Coasian Dynamics and Endogenous Learning

Tuomas Laiho and Julia Salmi*

March 31, 2020

Abstract

We study social learning in markets where a monopolistic seller and potential buyers learn from past buyers' reviews and experiences. Buyers' ability to time their purchases allows them to free-ride on information from other buyers. Our main result is that stationary equilibria are inefficient because of too little and too slow sales. We also show that the seller's commitment power may improve welfare. In addition, we analyze how good news and bad news learning affect the results.

JEL classification: D42, D83

Keywords: Social learning, experimentation, Coase conjecture, durable goods

1 Introduction

In many markets, the seller and the potential buyers of a product get valuable information from the experiences of past buyers and reviewers. Learning from

*Laiho: Department of Economics, University of Oslo, Moltke Moes vei 31, 0851 Oslo, Norway, tuomasl@econ.uio.no. Salmi: Department of Economics, University College London, j.salmi@ucl.ac.uk. We are very grateful to Pauli Murto and Juuso Välimäki for their support and encouragement throughout this project. We also thank Dino Gerardi, Bård Harstad, Daniel Hauser, Johannes Hörner, Klaus Kultti, Konrad Mierendorff, Juuso Toikka, Andrey Zhukov and seminar audiences at the AMFEA 2018, EARIE 2018, ESWM 2018, Helsinki GSE, NHH, SAET 2019, UCL and University of Oslo for their helpful comments. Salmi acknowledges financial support from Emil Aaltonen Foundation, OP Group Research Foundation, Yrjö Jahnsson Foundation, and the ERC grant no. 682417 through PI Vasiliki Skreta.

past outcomes is particularly important when buyers have the ability to time their purchases as is the case with durable or experience goods. The response from past consumers and reviewers plays a critical role in how much and how long a product is sold.[1] Social media and many different websites and apps for reviewing products have arguably made this dissemination of information more prominent than ever before. So far the existing literature has neglected this *endogenous* source of information when buyers are forward-looking and can wait for more information.

In this paper, we study pricing in a market where both a monopolistic seller and buyers learn product quality over time as buyers purchase the product. We ask how endogenous learning and buyers' ability to wait for more information affect the profit maximizing schedule of sales. Do the particulars of the learning process matter for the sales path? How does the seller's commitment power affect efficiency? To answer these questions, we solve the seller's profit maximizing schedule of sales with and without commitment power.

Buyers' dynamic incentives are an important aspect of the seller's problem. Firstly, buyers may want to postpone their purchases in order to free-ride on the information generated by other buyers. Secondly, if buyers expect prices to be low in the future they are unwilling to pay a high price today. The second part is why Coasian dynamics are central to the problem: if the seller is unable to resist future price decreases, she cannot capture the value of information from the buyers. We show that the combination of endogenous learning and forward-looking buyers distorts the seller's entry and exit decisions, and – surprisingly – causes *inefficient* delays. We also show that these inefficiencies can be alleviated if the seller can commit to the schedule of sales.

In our model, a monopolist is selling a product of uncertain quality to a fixed mass of buyers who have unit demand and decide when to purchase. Buyers' valuation for the product depends on their private type and on the unknown quality of the product. We model learning about product quality by assuming

[1] In the realm of experience goods, plays, musicals or movies often have their runs shortened or cancelled altogether if initial reviews are unfavorable or if the show does not garner enough demand. For example, between 1996-2000 more than half of the shows on Broadway were shown less than 10 times (see Simonoff and Ma (2003)). There are also many notable durable goods which have failed despite high hopes such as Samsung's Note 7, Apple's Newton or AT&T's Picturephone.

that a perfectly revealing signal (news) arrives according to a Poisson arrival rate that is proportional to the current rate of sales. Pure good news and pure bad news learning are special cases of this model. A natural interpretation for the learning process is that sales generate a probability that the product is reviewed by either a professional reviewer or a random consumer. While our model is set up in discrete time, we characterize the monopolist's behavior in the continuous time limit in which she can flexibly change the price. The key novelty is combining the buyers' ability to time their purchases with the endogenous arrival of information.

The (constrained) efficient benchmark is simple in our model: starting from the highest type, the buyers purchase without a delay until the marginal cost outweighs the social value of a purchase, which consists of both the value of consumption and the value of information. Our main interest, however, is in the stationary Markov perfect equilibria in which the profit maximizing seller optimizes the sales schedule in each state. The payoff relevant state consists of the distribution of the remaining buyers in the market and the belief that the quality is high. The profit maximizing seller sells until the marginal cost outweighs the value of consumption and the value of information for herself. Because the latter is always less than the total value of information, the sellers' exit and entry decisions are necessarily inefficient. The seller may also find it optimal to delay sales.

Delays in stationary equilibria arise purely from the seller's need to control buyers' willingness to pay, because we have shut down the seller's informational incentives to delay sales by assuming that learning happens instantaneously rather than gradually. That is, purchases generate news today but not tomorrow.[2] Delaying sales allows the seller to diminish the buyers' value of waiting by slowing down both the information flow and the possible decrease in the price. Thus, the slower the sales, the higher is the buyers' willingness to pay today. Delays occur when selling fast implies that the buyers' value of waiting is so high that they effectively demand more surplus than what is available. The seller then delays sales so that she is just indifferent between selling and exiting the market, which makes delays sequentially rational. This indifference implies that the stationary equilibria are unique only up to the length of these delays.

[2] See Laiho, Murto and Salmi (2020) on experimentation problems with gradual learning.

Interestingly, delays only occur when learning happens at least partially through bad news. An intuitive explanation for this result is that bad news provides decision relevant information immediately for a buyer who would otherwise purchase and thus bad news learning increases their value of waiting. Good news, on the other hand, provides decision relevant information only if a buyer (or the seller) is not willing to purchase (sell). Bad news and good news learning therefore affect decisions at different margins.

Besides delays, informational free-riding also explains why the seller sells too little compared to the efficient solution. Efficiency calls for the early prices to be low enough to take into account the information the early buyers produce. The seller's willingness to do so, however, depends on her share of the total surplus which is always less than one because the buyers must be compensated for their ability to wait for more information. This distorts the seller's entry and exit decisions. An important part of this is that in stationary equilibria the seller's value is limited by her inability to resist future price decreases. We prove that a version of the Coase conjecture holds for stationary equilibria. To ensure that our results do not rely on the monopolist making zero profits after uncertainty is resolved, we employ the (realistic) assumption that she faces a capacity constraint which limits the maximal rate of production.[3] Our novel proof enables the analysis with a capacity constraint and it also extends to other learning processes besides the exponential bandit model we use.

