret in the worst-case scenario, i.e. against the worst-case type.

I study two versions of the problem varying the seller's commitment power. When the seller has commitment, he chooses a mechanism to minimise the worst-case regret described above subject to the buyer's best response. The optimal mechanism exhibits regret equalisation across time: the seller's regret against the worst-case type who buys at time t equals regret against the worst-case type who buys at time $t + 1$, and also equals regret against the worst-case type who does not buy the good (provided such types exist). Regret equalisation implies that the optimal mechanism is time-consistent: at no point can the seller benefit from replacing it with a different mechanism. This feature is particular to the regret objective and stands in contrast with many standard dynamic settings, where optimal mechanisms induce suboptimal future outcomes in order to improve incentives at earlier stages.

To study the case where the seller lacks commitment, I define a solution concept called Perfect Regret Equilibrium (PRE). It mirrors Perfect Bayesian Equilibrium (PBE) except for the seller's objective. In a PRE the seller uses the history of prices and the buyer's equilibrium strategy to infer which types have purchased the good, akin to updating beliefs in a PBE. He then minimises worst-case regret against the

[1] Many extensions to multiple goods and/or multiple buyers have been studied. See Carroll (2017); Du (2018); Brooks and Du (2021); Che (2022); He and Li (2022); Suzdaltsev (2022a,c); Zhang (2023) for worst-case profit maximisation and Koçyiğit, Rujeerapaiboon, and Kuhn (2022); Zhang (2022); Koçyiğit, Kuhn, and Rujeerapaiboon (2024) for worst-case regret minimisation.

remaining types. Regret against a given type v is backward-looking: at any history from time t, he compares his lifetime equilibrium profits against v (calculated from time 0) to the best lifetime profits he could achieve against v by changing his strategy from time 0. Hence, the seller's strategy has less impact on regret as the game goes on, leading to nonstationary equilibrium behaviour. The advantage of this formulation of regret is that the seller's preferences are dynamically consistent, which roughly means that an optimal strategy remains optimal at subsequent histories holding the buyer's strategy fixed (Proposition 1). This is a desirable but elusive feature of dynamic models with nonstandard preferences (Carroll, 2019). When preferences lack dynamic consistency, there may not exist an equilibrium where strategies are immune to all possible deviations. For instance, Schlag and Zapechelnyuk (2024) consider another regret-based solution concept for dynamic games where equilibrium strategies only need to be immune to one-shot deviations, even though the one-shot deviation principle does not apply (see Section 3.5 for a detailed discussion). PRE is more strict because it requires optimality against all deviations. Yet, the one-shot deviation principle applies similarly to standard dynamic games (Proposition 2).

The main result of the paper characterises a PRE in the "gap case" $\underline{v} > 0$ (Theorem 1). The equilibrium is unique up to buyer indifference. It shares a lot of structure with the commitment solution due to the time consistency of the latter. In both cases the seller's preferences depend endogenously on the best counterfactual payoffs against each type. The price path of a PRE matches the optimal mechanism for a seller whose preferences are determined by the PRE counterfactual payoffs. It differs from the optimal mechanism only because the PRE counterfactual payoffs differ from the counterfactual payoffs under commitment.

Equilibrium sales in a PRE are qualitatively different from their PBE counterparts. Since regret is measured against a time-0 benchmark, sales at late stages of the game are less relevant due to discounting. In the long run regret against each type approaches the best counterfactual payoff, which is increasing in the type. The seller, therefore, prioritises extracting surplus from the high types. Consequently, low types may never purchase the good. This contrasts the Bayesian model, where the good is sold to all types in finitely many periods (Fudenberg, Levine, and Tirole, 1985).

The differences with the Bayesian model vanish as offers become frequent. Gul, Sonnenschein, and Wilson (1986) showed that in the limiting PBE outcome the good sells immediately at a price of $\underline{v}$, confirming a conjecture of Coase (1972).[2] This is attributed to the time inconsistency of the seller's optimal mechanism. The Coase conjecture also holds in my model with regret, despite the time consistency of the seller's optimal mechanism (Theorem 2). The result is, instead, driven by the time inconsistency of the seller's optimal *counterfactual* mechanism against each type, which produces the best counterfactual payoffs under commitment. My result contributes to a large literature investigating the sensitivity of the Coase conjecture to different specifications of the model.[3]

The rest of the paper proceeds as follows. Section 2 illustrates the results of the paper in a simple setting with two periods. Section 3 presents the general model where time is infinite, defines the PRE solution concept, and discusses related literature. Section 4 describes the commitment outcome and Section 5 presents the PRE characterisation. Extensions are discussed in Section 6. Omitted proofs appear in the Appendix.

2 Two-period example

This section presents a two-period version of the model studied in the rest of the paper. A buyer wants to purchase a single unit of a good from a seller. The buyer's type is his value v, which is unknown to the seller. The seller only knows that $v \in [\underline{v}, \bar{v}]$, where $\bar{v} \geq \underline{v} \geq 0$. The seller offers prices for the good at each time $t = 0, 1$. The buyer maximises utility. He gets utility $\delta^t(v - p)$ from purchasing the good at time t for a price p, and he gets 0 otherwise. Assume that the buyer purchases the good when indifferent.

2.1 Commitment

Suppose the seller can commit to a mechanism consisting of prices p_0, p_1 posted at times $t = 0, 1$, respectively. Given a mechanism (p_0, p_1), let B_0, B_1 denote the sets of types who purchase the good at times $t = 0, 1$, respectively, and let B_∞ denote

[2] See also Sobel and Takahashi (1983); Fudenberg, Levine, and Tirole (1985); Ausubel and Deneckere (1989).

[3] See McAfee and Wiseman (2008); Kim (2009); Fuchs and Skrzypacz (2010); Board and Pycia (2014); Ortner (2017); Nava and Schiraldi (2019); Chang and Lee (2022); Groseclose (2022); Brzustowski, Georgiadis-Harris, and Szentes (2023); **?**); Doval and Skreta (2024).

the set of types who do not purchase the good. Let $B_0 = [v_1, \bar{v}]$, $B_1 = [v_2, v_1)$, and $B_\infty = [\underline{v}, v_2)$, as illustrated in Figure 1. The seller's payoff (profits) against type v is

$$u_S(v) = \begin{cases} p_0 & \text{if } v \in B_0 \\ \delta p_1 & \text{if } v \in B_1 \\ 0 & \text{if } v \in B_\infty \end{cases}.$$

If the seller wants to maximise profits against a particular type v, he can set prices $p_0 = p_1 = v$. Type v will then buy at $t = 0$ resulting in the payoff $C(v) = v$. This is called the best counterfactual payoff against v.

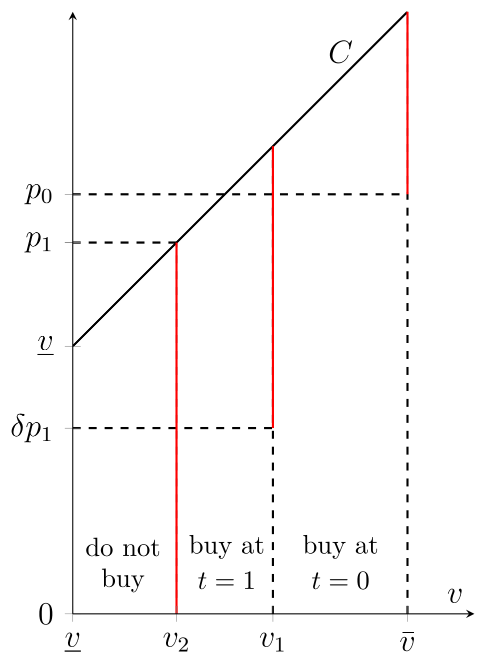

Figure 1: A mechanism

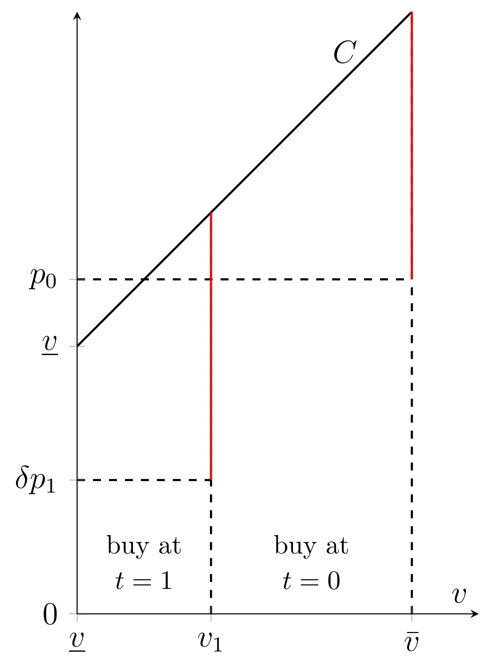

Figure 2: Optimal mechanism

The seller's regret $R(v) = C(v) - u_S(v)$ against type v is the difference between the seller's payoff against v and the best payoff he could obtain against v. In general, no mechanism can maximise profits against all types, so regret against some types will be positive. The seller wants to guarantee a level of regret irrespective of the buyer's type. His objective is to minimise regret against the worst-case type of buyer:

$$\min_{p_0, p_1} \sup_{v \in [\underline{v}, \bar{v}]} R(v). \tag{1}$$

Solution

Consider the seller's (worst-case) regret against (types in) B_0. The seller gets the same payoff against each of these types because all of them buy the good at $t = 0$ for a price of p_0. On the other hand, the best counterfactual payoff $C(v)$ is increasing in v. Hence, the seller's regret against B_0 equals his regret $R(\bar{v}) = C(\bar{v}) - p_0$ against the highest type in B_0, depicted on Figure 1 by a red vertical line at $\bar{v}$. Similarly, his regrets against B_1 and B_∞ equal $C(v_1) - \delta p_1$ and $C(v_2)$, respectively. They are depicted on Figure 1 by red vertical lines at v_1 and v_2. The seller's problem (1), therefore, amounts to minimising worst-case regret against the three groups of buyers B_0, B_1, B_∞, corresponding to the three red lines:

$$\min_{p_0,p_1} \max \left\{ \underbrace{C(\bar{v}) - p_0}_{\text{regret against } B_0}, \underbrace{(C(v_1) - \delta p_1)\mathbb{1}_{v_1 > \underline{v}}}_{\text{regret against } B_1}, \underbrace{C(v_2)\mathbb{1}_{v_2 > \underline{v}}}_{\text{regret against } B_\infty} \right\}, \tag{2}$$

where $\mathbb{1}$ is the indicator function. If $v_1 = \underline{v}$, then the good is sold to all types at $t = 0$, so $B_1 = \emptyset$ and the seller does not incur regret against types who buy at $t = 1$. Similarly, if $v_2 = \underline{v}$, then $B_\infty = \emptyset$ and the seller does not incur regret against types who do not purchase the good because there are no such types.

In what follows I show that the optimal mechanism equalises regret across all nonempty groups of buyer types. Fix a mechanism (p_0, p_1) and suppose $v_1 > \underline{v}$. The seller's problem (2) shows that p_1 does not affect regret against B_0. Hence, an optimal choice of p_1 minimises regret against the worst-case type in B_1 and B_∞:[4]

$$p_1 \in \operatorname*{argmin}_{p} \max \left\{ \underbrace{C(v_1) - \delta p}_{\text{regret against } B_1}, \underbrace{C(v_2)\mathbb{1}_{v_2 > \underline{v}}}_{\text{regret against } B_\infty} \right\}. \tag{3}$$

If $v_2 > \underline{v}$, then type v_2 is indifferent between accepting and rejecting p_1 at $t = 1$, so $p_1 = v_2$. Condition (3) then requires equalisation of regret against B_1 and B_∞. For example, if $C(v_1) - \delta p_1 < C(v_2)$, then the seller incurs higher regret against B_∞, as shown in Figure 1. A small reduction in p_1 would improve the objective in (3) by decreasing regret against B_∞, while ensuring it will not be overtaken by regret

[4] If regret against B_0 is sufficiently large, then multiple choices of p_1 may be optimal, some of which do not minimise regret against B_1 and B_∞. This does not arise in the optimal mechanism.

against B_1. Regret equalisation implies that $C(v_1) = C(v_2) + \delta v_2$. When this cannot be satisfied for any v_2, the good is sold to all types at $t = 1$ at price $p_1 = \underline{v}$, and $v_2 = \underline{v}$. Hence, in both cases $v_2 = f_1^C(v_1)$, where

$$f_1^C(v) = \sup\{v' \geq \underline{v} | C(v) \geq C(v') + \delta v'\} \tag{4}$$

with the convention $\sup \emptyset = \underline{v}$. Since regret against B_∞ equals regret against B_1 when $B_\infty \neq \emptyset$, problem (2) shows that the price in the initial period satisfies

$$p_0 \in \underset{p}{\operatorname{argmin}} \max \left\{ \underbrace{C(\bar{v}) - p}_{\text{regret against } B_0}, \underbrace{(C(v_1) - \delta p_1)\mathbb{1}_{v_1 > \underline{v}}}_{\text{regret against } B_1} \right\}. \tag{5}$$

The optimal price p_0 can now be derived similarly to p_1. If $v_1 > \underline{v}$, then type v_1 is indifferent between purchasing at $t = 0$ and $t = 1$, so $p_0 = (1 - \delta)v_1 + \delta p_1$. The optimal mechanism equalises regret against B_0 and B_1, i.e. $C(\bar{v}) - p_0 = C(v_1) - \delta p_1$. Substituting the indifference condition for v_1 results in $C(\bar{v}) = C(v_1) + (1 - \delta)v_1$. If the latter cannot be satisfied for any v_1, then $v_1 = \underline{v}$. Hence, $v_1 = f_0^C(\bar{v})$, where

$$f_0^C(v) = \sup\{v' \geq \underline{v} | C(v) \geq C(v') + (1 - \delta)v'\}. \tag{6}$$

Equations (4) and (6) describe the law of motion for the highest outstanding type in each period, referred to as *the state*. Substituting $C(v) = v$ gives $v_1 = \max\{\bar{v}/(2 - \delta), \underline{v}\}$ and $v_2 = \max\{v_1/(1 + \delta), \underline{v}\}$. The optimal prices are inferred from the types v_1, v_2 via the equations $p_0 = (1 - \delta)v_1 + \delta p_1$ and $p_1 = v_2$. For the choice of parameters in Figure 2 $v_1 = 6, v_2 = 4, p_0 = 5, p_1 = 4$.

Time consistency

Suppose the seller chooses the optimal mechanism (p_0^*, p_1^*). He sets price p_0^* at $t = 0$ and the buyer rejects it. Then at $t = 1$ he is given the option to abandon the mechanism by setting a price p_1. This is not anticipated by the buyer, who believes the seller is committed to (p_0^*, p_1^*). Can the seller benefit from choosing $p_1 \neq p_1^*$? The answer depends on his preferences at $t = 1$. I assume the seller minimises his worst-case regret against types who did not purchase the good at $t = 1$, given by $\sup_{v \in [\underline{v}, v_1)} R(v)$.[5] The interpretation is that the seller regrets the extra profits he could have made from the beginning of the game with knowledge of the buyer's type, as outlined in the introduction. It follows that p_1 satisfies (3) and, therefore, equals p_1^*. Hence, the optimal mechanism remains optimal from the perspective of $t = 1$. This property of the optimal mechanism is often termed *time consistency* following Kydland and Prescott (1977).

If the seller has an exogenously given counterfactual C that is strictly increasing and continuous, the resulting optimal mechanism, called C-optimal, is also time-consistent. This is crucial for the equilibrium analysis below.

2.2 No commitment

Now suppose the seller cannot commit to prices. A strategy for the seller is therefore a price p_0 set at $t = 0$, and a function $P_1(v)$ representing the price at $t = 1$ when the remaining types are believed to be $[\underline{v}, v)$. In equilibrium the seller follows the regret objectives at $t = 0$ and $t = 1$ outlined so far: His regret $R(v) = C(v) - \pi(v)$ against type v is defined analogously to the case of commitment, except that $C(v)$ is now the best payoff that can be obtained against v given the buyer's equilibrium strategy. At $t = 0$ the seller's equilibrium strategy minimises worst-case regret $\sup_{v \in [\underline{v}, \bar{v}]} R(v)$ against all types. At $t = 1$ the seller minimises worst-case regret $\sup_{v \in [\underline{v}, v')} R(v)$ against types $[\underline{v}, v')$ who have not purchased at $t = 0$ according to their equilibrium strategy given the price set at $t = 0$.

Let $[v_1, \bar{v}]$ be the types who buy at $t = 0$ and $[v_2, v_1)$ be the types who buy at $t = 1$ in equilibrium. Despite lack of commitment, the states v_1, v_2 follow the same laws of motion (4) and (6) as in the setting with commitment. This is a consequence of the time consistency of the C-optimal mechanism discussed above. At $t = 1$ the seller without commitment faces the same problem (3) as a seller with commitment who is given the option to abandon his previously chosen mechanism. Hence, the

[5] Alternatively, the seller could consider worst-case continuation regret in line with the solution concept of Schlag and Zapechelnyuk (2024) (see Section 3.5). In this case the optimal mechanism would not be time-consistent.

price p_1 set on the equilibrium path satisfies

$$p_1 \in \underset{p}{\operatorname{argmin}} \max \left\{ C(v_1) - \delta p, C(v_2) \mathbb{1}_{v_2 > \underline{v}} \right\}, \tag{7}$$

where $v_2 = p$ is the lowest type who buys the good, mirroring the analysis of (3). Thus, $p_1 = v_2 = f_1^C(v_1)$. The same analysis applies to any history at $t = 1$ where types $[\underline{v}, v)$ are present, resulting in an optimal price $P_1(v) = f_1^C(v)$. This completes the seller's strategy at $t = 1$. Similarly, the equilibrium price at $t = 0$ satisfies

$$p_0 \in \underset{p}{\operatorname{argmin}} \max \left\{ C(\bar{v}) - p, (C(v_1) - \delta P_1(v_1)) \mathbb{1}_{v_1 > \underline{v}} \right\}, \tag{8}$$

mirroring (5). Hence, $v_1 = f_0^C(\bar{v})$ and $p_0 = (1-\delta)v_1 + \delta p_1 = (1-\delta)v_1 + \delta v_2$, as in the case of commitment.

Despite following the same law of motion as the optimal mechanism, the equilibrium states v_1, v_2 differ from their counterparts under commitment. This is because the best counterfactual payoff the seller can obtain against type v in equilibrium (without commitment) is lower than the best commitment payoff. The equilibrium counterfactual C is characterised by a fixed point. Given C, the seller's strategy is determined as described above. This, in turn, determines the buyer's strategy as the best response to the pricing of the seller. Finally, the buyer's strategy determines the best payoffs the seller can obtain against each type, and they should match the counterfactual payoffs in C.

I now use the above steps to verify that the equilibrium counterfactual for the parameters $\underline{v} = 4, \bar{v} = 9, \delta = 0.5$ considered so far is given by[6]

$$C(v) = \begin{cases} \frac{v}{2} + 2 & \text{if } v < 8 \\ v - 2 & \text{if } v \geq 8 \end{cases}.$$

[6] In general, C can have a different number of linear segments.

It follows that the seller's strategy at $t = 1$ satisfies[7]

$$P_1(v) = f_1^C(v) = \begin{cases} 4 & \text{if } v < 8 \\ v - 4 & \text{if } v \geq 8 \end{cases}.$$

This can be used to determine equilibrium states and prices $v_1 = 5, v_2 = 4, p_0 = 4.5, p_1 = 4$.

Now consider the buyer's best response to the seller's strategy. At $t = 1$ type v accepts any price $p \leq v$. At $t = 0$ type v accepts a price p iff $v - p \geq \delta(v - P_1(v_1(p)))$, where $v_1(p)$ is the lowest type who accepts. It follows that v accepts iff $(1 - \delta)v \geq p - \delta P_1(v)$.

The best lifetime payoff the seller can obtain against type v given the buyer's strategy results from selling the good to v at $t = 0$ or $t = 1$. The best payoff from selling at $t = 1$ is δv. The best payoff from selling at $t = 0$ is $(1 - \delta)v + \delta P_1(v)$. Since $\delta = 0.5$, selling at $t = 0$ is more profitable. It can be shown that $C(v) = (1 - \delta)v + \delta P_1(v)$, completing the verification of the equilibrium counterfactual C.

Figure 3 summarises the solutions with and without commitment. Red vertical lines depict regret against types who purchase at $t = 0$ and $t = 1$ in the case of commitment, as in Figure 2. Blue vertical lines depict these regrets in the equilibrium without commitment. The function C^{com} of best counterfactual payoffs under commitment is steeper than the the function C^{eqm} of best counterfactual payoffs in equilibrium. It follows from the laws of motion (4) and (6) that commitment produces higher states and, therefore, results in slower sales. Indeed, at $t = 0$ types above 6 buy when there is commitment, while types above 5 buy in equilibrium. Both the commitment and equilibrium prices at $t = 0$ satisfy $p_0 = (1 - \delta)v_1 + \delta v_2$ for their respective states v_1, v_2. Hence, the equilibrium price of 4.5 is lower than the commitment price of 5.

[7] If $v < 8$, then $C(v) < 6 = C(\underline{v}) + \delta \underline{v}$, so $f_1^C(v) = \underline{v} = 4$. If $v \geq 8$, then $C(v) = v - 2 = \frac{v-4}{2} + 2 + \frac{v-4}{2} = C(v - 4) + \delta(v - 4)$, so $f_1^C(v) = v - 4$.

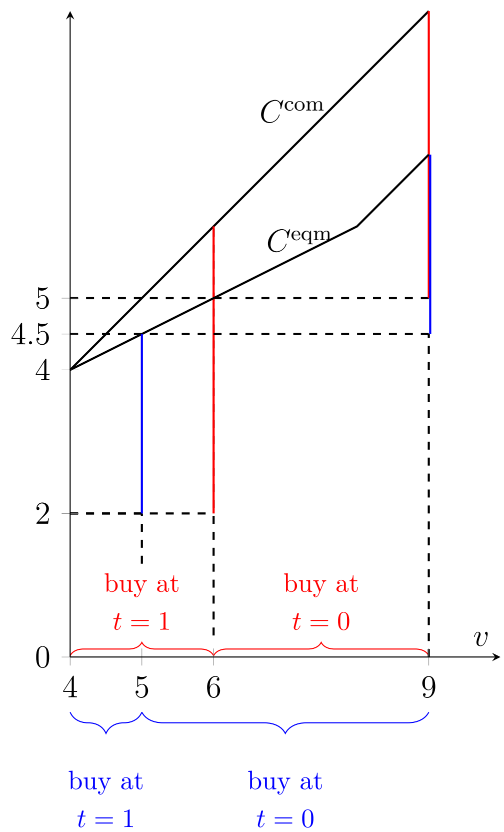

Figure 3: The solutions with and without commitment

3 Model

A monopolist seller has a single unit of a durable good. A buyer has value $v \in \mathcal{V} = [\underline{v}, \bar{v}]$ which is known to himself but not to the seller, where $\bar{v} \geq \underline{v} \geq 0$. In what follows I will refer to the buyer with value v as type v or simply v.