To analyze how Coasian dynamics affect the inefficiencies of the stationary equilibria, we characterize the full commitment solution to the seller's problem. The commitment monopolist always sells at full capacity and thus does not inefficiently delay sales. An important reason for this is that she can limit buyers' informational free-riding by committing to reduce sales in the future. In addition to selling faster, we also show that the commitment monopolist may sell more than the equilibrium monopolist. A key insight to this result lies in the Coase conjecture itself: only the monopolist with commitment power can charge high prices, and thus capture a larger share of the surplus, after the initial uncertainty is resolved. Consequently, she has stronger incentives to sell the product when

[3] Our main results also hold in the limit when the capacity goes to infinity.

uncertainty is present.

Either inefficiency – too small or too slow sales – may make the consumer surplus greater in the commitment solution. An important conclusion from our analysis is that the lack of commitment power affects welfare differently in new markets than it does in mature markets. When consumers are still unsure about a product, commitment power strengthens a firm's incentives to keep selling the product. This is because the cost of information must be borne by the seller, and the seller agrees to do so only if she can recoup the cost through future profits. Coasian dynamics in turn imply that commitment power helps to achieve this. In contrast, when a product is well established, commitment power only serves to reduce consumer surplus.

Related literature

This paper contributes to the literature on durable goods and bargaining and to the literature on social learning and experimentation. The solutions to the complete information version of our model are in line with the classic durable goods papers (see Stokey (1979), Stokey (1981) and Bulow (1982) for early papers).[4] Deneckere and Liang (2006) analyze a related bargaining model with interdependent valuations in which the (uninformed) buyer makes offers to the (informed) seller. Delays in their model, like in ours, stem from the need to diminish the value of waiting. Fuchs and Skrzypacz (2013) analyze the same model when the gap between the lowest buyer's valuation and the seller's cost goes to zero. Board (2008) characterizes the commitment solution in the case when demand varies exogenously over time. Fuchs and Skrzypacz (2010) analyze a bargaining model that is analogous to durable goods models with a continuum of types. The novelty is a stochastic arrival of an outside option, which creates incentives to delay trade. Lomys (2018) analyzes a similar bargaining model where the parties exogenously learn about the existence of the outside option and shows that the presence of exogenous learning leads to efficient delays while interdependency in values leads

[4]For a discussion on commitment power and the Coase conjecture, see Ausubel and Deneckere (1989), Sobel (1991), von der Fehr and Kühn (1995), Mcafee and Wiseman (2008) and Ortner (2017).

to inefficient delays in the stationary equilibrium.

The most related papers studying social learning are Kremer, Mansour and Perry (2014), Che and Hörner (2017) and Frick and Ishii (2016). All three papers look at situations where prices cannot adjust to facilitate experimentation. Kremer, Mansour and Perry (2014) and Che and Hörner (2017) study social learning from the point of view of a social planner who designs a recommendation system for the consumers. Both papers show that the planner's solution generally exhibits regions of over-recommendation to induce consumers to experiment. Over-recommendation arises from the same reasons that cause the monopolist to sell at a price below the marginal cost in our model. Frick and Ishii (2016) analyze similar environment with a fixed price and without a monopolist. Because of informational free-riding, some buyers may choose to wait in the equilibrium.

A number of papers have looked at pricing and learning. Bergemann and Välimäki (1996 and 2006) and Weng (2015) analyze idiosyncratic learning and pricing. Bose, Orosel, Ottaviani and Vesterlund (2006) look at monopoly pricing in a herding model with exogenous timing of purchases. Bonatti (2011) studies menu pricing in the presence of social learning. In these papers, buyers do not make timing decisions. Therefore, unlike in our paper, the solution does not have to take into account the buyers' incentives to delay purchases.

The rest of the paper is organized as follows. In Section 2, we introduce the model. In Section 3, we analyze the buyers' problem. In Section 4, we characterize behavior in stationary equilibria. In Section 5, we solve the commitment solution and compare it to stationary equilibria.

2 Model

2.1 Actions and payoffs

A monopolistic seller (she) is selling a product of unknown quality to a mass of buyers over time. Time is discrete, $t \in \{0, \Delta, 2\Delta, \dots\}$, with period length Δ. At the beginning of each period, the monopolist decides how much she sells, $\lambda_t \Delta$, constrained by the capacity ρ at her disposal: $\lambda_t\Delta \leq \rho\Delta$.[5] Each buyer (he) chooses when to purchase, has a unit demand and exits the market upon purchasing the product. The price at each period, p_t, is determined so that the market clears: a mass $\lambda_t\Delta$ of buyers is willing to purchase at price p_t.[6]

Let $\omega \in \{1, 0\}$ denote the quality of the product, where 1 denotes high quality and 0 low quality. Neither the seller nor the buyers observe ω, but they share a common belief: $\Pr(\omega = 1) = x_t$. Each buyer has a (private) type $\theta \in [\underline{\theta}, \bar{\theta}]$ distributed according to F, which has a continuously differentiable density function f and an increasing hazard rate. Buyer i's valuation for the product is $\theta_i\omega$, which implies that his expected valuation equals $\theta_i x_t$. All players are risk neutral. We normalize the value of the buyers' outside option to zero and the mass of buyers to one.

The monopolist faces a constant marginal cost of production, $c > 0$. We assume the 'no-gap' case: $c > \underline{\theta}$, so that the seller never wants to sell to the full distribution of buyers. Lastly, we assume that both the seller and the buyers discount the future with the same discount rate $r > 0$. For the most part, our analysis focuses on the continuous time limit as the period length goes to zero, $\Delta \to 0$.

2.2 Learning

The seller and the buyers observe public signals (news) about the product quality over time as buyers purchase the product. We model this learning process with an exponential bandit similar to those analyzed by Keller, Rady and Cripps (2005) and by Keller and Rady (2015). Central to our model of learning is that signals arrive proportionally to the amount of sales, so that the more the monopolist sells the more she and the buyers learn. For tractability, we focus on signals that are *perfectly revealing*, meaning that a single signal is conclusive evidence that the

[5]The capacity constraint implies that in the continuous time limit the price can react to information immediately because the seller is selling at a finite rate. All the main results hold for the limit as the capacity goes to infinity, $\rho \to \infty$.

[6]Choosing quantities is equivalent to choosing prices in our setup where the monopolist would never want to set a price that leads to rationing: higher types have a (weakly) higher willingness to pay in both states of the world and hence it is cheaper to use them for experimentation.

product is either high or low quality.

Suppose that the sales in period t equal $\lambda_t\Delta$. We can then describe the news process as follows. At the beginning of the next period, a perfectly revealing public signal about the product quality being *high* arrives with probability $1 - e^{-\omega\delta_G\lambda_t\Delta}$, where $\delta_G \geq 0$ denotes the arrival rate of the signal revealing high quality (good news) per a unit of sales. Similarly, a perfectly revealing signal about the quality being *low* arrives with probability $1 - e^{-(1-\omega)\delta_B\lambda_t\Delta}$, where $\delta_B \geq 0$ denotes the arrival rate of the signal revealing low quality (bad news).