3.1 Timing and strategies

The game unfolds in discrete time $t = 0, 1, \ldots$. Each period the seller offers a price, then the buyer accepts or rejects it. The prices belong to a bounded set $\mathcal{P}$ that contains $\mathcal{V}$.[8] If the buyer accepts, the game ends; otherwise, the game proceeds to the next period. A history of length t is a sequence $(p_0, p_1, \ldots, p_{t-1}) \in \mathcal{P}^t$ of prices

[8]Boundedness of $\mathcal{P}$ is required only for the proof of the one-shot deviation principle (Proposition 2). The remaining results continue to hold when $\mathcal{P}$ is unbounded, but the proofs of Lemma 6 and Proposition 5 need to be modified (see Kostadinov (2024)).

offered up to time t that have been rejected by the buyer. Let $\mathcal{H}^t$ be the space of histories of length $t \geq 1$, $\mathcal{H}^0 = \{h^0\}$ be the singleton set containing the initial history, and $\mathcal{H} = \cup_{t=0}^{\infty} \mathcal{H}^t$ be the space of all histories.

A strategy for the seller is a function $\sigma_S : \mathcal{H} \to \mathcal{P}$ mapping histories to price offers. A strategy for the buyer is a function $\sigma_B : \mathcal{V} \times \mathcal{H} \times \mathcal{P} \to \{A, R\}$ where $\sigma_B(v, h, p) = A(R)$ means that type v accepts (rejects) price p offered by the seller following history h. Let Σ_S and Σ_B be the respective strategy spaces for the seller and the buyer.

A strategy σ_B for the buyer can be used to infer the set of types $V[h|\sigma_B]$ present at history h, i.e. the types who reject all prices along h according to σ_B. Formally, $V[h^0|\sigma_B] = \mathcal{V}$ and

$$V[h|\sigma_B] = \{v \in \mathcal{V} | \sigma_B(v, (p_0, .., p_{s-1}), p_s) = R \text{ for all } s < t - 1\} \tag{9}$$

for histories $h = (p_0, .., p_{t-1}) \in \mathcal{H}^t$ such that $t > 0$ and the RHS of (9) is nonempty, where $(p_0, .., p_{s-1}) = h^0$ when $s = 0$. Any other history h' cannot be reached if the buyer is following σ_B. I set $V[h'|\sigma_B] = \{\underline{v}\}$ for all such histories. Let $\bar{V}[h|\sigma_B] := \sup V[h|\sigma_B]$ be the highest buyer type present at history h.

It will be useful to define histories $h^s[h|\sigma_S]$ on the path of a seller strategy σ_S following a history h. For any history $h \in \mathcal{H}^t$ let $h^t[h|\sigma_S] = h$ and $h^{s+1}[h|\sigma_S] = (h^s[h|\sigma_S], \sigma_S(h^s[h|\sigma_S]))$ for all $s \geq t$.

3.2 Payoffs and regret

Players share a common discount factor $\delta \in (0, 1)$. Suppose they follow a strategy profile $\sigma = (\sigma_S, \sigma_B)$. The seller's lifetime payoff given history h when the buyer's type is v is given by

$$u_S(v, h|\sigma) = \delta^t p_t,$$

where $t = \inf\{s \geq t | \sigma_B(v, h^s[h|\sigma_S], \sigma_S(h^s[h|\sigma_S])) = A\}$ is the first time following h when the good is sold to type v and $p_t = \sigma_S(h^t[h|\sigma_S])$ is the corresponding sales price. If such a time t does not exist, then the good is never sold and $u_S(v, h|\sigma) = 0$.

Type v's payoff given history h and an offer p is

$$u_B(v, h, p|\sigma) = \delta^t(v - p_t),$$

where t and p_t are defined analogously except for replacing σ_S with the strategy σ_S^p that differs only in $\sigma_S^p(h) = p$.

The highest lifetime payoff the seller can obtain against type v is

$$\bar{u}_S(v|\sigma_B) = \sup_{\sigma'_S \in \Sigma_S} u_S(v, h^0|\sigma'_S, \sigma_B).$$

Suppose the seller has a payoff benchmark against every type given by a function $C : \mathcal{V} \to \mathbb{R}$. Regret against v at any history h is equal to the difference between the benchmark $C(v)$ and the seller's lifetime payoff against v given h:

$$R(v, h|\sigma, C) = C(v) - u_S(v, h|\sigma).$$

The seller's (worst-case) regret against (types in) $\hat{V}$ at h is given by $\sup_{v \in \hat{V}} R(v, h|\sigma, C)$. When $\hat{V} = V[h|\sigma_B]$ this corresponds to the seller's regret against all types believed to be present at h, denoted

$$R(h|\sigma, C) = \sup_{v \in V[h|\sigma_B]} R(v, h|\sigma, C). \tag{10}$$

3.3 Equilibrium

I now propose an equilibrium concept modelled after Perfect Bayesian Equilibrium (PBE) with the difference that the seller minimises lifetime regret instead of maximising payoff.

Definition 1. *A strategy profile σ and a counterfactual function $C : \mathcal{V} \to \mathbb{R}$ form a Perfect Regret Equilibrium (PRE) if the following conditions hold:*

(E1) $\sigma_B \in \operatorname{argmax}_{\sigma'_B \in \Sigma_B} u_B(v, h, p|\sigma_S, \sigma'_B)$ *for all* $v \in \mathcal{V}, h \in \mathcal{H}, p \in \mathcal{P}$

(E2) $\sigma_S \in \operatorname{argmin}_{\sigma'_S \in \Sigma_S} R(h|\sigma'_S, \sigma_B, C)$ *for all* $h \in \mathcal{H}$

(E3) $C(v) = \bar{u}_S(v|\sigma_B)$.

The general principle is as follows. At any history h each player considers all possible realisations x of his opponents' private information drawing inference from past actions and equilibrium strategies. When contemplating a strategy for the rest of the game (each type of) each player computes two payoffs for each realisation x – (i) the best lifetime payoff he could attain if he could return to time 0 and choose a strategy for the entire game with knowledge of x, and (ii) the lifetime payoff based on all actions taken up to h and employing the chosen strategy from h onwards. The player chooses a strategy that minimises worst-case lifetime regret defined as the difference between (i) and (ii), assuming the worst realisation x and assuming that other players follow their equilibrium strategies.

I now demonstrate how this principle is applied to produce Definition 1. Since the seller is uninformed, there is a single realisation x_S of his private information. The buyer's best counterfactual payoff described by (i) above is therefore independent of his strategy, as the worst-case realisation of the seller's private information is always x_S. Hence, to minimise lifetime regret, the buyer maximises payoff (ii) or, equivalently, maximises the continuation payoff from h, exactly as he would in a PBE. This produces condition ($E1$).

Unlike the seller, the buyer has private information about his value v. Hence, at each history h the seller minimises his lifetime regret in (10) against the worst-case type v from the types $V[h|\sigma_B]$ believed to be present at h according to the buyer's strategy.[9] This produces condition ($E2$).

The final condition ($E3$) requires that the benchmark C against which the seller measures regret is given by the payoff (i) above, i.e. $C(v)$ is the best payoff he could obtain counterfactually knowing that the buyer's type is v. Henceforth, I refer to C as the counterfactual (function) and to $C(v)$ as the (best) counterfactual payoffs.

Remark 1. PRE can be interpreted as a solution concept for a game with multiple priors. Let Π be the set of all probability distributions over $\mathcal{V}$ and suppose the seller's prior beliefs are given by the set $\Pi_d = \{\pi \in \Pi | \pi(v) = 1 \text{ for some } v \in \mathcal{V}\}$ of all degenerate distributions. Then $\mu(h|\sigma_B) = \{\pi \in \Pi | \pi(v) = 1 \text{ for some } v \in V[h|\sigma_B]\}$ are the posterior beliefs following history h obtained by updating each prior by Bayes'

[9] Recall that $V[h|\sigma_B] = \{\underline{v}\}$ when h is inconsistent with σ_B, i.e. the seller infers he is facing type $\underline{v}$. This has no impact on behaviour on the equilibrium path by the same argument used for the Bayesian model: any other belief would cause the seller to charge a higher price at h, preserving the buyer's incentives to accept at prior to h.

rule using the buyer's strategy σ_B. Hence, $(E2)$ is equivalent to the (sequential) optimality of the seller's strategy under beliefs $\{\mu(h|\sigma_B)\}_{h\in\mathcal{H}}$ in the following sense:

$$\sigma_S \in \underset{\sigma'_S\in\Sigma_S}{\text{argmin}} \sup_{\pi\in\mu(h|\sigma_B)} \int R(v,h|\sigma'_S,\sigma_B,C)d\pi(v) \text{ for all } h\in\mathcal{H}. \tag{11}$$

Suppose the set of prior beliefs contains Π_d and let $\{\nu(h|\sigma_B)\}_{h\in\mathcal{H}}$ be the corresponding posterior beliefs. Optimality under $\{\nu(h|\sigma_B)\}_{h\in\mathcal{H}}$ is equivalent to (11) because $\nu(h|\sigma_B)\supseteq\mu(h|\sigma_B)$ and each distribution in $\nu(h|\sigma_B)$ is supported on $V[h|\sigma_B]$ so the worst-case distribution in $\nu(h|\sigma_B)$ is degenerate. Hence, PRE is equivalent to a solution concept where the seller entertains any set of priors containing Π^d, e.g. the set of all possible priors.

3.4 Properties of PRE

The most unusual feature of the seller's preferences is their backward-looking nature. Regret is based on what the seller could have done at earlier stages of the game, even though his strategy can only influence the future. Therefore, his preferences violate a common formulation of consequentialism which states that preferences should not depend on past actions.

Another nonstandard feature of the seller's preferences is that they are *not* exogenously given. They are defined only in the context of the counterfactual payoffs in C, which are determined endogenously by the buyer's equilibrium strategy. Nevertheless, given any counterfactual C the preferences satisfy two desirable properties. First, they are dynamically consistent in the following sense adapted from Epstein and Schneider (2003).

Proposition 1. *Fix a time t and let $h^t\in\mathcal{H}^t$, $C:\mathcal{V}\to\mathbb{R}$, $\sigma_B\in\Sigma_B$, $\sigma_S,\sigma'_S\in\Sigma_S$ with $\sigma_S(h)=\sigma'_S(h)$ for all $h\in\mathcal{H}^s$ with $s\leq t$. If $R(h^t,\sigma_S(h^t)|\sigma_S,\sigma_B,C)\leq R(h^t,\sigma_S(h^t)|\sigma'_S,\sigma_B,C)$, then $R(h^t|\sigma_S,\sigma_B,C)\leq R(h^t|\sigma'_S,\sigma_B,C)$.*

Proposition 1 is interpreted as follows. Fix a strategy σ_B for the buyer. If the seller finds it suboptimal to switch from σ_S to σ'_S at the time-$t+1$ history $(h^t,\sigma_S(h^t))$, then this switch is also suboptimal when evaluated at the time-t history h^t. Hence, the choice between σ_S and σ'_S is consistent across time. A useful heuristic is that "the optimal strategy at t remains optimal at $t+1$" but this is a stronger statement that does not hold in general because there may not exist an optimal strategy at $t+1$, or there may be multiple optimal strategies at t, some of which are not optimal at $t+1$.

The seller's preferences also satisfy the one-shot deviation principle, which states that a strategy σ_S is optimal iff there is no better strategy σ'_S that differs from σ_S at a single history.

Proposition 2. *(σ_S, σ_B, C) satisfies (E2) iff $R(h|\sigma_S, \sigma_B, C) \leq R(h|\sigma'_S, \sigma_B, C)$ for all $h \in \mathcal{H}$ and σ'_S with $\sigma_S(h') = \sigma'_S(h')$ for all $h' \neq h$.*

In standard dynamic games players maximise discounted expected utility. These preferences are dynamically consistent and the one-shot deviation principle applies.[10] Moreover, there exists a stationary recursive characterisation of equilibrium (Abreu, Pearce, and Stacchetti, 1990). For example, in the standard durable-good monopoly model the sets of (continuation) PBE from all histories where the seller holds a given belief about the buyer's type are identical. This is not true for PRE because regret is based on lifetime payoffs computed from time 0, so the impact of the seller's strategy diminishes over time. This time dependence results in nonstationary equilibrium behaviour described in Section 5.2.

3.5 Relation to the literature

Decision rules based on worst-case regret originate in Savage (1951). The first application of this criterion to games of incomplete information is in Linhart and Radner (1989).[11] They consider a sealed-bid mechanism for bilateral trade and assume each player minimises regret against the worst-case strategy of the worst-case type of his opponent.[12] There are no constraints on the strategies under consideration except that the buyer does not bid above his value and the seller does not bid below his cost. In contrast, PRE posits that the seller minimises regret against the *equilibrium* strategy of the worst-case buyer type. This feature of PRE appears in the solution concept of Hyafil and Boutilier (2004) for static games.[13] It also features in a number

[10] Since the buyer's preferences in my model amount to maximising discounted expected utility, they are also dynamically consistent and satisfy the one-shot deviation principle.

[11] Renou and Schlag (2010) and Halpern and Pass (2012) offer regret-based solution concepts for static games of complete information.

[12] See Shafer (2020) for an extension to multiple buyers and sellers.

[13] **?** adopts a similar approach from the interim perspective: regret-minimising bidders in a first-price auction correctly anticipate summary statistics of the distribution of their opponents' bids across all possible types.

of static settings where a mechanism designer minimises regret against the worst-case types of the agents given optimal behaviour by the latter (Hurwicz and Shapiro, 1978; Bergemann and Schlag, 2008, 2011; Beviá and Corchón, 2022; Koçyiğit, Rujeerapaiboon, and Kuhn, 2022; Koçyiğit, Kuhn, and Rujeerapaiboon, 2024; Suzdaltsev, 2022b; Zhang, 2022; Guo and Schmaya, 2023, 2025).

A vast literature originating in Hart and Mas-Colell (2000, 2001) studies the implications of regret objectives for the long-run distribution of action profiles in repeated games.[14] In these papers players consider the counterfactual improvement in payoff from changing their strategy from time 0, similarly to PRE. However, this counterfactual gain is computed under the assumption that the actions of their opponents observed so far would remain the same.[15] In a PRE, instead, the seller takes into account that switching to an alternative strategy would change the buyer's behaviour.

Caldentey, Liu, and Lobel (2017) study a continuous-time model of durable-good monopoly where the seller minimises regret and has commitment. Section 4.2 shows that the optimal mechanism in the frequent-action limit of my model matches their solution.

PRE contributes to the literature on dynamic games with ambiguity aversion. Ellis (2018) has shown that departing from standard preferences creates a conflict between dynamic consistency and consequentialism, provided a certain common prior assumption is maintained. This has led to two approaches.

One branch of the literature insists on dynamic consistency. Hanany, Klibanoff, and Mukerji (2020) propose a solution concept similar to PBE for players with smooth ambiguity preferences who hold multiple prior beliefs (Klibanoff, Marinacci, and Mukerji, 2005).[16] These preferences are dynamically consistent when beliefs are updated in a specific way that violates consequentialism. This approach preserves the one-shot deviation principle. PRE exhibits the same tradeoff – dynamic consistency and the one-shot deviation principle hold (Propositions 1 and 2), while consequentialism is sacrificed.

It is also possible to obtain dynamic consistency in specific games games with information structure satisfying a rectangularity property (Epstein and Schneider,

[14] See Chapter 7 of Cesa-Bianchi and Lugosi (2006) for a survey.

[15] Schlag and Zapechelnyuk (2012) show that relaxing this assumption has a substantial impact on the results of this literature.

[16] de Castro and Galvao (2019) and de Castro, Galvao, and da Siva Nunes (2025) apply a similar approach to a single agent with quantile preferences.

2003). Pahlke (2022) defines a solution concept for players with minmax preferences (Gilboa and Schmeidler, 1989). While these preferences are dynamically inconsistent in general, Pahlke (2022) obtains dynamic consistency by modifying ex-ante beliefs. Liu (2024) studies dynamic contracts where the principal faces ambiguity about the actions available to the agent but, unlike in standard models, the agent is myopic. **?** study durable-good monopoly where the seller faces ambiguity about the dynamic process governing how the buyer learns about his value. They restrict attention to processes that make the seller's preferences dynamically consistent.

A second branch of the literature works with dynamically inconsistent preferences, while preserving consequentialism (Bose and Daripa, 2009; Bose and Renou, 2014; Auster and Kellner, 2022; **?**).[17] The solution concepts in this literature require that strategies are immune to one-shot deviations, or satisfy the closely related notion of consistent planning (Strotz, 1955). However, the one-shot deviation principle does not apply due to lack of dynamic consistency. Hence, in equilibrium players can profit from multi-stage deviations.[18]

The "dynamic inconsistency" approach above is also adopted in a regret-based solution concept called Perfect Compromise Equilibrium (PCE) proposed by Schlag and Zapechelnyuk (2024).[19] Both PRE and PCE posit that at every history players minimise regret against the worst-case type of their opponent given their equilibrium strategy. The difference lies in the formulation of regret. In a PCE a player's regret against type v of his opponent equals the difference between his equilibrium *continuation* payoff against v and the best *continuation* payoff he could obtain against v *from the current history*. He computes payoffs assuming not only that his opponent will follow their equilibrium strategy but also that he himself will follow his equilibrium strategy in subsequent periods. This has two important implications. First, attention is restricted to *one-shot deviations* when computing the best continuation payoff against v. Second, the optimality of equilibrium strategies is checked only against *one-shot deviations*. In a PRE regret against v equals the difference between the *lifetime* equilibrium payoff against v and the best *lifetime* payoff against v. *All deviations* are considered when computing the best payoff against v, and equilibrium

[17] Auster, Che, and Mierendorff (2024) apply this approach to a single-agent stopping problem.

[18] In special settings there may be no profitable multi-stage deviations, despite dynamic inconsistency. Malladi (2020) shows this in the context of a single-agent search problem.

[19] Hayashi (2009) uses the same regret objective in a single-agent stopping problem. In **?** I apply PCE to a dynamic setting with uncertainty about future stage games.

strategies are optimal against *all deviations*. These differences are further illustrated in the Appendix, which contains a derivation of the PCE in the two-period setting from Section 2.

4 The commitment solution

In this section I assume the seller can commit to a strategy. This amounts to the design of a mechanism consisting of a sequence of prices $(p_t)_{t=0}^{\infty}$. The seller minimises regret at the initial history h^0 subject to the buyer's optimal behaviour. It will be useful to study the general problem of finding a regret-minimising mechanism given an arbitrary function $C : \mathcal{V} \to \mathbb{R}$ of counterfactual payoffs. Such a mechanism is called C-optimal.

Definition 2. *Let $C : \mathcal{V} \to \mathbb{R}$. A mechanism $(p_t)_{t=0}^{\infty}$ is C-optimal if it solves*

$$\begin{aligned} &\min_{(p_t)_{t=0}^{\infty}, \sigma \in \Sigma_S \times \Sigma_B} \quad R(h^0|\sigma, C) \\ &s.t\ \sigma_S(h^t) = p_t \text{ for all } t, h^t \in \mathcal{H}^t \\ &(\sigma, C) \text{ satisfies } (E1) \end{aligned}$$

An optimal mechanism minimises regret given the best counterfactual payoffs that can be obtained with commitment. These payoffs are given by the identity function $C(v) = v$ because the seller can extract all surplus from a given type v by committing to price at v in perpetuity.

Definition 3. *A mechanism is optimal if it is C-optimal for $C(v) = v$.*

I now outline the structure of C-optimal mechanisms when C belongs to the following class $\mathcal{C}$, which includes the identity function.

Definition 4. *$\mathcal{C}$ denotes the class of continuous functions $C : \mathcal{V} \to \mathbb{R}$ that are increasing at a nonvanishing rate, i.e. there exists $\lambda > 0$ such that $C(v') - C(v) \geq \lambda(v' - v)$ for all $v' > v$.*

The solution resembles its counterpart from the two-period example in Section 2. Optimal prices depend on time t and a state variable v_t representing the highest outstanding buyer type at time t (in the supremum sense) with the convention $\sup \emptyset =$

$\underline{v}$. The state is strictly decreasing over time until a (possibly infinite) time T with $v_{T+1} = \underline{v} < v_T$ when the good is sold to all remaining types. Hence, the set of types B_t who buy the good at time t is nonempty for all $t < T$ and satisfies $(v_{t+1}, v_t) \subseteq B_t \subseteq [v_{t+1}, v_t]$. Let $B_\infty = \mathcal{V} \setminus \cup_{t=0}^{\infty} B_t$ be the set of types who never buy the good.

The state evolves according to the law of motion $v_{t+1} = f_t^C(v_t)$ where

$$f_t^C(v) := \sup\{v' < v | C(v) \geq C(v') + \delta^t(1-\delta)v'\} \tag{12}$$

similarly to (4) and (6) from the two-period example. Let

$$f_{s\to t}^C(v) := f_s^C \circ f_{s-1}^C \circ ... \circ f_{t+1}^C \circ f_t^C(v) \tag{13}$$

for all $s \geq t$ with the convention $f_{t-1\to t}^C(v) = v$. Hence, a C-optimal mechanism results in state $v_t = f_{t-1\to 0}^C(\bar{v})$ at every time t. Note that $f_t^C(\underline{v}) = \underline{v}$ due to $\sup \emptyset = \underline{v}$, so $v_{t+1} = \underline{v}$ whenever $v_t = \underline{v}$.

The price at time t and state v is given by

$$P_t^C(v) = (1-\delta) \sum_{s=t}^{\infty} \delta^{s-t} f_{s\to t}^C(v). \tag{14}$$

A C-optimal mechanism sets price $P_t^C(v_t)$ at each time t when sales are made. These prices ensure that the state follows the law of motion (12) when the buyer acts optimally. In particular, (14) implies

$$v_{t+1} - P_t^C(v_t) = \delta(v_{t+1} - P_{t+1}^C(v_{t+1})), \tag{15}$$

which means that type v_{t+1} is indifferent between buying the good at times t and $t+1$. Types above v_{t+1} strictly prefer to buy at t, while types below v_{t+1} strictly prefer not to buy at t, resulting in state v_{t+1} at time $t+1$.

The following proposition summarises the characterisation of C-optimal mechanisms.

Proposition 3. *Let $C \in \mathcal{C}$. Then $(p_t^*)_{t=0}^{\infty} = (P_t^C(f_{t-1\to 0}^C(\bar{v})))_{t=0}^{\infty}$ is a C-optimal mechanism. Any C-optimal mechanism $(p_t)_{t=0}^{\infty}$ has $p_t = p_t^*$ for all t such that $f_{t-1\to 0}^C(\bar{v}) > \underline{v}$.*

A C-optimal mechanism equalises regret across time similarly to the two-period setting. To see this, note that the seller obtains payoff $\delta^t P_t^C(v_t)$ against all types in B_t.

The worst-case type in B_t is therefore $v_t = \sup B_t$ because C is strictly increasing. It follows that regret against B_t equals $C(v_t) - \delta^t P_t^C(v_t)$. If $v_{t+1} > \underline{v}$, then the continuity of C and the law of motion (12) imply that

$$C(v_t) = C(v_{t+1}) + \delta^t(1-\delta)v_{t+1}.$$

It follows from (15) that

$$C(v_t) - \delta^t P_t^C(v_t) = C(v_{t+1}) - \delta^{t+1} P_{t+1}^C(v_{t+1}),$$

i.e. the optimal mechanism equalises worst-case regret against B_t and B_{t+1}. If sales continue indefinitely, i.e. $B_\infty \neq \emptyset$, then regret against B_t also equals regret against B_∞. This is because regret against B_t equals regret against B_s for arbitrarily large s. The latter equals $C(v_s) - \delta^s P_s^C(v_s)$ and, therefore, converges to $C(v_\infty)$, where $v_\infty = \lim_{s\to\infty} v_s$. On the other hand, regret against B_∞ is given by $C(v_\infty)$ because no sales are made to v_∞ and lower types.

The argument for the optimality of regret equalisation is similar to the two-period example. The seller's regret can be rewritten as follows:

$$R(h^0|\sigma, C) \equiv \sup_{v\in\mathcal{V}} R(h^0, v|\sigma, C) = \sup_{t\in\{0,1,..,\infty\}:B_t\neq\emptyset} \quad \sup_{v\in B_t} R(h^0, v|\sigma, C).$$

The optimal mechanism equates regret across each nonempty group of buyers $B_0, B_1, .., B_\infty$ because decreasing regret against one group will increase regret against another. For example, the seller can decrease regret against B_0 by charging a higher price p_0. However, this would make more types wait until $t = 1$ to purchase the good, thereby increasing regret against B_1. If the seller tries to compensate for this by raising all prices, then there will be more types who never buy the good, thereby increasing regret against B_∞.

4.1 Time Consistency

The regret equalisation of the optimal mechanism implies a powerful time consistency property echoing the two-period setting. Suppose the buyer believes the seller is committed to an optimal mechanism $(p_t^*)_{t=0}^\infty$. The seller follows the mechanism until a history $h \in \mathcal{H}^t$. At this point he is (unexpectedly) given the opportunity to commit to a different sequence of prices for the remainder of the game. A *continuation-optimal* mechanism from h minimises regret against the remaining types $V[h|\sigma_B]$, aligning with his preferences in a PRE given by condition $(E2)$. This amounts to minimising regret against the worst-case (nonempty) group of types $B_t, B_{t+1}, .., B_\infty$, which is already achieved by the optimal mechanism due to regret equalisation. Hence, an (ex-ante) optimal mechanism is continuation-optimal from any on-path history h. This property is often called *time consistency* following Kydland and Prescott (1977).

Optimal mechanisms in standard settings often lack time consistency. If the durable-good monopolist is a Bayesian expected-utility maximiser, then he would commit to set the static monopoly price in each period (Stokey, 1979). The buyer's best response is to purchase in the initial period if his value exceeds the price, and to never purchase otherwise. Therefore, the seller would benefit from lowering prices after the initial sales are made, making the optimal mechanism time-inconsistent.

Proposition 4 states the time consistency result for my model and generalises it to any counterfactual $C \in \mathcal{C}$ by showing that the C-optimal mechanism (p_t^*) described in Proposition 3 remains continuation-optimal from every history on its path (given C). Proposition 4 also describes continuation-optimal mechanisms from any off-path history $h \in \mathcal{H}^t$ in state v, where the seller prices at $P_s^C(f_{s-1\to t}^C(v))$ at any time $s \geq t$. This is used in Section 5 to characterise PRE, where the seller lacks commitment.

Proposition 4. *Suppose that* (σ, C) *satisfies* $(E1)$ *and* $C \in \mathcal{C}$. *Fix a time* t, *a history* $h \in \mathcal{H}^t$, *and let* $v = \bar{V}[h|\sigma_B]$. *Then*

$$R(h|\sigma, C) \geq C(v) - \delta^t P_t^C(v). \tag{16}$$

In addition,

- *If* $v > 0$, *then* (16) *holds at equality only if* $\sigma_S(h^s[h|\sigma_S]) = P_s^C(f_{s-1\to t}^C(v))$ *for all* $s \geq t$ *such that* $f_{s-1\to t}^C(v) > \underline{v}$ *or* $s = t$.
- *If* $\sigma_S(h^s[h|\sigma_S]) = P_s^C(f_{s-1\to t}^C(v))$ *for all* $s \geq t$, *then* (16) *holds at equality.*

The time consistency of the optimal mechanism does *not* follow from the dynamic consistency of the seller's preferences (Proposition 1). These two consistency notions require robustness of optimal choice to different forms of commitment. Dynamic consistency implies that when the buyer's strategy is fixed the seller does not benefit from committing to his own strategy. This type of commitment is beneficial only if the seller is worried that his preferences at subsequent histories will not align with his optimal choice from today's perspective. On the other hand, the seller's strategy is time-consistent if he cannot benefit from committing to a strategy and forcing the buyer to best-respond to it. This type of commitment is more powerful because it allows the seller to influence the buyer's strategy by acting as a Stackelberg leader.[20] Note that the expected utility preferences are dynamically consistent, yet they result in a time-inconsistent optimal mechanism, as discussed above.

4.2 Frequent offers

I now use Proposition 3 to give a more precise characterisation of the optimal mechanism, where counterfactual payoffs are given by $C(v) = v$. In this case the law of motion (12) simplifies to

$$f_t^C(v) = \max\left\{\frac{v}{1+\delta^t(1-\delta)}, \underline{v}\right\}. \tag{17}$$

This can be used to obtain a closed form for the states $v_t = f^C_{t-1\to 0}(\bar{v})$ and prices $p_t = (1-\delta)\sum_{s\geq t}\delta^{s-t}v_{s+1}$.

It is convenient to study the limiting case where the monopolist can make arbitrarily frequent offers.[21] Let $\Delta > 0$ denote the units of real time between consecutive periods. If players discount at a rate r, then $\delta = e^{-r\Delta}$. As offers become arbitrarily frequent, i.e. $\Delta \to 0$, the law of motion (17) converges to

$$\dot{v}_t = -re^{-rt}v_t$$

for $v_t > \underline{v}$, where t measures units of real time, not periods. The solution is

$$v_t = \max\left\{\frac{\bar{v}}{e}\exp(\exp(-rt)), \underline{v}\right\}.$$

Let $T = \sup\{t \geq 0 | v_t > \underline{v}\}$ be the time it takes the monopolist to sell to all types. If $\underline{v} > \bar{v}/e$, then $T = -\ln(1+\ln(\underline{v}/\bar{v}))$ and all types are served. The seller's regret is

[20] In single-agent decision problems time consistency is equivalent to dynamic consistency, so the terms are used interchangeably.

[21] Frequent offers are in the monopolist's interest because commitment allows him to replicate any pricing policy with less frequent offers.

given by $R^* = -\underline{v}\ln(\underline{v}/\bar{v})/r$. If $\underline{v} \leq \bar{v}/e$, then $T = \infty$ and sales continue indefinitely. In this case the seller's regret is $R^* = \bar{v}/e$. The resulting prices at any time $t \leq T$ are given by

$$p_t = \frac{v_t - R^*}{e^{-rt}}.$$

Figure 4 and Figure 5 plot the solution holding r and $\bar{v}$ constant while varying $\underline{v}$. When $\underline{v}$ is high, there is little ambiguity about the buyer's type and the good is sold in finite time. When $\underline{v}$ is low, there is more ambiguity and some types are not served.[22]

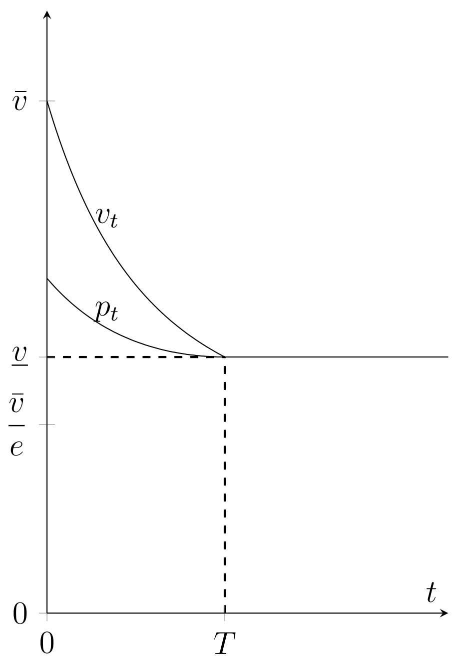

Figure 4: $T < \infty$

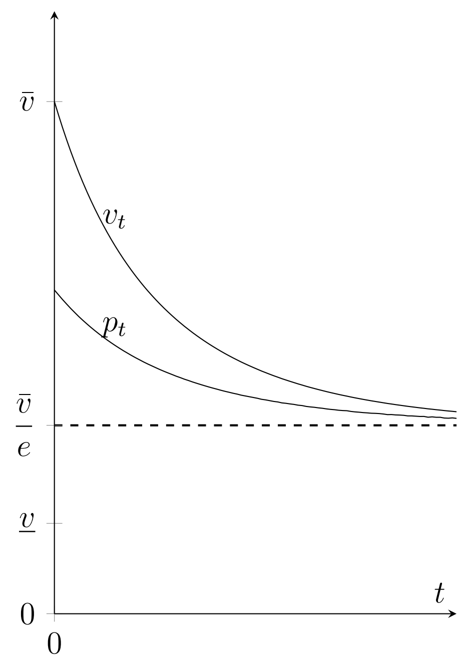

Figure 5: $T = \infty$

The limiting optimal mechanism described above matches the optimal mechanism in continuous time, described in Corollary 2 in Caldentey, Liu, and Lobel (2017). The dynamic pricing of the seller can also be interpreted as a menu of contracts $(x(v), p(v))_{v \in \mathcal{V}}$ in a static setting where $x(v)$ and $p(v)$ respectively denote the probability of selling the good and the transfer associated with contract v. Lower prices are paired with lower probability of obtaining the good to match the discounting in

[22]The parameters for Figure 4 are $\underline{v} = 0.85, \bar{v} = 1.7, r = 2$ and the parameters for Figure 5 are $\underline{v} = 0.3, \bar{v} = 1.7, r = 2$.

the dynamic model. This resulting menu is given by

$$x(v) = \begin{cases} 0 & \text{if } v \leq \frac{\bar{v}}{e} \\ \ln(v/\bar{v}) + 1 & \text{if } v \geq \frac{\bar{v}}{e} \end{cases} \qquad p(v) = \begin{cases} 0 & \text{if } v \leq \frac{\bar{v}}{e} \\ v - R^* & \text{if } v \geq \frac{\bar{v}}{e} \end{cases}.$$

It matches the optimal static mechanism in this setting described in Corollary 1 in Bergemann and Schlag (2005). In the Bayesian model optimal dynamic pricing also matches the optimal static mechanism, but the seller charges the same price to every type who is served. Stokey (1979) attributes this to the equal rates of time preference of both parties, who share a discount factor. In this paper discount factors are also equal, but the regret-based preferences of the seller are nonstationary, while the buyer's preferences for maximising discounted expected utility are stationary. This difference in time preference produces price discrimination in the optimal mechanism.

5 Equilibrium Characterisation

I now study Perfect Regret Equilibrium in order to characterise the outcome when the seller cannot commit to future prices. Most of the results apply to the gap case $\underline{v} > 0$, where gains from trade with all types are positive. Discussion of the no-gap case $\underline{v} = 0$ is deferred to Section 5.5.

The main result of the paper, Theorem 1 below, establishes the existence of a PRE. The equilibrium counterfactual function, buyer payoffs, and seller regret are uniquely determined. The equilibrium strategies at each history are unique up to a single buyer type who is indifferent between accepting and rejecting.

Definition 5. *Strategies $\sigma = (\sigma_S, \sigma_B)$ are essentially equivalent to strategies $\sigma' = (\sigma'_S, \sigma'_B)$ given counterfactual $C : \mathcal{V} \to \mathbb{R}$, denoted $\sigma \overset{C}{=} \sigma'$, if for any history $h \in \mathcal{H}$ and price $p \in \mathcal{P}$*

- $\sigma_S(h) = \sigma'_S(h)$
- $\sigma_B(v, h, p) = \sigma'_B(v, h, p)$ *for all but a single type $v \in \mathcal{V}$*
- $R(h|\sigma, C) = R(h|\sigma', C)$
- $u_B(v, h, p|\sigma) = u_B(v, h, p|\sigma')$ *for all $v \in \mathcal{V}$.*

Definition 6. *A PRE (σ, C) is essentially unique if every PRE (σ', C') satisfies $C = C'$ and $\sigma \overset{C}{=} \sigma'$.*

Theorem 1. *If $\underline{v} > 0$, there exists an essentially unique PRE.*

Given the equilibrium counterfactual C, the equilibrium strategies σ are essentially equivalent to the profile σ^C defined below.

Definition 7. *For any $C \in \mathcal{C}$, the strategies $\sigma^C = (\sigma_S^C, \sigma_B^C)$ are given by*

- $\sigma_B^C(v, h, p) = A$ *iff* $p \leq (1-\delta)v + \delta P_{t+1}^C(\min\{v, \bar{V}[h|\sigma_B^C]\})$
- $\sigma_S^C(h) = P_t^C(\bar{V}[h|\sigma_B^C])$

for all $t, h \in \mathcal{H}^t, v \in \mathcal{V}, p \in \mathcal{P}$.

There is a deep connection between the PRE (σ^C, C) and the commitment outcome given the same counterfactual C, similarly to the two-period example. The following lemma shows that if σ^C is played, then states follow the law of motion and the seller follows a continuation-optimal mechanism from any history h given C, as described in Proposition 4. The seller's regret at any history equals regret against the highest type who purchases the good immediately, thereby attaining the lower bound in (16). In particular, setting $h = h^0$ shows that states, prices and regret on the equilibrium path match their counterparts from the C-optimal mechanism described in Proposition 3.

Lemma 1. *Let $C \in \mathcal{C}$. The strategies σ^C are uniquely defined and satisfy*

- $\bar{V}\left[h^s|\sigma_B^C\right] = f_{s-1\to t}^C(v)$
- $\sigma_S^C(h^s) = P_s^C(f_{s-1\to t}^C(v))$
- $R(h|\sigma^C, C) = C(v) - \delta^t P_t^C(v)$

for all times $s \geq t$ and histories $h \in \mathcal{H}^t, h^s = h^s[h|\sigma_S^C]$ with $v = \bar{V}[h|\sigma_B^C]$.

The connection between the outcomes with and without commitment extends to any counterfactual $C \in \mathcal{C}$. Proposition 5 below states that if the seller's preferences are based on an exogenous $C \in \mathcal{C}$, then the optimality of the strategies embodied in equilibrium conditions $(E1)$ and $(E2)$ isolates an essentially unique profile σ^C. Hence, states and prices match their counterparts from any C-optimal mechanism even when $(E3)$ does not hold.

Definition 8. *(σ, C) is a quasi-PRE if it satisfies $(E1)$ and $(E2)$.*

Proposition 5. *Let $C \in \mathcal{C}$ and $\underline{v} > 0$. Then (σ, C) is a quasi-PRE iff $\sigma \overset{C}{=} \sigma^C$.*

The existence of a quasi-PRE (σ^C, C) follows from the time consistency of C-optimal mechanisms described in Proposition 4. The latter implies that the seller cannot benefit from abandoning σ^C at any history and committing to an alternative strategy for the rest of the game, which is closely related to the optimality of the seller's strategy σ_S^C in the setting without commitment. The essential uniqueness of the quasi-PRE strategies σ^C is obtained by showing that the seller can guarantee the lower bound on regret in (16), thereby uniquely determining equilibrium regret at every history (Lemma 8 in the Appendix).

The PRE characterisation is completed as follows. First, any PRE (σ, C) has $C \in \mathcal{C}$ (Lemma 12 in the Appendix). Hence, Proposition 5 implies that equilibrium strategies are essentially equivalent to σ^C. The equilibrium counterfactual C is characterised by fixed-point property (18) below. Given C, the buyer's strategy σ_B^C determines the best payoff the seller can obtain against any given type v. This amounts to $(1-\delta)v + \delta P_1^C(v)$, obtained by selling to v at $t = 0$ at the highest price he would accept. Equilibrium condition $(E3)$ requires that this payoff matches the counterfactual $C(v)$, i.e.

$$C(v) = (1-\delta)v + \delta P_1^C(v). \tag{18}$$

The proof of Theorem 1 is completed by showing that there exists a unique counterfactual $C \in \mathcal{C}$ that satisfies (18) (Lemma 15 in the Appendix).

5.1 Comparison with optimal mechanism

Figure 6 shows the best counterfactual payoffs $C^{\text{com}}(v) = v$ under commitment and the best counterfactual payoffs C^{eqm} in any PRE.[23] The solutions with and without commitment share the same structure, differing only in their respective counterfactuals. Indeed, if there were a PRE (σ, C^{com}), then Proposition 5 and Lemma 1 would imply that equilibrium prices match the optimal mechanism. Conversely, if the seller had commitment and the counterfactual payoffs were given by C^{eqm}, then the resulting C^{eqm}-optimal mechanism would match the price path of the PRE $(\sigma^{C^{\text{eqm}}}, C^{\text{eqm}})$.

[23] The parameters for Figure 6 are $\underline{v} = 1, \bar{v} = 10, \delta = 0.8$.

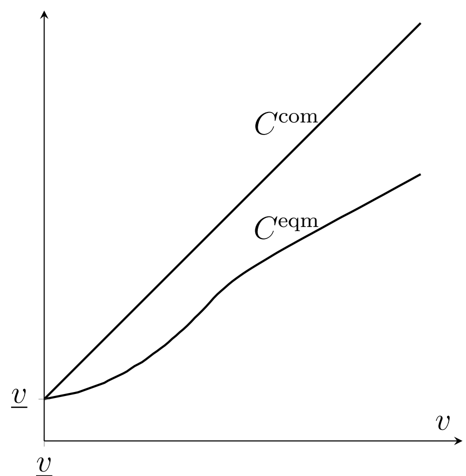

Figure 6: Counterfactuals with and without commitment

In the two-period example it was possible to show that C^{eqm} is flatter than C^{com}, which implies that prices are lower and sales are faster in equilibrium.[24] However, the rates of increase of C^{eqm} and C^{com} in the infinite-horizon setting are not necessarily ranked.[25] Numerical examples indicate it is possible that more types buy the good in the initial period under commitment, i.e. $f_0^{C^{\text{com}}}(\bar{v}) < f_0^{C^{\text{eqm}}}(\bar{v})$.

5.2 Nonstationarity and Rationing

All PRE exhibit nonstationary behaviour illustrated in Figure 7 below. Consider histories $h^t \in \mathcal{H}^t$ and $h^{t+1} \in \mathcal{H}^{t+1}$, both of which are in the same state v. The law of motion (12) and the monotonicity of the equilibrium counterfactual C imply that $f_t^C(v) < f_{t+1}^C(v)$ (unless both equal $\underline{v}$). Hence, less types buy the good at h^{t+1} than at h^t. More generally, starting from h^{t+1} results in slower sales in all subsequent periods in the sense that $f_{t+k+1\to t+1}^C(v)$ (the state k periods from h^{t+1}) is higher than $f_{t+k\to t}^C(v)$ (the state k periods from h^t). This results in a higher price $P_{t+1}^C(v)$ at h^{t+1} compared to the price $P_t^C(v)$ charged at h^t.

[24]This feature of the two-period model does not depend on parameters.
[25]An analytical example is available upon request.

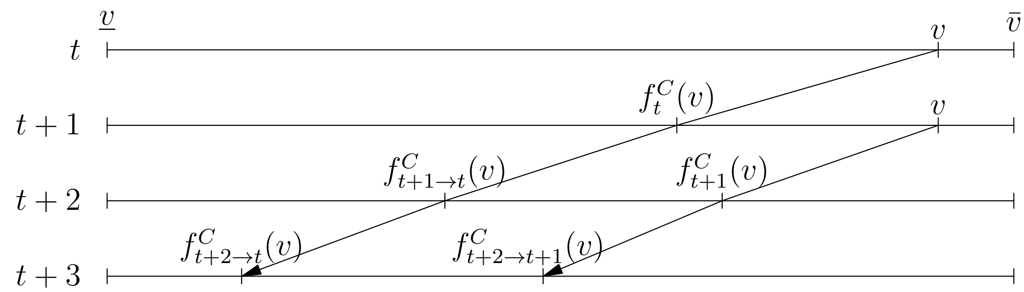

Figure 7: Nonstationarity of law of motion

The dependence of equilibrium prices on time for a fixed state distinguishes PRE from PBE, where prices depend only on the state. The difference is due to the time-0 benchmark for regret, which makes the seller's preferences nonstationary. As time goes on, the best counterfactual payoff $C(v)$ against a given type v is unchanged, but the seller's (lifetime) payoff from selling to v at a given price decreases due to discounting. In equilibrium higher types accept higher prices, so payoffs against them from histories in the same state decrease faster over time in comparison to payoffs against lower types. This makes the seller prioritise selling to high types more heavily as time goes on, resulting in higher prices, i.e. $P_{t+1}^C(v) > P_t^C(v)$.

A similar slowdown of sales is observed on the equilibrium path. If the good is not sold after sufficiently many periods, equilibrium payoffs against all remaining types become arbitrarily close to zero due to discounting, even if the seller were to extract full surplus from each of them. Therefore, regret against each type v converges to $C(v)$. The strict monotonicity of C implies that the seller's regret against sufficiently high types exceeds regret against the remaining types regardless of his continuation strategy. Hence, worst-case regret minimisation causes the seller to focus entirely on selling to high types. The resulting exclusion of low types is often termed "rationing".

Proposition 6. *There exists $N > 0$ such that every PRE (σ, C) satisfies $\lim_{t\to\infty} f_{t\to 0}^C(\bar{v}) > \underline{v}$ whenever $\underline{v} > 0$ and $\bar{v}/\underline{v} > N$.*

Section 4.2 showed that rationing can arise in an optimal mechanism. The possibility of rationing in equilibrium, however, creates a qualitative difference between my model and the Bayesian model.[26] In a PBE the good is sold to all types in finitely

[26]More specific comparisons between PRE and PBE are difficult to make unless a specific prior over the buyer's types is considered for the latter. For example, PBE prices can be higher than PRE prices if there is a lot of mass near the top of the type space and vice versa.

many periods.[27] Proposition 6 shows that this is not true in a PRE when the set of types is sufficiently large.

5.3 Effect of ambiguity

I now consider the effect of increased ambiguity, modelled as an expansion of the set of types $\mathcal{V}$ the seller considers to be possible.

Proposition 7. *Suppose $\underline{v} > 0$. Let (σ, C) be a PRE, and let $v_t = f^C_{t-1\to 0}(\bar{v})$ and $p_t = P^C_t(v_t)$ be the equilibrium states and prices at each time t. Consider an alternative game with $\mathcal{V} = [\alpha\underline{v}, \beta\bar{v}]$ for some $\alpha > 0, \beta > 0$. Let $(\hat{\sigma}, \hat{C})$ be a PRE of the alternative game with states $(\hat{v}_t)_{t=0}^{\infty}$ and prices $(\hat{p}_t)_{t=0}^{\infty}$.*

(i) If $\alpha = \beta$, then $\hat{v}_t = \beta v_t$ and $\hat{p}_t = \beta p_t$.

(ii) If $\alpha = 1 < \beta$, then $\hat{v}_t \geq v_t$ and $\hat{p}_t \geq p_t$.

Part (i) of Proposition 7 states that if the set of types is scaled by a constant β, then equilibrium states and prices scale by the same factor. Hence, if ambiguity increases while holding $\bar{v}/\underline{v}$ constant, then prices increase and the same proportion of types buys the good in each period.