If no news arrives, market participants update their beliefs by using the Bayes' rule based on the arrival rates δ_G and δ_B. Given a prior x_0, we can write the no news posterior (belief) x_t absent news as

$$x_t = \frac{x_0 e^{-\delta_G Q_t}}{x_0 e^{-\delta_G Q_t} + (1 - x_0)e^{-\delta_B Q_t}},$$

where $Q_t = \sum_{k=0}^{t-\Delta} \lambda_k \Delta$ is the total amount of sales prior to time t. Notice that absent news the belief, x_t, is decreasing in sales Q_t when $\delta_G > \delta_B$, i.e. in the *good news environment*, and increasing when when $\delta_G < \delta_B$, i.e. in the *bad news environment.*

3 Buyers' problem

3.1 Optimal purchases

We start by analyzing an individual buyer's problem with the aim of characterizing prices and learning dynamics in the market. Each buyer decides when to buy the product based on his own valuation, current and expected future prices and the information available at each time period t. Formally, buyers' purchasing decisions are adapted to the filtration $\mathcal{F}_t$ generated by the news process, past sales and the amount the seller sells today. Suppose first that prices follow a stochastic process $P = \{p_t\}$ adapted to $\mathcal{F}_t$. Each individual buyer is then a price taker and considers P as given. Thus, a buyer's problem is to choose an optimal stopping time τ^s adapted to $\mathcal{F}_t$ such that it maximizes his payoff:

$$\sup_{\tau^s} \mathbb{E}(e^{-r\tau^s}(\theta\omega - p_\tau)).$$

Using the properties of the news process, we can describe the information available to the buyers with the help of three random times: σ when any news arrives, σ_G when good news arrives and σ_B when bad news arrives.[7] It is clear that after bad news no buyer purchases so we can focus on behavior in the no news ($\sigma > t$) and in the good news ($\sigma_G \leq t$) cases. Let $\tau(\theta)$ denote the time when a buyer of type θ buys if no news has arrived and let $\tau^G(\theta; k)$ denote the time he buys if good news arrives at time k. In addition, let infinite purchasing times denote the situation when a buyer never buys. We then have a familiar comparative statics result:

Lemma 1 (Skimming property)**.** *Optimal purchasing times are decreasing in type:* $\tau(\theta') \leq \tau(\theta)$ *and* $\tau^G(\theta'; k) \leq \tau^G(\theta; k)$ *for all* $\theta' > \theta$ *and* k*.*

Lemma 1 follows directly from the Topkis's theorem once we observe that the buyer's discounted stopping payoff has strictly decreasing differences in type and purchasing time and is supermodular in purchasing time. Intuitively, higher types must be willing to buy first because waiting implies they are forgoing a larger expected payoff.

3.2 Learning dynamics

Next, we characterize how buyers' behavior affects the evolution of the belief in the market. The skimming property implies that knowing the current highest type, θ_t is enough to pin down the distribution of remaining buyers because all higher types must have already bought. Thus, the evolution of the current highest type θ_t is described by

$$F(\theta_t) - F(\theta_{t+\Delta}) = \lambda_t \Delta.$$

Knowing current highest type is also sufficient information for the evolution of the belief absent news because the total sales must equal $1 - F(\theta_t)$. Thus, the no news posterior is a function of the current highest type θ_t and can be defined as

$$x(\theta_t; x_0) := \frac{x_0 e^{-\delta_G(1-F(\theta_t))}}{x_0 e^{-\delta_G(1-F(\theta_t))} + (1-x_0)e^{-\delta_B(1-F(\theta_t))}}.$$

[7]Formally, let $Y_t \in \{-1, 0, 1\}$ be the news process with -1 denoting bad news, 0 no news and 1 good news. Then $\sigma_G = \inf\{t \geq 0 : Y_t = 1\}$, $\sigma_B = \inf\{t \geq 0 : Y_t = -1\}$ and $\sigma = \inf\{\sigma_G, \sigma_B\}$.

In addition, the probability that news does not arrive between sales to types between θ_t and θ_k ($\sigma > k$ conditional on $\sigma > t$) can be written as

$$\beta(\theta_t, \theta_k; x_0) := x(\theta_t; x_0) e^{-\delta_G(F(\theta_t)-F(\theta_k))} + (1 - x(\theta_t; x_0)) e^{-\delta_B(F(\theta_t)-F(\theta_k))}.$$

We often shorthand $x(\theta_t; x_0)$ to $x(\theta_t)$ and $\beta(\theta_t, \theta_k; x_0)$ to $\beta(\theta_t, \theta_k)$ for notational convenience.

3.3 Price dynamics

The market clearing price, p_t, is such that the next period's current highest type, $\theta_{t+\Delta}$, is indifferent between buying and waiting:

$$\theta_{t+\Delta} x_t - p_t = W_\Delta(\theta_{t+\Delta}; \theta_t, x_0), \tag{1}$$

where $W_\Delta(\theta_{t+\Delta}; \theta_t, x_0)$ is the value of the next period's highest type given that the current highest type is θ_t. Before defining the value, notice that the skimming property implies that any strategy set by the seller defines a sequence of highest remaining types $\{\theta_t\}$ which induces corresponding purchasing times τ and τ^G for each type. We can then define the buyers' value after good news and before news the following way.

Suppose first that good news arrives in period k when θ_k is the highest type ($\sigma_G = k$) and that prices follow a path $\{p_t^G\}$ thereafter. The value of a buyer of type θ is then

$$W_\Delta^G(\theta; \theta_t, \theta_k) := \sup_{\tau^G \geq t} \{e^{-r(\tau^G - t)}(\theta - p_{\tau^G}^G)\}, \tag{2}$$

where t is current time and θ_t is the current highest type.