Part (ii) considers the case where the increase in ambiguity arises from adding types higher than $\bar{v}$ to the set $\mathcal{V}$. This results in slower sales and higher prices. A key part of the argument is that continuation strategies after a type $v \in [\underline{v}, \bar{v}]$ accepts are identical to their counterparts in the original game with $\mathcal{V} = [\underline{v}, \bar{v}]$. Hence, counterfactual payoffs against types in $[\underline{v}, \bar{v}]$ remain unchanged and are lower than counterfactual payoffs against types higher than $\bar{v}$. The original prices $(p_t)_{t=0}^{\infty}$ would therefore produce higher regret against types who buy at time 0 than regret against the remaining types, thereby violating regret equalisation. Hence, the seller sets a higher initial price $\hat{p}_0 \geq p_0$ and similar arguments apply to subsequent prices.

Ambiguity can also be expanded by adding types lower than $\underline{v}$, resulting in a new set of types $\mathcal{V} = [\alpha\underline{v}, \bar{v}]$ with $\alpha < 1$. Unlike in part (ii), this would change counterfactual payoffs against the original types $[\underline{v}, \bar{v}]$. For example, after types $[\underline{v}, \bar{v}]$ purchase the good, prices will continue dropping to serve the newly introduced lower

[27] See Lemma 3 in Fudenberg, Levine, and Tirole (1985).

types. Hence, equilibrium strategies change and so do counterfactual payoffs.[28] Consequently, the new game can have higher or lower prices depending on parameters.

5.4 Frequent offers

Coase (1972) conjectured that lack of commitment causes the monopolist to charge prices close to the lowest value $\underline{v}$. Gul, Sonnenschein, and Wilson (1986) formalised the Coase conjecture in the Bayesian model by demonstrating that prices approach $\underline{v}$ when the seller makes sufficiently frequent offers, which corresponds to the limit as players become arbitrarily patient. The following result shows that the Coase conjecture also holds in my model with regret.

Theorem 2. *If $\underline{v} > 0$, the prices offered by the seller in any PRE converge uniformly to $\underline{v}$ as $\delta \to 1$.*

In the Bayesian model the Coase conjecture is attributed to the time inconsistency of the optimal mechanism. If the seller could commit, he would charge the static monopoly price in perpetuity and, therefore, sales would occur only in the initial period (Stokey, 1979). In any subsequent period it would be profitable to abandon the optimal mechanism by lowering prices in order to generate sales. In equilibrium the seller cannot commit not to act on these incentives. Therefore, the buyer anticipates that prices will be lowered, causing the seller to charge low prices in the first place.

The intuition from the Bayesian model does not translate directly to my setting because in any PRE (σ, C) the seller follows the C-optimal mechanism, which is time-consistent (Proposition 4). The Coase conjecture is, instead, driven by the time inconsistency of the seller's optimal *counterfactual* mechanism against type $v > \underline{v}$, which prices at v in perpetuity. This mechanism forms the benchmark $C(v)$ for the seller's regret against v in the case of commitment. However, if the seller prices in this manner when he lacks commitment, type v would not buy the good in anticipation of future discounts. This ends up lowering the best counterfactual payoff $C(v)$ to $\underline{v}$ in the limit as $\delta \to 1$. Since all counterfactual payoffs are close to $\underline{v}$, the monopolist lacks motivation to extract higher revenue from higher types and, therefore, sells the good quickly, validating the Coase conjecture.

[28] Formally, the proof of part (i) shows that $\hat{C}(\alpha v) = \alpha C(v)$ for all $v \in [\underline{v}, \bar{v}]$ but C is not linear, so $\hat{C}(\alpha v) \neq C(\alpha v)$.

In the Bayesian model the good is sold in a number of periods bounded by some finite T as $\delta \to 1$.[29] This implies that the buyer needs to wait at most T periods to purchase the good at price $\underline{v}$. The cost of waiting becomes negligible as $\delta \to 1$, driving equilibrium prices down to $\underline{v}$. In a PRE, however, the good may not be sold to some types when $\bar{v}$ is high (Proposition 6). Nevertheless, the proof of Theorem 2 shows that the good is sold in finitely many periods in the limit as $\delta \to 1$ (holding $\mathcal{V}$ constant), so prices converge to $\underline{v}$.

5.5 The no-gap case

The PRE results presented so far were for the gap case $\underline{v} > 0$. The no-gap case $\underline{v} = 0$ presents some additional hurdles. First, the time consistency result (Proposition 4) does not necessarily hold because it requires that the equilibrium counterfactual C belongs to the class $\mathcal{C}$. In the gap case this is proved via upward induction (Lemma 12 in the Appendix). The argument shows that the desired properties of C hold on some interval $[\underline{v}, v^*]$ where the next-period state is $\underline{v}$. The properties are then extended to a larger interval $[\underline{v}, v^* + \varepsilon]$ and so on. This method cannot be applied to the no-gap case because on the equilibrium path type $\underline{v} = 0$ never buys the good, so the state never reaches $\underline{v}$. Similar difficulties produce equilibrium multiplicity in the Bayesian model (Ausubel and Deneckere, 1989).

Secondly, even if I restrict attention to PRE (σ, C) with counterfactual $C \in \mathcal{C}$, it is unclear whether C is uniquely determined. This is again because the corresponding proof for the gap case uses an upward induction argument (Lemma 15 in the Appendix). On the positive side, the prices in any PRE (σ, C) with $C \in \mathcal{C}$ match the prices of a C-optimal mechanism.[30] These prices converge uniformly to $\underline{v} = 0$ as $\delta \to 1$ by an argument similar to the proof of Theorem 2. One such equilibrium is characterised below. It is remarkably simple in comparison to the gap case due to the linearity of the counterfactual C, which implies prices are linear in the state. In addition, equilibrium prices are lower than their counterparts in the optimal mechanism because the slope of C is smaller than one (see Section 5.1).

Proposition 8. *If $\underline{v} = 0$, there exists a PRE (σ^C, C) with $C(v) = \alpha v$ and $\alpha \in (0, 1)$.*

[29] See Lemma 3 in Fudenberg, Levine, and Tirole (1985).

[30] This can be shown similarly to Proposition 5. In fact, the equilibrium strategies are essentially equivalent to σ^C except that in state $\underline{v} = 0$ the seller may set a positive price because he is indifferent between making no sales and selling to type 0 at a price of 0.

6 Extensions

6.1 Mixed strategies

The analysis so far assumed that both players use pure strategies. Regret against type v at history h was defined as the difference between the best lifetime payoff against v and the equilibrium payoff against v given history h. The same definition can be extended to mixed strategies by replacing the lifetime payoffs above with the corresponding expected lifetime payoffs under the equilibrium strategies.

Bergemann and Schlag (2008) have shown that the seller can benefit from adopting a mixed strategy in the static setting. The same is true in the dynamic setting of this paper. Recall that with pure strategies regret against the worst-case type who buys at t equals regret against the worst-case type who buys at $t+1$. However, regret differs across types who buy at t because they all pay the same price. Random prices allow for further separation. Higher types find a larger set of prices (in the support of the mixed strategy) acceptable and, therefore, buy the good with higher probability. I conjecture that this will result in regret equalisation across all types who are not rationed, which features in the static setting (Bergemann and Schlag, 2008).

The introduction of mixed strategies for the buyer depends on how the expected equilibrium payoff of the seller given h is computed. It may be intuitive to consider the expected lifetime payoff conditional on the actions in history h when the equilibrium strategies are followed from h onwards. The problem with this definition is that all interior acceptance probabilities are treated identically by the seller's objective. Indeed, it is irrelevant whether type v accepts with probability 0.99 or 0.01 because in both cases v belongs to the set $V[h|\sigma_B]$ of remaining types at the subsequent history h. This can be rectified by accounting for the buyer's randomisation at histories prior to h. For example, if v accepts p with probability 0.99 at the initial history, the seller's equilibrium lifetime payoff against v given the subsequent history $h = (p) \in \mathcal{H}^1$ is a convex combination with weight 0.99 on the lifetime payoff he would have received if v accepted p at time 0 (equal to p), and weight 0.01 on the lifetime payoff conditional on the actions in h (in particular, type v's rejection of p at time 0) and the equilibrium strategies from h onwards. Defining lifetime payoffs in this manner preserves the dynamic consistency of the seller's preferences, allowing for analysis similar to the baseline model. The PRE characterisation is largely unaffected: the only change is that an indifferent type $v > \underline{v}$ may mix between accepting and rejecting, which has no effect on the seller's regret due to the continuum of types.

6.2 Finite number of types

When there are finitely many buyer types a PRE in pure strategies may not exist. The reasons are similar to the Bayesian model (Fudenberg and Tirole, 1991; **?**). If a high type v_H accepts a price with probability 1, there may be a sharp drop in price in the following period as the good is sold to lower types, which makes it suboptimal for v_H to accept in the first place. Hence, the buyer needs to mix in equilibrium. This, in turn, may require the seller to mix so that expected prices support the buyer's indifference. Introducing mixing along the lines of Section 6.1 above restores equilibrium existence in a two-period model with two types.[31] I conjecture that the analysis can be extended to the infinite-horizon setting.

6.3 Other Objectives

In this section I consider two alternative objectives for the seller: maxmin utility (Gilboa and Schmeidler, 1989) and smooth ambiguity preferences (Klibanoff, Marinacci, and Mukerji, 2005). I assume $\underline{v} > 0$ and show that in both cases the good is sold to all types within a finite number of periods. Hence, similarly to PBE, these alternative objectives cannot result in rationing, thereby distinguishing them from PRE (Proposition 6). Moreover, the number of periods it takes to sell the good is bounded uniformly in δ, so the Coase conjecture holds under both objectives.

Suppose the seller entertains all degenerate priors supported on $\mathcal{V}$ as in Remark 1. At each history he maximises utility under the worst-case belief that has not been contradicted by the buyer's equilibrium strategy. A maxmin equilibrium can then be defined as a strategy profile σ that satisfies $(E1)$ and

$$\sigma_S \in \underset{\sigma'_S \in \Sigma_S}{\operatorname{argmax}} \inf_{v \in V[h|\sigma_B]} u_S(v, h|\sigma'_S, \sigma_B) \text{ for all } h \in \mathcal{H}.$$

Note that any price accepted by $\underline{v}$ is also accepted by all other types. Type $\underline{v}$ is therefore the worst-case type at any history. Hence, the maxmin objective boils down to maximising the payoff against $\underline{v}$. A standard argument shows that $\underline{v}$ accepts all prices lower than $\underline{v}$ and rejects all higher prices. Hence, there exists a unique maxmin

[31] This example is available upon request.

equilibrium where the monopolist sells the good to all types immediately at a price of $\underline{v}$.

Hanany, Klibanoff, and Mukerji (2020) propose a solution concept for players with smooth ambiguity preferences (Klibanoff, Marinacci, and Mukerji, 2005). It can be adapted to my setting as follows. The seller has a finite set of priors $\pi_1, .., \pi_K$ supported on $\mathcal{V}$ and each of them admits a density function that is bounded away from zero and bounded above. There is a prior distribution μ_{h^0} over $\{\pi_1, .., \pi_K\}$. A strategy profile σ is sequentially optimal if it satisfies $(E1)$ and

$$\sigma_S \in \underset{\sigma'_S \in \Sigma_S}{\operatorname{argmax}} \sum_{\pi \in \Delta V} \phi(\mathbb{E}_\pi[u_S(v, h|\sigma'_S, \sigma_B)])\mu_h(\pi) \text{ for all } h \in \mathcal{H},$$

where ϕ is a strictly increasing concave function reflecting ambiguity aversion and μ_h is updated from μ_{h^0} using the smooth rule (Klibanoff, Marinacci, and Mukerji, 2005). The updated distribution satisfies $\mu_h(\pi) > 0$ only if π is the Bayes update of some prior π_k using σ_B.[32,33] Following the argument for Lemma 3 in Fudenberg, Levine, and Tirole (1985), there exists $\varepsilon_k > 0$ and a sufficiently large N_k independent of δ such that every N_k periods the highest type in the support of the Bayes update of π_k decreases by at least ε_k. (Otherwise, the bounded density of π_k would imply that the monopolist prefers to sell to all types at $\underline{v}$ immediately.) Since there are finitely many priors $\pi_1, .., \pi_K$, it follows that the good is sold within a finite number of periods independent of δ.

References

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

Auster, S., Y.-K. Che, and K. Mierendorff (2024): "Prolonged Learning and Hasty Stopping: The Wald Problem with Ambiguity," *American Economic Review*, 114, 426–461.

[32]If a history h is inconsistent with σ_B for any prior π_k, it is without loss of generality to set $\mu_h(\pi_{\underline{v}}) = 1$, where $\pi_{\underline{v}}(\underline{v}) = 1$.

[33]The remaining aspects of updating via the smooth rule are inessential for the argument.

Auster, S. and C. Kellner (2022): "Robust bidding and revenue in descending price auctions," *Journal of Economic Theory*, 199, 105072.

Ausubel, L. M. and R. J. Deneckere (1989): "Reputation in bargaining and durable goods monopoly," *Econometrica*, 57, 511–531.

Bergemann, D. and K. H. Schlag (2005): "Robust monopoly pricing: The case of regret," European University Institute Working Paper No.2005/10.

——— (2008): "Pricing without priors," *Journal of the European Economic Association*, 6, 560–569.

——— (2011): "Robust monopoly pricing," *Journal of Economic Theory*, 146, 2527–2543.

Beviá, C. and L. Corchón (2022): "Contests with dominant strategies," *Economic Theory*, 74, 1–19.

Board, S. and M. Pycia (2014): "Outside options and the failure of the Coase conjecture," *American Economic Review*, 104, 656–671.

Bose, S. and A. Daripa (2009): "A dynamic mechanism and surplus extraction under ambiguity," *Journal of Economic theory*, 144, 2084–2114.

Bose, S. and L. Renou (2014): "Mechanism design with ambiguous communication devices," *Econometrica*, 82, 1853–1872.

Brooks, B. and S. Du (2021): "Maxmin auction design with known expected values," Working Paper.

Brzustowski, T., A. Georgiadis-Harris, and B. Szentes (2023): "Smart contracts and the coase conjecture," *American Economic Review*, 113, 1334–1359.

Caldentey, R., Y. Liu, and I. Lobel (2017): "Intertemporal pricing under minimax regret," *Operations Research*, 65, 104–129.

Carrasco, V., V. F. Luz, N. Kos, M. Messner, P. Monteiro, and H. Moreira (2018): "Optimal selling mechanisms under moment conditions," *Journal of Economic Theory*, 177, 245–279.

Carrasco, V., V. F. Luz, P. K. Monteiro, and H. Moreira (2019): "Robust mechanisms: the curvature case," *Economic Theory*, 68, 203–222.

Carroll, G. (2017): "Robustness and separation in multidimensional screening," *Econometrica*, 85, 453–488.

——— (2019): "Robustness in mechanism design and contracting," *Annual Review of Economics*, 11, 139–166.

Cesa-Bianchi, N. and G. Lugosi (2006): *Prediction, learning, and games*, Cambridge university press.

Chang, D. and J. J. Lee (2022): "Price skimming: Commitment and delay in bargaining with outside option," *Journal of Economic Theory*, 205, 105528.

Che, E. (2022): "Robustly optimal auction design under mean constraints," in *Proceedings of the 23rd ACM Conference on Economics and Computation*, 153–181.

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

de Castro, L. and A. Galvao (2019): "Dynamic quantile models of rational behavior," *Econometrica*, 87, 1893–1939.

de Castro, L., A. Galvao, and D. da Siva Nunes (2025): "Dynamic economics with quantile preferences," *Theoretical Economics*, 20, 353–425.

Doval, L. and V. Skreta (2024): "Optimal mechanism for the sale of a durable good," *Theoretical Economics*, 19, 865–915.

Du, S. (2018): "Robust mechanisms under common valuation," *Econometrica*, 86, 1569–1588.

Ellis, A. (2018): "On dynamic consistency in ambiguous games," *Games and Economic Behavior*, 111, 241–249.

Epstein, L. G. and M. Schneider (2003): "Recursive multiple-priors," *Journal of Economic Theory*, 113, 1–31.

Fuchs, W. and A. Skrzypacz (2010): "Bargaining with arrival of new traders," *American Economic Review*, 100, 802–836.

Fudenberg, D., D. Levine, and J. Tirole (1985): "Infinite horizon models of bargaining with one-sided uncertainty," in *Game Theoretic Models of Bargaining*, Cambridge University Press, vol. 73, 79.

Fudenberg, D. and J. Tirole (1991): *Game Theory*, MIT Press.

Gilboa, I. and D. Schmeidler (1989): "Maxmin expected utility with non-unique prior," *Journal of mathematical economics*, 18, 141–153.

Groseclose, T. (2022): "The Coase Conjecture When the Monopolist and Customers Have Different Discount Rates," Working Paper.

Gul, F., H. Sonnenschein, and R. Wilson (1986): "Foundations of dynamic monopoly and the Coase conjecture," *Journal of Economic Theory*, 39, 155–190.

Guo, Y. and E. Schmaya (2023): "Regret-Minimizing Project Choice," *Econometrica*, 91, 1567–1593.

——— (2025): "Robust Monopoly Regulation," *American Economic Review*, 115, 559–634.

Halpern, J. Y. and R. Pass (2012): "Iterated regret minimization: A new solution concept," *Games and Economic Behavior*, 74, 184–207.

Hanany, E., P. Klibanoff, and S. Mukerji (2020): "Incomplete information games with ambiguity averse players," *American Economic Journal: Microeconomics*, 12, 135–187.

Hart, S. and A. Mas-Colell (2000): "A simple adaptive procedure leading to correlated equilibrium," *Econometrica*, 68, 1127–1150.

——— (2001): "A general class of adaptive strategies," *Journal of Economic Theory*, 98, 26–54.

Hayashi, T. (2009): "Stopping with anticipated regret," *Journal of Mathematical Economics*, 45, 479–490.

He, W. and J. Li (2022): "Correlation-robust auction design," *Journal of Economic Theory*, 200, 105403.

Hurwicz, L. and L. Shapiro (1978): "Incentive structures maximizing residual gain under incomplete information," *The Bell Journal of Economics*, 180–191.

Hyafil, N. and C. Boutilier (2004): "Regret minimizing equilibria and mechanisms for games with strict type uncertainty," in *Proceedings of the 20th conference on Uncertainty in artificial intelligence*, 268–277.

Kim, K. (2009): "The Coase conjecture with incomplete information on the monopolist's commitment," *Theoretical Economics*, 4, 17–44.

Klibanoff, P., M. Marinacci, and S. Mukerji (2005): "A smooth model of decision making under ambiguity," *Econometrica*, 73, 1849–1892.

Koçyiğit, Ç., D. Kuhn, and N. Rujeerapaiboon (2024): "Regret Minimization and Separation in Multi-Bidder, Multi-Item Auctions," *INFORMS Journal on Computing*.

Koçyiğit, Ç., N. Rujeerapaiboon, and D. Kuhn (2022): "Robust multidimensional pricing: Separation without regret," *Mathematical Programming*, 1–34.

Kostadinov, R. (2024): "Regret in Durable-Good Monopoly," Working Paper 2024-02, Department of Economics, McMaster University.

Kydland, F. E. and E. C. Prescott (1977): "Rules rather than discretion: The inconsistency of optimal plans," *Journal of political economy*, 85, 473–491.

Libgober, J. and X. Mu (2021): "Informational robustness in intertemporal pricing," *The Review of Economic Studies*, 88, 1224–1252.

Linhart, P. B. and R. Radner (1989): "Minimax-regret strategies for bargaining over several variables," *Journal of Economic Theory*, 48, 152–178.

Liu, C. (2024): "Robust Contracts with Exploration," *arXiv preprint arXiv:2212.00157*.

Malladi, S. (2020): "Searching in the Dark and Learning Where to Look," Working Paper.

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, 1930–1968.

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

Pahlke, M. (2022): “Dynamic consistency in incomplete information games with multiple priors,” *Games and Economic Behavior*, 133, 85–108.

Renou, L. and K. H. Schlag (2010): “Minimax regret and strategic uncertainty,” *Journal of Economic Theory*, 145, 264–286.

Savage, L. (1951): “The theory of statistical decision,” *Journal of the American Statistical association*, 46, 55–67.

Schlag, K. and A. Zapechelnyuk (2012): “On the impossibility of achieving no regrets in repeated games,” *Journal of Economic Behavior & Organization*, 81, 153–158.

——— (2024): “Compromise, Don’t Optimize: Generalizing Perfect Bayesian Equilibrium to Allow for Ambiguity,” *Journal of Political Economy Microeconomics*, 2.

Shafer, R. C. (2020): “Minimax regret and failure to converge to efficiency in large markets,” *Games and Economic Behavior*, 124, 281–287.

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*, 93, 355–371.

Strotz, R. (1955): “Myopia and Inconsistency in Dynamic Utility Maximization,” *Review of Economic Studies*, 23, 165–180.

Suzdaltsev, A. (2022a): “Distributionally robust pricing in independent private value auctions,” *Journal of Economic Theory*, 206, 105555.

——— (2022b): “Minimax-regret Contracts and Optimal Innovation Prizes,” Working Paper.

——— (2022c): "An Optimal Distributionally Robust Auction," Working Paper.

Zhang, W. (2022): "Auctioning multiple goods without priors," *arXiv preprint arXiv:2204.13726*.

——— (2023): "Correlation-Robust Optimal Auctions," Working Paper.

Appendix

Proof of Proposition 1

Let $B_t = V[h^t|\sigma_B]$ and $B_{t+1} = V[h^t, \sigma_S(h^t)|\sigma_B]$. Then

$$\begin{aligned} R(h^t|\sigma_S, \sigma_B, C) &= \sup_{v \in B_t} R(v, h^t|\sigma_S, \sigma_B, C) \\ &= \max\left\{ \sup_{v \in B_t \setminus B_{t+1}} R(v, h^t|\sigma_S, \sigma_B, C), \sup_{v \in B_{t+1}} R\left(v, (h^t, \sigma_S(h^t))|\sigma_S, \sigma_B, C\right) \right\} \\ &\leq \max\left\{ \sup_{v \in B_t \setminus B_{t+1}} R(v, h^t|\sigma'_S, \sigma_B, C), \sup_{v \in B_{t+1}} R\left(v, (h^t, \sigma_S(h^t))|\sigma'_S, \sigma_B, C\right) \right\} \\ &= R(h^t|\sigma'_S, \sigma_B, C), \end{aligned}$$

where the inequality follows from the hypothesis of the proposition because $R(v, h^t|\sigma'_S, \sigma_B, C) = R(v, h^t|\sigma_S, \sigma_B, C) = C(v) - \delta^t \sigma_S(h^t)$ for all $v \in B_t \setminus B_{t+1}$.

Proof of Proposition 2

The following notation for strategy profiles is adopted throughout the proof: $\sigma = (\sigma_S, \sigma_B), \sigma' = (\sigma'_S, \sigma_B), \hat{\sigma} = (\hat{\sigma}_S, \sigma_B), \hat{\sigma}^t = (\hat{\sigma}^t_S, \sigma_B)$.

It suffices to show the "if" direction. To this end, suppose that σ_S is optimal against one-shot deviations in the sense that $R(h|\sigma, C) \leq R(h|\sigma', C)$ for all $h \in \mathcal{H}$ and σ'_S with $\sigma'_S(h') = \sigma_S(h')$ for all $h' \neq h$. Suppose, towards a contradiction, that $R(h|\hat{\sigma}, C) < R(h|\sigma, C)$ for some time t, history $h \in \mathcal{H}^t$, and strategy $\hat{\sigma}_S$. Let $h^s = h^s[h|\hat{\sigma}_S]$ for all $s \geq t$, and let $T = \sup\{s \geq t | \hat{\sigma}_S(h^s) \neq \sigma_S(h^s)\}$ be the last time $\hat{\sigma}_S$ differs from σ_S on the path of $\hat{\sigma}_S$ from h. For each $s \geq t$ consider the strategy $\hat{\sigma}^s_S$ that follows $\hat{\sigma}_S$ at time s and earlier times, and follows σ after time s.