Suppose next that no news has arrived by time t ($\sigma > t$) and that prices follow a path $\{p_t\}$ absent news. The value of a buyer of type θ is then

$$\begin{aligned} W_\Delta(\theta; \theta_t, x_0) := \sup_{\tau \geq t} \Big\{ & e^{-r(\tau - t)} \beta(\theta_t, \theta_\tau)(\theta x_\tau - p_\tau) \\ & + \sum_{k=t+\Delta}^{\tau} e^{-r(k-t)} \Pr(\sigma_G = k | \sigma > t) W_\Delta^G(\theta; \theta_k, \theta_k) \Big\}, \end{aligned} \tag{3}$$

where $\Pr(\sigma_G = k|\sigma > t)$ is the probability of good news arriving at time $k \geq t$ ($\sigma_G = k$) conditional on no news having arrived by time t:

$$\begin{aligned}\Pr(\sigma_G = k|\sigma > t) &= \Pr(\sigma_G = k|\sigma > k - \Delta)Pr(\sigma > k - \Delta|\sigma > t) \\ &= \beta(\theta_t, \theta_{k-\Delta})(1 - e^{-\delta_G(F(\theta_{k-\Delta})-F(\theta_k))})x(\theta_{k-\Delta}).\end{aligned}$$

4 Stationary equilibria

4.1 Discrete time

We are mostly interested in the seller's behavior when she cannot commit to a strategy beforehand. We focus on stationary Markov perfect equilibria in which equilibrium strategies depend only on payoff relevant variables. We define the game in discrete time and first analyze discrete time necessary conditions for equilibria. One of our results is to show that a suitable form of the Coase conjecture holds: the seller either sells at full rate or her value is close to zero. After analyzing the discrete time necessary conditions, we take the continuous time limit and characterize equilibrium properties that arise from the necessary conditions in the limit.

4.1.1 Equilibrium and strategies

The payoff relevant information is fully characterized by the current highest type in the market, θ_t, and the current belief, x_t (from Lemma 1). We refer to (θ_t, x_t) as the state and let $S := [\underline{\theta}, \overline{\theta}] \times [0, 1]$ denote the set of all possible states. We define a stationary equilibrium as a pair of functions $(\Lambda_\Delta, P_\Delta)$. The function $\Lambda_\Delta : S \to [0, \rho]$ maps states to rates of sales and the function $P_\Delta : \Theta \times S \to \mathbb{R}$ specifies the highest price each type is willing to pay given the current state.

Together, the rate of sales and the acceptance price specify the market price: $p(\theta_t, x_t) = P_\Delta(\theta_{t+\Delta}; \theta_t, x_t)$, where $\theta_{t+\Delta}$ solves $F(\theta_t) - F(\theta_{t+\Delta}) = \Lambda_\Delta(\theta_t, x_t)\Delta$. Therefore, a profile $(\Lambda_\Delta, P_\Delta)$ uniquely pins down the path of sales and prices. We defined the buyer's problem in Section 3. Notice that the market clearing equation (1) induces a price that the buyers are willing to accept for a given profile $(\Lambda_\Delta, P_\Delta)$.

Given a price function P_Δ, we can write the seller's problem recursively as

$$V_\Delta(\theta_t, x_t) = \max_{\lambda\in[0,\rho]} \Big\{(P_\Delta(\theta_{t+\Delta}; \theta_t, x_t) - c)\lambda\Delta + e^{-r\Delta}\beta(\theta_t, \theta_{t+\Delta})V_\Delta(x_{t+\Delta}, \theta_{t+\Delta}) + e^{-r\Delta}\Pr(\sigma_G = t + \Delta|\sigma > t)V^G_\Delta(\theta_{t+\Delta})\Big\}, \tag{4}$$

where the probabilities $\beta(\theta_t, \theta_{t+\Delta})$ and $\Pr(\sigma_G = t + \Delta|\sigma > t)$ are as specified in Section 3. Similarly, the seller's problem after good news is

$$V^G_\Delta(\theta_t) = \max_{\lambda\in[0,\rho]} \Big\{(P(\theta_{t+\Delta}; \theta_t, 1) - c)\lambda\Delta + e^{-r\Delta}V^G_\Delta(\theta_{t+\Delta})\Big\}.$$

A stationary equilibrium is a profile $(\Lambda_\Delta, P_\Delta)$ such that the quantities and the prices are sequentially rational.[8]

Definition 1. *A stationary Markov perfect equilibrium is characterized by a sales strategy Λ_Δ and a price function P_Δ such that for each state $(\theta_t, x_t) \in S$:*

- $\Lambda_\Delta(\theta_t, x_t)$ *maximizes* (4)
- *price* $P_\Delta(\theta; \theta_t, x_t)$ *satisfies* (1),
- *the belief equals 1 after good news, 0 after bad news, and the no news belief evolves according to:*

$$x_{t+\Delta} = \frac{x_t e^{-\delta_G\lambda_t\Delta}}{x_t e^{-\delta_G\lambda_t\Delta} + (1 - x_t)e^{-\delta_B\lambda_t\Delta}},$$

- *the highest type in the market evolves according to:*

$$F(\theta_t) - F(\theta_{t+\Delta}) = \lambda_t\Delta.$$

When product quality is known, the existence of a stationary equilibrium can be established by arguing that the seller's payoff is increasing in the buyer's type,

[8]The bargaining literature looks at weak Markov perfect equilibria where the seller makes price offers, which are allowed to depend on the current state and previous period's price (see Fudenberg, Levine and Tirole (1985) and Ausubel and Deneckere (1989)). The price setting problem reduces to the same system of equations as the quantity choice problem in our model. The use of quantities instead of prices enables us to bypass the question of deviations to out of equilibrium prices and therefore our equilibrium is a strong Markov equilibrium in the sense that equilibrium strategies depend only on the payoff relevant state.

which guarantees that the game ends in a finite number of periods.[9] In our environment, the seller's expected payoff can be larger from low types than from high types because later sales happen conditional on information from earlier sales. Deneckere and Liang (2006) face a similar problem when analyzing bargaining with interdependent values: the monotonicity breaks down when the buyer's value is positively correlated with the seller's cost. Their proof for the existence of a stationary equilibrium extends to our model.

4.1.2 Discrete time necessary conditions

We start by characterizing the necessary conditions for a stationary equilibrium in discrete time. To begin, note that it is enough for the seller to know the current highest type and whether good or bad news has arrived to keep track of the state. Thus, the problem reduces from two state variables to one state variable.

We first define an important function for the analysis: the expected *marginal joint payoff* per type, m_Δ:

$$m_\Delta(\theta; x_0) := (\theta + \delta_G e^{-r\Delta} V_\Delta^G(\theta)) x(\theta; x_0) - c.$$

The function m_Δ describes the seller's payoff in the case she sells to type θ and then exits absent good news. It measures the amount of surplus available at each type and hence the name joint payoff.[10] For completeness, let $m_\Delta(\theta; 1) := \theta - c$. Notice that if the monopolist is not expected to sell, the buyer's value is zero and the seller captures the entire marginal joint payoff. Therefore, the seller has a profitable deviation if the buyers expect her to stop selling when the marginal joint payoff is positive. This yields:

Lemma 2. *The monopolist sells in state $(\theta, x(\theta; x_0))$ if $m_\Delta(\theta; x_0) > 0$.*

We are especially interested in equilibria with delays so that sales are below capacity. To analyze this in more detail, we define that the monopolist sells at

[9] See Fudenberg, Levine and Tirole (1985), Gul, Sonnenschein and Wilson (1986) and Ausubel and Deneckere (1989).

[10] The marginal joint payoff is simply the sellers marginal payoff, price minus the marginal cost and the good news continuation value, plus the buyers' payoff, the valuation minus the price.

almost full rate if her excess capacity over the next K periods is less than her per period capacity.

Definition 2. *Sales take place at almost full rate K-close to θ if for all θ' such that $F(\theta) - F(\theta') \in (0, \frac{K}{\rho}\Delta)$,*

$$\tau_\Delta(\theta') - \tau_\Delta(\theta) \leq \frac{F(\theta) - F(\theta') + \rho\Delta}{\rho}.$$

Using this definition, we establish that a suitable form of the Coase conjecture holds for our model.

Theorem 1. *(Coase conjecture) For all $\epsilon > 0$ and for all $K > 0$, there exists $\bar{\Delta}$ such that if $\Delta < \bar{\Delta}$, at least one of the following holds for all stationary equilibria and all θ: (i) sales take place at almost full rate K-close to θ or (ii) $V_\Delta(\theta; x_0) \leq \epsilon$.*

Let the period length be small. Theorem 1 says that if the seller does not want to speed up sales, her value must be small. The intuition for this is that whenever the seller gets a positive value from the next type, she faces an irresistible temptation to speed up sales to sell also to that type. In equilibrium, she ends up selling until either the capacity constraint binds or the value from sales is close to zero. It is easy to see that without the capacity constraint this implies that the monopolist always speeds up her sales until her value is close to zero – the usual form of the Coase conjecture. Essentially, the equilibrium monopolist suffers from Coasian dynamics: she lowers the price until she has sold to all types from which she can obtain positive marginal value.

More explicitly, Theorem 1 establishes a connection between the purchasing times in two states which can be reached in a finite number of periods (K). The reason we restrict to states close to type θ is that without the constraint we can only establish that the monopolist's value must be low for some type in the interval between types θ and θ' but not necessarily at type θ. Because the result holds for any K, the condition (i) in Theorem 1 is equivalent to no delays in sales when the period length shrinks.

The complete proof of Theorem 1 is laborious and can be found in Appendix B.1. The capacity constraint in particular complicates the proof because the usual arguments for the Coase conjecture build on deviations which double the rate of sales. Such a deviation may violate the capacity constraint in our model. With the help of additional continuity results, we prove Theorem 1 by looking at multi-period deviations. The main intuition remains the same as in other proofs for the Coase conjecture: if the seller gets positive value in some future equilibrium state, she wants to get there as fast as possible.

Our proof for Theorem 1 extends to any environment in which the *distribution* of net surplus, consumer's valuation minus the production cost, depends on the cumulative past sales but not directly on time. This includes a model of learning where each new sale generates a normally distributed signal and a model where the marginal cost of production depends, deterministically or stochastically, on past sales.

4.2 Continuous time: necessary conditions

For the rest of Section 4, we focus on characterizing the equilibrium behavior in the continuous time limit ($\Delta \to 0$). We first take the limit of the discrete time necessary conditions and then use these conditions to characterize equilibrium properties in the limit. We start by defining two functions: the expected *marginal joint payoff* per type, m, and the seller's expected *marginal payoff* per type, π.

$$\begin{aligned} m(\theta; x_0) &:= (\theta + \delta_G V^G(\theta))x(\theta; x_0) - c, \\ \pi(\theta; x_0) &:= (\theta + \delta_G V^G(\theta))x(\theta; x_0) - c - W(\theta; \theta, x_0). \end{aligned}$$

The function m is simply the continuous time limit of the marginal joint payoff, m_Δ, and its interpretation is the same as in discrete time. The marginal payoff function π keeps track of the seller's share of the joint payoff. We thus have $m(\theta) \geq \pi(\theta)$ always.

The continuation play affects the marginal payoff both through the value of good news, V^G, and through the buyers' value, W, which are the continuous time limits of their discrete time versions and which we define next. First notice that any sales path (Λ) induces a corresponding schedule of purchasing times for each type. Given the sales paths induced by the seller's and the buyers' equilibrium strategies, the seller's value is equal to an integral over her marginal payoff. In the continuous time limit, this gives

$$V(\theta_t; x_0) = \int_{\underline{\theta}}^{\theta_t} e^{-r(\tau(\theta)-t)}\beta(\theta, \theta_t)\pi(\theta; x_0)f(\theta)d\theta. \tag{5}$$

For the buyers' value, we first use the optimality of the purchasing times and the envelope theorem to write the value as an integral over type. In the continuous time limit, we then have (Appendix A):

$$\begin{aligned} W(\theta; \theta_t, x_0) =& x(\theta) \int_{\underline{\theta}}^{\theta} e^{-r(\tau(s)-\tau(\theta))-\delta_G(F(\theta_t)-F(s))} ds \\ &+ x(\theta) \int_{\underline{\theta}}^{\theta} \int_{s}^{\theta_t} e^{-r(\tau^G(s;z)-\tau(\theta))-\delta_G(F(\theta_t)-F(z))} \delta_G f(z) dz ds. \end{aligned} \tag{6}$$

We can now state the necessary conditions for any stationary equilibrium in the continuous time limit. First note that both the seller's and the buyers' values have to be weakly positive. Combining this with the continuous time limits of Lemma 2 and Theorem 1 then yields the following (Appendix B.2).

Lemma 3. *In the continuous time limit, any stationary equilibrium satisfies:*

(i) If $V(\theta; x_0) > 0$, the monopolist must sell at full rate: $\lambda(\theta) = \rho$.

(ii) If $V(\theta; x_0) = 0$ and $m(\theta; x_0) > 0$, there must be delay in sales so that $\lim_{s\to\theta^+} \pi(s; x_0) \geq 0$.

(iii) If $V(\theta; x_0) = 0$ and $m(\theta; x_0) \leq 0$, the monopolist does not sell: $\lambda(\theta) = 0$.

(iv) The monopolist's value function V is continuous in the current highest type θ.

We use the necessary conditions in Lemma 3 as our solution concept in the continuous time limit. Together with equations (5) and (6) they pin down the equilibrium sales path. To see whether a schedule of purchasing times can be supported as a stationary equilibrium, we need to check whether the necessary conditions are satisfied everywhere on the no news and after good news paths.

4.3 Continuous time: equilibrium after good news

We first characterize equilibrium behavior after good news. Suppose good news arrives when θ_k is the current highest type ($\sigma_G = k$). The game is then of complete information and the Coase conjecture implies that the monopolist sells at full rate to all types whose valuation is greater than the marginal cost. This follows immediately from Lemma 3 once we argue that the capacity constraint ensures that the seller always gets a strictly positive share of the social surplus (Appendix B.2).