Case 1: Suppose $T < \infty$. Then

$$R(h^s|\hat{\sigma}^s, C) \geq R(h^s|\sigma, C) \quad \text{for all } t \leq s \leq T \tag{19}$$

because σ_S is optimal against one-shot deviations and $R(h^s|\hat{\sigma}^s, C) = R(h^s|\sigma', C)$ for the one-shot deviation σ'_S with $\sigma'_S(h^s) = \hat{\sigma}_S(h^s)$. Moreover,

$$R(h^s|\hat{\sigma}, C) \geq R(h^s|\sigma, C) \quad \text{for all } t \leq s \leq T. \tag{20}$$

For $s = T$ this follows from (19) because $R(h^T|\hat{\sigma}^T, C) = R(h^T|\hat{\sigma}, C)$ by definition of T. If $R(h^{s+1}|\hat{\sigma}, C) \geq R(h^{s+1}|\sigma, C)$, then Proposition 1 and (19) imply that

$$R(h^s|\hat{\sigma}, C) \geq R(h^s|\hat{\sigma}^s, C) \geq R(h^s|\sigma, C).$$

Hence, (20) holds in contradiction to $R(h|\hat{\sigma}, C) < R(h|\sigma, C)$, since $h = h^t$.

Case 2: Suppose $T = \infty$. If $\hat{\sigma}_S$ is followed from h onwards, then for each type v and time s one of the following holds:

- Type v buys the good prior to time s, so $R(v, h|\hat{\sigma}, C) = R(v, h|\hat{\sigma}^s, C)$.
- Type v does not buy the good, or buys after time s, so

$$R(v, h|\hat{\sigma}, C) \geq C(v) - \delta^s \sup \mathcal{P}$$
$$\text{and } R(v, h|\hat{\sigma}^s, C) \leq C(v) - \delta^s \min\{\inf \mathcal{P}, 0\}.$$

In both cases there exists a constant κ independent of v and s such that $R(v, h|\hat{\sigma}^s, C) \leq R(v, h|\hat{\sigma}, C) + \delta^s \kappa$. Since $R(h|\hat{\sigma}, C) < R(h|\sigma, C)$, it follows that $R(h|\hat{\sigma}^s, C) < R(h|\sigma, C)$ for sufficiently large s. Since $\hat{\sigma}^s_S$ follows σ_S after time s, the optimality of σ_S against one-shot deviations yields a contradiction by the argument for Case 1.

PCE in the two-period setting

Consider the two-period example from Section 2 and the parameters $\underline{v} = 4, \bar{v} = 9, \delta = 0.5$ considered therein. A PCE for the parameters is found as follows. At $t = 1$ each type $v \geq p$ accepts price p. Hence, the best continuation payoff the seller can obtain against v is $C_1(v) = v$. Therefore, the regret-minimising price $P_1(v)$ at a history at

$t = 1$ in state v solves

$$P_1(v) \in \operatorname*{argmin}_{p_1} \left\{C_1(v) - p_1, C_1(v_2)\mathbb{1}_{v_2 > \underline{v}}\right\}, \tag{21}$$

where $v_2 = p_1$. Note that, unlike in (7), p_1 is not discounted because regret is computed from the perspective of $t = 1$. It follows that

$$P_1(v) = \begin{cases} 4 & \text{if } v < 8 \\ \frac{v}{2} & \text{if } v \geq 8 \end{cases}$$

and regret is equalised across time (unless the good is sold to all types), similarly to PRE. At $t = 0$ the buyer anticipates that the seller will follow P_1 at $t = 0$. Similarly to the PRE example, it can be shown that the best payoff against type v is

$$C_0(v) = (1-\delta)v + \delta P_1(v) = \begin{cases} \frac{1}{2}v + 2 & \text{if } v < 8 \\ \frac{3}{4}v & \text{if } v \geq 8 \end{cases}.$$

Note that C_0 differs from C_1 because they represent the best *continuation* payoffs from $t = 0$ and $t = 1$ respectively instead of the best *lifetime* payoffs considered in PRE, which are always computed from time 0. The PCE characterisation is completed by the price p_0 at $t = 0$, which solves

$$p_0 \in \operatorname*{argmin}_{p} \max \left\{C_0(\bar{v}) - p, (C_0(v_1) - \delta P_1(v_1))\mathbb{1}_{v_1 > \underline{v}}, C_0(v_2)\mathbb{1}_{v_2 > \underline{v}}\right\}, \tag{22}$$

with $p = (1-\delta)v_1 + \delta P_1(v_1)$ and $v_2 = \max\{v_1/2, \underline{v}\}$. This is obtained similarly to (8), except for the inclusion of the regret $C_0(v_2)\mathbb{1}_{v_2 > \underline{v}}$ against types who never purchase the good. Unlike in PRE, this regret may not equal the regret $(C_0(v_1) - \delta P_1(v_1))\mathbb{1}_{v_1 > \underline{v}}$ against types who purchase at $t = 1$, despite the regret equalisation $C_1(v_1) - P_1(v_1) = C_1(v_2)$ observed earlier. Nevertheless, regret at $t = 1$ is higher for every v_1[34], so

$$p_0 \in \operatorname*{argmin}_{p} \max \left\{C(\bar{v}) - p, (C(v_1) - \delta P_1(v_1))\mathbb{1}_{v_1 > \underline{v}}\right\}.$$

[34] If $P_1(v_1) = \underline{v}$, then $C_0(v_1) - \delta P_1(v_1) \geq C(v_2)\mathbb{1}_{v_2 > \underline{v}}$ obtains trivially. Otherwise, $v > 8$, so $C_0(v_1) - \delta P_1(v_1) = \frac{v_1}{2} > \frac{v_1}{4} + 2 = C_0(\frac{v_1}{2})$.

This mirrors (8) from the PRE example, so $p_0 = f_0^{C_0}(\bar{v})$. The solution has $v_1 = 4.75, v_2 = 4, p_0 = 4.375, p_1 = 4$.

In equilibrium regret against types who buy at $t = 0$ equals regret against types who buy at $t = 1$, which is higher than regret against types who never buy the good. Since all types buy at $t = 0$ or $t = 1$, there is no multi-stage deviation that lowers the seller's regret at $t = 0$. However, if $\bar{v}$ is large enough some types will not buy the good and regret could be improved by increasing prices at $t = 0$ and $t = 1$ simultaneously. Such a multi-stage deviation, however, is not allowed in a PCE.

Proofs of Proposition 3 and Proposition 4

Proposition 3 follows from Proposition 4 applied to history h^0, noting that the existence of an optimal mechanism follows from the existence of a strategy for the buyer that is a best reponse to the mechanism (p_t^*) in the sense of $(E1)$. Such a strategy can be constructed using Lemma 4(i) below.

The proof of Proposition 4 requires some preliminary results. Lemma 2 below lists basic properties of the laws of motion (12), (13), and the prices in (14).

Lemma 2. *If $C \in \mathcal{C}$, then for all times $s \geq t$*

(i) f_t^C, $f_{s\to t}^C$, and P_t^C are increasing and continuous.

(ii) $0 < f_t^C(v) < v$ for all $v > \underline{v}$ and $f_t^C(\underline{v}) = \underline{v}$.

Proof. Consider the function $\psi : \mathcal{V} \to \mathbb{R}$ given by $\psi(v) = C(v) + \delta^t(1-\delta)v$. It is strictly increasing and continuous because $C \in \mathcal{C}$. It follows that $C(v) = \psi(f_t^C(v))$ if $C(v) \geq \psi(\underline{v})$ and $f_t^C(v) = \underline{v}$ if $C(v) \leq \psi(\underline{v})$. Since ψ is strictly increasing, it has a continuous and strictly increasing inverse $\psi^{-1} : [\psi(\underline{v}), \psi(\bar{v})] \to \mathbb{R}$. Define $\varphi : [C(\underline{v}), C(\bar{v})] \to \mathbb{R}$ as follows:

$$\varphi(x) = \begin{cases} \psi^{-1}(x) & \text{if } x \geq \psi(\underline{v}) \\ \underline{v} & \text{if } x < \psi(\underline{v}) \end{cases}.$$

Note that φ is well-defined because $C(\bar{v}) \leq \psi(\bar{v})$. It is increasing and continuous because ψ^{-1} is increasing and continuous. Hence, $f_t^C = \varphi \circ C$ is increasing and continuous. Since t is arbitrary, an inductive argument establishes that $f_{s\to t}^C$ is increasing and continuous for all $s \geq t$. Hence, the function $\alpha : \mathcal{V} \to \mathcal{V}^\infty$ given by

$v \mapsto (f^C_{s\to t}(v))_{s=t}^\infty$ is increasing and continuous. It follows that $P^C_t = \beta \circ \alpha$ is increasing and continuous, where $\beta : (v_s)_{s=t+1}^\infty \mapsto (1-\delta)\sum_{s=t}^\infty \delta^{s-t} v_{s+1}$.

Towards part (ii), suppose $v > \underline{v}$. If $f^C_t(v) = 0$, then $\underline{v} = 0$ and $C(v) < \psi(v')$ for all $v' > 0$, which implies $C(v) \leq C(0)$ in contradiction to the strict monotonicity of C. Hence, $f^C_t(v) > 0$. Since $C(v) < \psi(v)$, the continuity of ψ implies that $C(v) < \psi(\hat{v})$ for some $\hat{v} < v$. It follows from the monotonicity of ψ that $f^C_t(v) \leq \hat{v} < v$. Finally, $f^C_t(\underline{v}) = \underline{v}$ holds by definition. □

Lemma 3(i) below formalises the regret equalisation of the optimal mechanism. Lemma 3(ii) implies that the states in an optimal mechanism follow the law of motion.

Lemma 3. *The following hold for any time t, type $v \in \mathcal{V}$, and $C \in \mathcal{C}$:*

(i) $C(v) - \delta^t P^C_t(v) \leq C(f^C_{s-1\to t}(v)) - \delta^s P^C_s(f^C_{s-1\to t}(v))$ for all $s > t$. This holds at equality when $f^C_{s-1\to t}(v) > \underline{v}$.

(ii) Any decreasing sequence $(v_s)_{s\geq t} \in \mathcal{V}^\infty$ with $v_t = v$ and

$$\sup_{s\geq t: v_s > \underline{v} \text{ or } s=t} \quad C(v_s) - \delta^s(1-\delta)\sum_{k=0}^\infty \delta^k v_{s+k+1} \leq C(v) - \delta^t P^C_t(v) \tag{23}$$

has $v_s = f^C_{s-1\to t}(v)$ for all $s > t$.

Proof. Throughout the proof s refers to a time $s \geq t$ unless specified otherwise.

Part (i): The law of motion (13) implies that $f^C_{s\to t}(v) = f^C_s(f^C_{s-1\to t}(v))$. Hence,

$$P^C_s(f^C_{s-1\to t}(v)) = (1-\delta) f^C_{s\to t}(v) + \delta P^C_{s+1}(f^C_{s\to t}(v)). \tag{24}$$

Moreover, (12) and the continuity of C imply

$$C(f^C_{s-1\to t}(v)) \leq C(f^C_{s\to t}(v)) + \delta^s(1-\delta) f^C_{s\to t}(v), \tag{25}$$

and this holds at equality when $f^C_{s\to t}(v) > \underline{v}$. It follows from (24) and (25) that

$$C(f^C_{s-1\to t}(v)) - \delta^s P^C_s(f^C_{s-1\to t}(v)) \leq C(f^C_{s\to t}(v)) - \delta^{s+1} P^C_{s+1}(f^C_{s\to t}(v)),$$

and this holds at equality when $f^C_{s\to t}(v) > \underline{v}$. Part (i) now follows because $(f^C_{s-1\to t}(v))_{s\geq t}$ is decreasing and $f^C_{t-1\to t}(v) \equiv v$.

Part (ii): Consider any decreasing sequence $(v_s)_{s\geq t}$ with $v_t = v$ that satisfies (23). If $v = \underline{v}$, then the claim holds trivially. I henceforth assume $v > \underline{v}$. Let

$$p_s = (1-\delta)\sum_{k=0}^{\infty}\delta^k v_{s+k+1} \tag{26}$$

and $\Delta_s = v_s - f^C_{s-1\to t}(v)$. Hence, $p_s - P^C_s(f^C_{s-1\to t}(v)) = (1-\delta)\sum_{k=0}^{\infty}\delta^k\Delta_{s+k+1}$. Since $\Delta_t = 0$ by definition, part (ii) amounts to showing that $\Delta_s = 0$ for all $s > t$.

Case 1: Suppose there exists T such that $\Delta_s \leq 0$ for all $s > T$. Then $p_T \leq P^C_T(f^C_{T-1\to t}(v))$. Suppose, towards a contradiction, that $\Delta_T > 0$. Then the strict monotonicity of C implies that $C(v_T) > C(f^C_{T-1\to t}(v))$ and, consequently,

$$C(v_T) - \delta^T p_T > C(f^C_{T-1\to t}(v)) - \delta^T P^C_T(f^C_{T-1\to t}(v)) \geq C(v) - \delta^t P^C_t(v),$$

where the last inequality follows from part (i). This contradicts (23) because $\Delta_T > 0$ implies $v_T > \underline{v}$. Hence, $\Delta_T \leq 0$. An inductive argument establishes that $\Delta_s \leq 0$ for all s.

It is now possible to show that $\Delta_s = 0$ for all $s > t$. To see this, assume the contrary and consider the first time $s \geq t$ such that $\Delta_{s+1} < 0$. Since $\Delta_t = 0$, it follows that $\Delta_s = 0$ and, consequently $C(v_s) = C(f^C_{s-1\to t}(v))$. Moreover, $\Delta_{s+1} < 0$ implies $p_s < P^C_s(f^C_{s-1\to t}(v))$ because $\Delta_{s+k} \leq 0$ for all $k > 0$. Hence,

$$C(v_s) - \delta^s p_s > C(f^C_{s-1\to t}(v)) - \delta^s P^C_s(f^C_{s-1\to t}(v)).$$

Note that $v_s > \underline{v}$; otherwise $v_s = f^C_{s-1\to t}(v) = \underline{v}$, which implies $\Delta_{s+1} = 0$ because $\underline{v} \leq v_{s+1} \leq v_s = \underline{v}$ and $f^C_{s\to t}(v) = f^C_s(f^C_{s-1\to t}(v)) = f^C_s(\underline{v}) = \underline{v}$ by Lemma 2(ii). Hence, the above contradicts (23) due to part (i).

Case 2: Suppose that $\Delta_s > 0$ for infinitely many s. Note that (v_s) is decreasing by assumption and $(f^C_{s-1\to t}(v))$ is decreasing by definition, so both sequences converge. Hence, (Δ_s) converges.

Since $C \in \mathcal{C}$, there exists a lower bound λ on its slope. Let T satisfy $\lambda > 2\delta^T$ and let $s > T$ with $\Delta_s > 0$. It follows that

$$\lambda\Delta_s \leq C(v_s) - C(f^C_{s-1\to t}(v)) \leq \delta^s(p_s - P^C_s(f^C_{s-1\to t}(v))) \\ = \delta^s(1-\delta)\sum_{k=0}^{\infty}\delta^k\Delta_{s+k+1} < \lambda(1-\delta)\sum_{k=0}^{\infty}\delta^k\frac{\Delta_{s+k+1}}{2},$$

where the second inequality follows from (23) and part (i) because $v_s > \underline{v}$ due to $\Delta_s > 0$. Hence, $\Delta_{s+k_1} > 2\Delta_s$ for some $k_1 \geq 1$. The same argument can be applied to time $s + k_1$ to find a time $s + k_2 > s + k_1$ such that $\Delta_{s+k_2} > 2\Delta_{s+k_1}$. An inductive argument establishes that (Δ_s) is unbounded, contradicting its convergence. Hence, the assumption of case 1 holds, under which it was shown that $\Delta_s = 0$ for all $s > t$, as required. □

Lemma 4 characterises the buyer's best response to the continuation-optimal mechanism from any history.

Lemma 4. *Let $C \in \mathcal{C}$.*

(i) (σ, C) satisfies $(E1)$ iff

$$\sigma_B(v,h,p) = \begin{cases} A & \text{if } v - p > \sup_{s>t}\delta^{s-t}(v - \sigma_S(h^s[h,p|\sigma_S]) \\ R & \text{if } v - p < \sup_{s>t}\delta^{s-t}(v - \sigma_S(h^s[h,p|\sigma_S]) \end{cases}$$

for all $v \in \mathcal{V}, t, h \in \mathcal{H}^t, p \in \mathcal{P}$.

(ii) If (σ, C) satisfies $(E1)$ and $\sigma_S(h^s[h|\sigma_S]) = P^C_s(f^C_{s-1\to t}(v))$ for some t, $h \in \mathcal{H}^t$, and all $s \geq t$ where $v = \bar{V}[h|\sigma_B]$, then $\bar{V}[h^s[h|\sigma_S]|\sigma_B] = f^C_{s-1\to t}(v)$ for all $s \geq t$.

Proof. Suppose the seller offers p following $h \in \mathcal{H}^t$ and follows σ_S thereafter. Type v obtains $\delta^t(v-p)$ if he accepts p following h and $\delta^s(v - \sigma_S(h^s[h,p|\sigma_S])$ if he accepts at any subsequent time $s > t$. If he rejects all subsequent offers, he receives a payoff of 0, which equals the limit of $\delta^s(v - \sigma_S(h^s[h,p|\sigma_S])$ as $s \to \infty$ because $\mathcal{P}$ is bounded. The desired characterisation of $(E1)$ follows.

Towards the second part, type v prefers accepting $P^C_s(f^C_{s-1\to t}(v))$ at time s to accepting $P^C_{s+1}(f^C_{s\to t}(v))$ at time $s+1$ when $v - P^C_s(f^C_{s-1\to t}(v)) \geq \delta(v - P^C_{s+1}(f^C_{s\to t}(v)))$, which is equivalent to $v \geq f^C_{s\to t}(v)$ by (13) and (14). Type v prefers accepting at $s+1$ to accepting at s when $v \leq f^C_{s\to t}(v)$ and is indifferent when $v = f^C_{s\to t}(v)$.

Lemma 2(ii) implies that $f^C_{s\to t}(v) \geq f^C_{s+1\to t}(v)$. Hence, types $v \geq f^C_t(v)$ prefer accepting at t to any other strategy, and this preference is strict when $v > f^C_t(v)$. Hence, $\bar{V}[h^{t+1}[h|\sigma_S]|\sigma_B] = f^C_t(v)$. The rest of the claim follows inductively. □

Lemma 5 states that higher types purchase earlier than lower types, known as the "skimming property".

Lemma 5. *If (σ, C) satisfies $(E1)$ and $\sigma_B(v, h, p) = A$ for some $v \in \mathcal{V}$, $h \in \mathcal{H}$, $p \in \mathcal{P}$, then $\sigma_B(v', h, p) = A$ for all $v' > v$.*

Proof. Let $p_1, p_2, ...$ be the prices set by the seller following history h and price p. Since type v accepts, it follows that $v - p \geq \delta^t(v - p_t)$ for all $t \geq 1$, which is equivalent to $(1 - \delta^t)v \geq p - \delta^t p_t$. Hence, $(1 - \delta^t)v' \geq p - \delta^t p_t$ for any $v' > v$, so v' also accepts p following history h. □

To prove Proposition 4, consider $C \in \mathcal{C}$ and strategies σ such that (σ, C) satisfies $(E1)$. Fix a time t, a history $h \in \mathcal{H}^t$, and let $v = \bar{V}[h|\sigma_B]$. Let $h^s = h^s[h|\sigma_S]$, $v_s = \bar{V}[h^s|\sigma_B]$, and $p_s = \sigma_S(h^s)$ for all $s \geq t$, noting that $v_t = v$. Let $\mathcal{T}$ be the set of times $s \geq t$ when a sale occurs given that the players follow σ from h onwards, i.e.

$$\mathcal{T} = \{s \geq t | \sigma_B(v', h^s, p_s) = A \text{ for some } v' \in V[h^s|\sigma_B]\}.$$

Lemma 5 implies that $V[h^{s+1}|\sigma_B]$ equals either $[\underline{v}, v_{s+1}]$ or $[\underline{v}, v_{s+1})$ for all $s \in \mathcal{T}$.

If $\mathcal{T} = \emptyset$, then $R(h|\sigma, C) = C(v)$ by monotonicity and continuity of C because no type lower than v buys the good. Hence, (16) holds; it also holds at a strict inequality when $v > 0$ because Lemma 2(ii) implies $P^C_t(v) > 0$.

If $\mathcal{T} \neq \emptyset$, let $\bar{\tau} = \sup \mathcal{T}$. The following upper bound on prices applies at all times $s \in \mathcal{T}$:

$$p_s \leq (1 - \delta) \sum_{k=0}^{\infty} \delta^k v_{s+k+1}. \tag{27}$$

To see this, suppose first that $s = \bar{\tau}$ is finite. Hence, no further purchases occur, so $v_{s+1} = v_{s+k}$ for all $k > 1$. Since no type accepts a price higher than his value $p_s \leq v_{s+1} = (1 - \delta) \sum_{k=0}^{\infty} \delta^k v_{s+k+1}$, as required. Now suppose $s < \bar{\tau}$ and let s' be the first time in $\mathcal{T}$ after s. It follows that $v_{s+1} = .. = v_{s'}$. Hence,

$$p_s \leq (1 - \delta^{s'-s})v_{s+1} + \delta^{s'-s}p_{s'} = (1-\delta)\sum_{k=0}^{s'-s-1} \delta^k v_{s+k+1} + \delta^{s'-s}p_{s'}$$

because v_{s+1} or types arbitrarily close to it prefer to accept p_s than to wait until s' and accept $p_{s'}$. The desired upper bound (27) follows by applying the same argument to $p_{s'} \in \mathcal{T}$ and prices at subsequent times in $\mathcal{T}$. I now show that

$$\sup_{s \geq t: v_s > \underline{v} \text{ or } s=t} C(v_s) - \delta^s(1-\delta)\sum_{k=0}^{\infty} \delta^k v_{s+k+1} \leq R(h|\sigma, C). \tag{28}$$

This follows from the following three cases.

(i) If $s \in \mathcal{T}$, then v_s or an arbitrarily close type accepts at time s, so (27) implies

$$R(h|\sigma, C) \geq C(v_s) - \delta^s(1-\delta)p_s \geq C(v_s) - \delta^s(1-\delta)\sum_{k=0}^{\infty} \delta^k v_{s+k+1}. \tag{29}$$

(ii) If $t \leq s < \bar{\tau}$ and $s \notin \mathcal{T}$, let s' be the first time in $\mathcal{T}$ after s. Then $v_s = v_{s+1} = .. = v_{s'}$, so (29) evaluated at $s' \in \mathcal{T}$ implies

$$C(v_s) - \delta^s(1-\delta)\sum_{k=0}^{\infty} \delta^k v_{s+k+1} \leq C(v_s) - \delta^s(1-\delta)\sum_{k=s'-s}^{\infty} \delta^k v_{s+k+1}$$
$$= C(v_{s'}) - \delta^{s'}(1-\delta)\sum_{k=0}^{\infty} \delta^k v_{s'+k+1} \leq R(h|\sigma, C).$$

(iii) If $s > \bar{\tau}$ and $v_s > \underline{v}$, then $0 < v_s = v_{s+1} = v_{s+2}..$ and types lower than v_s never accept. Hence,

$$R(h|\sigma, C) \geq C(v_s) > C(v_s) - \delta^s(1-\delta)\sum_{k=0}^{\infty} \delta^k v_{s+k+1}.$$

To complete the proof that (16) holds when $\mathcal{T} \neq \emptyset$, I consider the following two cases.

(a) If $v_s \neq f^C_{s-1\to t}(v)$ for some $s > t$, then (28) and Lemma 3(ii) applied to $(v_s)_{s\geq t}$ imply that (16) holds at a strict inequality.

(b) If $v_s = f^C_{s-1\to t}(v)$ for all $s > t$, then

$$P^C_s(f^C_{s-1\to t}(v)) = (1-\delta)\sum_{k=0}^{\infty}\delta^k v_{s+k+1}. \tag{30}$$

It follows from Lemma 3(i) that the LHS of (28) equals $C(v) - P^C_t(v)$, so (16) holds.

Suppose that $v > 0$ and (16) holds at equality. Then $\mathcal{T} \neq \emptyset$ and case (b) above holds. If $v = \underline{v}$, then $P^C_t(v) = \underline{v}$. Since type $\underline{v}$ does not accept prices higher than his value and $\underline{v} > 0$, (16) at equality implies that $p_t = \underline{v}$. Hence, $p_s = P^C_s(f^C_{s-1\to t}(v))$ for all $s \geq t$ with $s = t$ or $f^C_{s-1\to t}(v) > \underline{v}$ (the latter does not hold for any $s > t$ by Lemma 2(ii)). If, instead, $v > \underline{v}$, then Lemma 2 implies that $v_s = f^C_{s-1\to t}(v) > f^C_{s\to t}(v) = v_{s+1}$ for all $s \geq t$ with $s = t$ or $f^C_{s-1\to t}(v) > \underline{v}$, so all these times belong to $\mathcal{T}$ and, therefore, satisfy $p_s \leq P^C_s(v_s)$ due to (27) and (30). If $p_s < P^C_s(v_s)$ for such a time s, then the fact that type v_s or an arbitrarily close type purchases at time s implies that

$$R(h|\sigma, C) > C(v_s) - \delta^s P^C_s(v_s) = C(v) - \delta^t P^C_t(v),$$

where the equality follows from Lemma 3(i). This contradicts (16) holding at equality. Hence, $p_s = P^C_s(v_s)$, as required.

Towards the last point, suppose $p_s = P^C_s(f^C_{s-1\to t}(v))$ for all $s \geq t$. It follows from Lemma 4(ii) that $v_s = f^C_{s-1\to t}(v)$ for all $s \geq t$. By Lemma 2(ii) $v_s > v_{s+1}$ for all $s \geq t$ such that $v_s > \underline{v}$. Moreover, if $v_{s+1} = \underline{v}$, then Lemma 2(ii) implies that $\underline{v} > 0$, so all types accept $p_s = \underline{v}$ at time s because future prices equal $\underline{v}$. Hence, sales occur at all times $s \geq t$ with $v_s > \underline{v}$ or $s = t$. It follows from the continuity of C that

$$\begin{aligned} R(h|\sigma, C) &= \sup_{s \geq t: v_s > \underline{v} \text{ or } s=t} \; \sup_{v' \in V[h^s|\sigma_B]: \sigma_B(v', h^s, p_s) = A} R(v', h^s|\sigma, C) \\ &= \sup_{s \geq t: f^C_{s-1\to t}(v) > \underline{v} \text{ or } s=t} C(f^C_{s-1\to t}(v)) - \delta^s P^C_s(f^C_{s-1\to t}(v)) = C(v) - \delta^t P^C_t(v), \end{aligned}$$

where the last equality follows from Lemma 3(i). Hence, (16) holds at equality.

Bounds on equilibrium outcomes

This section contains preliminary results providing a partial characterisation of equilibrium strategies and regret. Lemma 6 shows that the buyer always accepts prices below $\underline{v}$, so the seller offers at least $\underline{v}$.

Lemma 6. *If (σ, C) is a quasi-PRE and C is bounded, then $\sigma_S(h) \geq \underline{v}$ and $\sigma_B(v, h, p) = A$ for all $v \in \mathcal{V}, h \in \mathcal{H}, p < \underline{v}$.*

Proof. Let $\underline{p} = \inf_{h \in \mathcal{H}} \sigma_S(h)$. Since $\mathcal{P}$ is bounded, $\underline{p}$ is finite. Suppose $\underline{p} < \underline{v}$ towards a contradiction. Consider a history where the seller offers $p < (1-\delta)\underline{v} + \delta\underline{p} \in (\underline{p}, \underline{v})$. Since $\underline{v} - p > \delta(\underline{v} - \underline{p})$, Lemma 5 implies that all types accept p. However, the seller can decrease regret against all types by offering a slightly higher price, which would be accepted due to the strict inequality above. It follows that $\underline{p} \geq \underline{v}$, i.e. the seller never offers prices below $\underline{v}$. This in turn implies that $\underline{v}$ accepts any price $p < \underline{v}$ at any history, and so do all other types by Lemma 5. □

In what follows let $\phi(v-)$ denote the left limit of an increasing function $\phi : \mathcal{V} \to \mathbb{R}$, where $\phi(\underline{v}-) = \phi(\underline{v})$. I use $C-$ to denote the function with $C-(v) = C(v-)$ for all v and define $\bar{C}-$ analogously.

The following results require a generalisation of the laws of motion and prices that were defined only for $C \in \mathcal{C}$. The law of motion f_t^C for an increasing function C is defined as in (12) but the composite law of motion

$$f_{s \to t}^C(v) := f_s^{C-} \circ f_{s-1}^{C-} \circ \ldots \circ f_{t+1}^{C-} \circ f_t^C(v) \tag{31}$$

for $s \geq t$ is different from (13). The left limit $C-$ is used in place of C for state transitions from $t+1$ onwards because C may be discontinuous at a threshold type such as $f_t^C(v)$ who accepts at time t, so that types $[\underline{v}, f_t^C(v))$ remain at $t+1$. Finally, the prices P_t^C are defined analogously to (14). When $C \in \mathcal{C}$, the definitions of $f_t^C, f_{s \to t}^C, P_t^C$ align because $C- = C$, making (31) equivalent to (13).

Lemma 7 lists basic properties of increasing counterfactuals, similarly to Lemma 2.

Lemma 7. *If $C : \mathcal{V} \to \mathbb{R}$ is increasing, then for all times $s \geq t$*

(i) $f_t^C, f_{s \to t}^C$, *and* P_t^C *are increasing*

(ii) $f_t^C \leq f_{t+1}^C$

(iii) $f^C_{s\to t}(v-) = f^{C-}_{s\to t}(v)$ *for all* $v \in \mathcal{V}$

Proof. The monotonicity of f^C_t and f^{C-}_t follows from the argument used in the proof of Lemma 2(i) (note that $C-$ is increasing). It follows that $f^C_{s\to t}$ and P^C_t are increasing.

If $C(v) \geq C(v') + \delta^t(1-\delta)v'$ then $C(v) \geq C(v') + \delta^{t+1}(1-\delta)v'$. Hence, $f^C_t \leq f^C_{t+1}$.

Towards part (iii) it can be shown that

$$\begin{aligned} f^C_t(v-) &= \sup\{v' < v | C(\hat{v}) \geq C(v') + \delta^t(1-\delta)v' \text{ for some } \hat{v} < v\} \\ &= \sup\{v' < v | C(v-) \geq C(v'-) + \delta^t(1-\delta)v'\} = f^{C-}_t(v) \end{aligned} \tag{32}$$

for all $v > \underline{v}$ (recalling the convention $\sup\emptyset = \underline{v}$). It follows that

$$\begin{aligned} f^C_{s\to t}(v-) &= \sup_{\hat{v}<v} f^{C-}_s \circ ... \circ f^{C-}_{t+1} \circ f^C_t(\hat{v}) \\ &= f^{C-}_s \circ ... \circ f^{C-}_{t+1}\left(\sup_{\hat{v}<v} f^C_t(\hat{v})\right) = f^{C-}_s \circ ... \circ f^{C-}_{t+1} \circ f^{C-}_t(v) = f^{C-}_{s\to t}(v), \end{aligned}$$

where the second equality follows from the monotonicity of $f^C_t, f^C_{t+1}, .., f^C_s$ and the third equality follows from (32). If $v = \underline{v}$, then $f^C_{s\to t}(v-) \equiv f^C_{s\to t}(\underline{v}) = \underline{v} = f^{C-}_{s\to t}(v)$. $\square$

Lemma 8 below obtains bounds on acceptable prices and on the seller's regret. These bounds are later shown to be tight in any PRE (σ, C). Lemma 8 places minimal assumptions on the counterfactual C, so that the bounds can be used to establish necessary conditions for PRE in Lemma 12.

Lemma 8. *Let* (σ, C) *be a quasi-PRE such that* C *is bounded. Let* $\bar{C}(v) = \sup_{v'\leq v} C(v)$ *for all* v*. Then for any time* t *and history* $h \in \mathcal{H}^t$

(i) $\sigma_B(v, h, p) = A$ *for all* $v \in V[h|\sigma_B], p < (1-\delta)v + \delta P^{\bar{C}}_{t+1}(v)$

(ii) $R(h|\sigma, C) \leq \begin{cases} \bar{C}(v) - \delta^t P^{\bar{C}}_t(v) & \text{if } V[h|\sigma_B] = [\underline{v}, v] \\ \bar{C}(v-) - \delta^t P^{\bar{C}}_t(v-) & \text{if } V[h|\sigma_B] = [\underline{v}, v) \end{cases}$

Proof. Let $T \geq 0$ and define the functions

$$\hat{P}_t(v) = (1-\delta)\sum_{s=t}^{\infty} \delta^{s-t}\left(f^{\bar{C}}_{s\to t}(v)\mathbb{1}_{s<T} + \underline{v}\mathbb{1}_{s\geq T}\right)$$

for all t, where $\mathbb{1}$ denotes the indicator function. Note that $\hat{P}_t(v) \in \mathcal{V} \subseteq \mathcal{P}$ for all t, v because $f^{\bar{C}}_{s\to t}(v) \in \mathcal{V}$ for all $s \geq t$ by definition. Since $\bar{C}$ is increasing, each $\hat{P}_t$ is increasing by Lemma 7(i). Hence,

$$\begin{aligned}\hat{P}_t(v-) &= (1-\delta)\sum_{s=t}^{\infty}\delta^{s-t}\left(\lim_{v'\nearrow v} f^{\bar{C}}_{s\to t}(v')\mathbb{1}_{s<T} + \underline{v}\mathbb{1}_{s\geq T}\right)\\ &= (1-\delta)\sum_{s=t}^{\infty}\delta^{s-t}\left(f^{\bar{C}-}_{s\to t}(v)\mathbb{1}_{s<T} + \underline{v}\mathbb{1}_{s\geq T}\right)\end{aligned}$$

for all v by Lemma 7(iii). It follows that

$$\begin{aligned}\hat{P}_t(v) &= (1-\delta)\left[f^{\bar{C}}_t(v) + \delta\sum_{s=t+1}^{\infty}\delta^{s-t}\left(f^{\bar{C}-}_{s\to t+1}(f^{\bar{C}}_t(v))\mathbb{1}_{s<T} + \underline{v}\mathbb{1}_{s\geq T}\right)\right]\\ &= (1-\delta)f^{\bar{C}}_t(v) + \delta\hat{P}_{t+1}(f^{\bar{C}}_t(v)-) \end{aligned} \tag{33}$$

and, similarly,

$$\hat{P}_t(v-) = (1-\delta)f^{\bar{C}-}_t(v) + \delta\hat{P}_{t+1}(f^{\bar{C}-}_t(v)-) \tag{34}$$

for all $t < T$. I now show that the following statements hold for all $t \leq T$:

$(i)_t$ $\sigma_B(v, h, p) = A$ for all $v \in V[h|\sigma_B]$, $h \in \mathcal{H}^t$, $p < (1-\delta)v + \delta\hat{P}_{t+1}(v)$

$(ii)_t$ For all $h \in \mathcal{H}^t$ and $\varepsilon > 0$ there exists σ^ε_S such that

$$R(h|\sigma^\varepsilon_S, \sigma_B, C) \leq \begin{cases}\bar{C}(v) - \delta^t\hat{P}_t(v) + \varepsilon & \text{if } V[h|\sigma_B] = [\underline{v}, v]\\ \bar{C}(v-) - \delta^t\hat{P}_t(v-) + \varepsilon & \text{if } V[h|\sigma_B] = [\underline{v}, v)\end{cases}$$

$(iii)_t$ For all $h \in \mathcal{H}^t$

$$\sigma_S(h) \geq \begin{cases}\hat{P}_t(v) & \text{if } V[h|\sigma_B] = [\underline{v}, v]\\ \hat{P}_t(v-) & \text{if } V[h|\sigma_B] = [\underline{v}, v)\end{cases}$$

Lemma 8(i) follows from statement $(i)_t$ because $\hat{P}_t \to P^{\bar{C}}_t$ as $T \to \infty$. Similarly, Lemma 8(ii) follows from statement $(ii)_t$ by considering large T and small ε.

The proof is by induction. Lemma 6 states that the seller never prices below $\underline{v}$. This implies $(iii)_T$ because $\hat{P}_T = \underline{v}$ and $(i)_T$ because the buyer does not expect prices lower than $\underline{v}$. Statement $(ii)_T$ follows from $(i)_T$ because prices below $\underline{v}$ are accepted by all types.

For the inductive step suppose the statements hold for times $T, T-1, ...t+1$. To show $(i)_t$, fix a history $h \in \mathcal{H}^t$, a type $v \in V[h|\sigma_B]$, and price $p < (1-\delta)v + \delta\hat{P}_{t+1}(v)$. Suppose, towards a contradiction, that $\sigma_B(v, h, p) = R$. Let $s > t$ be the first time when type v buys the good following (h, p) given the strategies σ, and let $h' = h^s[h, p|\sigma_S]$.

If $s \leq T$, Lemma 5 implies that either $V[h'|\sigma_B] = [\underline{v}, v']$ for some $v' \geq v$ or $V[h'|\sigma_B] = [\underline{v}, v')$ for some $v' > v$. In both cases statement $(iii)_s$ implies that $\sigma_S(h') \geq \hat{P}_s(v)$ due to the monotonicity of $\hat{P}_s$. Hence, a sufficient condition for the suboptimality of the buyer's strategy in the sense of $(E1)$ is that

$$\delta^t(v-p) > \delta^s(v - \hat{P}_s(v)) \tag{35}$$

whenever $t < s \leq T$. If $s = t+1$, then (35) holds by definition of p. If (35) holds for $t < s < T$, then it also holds for $s+1$ because

$$\begin{aligned}\delta^t(v-p) > \delta^s(v - \hat{P}_s(v)) &= \delta^s \left(v - (1-\delta) f_s^{\bar{C}}(v) - \delta\hat{P}_{s+1}(f_s^{\bar{C}}(v)-)\right) \\ &\geq \delta^s \left(v - (1-\delta)v - \delta\hat{P}_{s+1}(v)\right) = \delta^{s+1}(v - \hat{P}_{s+1}(v)),\end{aligned}$$

where the first line follows from (33) and the second line follows from $f_s^{\bar{C}}(v) \leq v$ and the monotonicity of $\hat{P}_{s+1}$. This contradicts the optimality of the buyer's strategy.

Note that (35) evaluated at $s = T$ implies

$$\delta^t(v-p) > \delta^T(v - \hat{P}_T(v)) = \delta^T(v - \underline{v}). \tag{36}$$

Suppose the first time s when type v accepts following (h, p) has $s > T$. The buyer's strategy is again suboptimal because Lemma 6 implies the buyer's payoff is no larger than $\delta^s(v - \underline{v})$ and (36) implies that $\delta^t(v-p) > \delta^s(v - \underline{v})$. Finally, if type v rejects all offers following (h, p) he obtains a payoff 0, which is less than $\delta^t(v-p)$ by (36). Hence, $\sigma_B(v, h, p) = A$, completing the proof of statement $(i)_t$.

To prove statement $(ii)_t$ fix $\varepsilon > 0$ and $h \in \mathcal{H}^t$ such that $V[h|\sigma_B] = [\underline{v}, v]$. Suppose the seller offers $\hat{P}_t(v) - \varepsilon$ at time t and $\hat{P}_s(f_{s-1\to t}^{\bar{C}}(v)-) - \varepsilon$ at any time s such that $t < s \leq T$. Then statement $(i)_t$ implies that all types $v' \geq f_t^{\bar{C}}(v)$ accept at time t, since

$$\hat{P}_t(v) = (1-\delta) f_t^{\bar{C}}(v) + \delta \hat{P}_{t+1}(f_t^{\bar{C}}(v)-) \leq (1-\delta)v' + \delta \hat{P}_{t+1}(v')$$

by (33) and the monotonicity of $\hat{P}_{t+1}$. Similarly, it follows from statements $(i)_{t+1}, .., (i)_T$ and (34) that every type $v' \geq f_{s\to t}^{\bar{C}}(v) = f_s^{\bar{C}-}(f_{s-1\to t}^{\bar{C}}(v))$ accepts no later than time $s > t$. Let τ be the first time s with $t \leq s \leq T$ such that $f_{s\to t}^{\bar{C}}(v) = \underline{v}$. If such a time does not exist, set $\tau = T$. Since $\hat{P}_T = \underline{v}$, it follows from Lemma 6 that all types accept by τ, so the seller's regret is upper-bounded by

$$\max\left\{\bar{C}(v) - \delta^t \hat{P}_t(v), \max_{t<s\leq\tau} \bar{C}(f_{s-1\to t}^{\bar{C}}(v)-) - \delta^s \hat{P}_s(f_{s-1\to t}^{\bar{C}}(v)-)\right\} + \varepsilon. \tag{37}$$

Since $\bar{C}$ is increasing, the law of motion (12) implies that $\bar{C}(v) \geq \bar{C}(f_t^{\bar{C}}(v)-) + \delta^t(1-\delta) f_t^{\bar{C}}(v)$ when $f_t^{\bar{C}}(v) > \underline{v}$, i.e. $\tau > t$, so it follows from (33) that

$$\bar{C}(v) - \delta^t \hat{P}_t(v) \geq \bar{C}(f_t^{\bar{C}}(v)-) - \delta^{t+1} \hat{P}_{t+1}(f_t^{\bar{C}}(v)-). \tag{38}$$

Similarly, (12) applied to $\bar{C}-$ and (34) imply that,

$$\bar{C}(f_{s-1\to t}^{\bar{C}}(v)-) - \delta^s \hat{P}_s(f_{s-1\to t}^{\bar{C}}(v)-) \geq \bar{C}(f_{s\to t}^{\bar{C}}(v)-) - \delta^{s+1} \hat{P}_{s+1}(f_{s\to t}^{\bar{C}}(v)-) \tag{39}$$

for $t < s < \tau$. It follows from (37), (38), and (39) that $R(h|\sigma, C) \leq \bar{C}(v) - \delta^t \hat{P}_t(v) + \varepsilon$.

The case of $V[h|\sigma_B] = [\underline{v}, v)$ follows from a similar argument: Suppose the seller sets price $\hat{P}_s(f_{s-1\to t}^{\bar{C}-}(v)-) - \varepsilon$ at any time s with $t \leq s \leq T$. Then every type $v' \geq f_{s\to t}^{\bar{C}-}(v)$ accepts no later than time s. It follows from (34) that $R(h|\sigma, C) \leq \bar{C}(v-) - \delta^t \hat{P}_t(v-) + \varepsilon$. This completes the proof of statement $(ii)_t$.

Towards statement $(iii)_t$ suppose $\sigma_S(h) < \hat{P}_t(v)$ for some $h \in \mathcal{H}^t$ with $V[h|\sigma_B] = [\underline{v}, v]$. It follows from statement $(i)_t$, (33), and $f_t^{\bar{C}}(v) \leq v$ that type v accepts $\sigma_S(h)$ following h, so $R(h|\sigma, C) > \bar{C}(v) - \delta^t \hat{P}_t(v)$ in contradiction to $(ii)_t$ and $(E2)$. If $\sigma_S(h) < \hat{P}_t(v-)$ and $V[h|\sigma_B] = [\underline{v}, v)$, then $v > 0$, so $\bar{C}(v-) < \bar{C}(v-) + \delta^t(1-\delta)v$ and, hence, $f_t^{\bar{C}-}(v) < v$ by the left continuity of $\bar{C}-$. It follows from $(i)_t$ and (34) that types arbitrarily close to v accept, so the seller's regret is larger than $\bar{C}(v-) - \delta^t \hat{P}_t(v-)$ in contradiction to $(ii)_t$. This demonstrates $(iii)_t$, thereby completing the proof. $\square$

Proof of Lemma 1

At the initial history $\bar{V}[h^0|\sigma_B^C] = \bar{v}$, which pins down $\sigma_B^C(h^0)$. This, in turn, determines $\bar{V}[h^1|\sigma_B^C]$ for each $h^1 \in \mathcal{H}^1$, thereby determining $\sigma_B^C(h^1)$. The rest of the buyer's strategy σ_B^C and states $\bar{V}[\cdot|\sigma_B^C]$ are uniquely determined by induction. The seller's strategy is well-defined because $P_t^C(v) \in \mathcal{V} \subseteq \mathcal{P}$ for any t, v, which follows from (14) and $f_{s\to t}^C(v) \in \mathcal{V}$. It is unique because the states $\bar{V}[\cdot|\sigma_B^C]$ are unique.

Towards the remaining parts, note that $\sigma_B^C(v', h, P_t^C(v)) = A$ for all $v' \geq f_t^C(v)$ because $P_t^C(v) = (1-\delta)f_t^C(v) + \delta P_{t+1}^C(f_t^C(v)) \leq (1-\delta)v' + \delta P_{t+1}^C(\min\{v', v\})$ due to $f_t^C(v) \leq v$ and the monotonicity of P_{t+1}^C (Lemma 2(i)). Moreover, $\sigma_B^C(v', h, P_t^C(v)) = R$ for all $v' < f_t^C(v)$ because $P_t^C(v) > (1-\delta)v' + \delta P_{t+1}^C(v')$ and $v' < v$. It follows inductively that $\bar{V}[h^s|\sigma_B^C] = f_{s-1\to t}^C(v)$ for all $s \geq t$. Hence, $\sigma_S^C(h^s) = P_s^C(\bar{V}[h^s|\sigma_B^C]) = P_s^C(f_{s-1\to t}^C(v))$. It follows from the argument for the last part of Proposition 4 that $R(h|\sigma^C, C) = C(v) - \delta^t P_t^C(v)$.

Proof of Proposition 5

Here and in the proof of Theorem 1 it will be useful to work with the restriction $C_{\hat{V}}$ of a counterfactual C to an interval $\hat{V} = [\underline{v}, \hat{v}]$ or $\hat{V} = [\underline{v}, \hat{v})$. In a slight abuse of notation, the class $\mathcal{C}$ will be expanded to include every function $C_{\hat{V}}$ that satisfies the properties in Definition 4 except for the requirement that its domain is $\mathcal{V}$. The same convention applies to prices $P_t^{C_{\hat{V}}}$. Lemma 2, Lemma 3, and Lemma 7 apply to the expanded class $\mathcal{C}$.