Proposition 1. *In the continuous time limit of any stationary equilibrium, the monopolist sells at full rate to all types greater than the marginal cost after good news. That is, after good news* $\lambda(\theta) = \rho$ *for all* $\theta \geq c$ *and* $\lambda(\theta) = 0$ *for all* $\theta < c$.

Proposition 1 pins down the seller's strategy and value after good news, V^G, for all stationary equilibria (Appendix B.2). An important implication of Proposition 1 is that the capacity constraint does not give any commitment power in addition to the mechanical effect of constraining the rate of sales. After good news, the equilibrium always coincides with the (constrained) efficient solution irrespective of the capacity constraint.

4.4 Continuous time: equilibrium before news

We next analyze stationary equilibria before news. First notice that the marginal joint payoff must be zero when the seller exits from the market (follows from Lemma 3, see Appendix B.2).

Proposition 2. *In the continuous time limit, all stationary equilibria are such that the monopolist sells to all types greater than* θ_E *absent news. The cutoff* θ_E *is the largest* θ *such that*

$$(\theta_E + \delta_G V^G(\theta_E))x(\theta_E; x_0) - c = 0 \tag{7}$$

and $V(\theta; x_0) = 0$. *After good news, the monopolist sells according to Proposition 1.*

The result in Proposition 2 can be strengthened if the marginal joint payoff is monotone over type, because then there is a unique type such that it equals zero and thus (7) is also a sufficient condition for the monopolist to exit. This is true, for example, when good news learning dominates, $\delta_G \geq \delta_B$, because then both the belief and the type are decreasing over sales. When bad news is more likely, however, low type buyers are more optimistic when they purchase than high type buyers at their earlier time of purchase and thus the marginal joint payoff might be greater for low types.

Independent of the news environment, Proposition 2 entails that the seller may sell to buyers whose valuation is below the marginal cost in order to generate information. When this happens, the price falls below the marginal cost and the seller makes short-term losses. The difference between the news environments is that in the good news environment this behavior is fully captured by the value of good news. In the bad news environment, the seller might also sell to types for which the marginal joint payoff is negative because she can earn positive profits from the more optimistic lower types.

Proposition 2 tells us when the monopolist exits the market in a stationary equilibrium but not how she sells the product. We know from the first condition in Lemma 3 that the monopolist sells at full rate as long as her value is strictly positive. However, it turns out that when learning is present, the monopolist's payoff from selling at full rate may often be negative. In particular, a positive marginal joint surplus does not guarantee that the seller's marginal payoff is positive unlike with complete information. This is because the opportunity to wait for more information increases the buyers' value and thus makes it more costly to the monopolist to sell to early buyers. Because in equilibrium the seller cannot reduce the buyers' value by committing to reduce future sales, she has to introduce delays.

Delays mean that the expected purchasing times are strictly larger than they would be if the monopolist used the full capacity.

Definition 3. *An equilibrium features delays when there exists θ such that the purchasing time absent news, $\tau(\theta)$, is finite and satisfies*

$$\tau(\theta) > \frac{1 - F(\theta)}{\rho} =: \tau^{\rho}(\theta).$$

To analyze delays in more detail, we first define the seller's full rate payoff. Suppose the current highest type is θ and the monopolist sells at full rate until type θ^*. Then her full rate payoff at type θ can be defined as

$$\Pi^\rho(\theta, \theta^*; x_0) = \int_{\theta^*}^{\theta} e^{-r(\tau^\rho(s) - \tau^\rho(\theta))} \beta(\theta, s) \Big(m(\theta) - W^\rho(\theta; \theta) \Big) f(s) ds,$$

where $W^\rho(\theta; \theta)$ is the buyers' value defined in (6) given full rate sales. For a given initial belief x_0, the full rate payoff depends only on the exit type, θ^*, and the model primitives. Let $\hat{\Theta}_E(x_0)$ denote the set of potential exit types: $\hat{\Theta}_E(x_0) \coloneqq \{\theta \in [\underline{\theta}, \overline{\theta}] : m(\theta; x_0) = 0\}$. Using the above definitions, we have the following result for checking when delays happen.

Lemma 4. *Suppose that the equilibrium monopolist enters the market. If for every potential exit type $\hat{\theta}_E \in \hat{\Theta}_E(x_0)$ there exists $\theta > \hat{\theta}_E$ such that $\Pi^\rho(\theta, \hat{\theta}_E; x_0) < 0$, then every stationary equilibrium features delays in the continuous time limit.*

The intuition behind Lemma 4 is that whenever selling at full rate results in a negative payoff, the rate of sales adjusts so that the seller is indifferent between continuing to sell and exiting (Lemma 3 (ii), see Appendix B.2).

To understand delays better consider the following example. Let capacity go to infinity, $\rho \to \infty$, and suppose learning happens purely through bad news, $\delta_B > \delta_G = 0$. Assume as well that the marginal joint payoff is positive for the highest type, $m(\overline{\theta}; x_0) > 0$, so that the seller enters in any equilibrium. If she then sells at full rate (purchasing times go to zero), from (6) we have that a buyer's value equals:

$$W^\infty(\theta; \theta, x_0) = x(\theta) \int_{\theta_E}^{\theta} ds = x(\theta)(\theta - \theta_E),$$

where θ_E is the exit type in the equilibrium. Consequently, the monopolist's marginal payoff equals

$$\pi^\infty(\theta; x_0) = \theta x(\theta) - W^\infty(\theta; \theta, x_0) - c = \theta_E x(\theta) - c < 0.$$

The marginal payoff is negative for all types greater than the exit type ($\theta_E = c/x(\theta_E)$). This follows because the law of motion implies that the belief is greater at the exit type, $x(\theta_E) > x(\theta)$. Essentially, full rate sales imply that due to their ability to wait, the buyers demand more surplus than what is available. Therefore, there must be delays in any stationary equilibrium. The next result generalizes the intuition in this example (Appendix B.2).

Theorem 2. *Suppose that $\delta_B > 0$ and that the initial belief $x_0 \in (0,1)$ is such that the monopolist enters. Then, in the continuous time limit, there exists $\overline{\rho} \in (0,\infty)$ such that for $\rho > \overline{\rho}$ all stationary equilibria feature delays.*

Delays depend on the bad news arrival rate because the arrival of bad news provides decision relevant information precisely when a buyer is willing to purchase the product. The buyers must be compensated in price for the possibility of bad news, which in turn leads to the monopolist slowing down sales in order to keep the price high.

While Theorem 2 tells us when delays happen, it does not tell how the seller implements them. This depends on where the full rate payoff is negative. Figure 1 illustrates delays when the delay happens at an interior type, $\hat{\theta} > \theta_E$. The figure on the left is what happens if the monopolist sells at full rate until exiting at type θ_E. In the figure, the buyers' value of waiting is so large that both the full rate marginal payoff, π^ρ, and the full rate payoff, Π^ρ, are negative for high types. However, because the marginal joint payoff is positive, it is not sequentially rational for the monopolist to exit when the current highest type is greater than the cutoff type, $\theta > \theta_E$. Therefore, the monopolist must slow down sales so that she still has incentives to sell. This is illustrated in the figure on the right, which shows what happens in equilibrium: the monopolist introduces a delay in sales at $\hat{\theta}$, which reduces the buyers' value of waiting so that the full rate payoff becomes positive before $\hat{\theta}$.

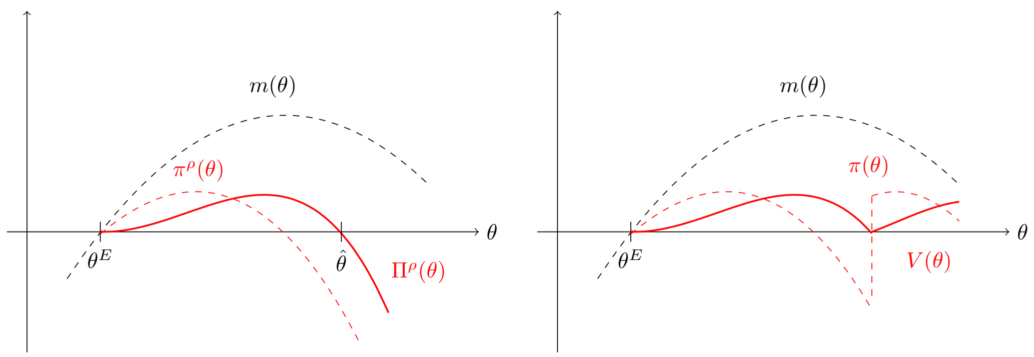

Figure 1: An illustration of an equilibrium with delays. The left-hand side depicts the marginal joint payoff together with the monopolist's payoff Π^ρ and marginal payoff π^ρ if she sells at full rate. The right-hand side depicts the corresponding equilibrium payoffs.

Notice that in Figure 1, and more generally, we have a point of discontinuity in the marginal payoff at the type at which the monopolist is indifferent ($\hat{\theta}$) because the monopolist is waiting before starting to sell again. The delay diminishes the buyers' value of waiting for types above the indifferent type and makes them willing to pay strictly more than after the delay. Notice as well that there can be several delays in an equilibrium. There is nothing qualitatively different when there are delays at several types, but, of course, both the seller's and the buyers' values must take future delays into account.

Because delays happen when the monopolist is indifferent between selling and waiting, the monopolist might be indifferent between different amounts of delay. In general, there is no closed form solution for the length of the delay and the equilibrium sales path. However, we can pin down the minimum length of delay because the monopolist's marginal payoff must be positive before delay (from Lemma 3, condition (ii)):

$$\lim_{\theta \to \hat{\theta}^+} \pi(\hat{\theta}; x_0) = (\hat{\theta} + \delta_G V^G(\hat{\theta}))x(\hat{\theta}) - c - e^{-rd} W(\hat{\theta}; \hat{\theta}, x_0) \geq 0, \tag{8}$$

where $d > 0$ is the length of the delay before the monopolist starts selling again. This lower bound ensures that the monopolist's payoff is non-negative close to the delay type. This type of waiting is Markovian in the sense that when the period length is still strictly positive, the current highest type keeps decreasing (the monopolist must sell a positive amount in discrete time because $m_\Delta(\theta) > 0$). Thus, sales grind to a halt only in the continuous time limit in which the seller waits for a strictly positive amount of time before selling again.[11]

So far we have discussed delays which happen at an interior type. What happens if the full rate payoff approaches the exit type from below?[12] It is clear that delaying sales in the way we described earlier does not work in this case, because we must have that the monopolist's payoff stays weakly positive even as the current highest type tends to the exit type, $\theta \to \theta_E^+$. Thus, in this case we must have that the expected rate of sales tends to zero as the marginal buyer tends to the exit type: $\lim_{\theta \to \theta_E^+} \lambda(\theta) = 0$. We show how to construct the rate of

[11]The delays in the interior case resemble those in Deneckere and Liang (2006) although, interestingly, they emerge from a different source.

[12]That is, the derivative of the marginal payoff is negative at the exit type: $\pi'(\theta_E) < 0$. Then, there is a neighborhood to the right of θ_E in which the seller's payoff is negative if she sells at full rate.

sales in this case in Appendix B.2.

As a final remark, we discuss how the capacity constraint affects equilibrium behavior. After good news, the capacity constraint simply makes the seller's value strictly positive as is evident from Proposition 1. It similarly limits the rate of sales before news but does not guarantee that the monopolist's payoff is necessarily positive from full rate sales as Theorem 2 makes clear. Because the rate of sales affects the price, the capacity constraint essentially restricts how low a price the seller can set today and thus it proxies for the seller's ability to commit to a high price in the short-term.

5 Comparison to commitment solution

5.1 Commitment solution

In this section, we solve the seller's problem with full commitment power. The commitment solution serves as an important benchmark for stationary equilibria. We characterize the solution directly in the continuous time limit.

The commitment monopolist maximizes her *ex-ante* payoff by designing a schedule of purchasing times for all types. Letting $\mathcal{T}$ denote a schedule of purchasing times for all types, we can write the commitment monopolist's problem

$$\max_{\mathcal{T}} \int_{\underline{\theta}}^{\theta_t} e^{-r(\tau(\theta)-\tau(\theta_t))}\beta(\theta,\theta_t)\Big(\theta x(\theta) - W(\theta;\theta) - c + \delta_G x(\theta)\Pi^G(\mathcal{T};\theta)\Big)f(\theta)d\theta,$$

subject to incentive compatibility and the capacity constraint. These constraints imply that the purchasing times are decreasing functions such that the difference in purchasing times between types θ and θ' is at least $(F(\theta) - F(\theta'))/\rho$.

The first step in solving the commitment monopolist's problem is to simplify her payoff by expressing it in terms of virtual surplus. This step resembles static mechanism design and papers analyzing dynamic problems with similar techniques (e.g. Board (2007)): we use the envelope theorem (Milgrom and Segal (2002)) to write a buyer's value as an integral over the type and then use Fubini's theorem to simplify the resulting expression for the monopolist's payoff (Appendix C).