Lemma 9 below pins down equilibrium prices and regret when the restriction of C to the remaining buyer types belongs to the (expanded) class $\mathcal{C}$.

Lemma 9. *Let (σ, C) be a quasi-PRE such that C is bounded. Fix t and $h \in \mathcal{H}^t$. Let $v^* = \bar{V}[h|\sigma_B]$ and $C^* = C-_{[\underline{v},v^*]}$. If $C_{V[h|\sigma_B]} \in \mathcal{C}$, then $R(h|\sigma, C) = C^*(v^*) - \delta^t P_t^{C^*}(v^*)$. Moreover, if $v^* > 0$, then $\sigma_S(h^s[h|\sigma_S]) = P_t^{C^*}(f_{s-1\to t}^{C^*}(v^*))$ for all $s \geq t$.*

Proof. Lemma 5 implies that such that $V[h|\sigma_B]$ equals either $[\underline{v}, v^*]$ or $[\underline{v}, v^*)$.

Suppose $V[h|\sigma_B] = [\underline{v}, v^*]$. Then $C_{[\underline{v},v^*]} \in \mathcal{C}$ is continuous, so $C^* = C-_{[\underline{v},v^*]} = C_{[\underline{v},v^*]} \in \mathcal{C}$. It follows that $C(v') = \sup_{v\leq v'} C(v)$ for all $v \leq v^*$, so Lemma 8(ii) implies that $R(h|\sigma, C) \leq C(v^*) - \delta^t P_t^C(v^*) = C^*(v^*) - \delta^t P_t^{C^*}(v^*)$.

Suppose $V[h|\sigma_B] = [\underline{v}, v^*)$. Then $C_{[\underline{v},v^*)}$ is continuous, so $C^*_{[\underline{v},v^*)} = C_{[\underline{v},v^*)} \in \mathcal{C}$. Since $C^*(v^*) = \lim_{v\nearrow v^*} C(v) = \lim_{v\nearrow v^*} C^*(v)$, it follows that $C^* \in \mathcal{C}$. Lemma 8(ii) implies that $R(h|\sigma, C) \leq C-(v^*) - \delta^t P_t^{C-}(v^*) = C^*(v^*) - \delta^t P_t^{C^*}(v^*)$.

In both cases (a) $C^* \in \mathcal{C}$, (b) $C^*_{V[h|\sigma_B]} = C_{V[h|\sigma_B]}$, and (c) $R(h|\sigma, C) \le C^*(v^*) - \delta^t P_t^{C^*}(v^*)$. Consider an extension of C^* to $\mathcal{V}$ with $C^*(v) = C^*(v^*) + v - v^*$ for all $v > v^*$. The extended C^* satisfies $C^* \in \mathcal{C}$ due to (a). Moreover, (σ, C^*) satisfies $(E1)$ because (σ, C) satisfies $(E1)$. Hence, Proposition 4 implies that $R(h|\sigma, C^*) \ge C^*(v^*) - \delta^t P_t^{C^*}(v^*)$. But $R(h|\sigma, C^*) = R(h|\sigma, C)$ due to (b), so (c) holds at equality, as required. If $v^* > 0$, it follows from Proposition 4 that $\sigma_S(h^s[h|\sigma_S]) = P_t^{C^*}(f^{C^*}_{s-1\to t}(v^*))$ for all $s \ge t$. □

The following lemma characterises the buyer's best response from a history h when the seller follows the continuation-optimal mechanism from h.

Lemma 10. *Let $C \in \mathcal{C}$. Suppose that $\sigma_S(h^s[h|\sigma_S]) = P_s^C(f^C_{s-1\to t}(\bar{V}[h|\sigma_B]))$ for any times $s \ge t$ and history $h \in \mathcal{H}^t$. Then (σ, C) satisfies $(E1)$ iff $\sigma_B(v, h, p) = \sigma_B^C(v, h, p)$ for all $v \in \mathcal{V}, t, h \in \mathcal{H}^t, p \in \mathcal{P}$ such that $v - p \ne \sup_{s>t} \delta^{s-t}(v - \sigma_S(h^s[h, p|\sigma_S])$.*

Proof. Fix a time t, history $h \in \mathcal{H}^t$, and price $p \in \mathcal{P}$. Let $v_t = \bar{V}[h|\sigma_B]$ and $v_{t+1} = \bar{V}[h, p|\sigma_B]$. For any $v, v' \in \mathcal{V}$ and $s > t$, let

$$p_s(v, v') = (1 - \delta^{s-t})v + \delta^{s-t} P_s^C(f^C_{s-1\to t+1}(v')).$$

Note that

$$\inf_{s>t} p_s(v, v') = p_{t+1}(v, v') \qquad \text{whenever } v \ge v' \tag{40}$$

because

$$\begin{aligned} p_s(v, v') &= (1 - \delta^{s-t})v + \delta^{s-t}[(1 - \delta) f^C_{s\to t+1}(v') + \delta P^C_{s+1}(f^C_{s\to t+1}(v'))] \\ &\le (1 - \delta^{s-t})v + \delta^{s-t}[(1 - \delta)v + \delta P^C_{s+1}(f^C_{s\to t+1}(v')) = p_{s+1}(v, v'). \end{aligned}$$

Lemma 4(i) implies that $(E1)$ is equivalent to

$$\sigma_B(v, h, p) = \begin{cases} A & \text{if } p < \inf_{s>t} p_s(v, v_{t+1}) \\ R & \text{if } p > \inf_{s>t} p_s(v, v_{t+1}) \end{cases} \tag{41}$$

for all $v \in \mathcal{V}, t, h \in \mathcal{H}^t, p \in \mathcal{P}$. On the other hand $\sigma_B^C(v, h, p) = A$ iff $p \le$

$p_{t+1}(v, \min\{v, v_t\})$. Consider the sets

$$V_A = \left\{v \middle| p \leq \inf_{s>t} p_s(v, v_{t+1})\right\}, \quad \hat{V}_A = \left\{v \middle| p \leq p_{t+1}(v, \min\{v, v_t\})\right\},$$
$$V_I = \left\{v \middle| p = \inf_{s>t} p_s(v, v_{t+1})\right\}, \quad \hat{V}_I = \left\{v \middle| p = p_{t+1}(v, \min\{v, v_t\})\right\}.$$

Lemma 2(i) implies that $p_s(\cdot, v')$ is strictly increasing and continuous for all $s > t$ and v'. Moreover, the sequence of functions $(p_s(\cdot, v'))_{s>t}$ converges uniformly to the identity function on $\mathcal{V}$ and is therefore equicontinuous. It follows that $\inf_{s>t} p_s(\cdot, v')$ is continuous for all v'. Hence, each of the sets V_A and $\hat{V}_A$ is either empty, or an interval $[v^*, \bar{v}]$ for some v^*, and $v \in V_I$ implies that $v = \inf V_A$. Moreover, $V_A = \hat{V}_A$ and $V_I = \hat{V}_I$ whenever

$$\inf_{s>t} p_s(v^*, v_{t+1}) = p_{t+1}(v^*, \min\{v^*, v_t\}), \tag{42}$$

where $v^* = \inf V_A$ or $v^* = \inf \hat{V}_A$ with $\inf \emptyset = \bar{v}$.

Towards the "only if" direction, suppose (σ, C) satisfies $(E1)$. Then $v_{t+1} = \min\{v^*, v_t\}$ by (41), where $v^* = \inf V_A$. Hence, $v^* \geq v_{t+1}$, so it follows from (40) and (42) that $V_A = \hat{V}_A$ and $V_I = \hat{V}_I$. It follows from (41) that type v accepts if $p < p_{t+1}(v, \min\{v, v_t\})$ and rejects if $p > p_{t+1}(v, \min\{v, v_t\})$. Hence, σ_B aligns with σ_B^C when p is offered following h, except for a type $v = v^*$ with $p = p_{t+1}(v, v_{t+1}) = v - \sup_{s>t} \delta^{s-t}(v - \sigma_S(h^s[h, p|\sigma_S]))$. The claim follows because t, h, p are arbitrary.

Towards the "if" direction, suppose $\sigma_B(v, h, p) = \sigma_B^C(v, h, p)$ for all $v \in \mathcal{V}$ such that $v - p \neq \sup_{s>t} \delta^{s-t}(v - \sigma_S(h^s[h, p|\sigma_S]))$, i.e. $p = \inf_{s>t} p_s(v, v_{t+1})$. Since V_I is singleton or empty, $v_{t+1} = \bar{V}[h, p|\sigma_B^C] = \min\{v^*, v_t\}$, where $v^* = \inf \hat{V}_A$. It follows from (40) and (42) that $V_A = \hat{V}_A$ and $V_I = \hat{V}_I$. Hence, (41) holds. Since t, h, p are arbitrary, it follows from Lemma 4(i) that (σ, C) satisfies $(E1)$. □

To prove Proposition 5, let (σ, C) be a quasi-PRE with $C \in \mathcal{C}$. Then C is bounded, so Lemma 9 and $\underline{v} > 0$ imply that $\sigma_S(h^s[h|\sigma_S]) = P_s^C(f_{s-1\to t}^C(\bar{V}[h|\sigma_B]))$ for all $s \geq t$ and $h \in \mathcal{H}^t$. It follows from Lemma 10 and Lemma 4(i) that σ_B equals σ_B^C except for one type at every history who is indifferent between accepting and rejecting. Hence, the buyer's payoff is uniquely determined. Lemma 9 also implies that $\sigma_S = \sigma_S^C$ and the seller's regret is $R(h|\sigma, C) = C(v) - \delta^t P_t^C(v)$ at any time-t history h in state v. It follows from Lemma 1 that the seller's regret is the same as the regret from the profile σ^C. Hence, $\sigma \overset{C}{=} \sigma^C$.

In the other direction, suppose $C \in \mathcal{C}$ and $\sigma \overset{C}{=} \sigma^C$. Lemma 1 implies that $\sigma_S(h^s[h|\sigma_S]) = P_s^C(f_{s-1\to t}^C(\bar{V}[h|\sigma_B]))$ for all $s \geq t$ and $h \in \mathcal{H}^t$. It follows from Lemma 10 and Lemma 4(i) that (σ, C) satisfies $(E1)$. It remains to show that the seller cannot profitably deviate from σ_S at any time-t history h in state v. By Proposition 2 it suffices to consider a one-shot deviation σ'_S, i.e. $\sigma'_S(h') = \sigma_S(h')$ for all $h' \neq h$. The buyer's incentives when facing σ'_S are the same as his incentives when facing σ_S following any history h' and price p except for histories h' preceding h. It follows from Lemma 4(i) that σ_B can be modified at these histories to produce a new strategy σ'_B such that $(\sigma'_S, \sigma'_B, C)$ satisfies $(E1)$. It follows from Proposition 4 that

$$R(h|\sigma'_S, \sigma_B, C) = R(h|\sigma'_S, \sigma'_B, C) \geq C(v) - \delta^t P_t^C(v).$$

On the other hand, $R(h|\sigma, C) = C(v) - \delta^t P_t^C(v)$ by Lemma 1. Hence, σ'_S is not a profitable deviation, thereby completing the proof.

Proof of Theorem 1

The proof uses the fact that equilibrium regret is increasing over time and in the state, as shown in Lemma 11 below.

Lemma 11. *The following hold for any time t, type $v^* \in \mathcal{V}$ and $C : [\underline{v}, v^*] \to \mathbb{R}$ with $C \in \mathcal{C}$:*

(i) $P_t^C \geq \delta P_{t+1}^C$ *and* $P_t^C(v) > \delta P_{t+1}^C(v)$ *for all* $v > 0$.

(ii) $C - \delta^t P_t^C$ *is strictly increasing.*

Proof. **Part (i):** If $v = \underline{v}$, then $P_t^C(v) = P_{t+1}^C(v) = \underline{v}$, so $P_t^C(v) > \delta P_{t+1}^C(v)$ if $v > 0$ and $P_t^C(v) = \delta P_{t+1}^C(v)$ if $v = 0$.

Suppose $v > \underline{v}$ and consider the sequence $(v_s)_{s\geq t}$ such that $v_t = v_{t+1} = v$ and $v_s = f_{s-1\to t+1}^C(v)$ for all $s > t+1$. Lemma 2(ii) implies that $v_{t+1} = v > f_t^C(v)$. It follows from the monotonicity of $f_{s-1\to t+1}^C$ (Lemma 2(i)) that

$$v_s = f_{s-1\to t+1}^C(v) \geq f_{s-1\to t+1}^C(f_t^C(v)) = f_{s-1\to t}^C(v)$$

for all $s > t+1$. Let $p_s = (1-\delta)\sum_{k=0}^{\infty}\delta^k v_{s+k+1}$. Hence,

$$p_t - P_t^C(v) = (1-\delta)\sum_{k=0}^{\infty}\delta^k(v_{t+k+1} - f_{t+k\to t}^C(v)) \geq 0$$

and, consequently, $C(v_t) - \delta^t p_t \leq C(v_t) - \delta^t P_t^C(v)$. The sequence $(v_s)_{s\geq t}$ is decreasing by Lemma 2(i) and satisfies $v_t = v$ and $v_{t+1} = v \neq f_t^C(v)$. It follows from Lemma 3(ii) that $C(v_s) - \delta^s p_s > C(v) - \delta^t P_t^C(v)$ for some $s > t$ with $v_s > \underline{v}$. On the other hand, Lemma 3(i) implies that

$$C(v) - \delta^{t+1}P_{t+1}^C(v) = C(f_{s-1\to t+1}^C(v)) - \delta^s P_s^C(f_{s-1\to t+1}^C(v)) = C(v_s) - \delta^s p_s,$$

so $\delta^t P_t^C(v) > \delta^{t+1}P_{t+1}^C(v)$, thereby completing the proof.

Part (ii): Suppose, towards a contradiction, that $\psi(v) := C(v) - \delta^t P_t^C(v)$ is not strictly increasing. Note that ψ is continuous because $C \in \mathcal{C}$ and P_t^C is continuous by Lemma 2(i). Hence, there exists $v \in (\underline{v}, v^*)$ such that for all $\varepsilon > 0$ there exists $v' \in (v, v+\varepsilon)$ with $\psi(v') \leq \psi(v)$. By Lemma 2(ii) there exists $v' > v$ such that $\psi(v') \leq \psi(v)$ and $f_t^C(v') < v$.

Let $v_t = v$ and $v_s = f_{s-1\to t}^C(v')$ for $s > t$. Note that $(v_s)_{s\geq t}$ is decreasing because $v_{t+1} = f_t^C(v') \leq v = v_t$. Then

$$\begin{aligned} C(v_s) - \delta^s(1-\delta)\sum_{k=0}^{\infty}\delta^k v_{s+k+1} &= C(v_s) - \delta^s P_s^C(f_{s-1\to t}^C(v')) \\ &\leq C(v') - \delta^t P_t^C(v') \leq C(v) - \delta^t P_t^C(v) \qquad (43) \end{aligned}$$

for all s such that $v_s > \underline{v}$ or $s = t$. The first inequality follows from $C(v_t) = C(v) < C(v')$ if $s = t$ and from Lemma 3(i) if $s > t$ and $v_s > \underline{v}$. The last inequality follows from $\psi(v') \leq \psi(v)$. It follows from (43) and Lemma 3(ii) that $f_{s-1\to t}^C(v') \equiv v_s = f_{s-1\to t}^C(v)$ for all $s > t$ and, consequently, $P_t^C(v) = P_t^C(v')$. This contradicts $\psi(v') \leq \psi(v)$ because $C(v') > C(v)$. Hence, ψ is strictly increasing, as required. □

Let $\mathcal{C}^*$ be the class of functions $C \in \mathcal{C}$ that satisfy the fixed-point property (18). (Recall that $\mathcal{C}$ is expanded to include functions with domain $[\underline{v}, v^*]$.) Theorem 1 follows from Lemma 12 and Lemma 15 below.

Lemma 12. *If $\underline{v} > 0$, then (σ, C) is a PRE iff $C : \mathcal{V} \to \mathbb{R}$, $C \in \mathcal{C}^*$, and $\sigma \stackrel{C}{=} \sigma^C$.*

Proof. **Proof of the "if" direction.** Let $C : \mathcal{V} \to \mathbb{R}$ satisfy $C \in \mathcal{C}^*$ and $\sigma \overset{C}{=} \sigma^C$. By Proposition 5 (σ, C) is a quasi-PRE. It remains to show that (σ, C) satisfies $(E3)$. Since $\sigma \overset{C}{=} \sigma^C$, any type v that has not accepted prior to history $h \in \mathcal{H}^t$ accepts prices below $(1-\delta)v + \delta P^C_{t+1}(v)$ at h and rejects higher prices. Hence,

$$\bar{u}_S(v|\sigma_B) = \sup_t \delta^t[(1-\delta)v + \delta P^C_{t+1}(v)].$$

It follows from Lemma 11(i) and $C \in \mathcal{C}$ that $\bar{u}_S(v|\sigma_B) = (1-\delta)v + P^C_1(v) = C(v)$, as required.

Proof of the "only if" direction. Let (σ, C) be a PRE. By Proposition 5 it suffices to show $C \in \mathcal{C}^*$. To this end, Lemma 6 implies that $C(\underline{v}) = \underline{v}$, so it suffices to show Lemma 13 and Lemma 14 below. □

Lemma 13. *If $\underline{v} > 0$, (σ, C) is a PRE, and $C_{[\underline{v}, v^*)} \in \mathcal{C}^*$, then $C_{[\underline{v}, v^*]} \in \mathcal{C}^*$.*

Lemma 14. *If $\underline{v} > 0$, (σ, C) is a PRE, and $C_{[\underline{v}, v^*]} \in \mathcal{C}^*$, then $C_{[\underline{v}, v^*+\varepsilon]} \in \mathcal{C}^*$ for some $\varepsilon > 0$.*

Note that C is bounded by $(E3)$ because the buyer never accepts prices higher than his value. Hence, any PRE (σ, C) satisfies the hypotheses of Lemma 6, Lemma 8, and Lemma 9, which are used in the proofs below.

Proof of Lemma 13. Fix a time t, a history $h^t \in \mathcal{H}^t$, and suppose v^* accepts price p following h^t, i.e. $\sigma_B(v^*, h^t, p) = A$. Let $h^{t+1} = (h^t, p)$ denote the subsequent history. By Lemma 5 $V[h^{t+1}|\sigma_B] \subseteq [\underline{v}, v^*)$. Since $C_{[\underline{v}, v^*)} \in \mathcal{C}^*$, it follows that $C^* := C-_{[\underline{v}, v^*]} \in \mathcal{C}^*$. Hence, Lemma 9 implies that $\sigma_S(h^{t+1}) \leq P^{C^*}_{t+1}(v^*)$ due to the monotonicity of $P^{C^*}_{t+1}$ (Lemma 2(i)). It follows from $\sigma_B(v^*, h^t, p) = A$ that $p \leq (1-\delta)v^* + \delta P^{C^*}_{t+1}(v^*)$. Hence,

$$C(v^*) = \bar{u}_S(v^*|\sigma_B) \leq \sup_t \delta^t[(1-\delta)v^* + \delta P^{C^*}_{t+1}(v^*)] = (1-\delta)v^* + \delta P^{C^*}_1(v^*) = C^*(v^*),$$

where the first equality holds by $(E3)$ and the second equality follows from Lemma 11(i). On the other hand, $\sup_{v' \leq v} C(v') = C^*(v)$ for all $v < v^*$, so Lemma 8(i) and Lemma 5 imply that v^* accepts any price below $(1-\delta)v + \delta P^{C^*}_1(v)$ at $t = 0$. Hence,

$$\bar{u}_S(v^*|\sigma_B) \geq \sup_{v < v^*} (1-\delta)v + \delta P^{C^*}_1(v) = (1-\delta)v^* + \delta P^{C^*}_1(v^*) = C^*(v^*)$$

by the monotonicity and continuity of $P_1^{C^*}$ (Lemma 2(i)). Hence, $C(v^*) = C^*(v^*)$, so $C_{[\underline{v},v^*]} = C^* \in \mathcal{C}^*$, as required. $\square$

Proof of Lemma 14. Let $\bar{C}(v) = \sup_{v' \leq v} C(v)$. It follows from $C_{[\underline{v},v^*]} \in \mathcal{C}^*$ that $\bar{C}_{[\underline{v},v^*]} = C_{[\underline{v},v^*]}$. Hence, Lemma 8(i) and Lemma 5 imply that any price below $(1-\delta)v^* + \delta P_1^C(v^*) = C(v^*)$ offered at time 0 is accepted by all types above v^*, so

$$\bar{C}(v) = \sup_{v^* < v' \leq v} C(v') \quad \text{for all } v > v^* \tag{44}$$

follows from $(E3)$. Consider $\varepsilon > 0$ and a time $T \geq 1$ such that the following hold:

$$\bar{C}(v^* + \varepsilon) < \bar{C}(v^*+) + \delta^T(1-\delta)^2\underline{v} \tag{45}$$

$$\delta^T(v^* + \varepsilon) < \underline{v}, \tag{46}$$

where $\bar{C}(v^*+)$ denotes the right limit of $\bar{C}$ at v^*. Consider a history $h \in \mathcal{H}^t$ such that $t \leq T$ and $V[h|\sigma_B] = [\underline{v}, v]$ where $v^* < v \leq v^* + \varepsilon$. Let $s + 1 = \inf\{\tau \geq t | \bar{V}[h^\tau[h|\sigma_S]|\sigma_B] \leq v^*\}$ be the first time the state is less than or equal to v^* following h. If s does not exist, then Lemma 5 implies there exists $v' > v^*$ such that following h all types in $(v^*, v']$ do not buy the good, so $R(h|\sigma, C) \geq \bar{C}(v^*+)$ by (44). This contradicts (45) because the seller can guarantee regret no larger than $\bar{C}(v^*+\varepsilon) - \delta^t\underline{v}$ at h by charging $\underline{v}$ (Lemma 6).

Let $h^s = h^s[h|\sigma_S], h^{s+1} = h^{s+1}[h|\sigma_S]$, and $v_{s+1} = \bar{V}[h^{s+1}|\sigma_B]$. Since $v_{s+1} \leq v^*$ and $\bar{C}_{[\underline{v},v^*]} = C_{[\underline{v},v^*]} \in \mathcal{C}^*$, it follows from Lemma 9 and $\underline{v} > 0$ that $\sigma_S(h^{s+1}) = P_{s+1}^{\bar{C}}(v_{s+1})$. It follows from the acceptance of price $\sigma_S(h^s)$ following h^s by type v_{s+1} (or an arbitrarily close type) that $\sigma_S(h^s) \leq (1-\delta)v_{s+1} + \delta P_{s+1}^{\bar{C}}(v_{s+1})$. If $v_{s+1} > \underline{v}$, then (44) implies that

$$\begin{aligned} R(h|\sigma, C) &\geq \max\{\bar{C}(v^*+) - \delta^s\sigma_S(h^s), R(h^{s+1}|\sigma, C)\} \\ &\geq \max\{\bar{C}(v^*+) - \delta^s[(1-\delta)v_{s+1} + \delta P_{s+1}^{\bar{C}}(v_{s+1})], \bar{C}(v_{s+1}) - \delta^{s+1}P_{s+1}^{\bar{C}}(v_{s+1})\}, \end{aligned} \tag{47}$$

where the second line follows from Lemma 9. If the seller, instead, offers a price lower than but arbitrarily close to $(1-\delta)v' + \delta P_{t+1}^{\bar{C}}(v')$ following h for some $v' \leq v^*$, then

v' accepts by Lemma 8(i). Hence,

$$R(h|\sigma, C) \leq \max\{\bar{C}(v) - \delta^t[(1-\delta)v' + \delta P^{\bar{C}}_{t+1}(v')], \bar{C}(v') - \delta^{t+1} P^{\bar{C}}_{t+1}(v')\}, \qquad (48)$$

where Lemma 8(ii) is used to upper bound the seller's continuation regret from h^{t+1}. Since $\underline{v} > 0$, Lemma 11(i) implies that $\delta^{t+1} P^{\bar{C}}_{t+1}(v_{s+1}) > \delta^{s+1} P^{\bar{C}}_{s+1}(v_{s+1})$ if $s > t$. It follows from (45), (47), and (48) with $v' = v_{s+1}$ that $s = t$, i.e. the monopolist sells to all types above v^* at h. Hence,

$$R(h|\sigma, C) = \max\{\bar{C}(v) - \delta^t[(1-\delta)v_{t+1} + \delta P^{\bar{C}}_{t+1}(v_{t+1})], \bar{C}(v_{t+1}) - \delta^{t+1} P^{\bar{C}}_{t+1}(v_{t+1})\}. \qquad (49)$$

Consider the two terms inside the maximum on the RHS of (48). Both terms are strictly increasing in v' because $P^{\bar{C}}_{t+1}$ is increasing (Lemma 7(i)) and $\bar{C} - \delta^{t+1} P^{\bar{C}}_{t+1}$ is strictly increasing on $[\underline{v}, v^*]$ where C equals $\bar{C}$ (Lemma 11(ii)). It follows from the law of motion (12) and the continuity of $\bar{C}_{[\underline{v}, v^*]}$ that the RHS of the regret guarantee (48) is minimised on $[\underline{v}, v^*]$ by $v' = f^{\bar{C}}_t(v)$. Hence, $v_{t+1} \leq v^*$ and (49) imply that $v_{t+1} = f^{\bar{C}}_t(v)$. If $\sigma_S(h) > (1-\delta)v_{t+1} + \delta P^{\bar{C}}_{t+1}(v_{t+1})$, then some types higher than v_{t+1} strictly prefer accepting $\sigma_S(h^{t+1}) = P^{\bar{C}}_{t+1}(v_{t+1})$ following h^{t+1} to accepting $\sigma_S(h)$ following h. Hence, $\sigma_S(h) \leq P^{\bar{C}}_t(v)$. The same conclusion holds when $v_{s+1} = \underline{v}$ because, similarly to (47) and (48), it can be shown that $\bar{C}(v) - \delta^t \underline{v} \geq R(h|\sigma, C) \geq \bar{C}(v^*+) - \delta^s \underline{v}$ and, therefore, $v_{t+1} = \underline{v}$ by (45). Similar arguments establish that $v_{t+1} = f^{\bar{C}-}_t(v)$ and $\sigma_S(h) \leq P^{\bar{C}-}_t(v) \leq P^{\bar{C}}_t(v)$ if $V[h|\sigma_B] = [\underline{v}, v)$, replacing $\bar{C}(v)$ with $\bar{C}(v-)$ in (48) and (49). Moreover, $\sigma_S(h) = P^{\bar{C}}_t(v)$ when $v \leq v^*$ due to Lemma 9 and $\underline{v} > 0$.

Consider the highest payoff the seller can obtain against type $v \in (v^*, v^* + \varepsilon]$. If v accepts a price $p > (1-\delta)v + \delta P^{\bar{C}}_{t+1}(v)$ following history $h \in \mathcal{H}^t$ with $t < T$, then the above analysis shows that the seller's next-period offer $\sigma_S(h, p)$ is no larger than $P^{\bar{C}}_{t+1}(v)$, since $\bar{V}(h, p|\sigma_B) \leq v$ by Lemma 5 and $P^{\bar{C}}_{t+1}$ is increasing by Lemma 7(i). This makes accepting p suboptimal for v. Hence, v rejects offers above $(1-\delta)v + \delta P^{\bar{C}}_{t+1}(v)$ following h and accepts lower offers by Lemma 8(i). It follows that

$$C(v) = \bar{u}_S(v|\sigma_B) = \max_{t<T}(1-\delta)v + \delta P^{\bar{C}}_{t+1}(v) = (1-\delta)v + \delta P^{\bar{C}}_1(v), \qquad (50)$$

where the first equality follows from ($E3$) and the second equality follows because it is not optimal to sell to v at time T or later due to (46) and Lemma 6. For the last equality it suffices to show that $P_t^{\bar{C}}(v) \geq \delta P_{t+1}^{\bar{C}}(v)$ for all $t < T$. Lemma 7(ii) and (45) imply that $f_t^{\bar{C}}(v) \leq f_{t+1}^{\bar{C}}(v) \leq v^*$. Suppose $f_t^{\bar{C}}(v) = f_{t+1}^{\bar{C}}(v)$. It follows that

$$P_t^{\bar{C}}(v) = (1-\delta)f_t^{\bar{C}}(v) + \delta P_{t+1}^{\bar{C}}(f_t^{\bar{C}}(v)) \geq \delta\left[(1-\delta)f_t^{\bar{C}}(v) + \delta P_{t+2}^{\bar{C}}(f_t^{\bar{C}}(v))\right] = \delta P_{t+1}^{\bar{C}}(v),$$

where the inequality follows from Lemma 11(i) because $\bar{C}_{[\underline{v},v^*]} \in \mathcal{C}$. Now suppose $f_t^{\bar{C}}(v) < f_{t+1}^{\bar{C}}(v)$. Let $\hat{C} : [\underline{v}, v] \to \mathbb{R}$ be the function that agrees with $\bar{C}$ on $[\underline{v}, f_{t+1}^{\bar{C}}(v)] \cup \{v\}$ and is linear elsewhere. The continuity of $\bar{C}$ on $[\underline{v}, v^*]$ and (12) imply that $\bar{C}(v) > \bar{C}(f_{t+1}^{\bar{C}}(v))$. Hence, $\hat{C} \in \mathcal{C}$. Since $f_t^{\bar{C}}(v) < f_{t+1}^{\bar{C}}(v)$ and $\hat{C}_{[\underline{v}, f_{t+1}^{\bar{C}}(v)]} = \bar{C}_{[\underline{v}, f_{t+1}^{\bar{C}}(v)]} \in \mathcal{C}$, it follows that $f_{s\to t}^{\hat{C}}(v) = f_{s\to t}^{\bar{C}}(v)$ for all $s \geq t$, so $P_t^{\hat{C}}(v) = P_t^{\bar{C}}(v)$. On the other hand, $f_{t+1}^{\hat{C}}(v) \geq f_{t+1}^{\bar{C}}(v)$ by construction, so Lemma 7(i) implies that $f_{s\to t+1}^{\hat{C}}(v) \geq f_{s\to t+1}^{\bar{C}}(v)$ for all $s \geq t+1$, so $P_{t+1}^{\hat{C}}(v) \geq P_{t+1}^{\bar{C}}(v)$. It follows from Lemma 11(i) applied to $\hat{C}$ that $P_t^{\bar{C}}(v) \geq \delta P_{t+1}^{\bar{C}}(v)$.

Equation (50) and the monotonicity of $P_1^{\bar{C}}$ imply that $C_{[\underline{v},v^*+\varepsilon]}$ is increasing and its slope is lower-bounded by $1-\delta$. Hence, $C_{[\underline{v},v^*+\varepsilon]} = \bar{C}_{[\underline{v},v^*+\varepsilon]}$, so (50) implies that

$$C(v) = (1-\delta)v + \delta P_1^C(v) = (1-\delta)v + \delta(1-\delta)f_1^C(v) + \delta^2 P_2^C(f_1^C(v)) \tag{51}$$

for all $v \in [\underline{v}, v^* + \varepsilon]$. To show $C_{[\underline{v},v^*+\varepsilon]} \in \mathcal{C}^*$ it remains to show that it is continuous. Since $f_1^C(v) \leq v^*$ for all $v \leq v^* + \varepsilon$ and P_2^C is continuous on $[\underline{v}, v^*]$, it suffices to show the continuity of f_1^C on $[\underline{v}, v^* + \varepsilon]$. To this end, the law of motion (12) implies that $C(v) \leq C(f_1^C(v)) + \delta(1-\delta)f_1^C(v)$, holding at equality if $f_1^C(v) > \underline{v}$. Hence, (51) implies that

$$(1-\delta)v \leq C(f_1^C(v)) - \delta^2 P_2^C(f_1^C(v)),$$

holding with equality if $f_1^C(v) > \underline{v}$. Lemma 2 and Lemma 11(ii) imply that $C - \delta^2 P_2^C$ is strictly increasing and continuous on $[\underline{v}, v^*]$. Hence, it has a continuous inverse $(C - \delta^2 P_2^C)^{-1}$ on $[C(\underline{v}) - \delta^2 P_2^C(\underline{v}), C(v^*) - \delta^2 P_2^C(v^*)]$. It follows that

$$f_1^C(v) = \begin{cases} (C - \delta^2 P_2^C)^{-1}((1-\delta)v) & \text{if } (1-\delta)v \geq C(\underline{v}) - \delta^2 P_2^C(\underline{v}) \\ \underline{v} & \text{if } (1-\delta)v \leq C(\underline{v}) - \delta^2 P_2^C(\underline{v}) \end{cases} \tag{52}$$

for all $v \in [\underline{v}, v^* + \varepsilon]$ and, hence, f_1^C is continuous on $[\underline{v}, v^* + \varepsilon]$, as required. $\square$

Lemma 15. *If $\underline{v} > 0$ there exists a unique function $C : \mathcal{V} \to \mathbb{R}$ with $C \in \mathcal{C}^*$.*

Proof. **Existence.** The function $C : \{\underline{v}\} \to \mathbb{R}$ given by $C(\underline{v}) = \underline{v}$ satisfies $C \in \mathcal{C}^*$. Suppose there exists $C : [\underline{v}, v^*) \to \mathbb{R}$ with $C \in \mathcal{C}^*$. Then C can be extended to $[\underline{v}, v^*]$ by setting $C(v^*) = C-(v^*)$. The resulting function $C : [\underline{v}, v^*] \to \mathbb{R}$ belongs to $\mathcal{C}^*$. Hence, it suffices to show that any $C : [\underline{v}, v^*] \to \mathbb{R}$ with $C \in \mathcal{C}^*$ and $C(\underline{v}) = \underline{v}$ can be extended to $[\underline{v}, v^* + \varepsilon]$, where $\varepsilon > 0$ is given by

$$C(v^*) = (1-\delta)(v^* + \varepsilon) + \delta^2 P_2^C(v^*). \tag{53}$$

It follows, analogously to (52) from the proof of Lemma 14 that the function $\phi : [\underline{v}, v^* + \varepsilon] \to [\underline{v}, v^*]$ given by

$$\phi(v) = \begin{cases} (C - \delta^2 P_2^C)^{-1}((1-\delta)v) & \text{if } (1-\delta)v \geq C(\underline{v}) - \delta^2 P_2^C(\underline{v}) \\ \underline{v} & \text{if } (1-\delta)v \leq C(\underline{v}) - \delta^2 P_2^C(\underline{v}) \end{cases}$$

is strictly increasing on $\phi^{-1}(\underline{v}, v^*]$ and continuous with $\phi(v^*) = f_1^C(v^*)$. Let $C^* : [\underline{v}, v^* + \varepsilon] \to \mathbb{R}$ satisfy $C^*(v) = C(v)$ for $v < v^*$ and

$$C^*(v) = \begin{cases} C(\phi(v)) + \delta(1-\delta)\phi(v) & \text{if } (1-\delta)v \geq C(\underline{v}) - \delta^2 P_2^C(\underline{v}) \\ (1-\delta)v + \delta\underline{v} & \text{if } (1-\delta)v \leq C(\underline{v}) - \delta^2 P_2^C(\underline{v}) \end{cases}$$

for $v \in [v^*, v^* + \varepsilon]$. If $\phi(v^*) > \underline{v}$, then $\phi(v^*) = f_1^C(v^*)$ implies that $C^*(v^*) = C(f_1^C(v^*)) + \delta(1-\delta)f_1^C(v^*) = C(v^*)$. If $\phi(v^*) = \underline{v}$, then $C^*(v^*) = (1-\delta)v^* + \delta\underline{v} = (1-\delta)v^* + \delta P_1^C(v^*) = C(v^*)$, since $C \in \mathcal{C}^*$. Hence, C^* is continuous at v^*. Moreover, if $(1-\delta)v = C(\underline{v}) - \delta^2 P_2^C(\underline{v})$, then $\phi(v) = \underline{v}$ and $C(\underline{v}) + \delta(1-\delta)\underline{v} = C(\underline{v}) + \delta P_1^C(\underline{v}) - \delta^2 P_2^C(\underline{v}) = (1-\delta)v + \delta\underline{v}$. Hence, the continuity and monotonicity of ϕ imply that C^* is continuous and increasing.

If $\phi(v) = \underline{v}$, then $C^*(v) = (1-\delta)v + \delta\underline{v} \leq C(\underline{v}) - \delta^2 P_2^C(\underline{v}) + \delta\underline{v} = C(\underline{v}) + \delta(1-\delta)\underline{v}$. Hence, (53) and the monotonicity of C^* imply that

$$\begin{aligned} \phi(v) &= \sup\{v' < v^* | C^*(v) \geq C(v') + \delta(1-\delta)v'\} \\ &= \sup\{v' < v^* | C^*(v) \geq C^*(v') + \delta(1-\delta)v'\} = f_1^{C^*}(v). \end{aligned}$$

If $\phi(v) = \underline{v}$, then $C^*(v) = (1-\delta)v + \delta P_1^C(v)$ follows from $P_1^C(v) = \underline{v}$. Otherwise, $C^*(v) = C(\phi(v)) - \delta^2 P_2^C(\phi(v)) + \delta(1-\delta)\phi(v) + \delta^2 P_2^C(\phi(v)) = (1-\delta)v + \delta P_1^{C^*}(v)$. Hence, C^* satisfies the fixed-point property (18) and its slope is lower-bounded by

$(1-\delta)$. It follows that $C^* \in \mathcal{C}^*$, as required.

Uniqueness. Suppose $C : \mathcal{V} \to \mathbb{R}$ and $C^* : \mathcal{V} \to \mathbb{R}$ with $C, C^* \in \mathcal{C}^*$ and $C \neq C^*$. Let $v^* = \sup\{v | C(v) = C^*(v)\}$. Since C and C^* are continuous and $C(\underline{v}) = C^*(\underline{v}) = \underline{v}$ by the fixed-point property (18), it follows that $C_{[\underline{v},v^*]} = C^*_{[\underline{v},v^*]}$. Lemma 2 implies that f_1^C and $f_1^{C^*}$ are continuous, $f_1^C(v^*) < v^*$, and $f_1^{C^*}(v^*) < v^*$. Hence, there exists $v > v^*$ such that $f_1^C(v) \leq v^*$, $f_1^{C^*}(v) \leq v^*$, and $C(v) \neq C^*(v)$, i.e. $f_1^C(v) \neq f_1^{C^*}(v)$. Assume $f_1^C(v) < f_1^{C^*}(v)$ without loss of generality. Then Lemma 3(i) implies that

$$C(v) - \delta P_1^C(v) \leq C(f_1^C(v)) - \delta^2 P_2^C(f_1^C(v)) \tag{54}$$

$$C^*(v) - \delta P_1^{C^*}(v) = C(f_1^{C^*}(v)) - \delta^2 P_2^C(f_1^{C^*}(v)). \tag{55}$$

The LHS of (54) and (55) equal $(1-\delta)v$ by the fixed-point property (18) because $C, C^* \in \mathcal{C}^*$. On the other hand, the RHS of (54) is smaller than the RHS of (55) because $C - \delta^2 P_2^C$ is strictly increasing by Lemma 11(ii). The resulting contradiction completes the proof. □

Proof of Proposition 6

The rate of decrease of $v_t = f_{t\to 0}^C(\bar{v})$ can be bounded as follows:

$$\begin{aligned}\delta^t(1-\delta)v_{t+1} \geq C(v_t) - C(v_{t+1}) &= (1-\delta)(v_t - v_{t+1}) + \delta(P_1^C(v_t) - P_1^C(v_{t+1})) \\ &\geq (1-\delta)(v_t - v_{t+1}),\end{aligned}$$

where the first inequality follows from the law of motion (12) and the continuity of C, the equality follows from the fixed-point property (18) due to Lemma 12, and the last inequality follows from the monotonicity of P_1^C (Lemma 2(i)). Hence, $v_{t+1} \geq v_t/(1 + \delta^t)$ for all t. The convergence of the series $\sum_{t=0}^\infty \ln(1+\delta^t)$ implies that $\prod_{t=0}^\infty 1/(1+\delta^t) = 1/N$ for some $N > 0$. Hence, $\lim_{t\to\infty} v_t \geq \bar{v}/N > \underline{v}$ whenever $\bar{v}/\underline{v} > N$, as required.

Proof of Proposition 7

For any times t and $s \geq t$ and $C^* : [\alpha\underline{v}, \beta\bar{v}] \to \mathbb{R}$, define $\hat{f}_t^{C^*}$, $\hat{f}_{s\to t}^{C^*}$, and $\hat{P}_t^{C^*}$ analogously to (12), (13), and (14) for the game with $\mathcal{V} = [\alpha\underline{v}, \beta\bar{v}]$.

Towards part (i), let $\alpha = \beta$ and $C^* : [\alpha\underline{v}, \beta\bar{v}] \to \mathbb{R}$ with $C^*(\beta v) = \beta C(v)$ for all

$v \in \mathcal{V}$. It follows that

$$\hat{f}_t^{C^*}(\beta v) = \sup\{\beta v' | C^*(\beta v) \geq C^*(\beta v') + \delta^t(1-\delta)\beta v'\}$$
$$= \beta \sup\{v' | C(v) \geq C(v') + \delta^t(1-\delta)v'\} = \beta f_t^C(v)$$

for all t and $v \in \mathcal{V}$. Hence, $\hat{f}_{s\to t}^{C^*}(\beta v) = \beta f_{s\to t}^C(v)$ for all $s \geq t$ and, consequently, $\hat{P}_t^{C^*}(\beta v) = \beta P_t^C(v)$. Since C satisfies the fixed-point property (18) by Lemma 12, it follows that

$$C^*(\beta v) = \beta C(v) = \beta[(1-\delta)v + \delta P_1^C(v)] = (1-\delta)\beta v + \delta \hat{P}_1^{C^*}(\beta v)$$

for all $v \in \mathcal{V}$. Hence, C^* satisfies (18) for the game with $\mathcal{V} = [\alpha \underline{v}, \beta \bar{v}]$. It follows from Lemma 12 and Lemma 15 that $\hat{C} = C^*$. Hence, $\hat{v}_t = \hat{f}_{t-1\to 0}^{\hat{C}}(\beta \bar{v}) = \beta f_{t-1\to 0}^C(\bar{v}) = \beta v_t$ and $\hat{p}_t = \hat{P}_t^{\hat{C}}(\hat{v}_t) = \beta P_t^C(v_t) = \beta p_t$, as required.

Towards part (ii), let $\alpha = 1 < \beta$. Lemma 15 implies that $\hat{C}_{[\underline{v},\bar{v}]} = C$. Hence, $v_t = f_{t-1\to 0}^C(\bar{v}) = \hat{f}_{t-1\to 0}^{\hat{C}}(\bar{v}) \leq \hat{f}_{t-1\to 0}^{\hat{C}}(\beta \bar{v}) = \hat{v}_t$ due to the monotonicity of $\hat{f}_{t-1\to 0}^{\hat{C}}$ (Lemma 2(i)). It follows from the monotonicity of $\hat{P}_t^{\hat{C}}$ that $P_t^C(v_t) = \hat{P}_t^{\hat{C}}(v_t) \leq \hat{P}_t^{\hat{C}}(\hat{v}_t)$, as required.

Proof of Theorem 2

Let (σ, C) be a PRE. By Lemma 12 $\sigma \overset{C}{=} \sigma^C$ and $C \in \mathcal{C}^*$. Hence,

$$R(h^0|\sigma, C) = C(\bar{v}) - P_0^C(\bar{v}) \leq C(\bar{v}) - \delta P_1^C(\bar{v}) = (1-\delta)\bar{v},$$

where the sequence of (in)equalities follows from Lemma 1, Lemma 11(i), and the fixed-point property (18) due to $C \in \mathcal{C}^*$, respectively. When δ is sufficiently high $R(h^0|\sigma, C) < \underline{v}$, so the good is sold to all types in equilibrium; otherwise, the seller's regret is at least $C(\underline{v})$ and $C(\underline{v}) = \underline{v}$ by (18). Let $T = \inf\{s | \sigma_B(\underline{v}, h^s[h^0|\sigma_S], \sigma_S(h^s[h^0|\sigma_S])) = A\}$ be the first time when $\underline{v}$ accepts on the equilibrium path and $h^T = h^T[h^0|\sigma_S]$. Since $P_T^C \geq \underline{v}$ and no type accepts a price higher than his value, $\sigma_S(h^T) = \underline{v}$. Hence,

$$(1-\delta)\bar{v} \geq R(h^0|\sigma, C) \geq R(h^T|\sigma, C) \geq C(\underline{v}) - \delta^T \sigma_S(h^T) = (1-\delta^T)\underline{v}$$
$$\bar{v} \geq (1 + \delta + .. + \delta^{T-1})\underline{v}.$$

Hence, for sufficiently high δ the good is sold to all types in no more than $N = \lceil \bar{v}/\underline{v} \rceil$ periods. It follows from $\sigma_S = \sigma_S^C$, Lemma 1, (14), and $f_{t-1\to 0}^C(\bar{v}) = \underline{v}$ for all $t > N$ that all prices offered on the equilibrium path are bounded above by $(1-\delta^N)\bar{v} + \delta^N \underline{v}$ for sufficiently high δ and, hence, converge uniformly to $\underline{v}$ as $\delta \to 1$.

Proof of Proposition 8

Let $C(v) = \alpha v$ for $\alpha \in (0,1)$. It follows that $C \in \mathcal{C}$. The arguments for the "if" directions of Proposition 5 and Lemma 12 translate directly to the case of $\underline{v} = 0$. Hence, (σ^C, C) is a PRE if C satisfies the fixed-point property (18). To this end, the law of motion (12) reduces to $f_t^C(v) = \frac{\alpha}{\alpha+\delta^t(1-\delta)} v$, so $P_t^C(v) = v \prod_{s=t}^{\infty} \frac{\alpha}{\alpha+\delta^s(1-\delta)}$. Hence, (18) holds iff $\varphi(\alpha) = 0$ where $\varphi : [0,1] \to \mathbb{R}$ is given by

$$\varphi(\alpha) = \alpha - (1-\delta) - \delta \prod_{t=1}^{\infty} \frac{\alpha}{\alpha + \delta^t(1-\delta)}.$$

Since φ is continuous, $\varphi(0) < 0$, and $\varphi(1) > 0$, it follows from the Intermediate Value Theorem that there exists $\alpha \in (0,1)$ such that C satisfies (18), as required.