Lemma 5. *Suppose the monopolist's strategy induces a schedule of incentive com-*

patible purchasing times $\mathcal{T}$. Then the monopolist's payoff absent news at type θ_t in the continuous time limit equals

$$\Pi(\mathcal{T};\theta_t,x_0)=\int_{\underline{\theta}}^{\theta_t} e^{-r(\tau(\theta)-\tau(\theta_t))}\beta(\theta_t,\theta)\phi(\mathcal{T},\theta;\theta_t,x_0)f(\theta)d\theta,$$

where $\phi(\mathcal{T},\theta;\theta_t,x_0)$ is the (marginal) no news virtual surplus defined as

$$\phi(\mathcal{T},\theta;\theta_t,x_0):=\left(\theta+\delta_G\Phi_G(\mathcal{T},\theta;\theta_t)-\frac{F(\theta_t)-F(\theta)}{f(\theta)}\right)x(\theta)-c,$$

and $\Phi_G(\mathcal{T},\theta;\theta_t)$ is the virtual payoff of good news arriving when θ is the highest type:

$$\Phi_G(\mathcal{T},\theta;\theta_t):=\int_{\underline{\theta}}^{\theta} e^{-r\tau^G(s;\theta)}\left(s-c-\frac{F(\theta_t)-F(s)}{f(s)}\right)ds.$$

The no news virtual surplus in Lemma 5, $\phi(\mathcal{T},\theta;\theta_t,x_0)$, tells us how much surplus the seller can extract from each type absent news. The virtual surplus depends on the allocation through the *virtual payoff* from sales after good news, $\Phi_G(\mathcal{T},\theta;\theta_t)$. The virtual payoff tells how much value the monopolist receives from good news taking into account the increase in buyers' rents (before and after good news) caused by sales after good news. The result in Lemma 5 is remarkable because it allows us to write a dynamic optimization problem simply as a static problem of allocating purchasing times to buyers.

Similar to Section 4, the analysis of the commitment monopolist's problem breaks naturally into two parts: what happens after good news and what happens absent news. In the interest of space we do not cover the after good news case separately because its solution follows from nearly identical steps to the case absent news. Absent news, the analysis boils down to analyzing the behavior of the (marginal) no news virtual surplus evaluated at the beginning of the game:

$$\phi^*(\theta;\overline{\theta},x_0):=\left(\theta+\delta_G\Phi_G^*(\theta)-\frac{1-F(\theta)}{f(\theta)}\right)x(\theta)-c, \tag{9}$$

where $\Phi_G^*(\theta)$ denotes the virtual value of good news given an optimal strategy after good news. The virtual surplus captures the marginal trade-offs the monopolist is facing: each sale generates a benefit in terms of revenue and potential arrival of good news but comes with a cost of increasing buyers' rents and the marginal cost. Similar to the marginal payoff, the virtual surplus can be non-monotonic in the bad news environment: lower types are more optimistic and thus might yield more virtual surplus for the monopolist. This means we cannot solve the seller's problem by simply finding a cutoff type at which the virtual surplus equals zero.

To simplify the problem we first establish that the commitment monopolist always wants to implement minimal purchasing times (sell at full rate) for types who purchase (Appendix C):[13]

Lemma 6. *Suppose the commitment monopolist finds it optimal to sell at least until $\theta_t = \theta_N$ absent news. Then, in the continuous time limit, selling at full rate is optimal for all $\theta > \theta_N$.*

Given Lemma 6, the remaining question is when the commitment monopolist exits. Local optimality of the exit decision is satisfied if the virtual surplus equals zero, but to account for global deviations, we need to check whether any amount of future sales can bring positive value to the monopolist. First, we introduce a new function that describes the virtual payoff that the monopolist gets if she continues to sell (at full rate) until type θ^* given current highest type θ_t:

$$\Phi^\rho(\theta_t, \theta^*; x_0) := \int_{\theta^*}^{\theta_t} e^{-r(\tau^\rho(s)-\tau^\rho(\theta_t))}\beta(\theta_t, s)\phi^*(s; \overline{\theta}, x_0)f(s)ds.$$

To check whether future sales can bring the monopolist positive value, we define the maximal virtual payoff from future sales, $\Psi(\theta_t; x_0)$, using the full rate virtual payoff. Let $\mathbb{A}(\theta_t)$ denote the set of all local maxima for $\Phi^\rho(\theta^*; \theta_t, x_0)$ over the exit type θ^*. We can then define the function $\Psi(\theta_t; x_0)$ as

$$\Psi(\theta_t; x_0) := \max_{\theta^* \in \mathbb{B}(\theta_t)} \Phi^\rho(\theta_t, \theta^*; x_0),$$

where $\mathbb{B}(\theta_t) := \{a \in \mathbb{A}(\theta_t) \cup \{\underline{\theta}\} : a < \theta_t\}$. Whenever $\Psi(\theta_t; x_0)$ is negative, the seller is willing to exit. Using the maximal virtual payoff we can characterize the commitment monopolist's optimal strategy as follows (full proof in Appendix C).

Proposition 3. *In the continuous time limit, the commitment solution is such that*

[13] The intuition behind Lemma 6 is straightforward: selling at a lower than the full rate means that the monopolist is simply discounting more. One way to see this is that the monopolist's payoff is linear in the discounting term $e^{-r\tau(\theta)}$. The commitment monopolist does not need to delay sales, because she can simply commit to selling less in the future.

(i) absent news the monopolist sells at full rate for all types greater than θ_N. The cutoff θ_N is the largest θ such that

$$\left(\theta + \delta_G \Phi_G^*(\theta) - \frac{1 - F(\theta)}{f(\theta)}\right) x(\theta; x_0) = c$$

and $\Psi(\theta; x_0) < 0$. That is, absent news $\lambda(\theta) = \rho$ for $\theta \geq \theta_N$ and $\lambda(\theta) = 0$ for $\theta < \theta_N$.

(ii) after good news the monopolist sells at full rate until

$$\theta_G - \frac{1 - F(\theta_G)}{f(\theta_G)} = c$$

for a unique $\theta_G \in [\underline{\theta}, \overline{\theta}]$ independent of past history. That is, $\lambda(\theta) = \rho$ for $\theta \geq \theta_G$ and $\lambda(\theta) = 0$ for $\theta < \theta_G$.

In the first part of Proposition 3, the sufficiency condition is needed to take into account the value of information as sales today may bring value tomorrow. The second part of Proposition 3 says that the timing of good news does not affect the monopolist's strategy after good news in any way: whenever good news arrives the final allocation always corresponds to the static monopoly allocation. This is because the commitment monopolist takes into account *all* the rents to buyers caused by selling after good news. These rents accrue both *before* good news arrives and *after* good news has arrived.

An important part of Proposition 3 is that we are able to transform the monopolist's dynamic optimization problem into a static mechanism design problem. This simplifies the analysis significantly. While the solution is thus similar to static mechanism design, there are also important differences stemming from the dynamic nature of the underlying problem. The most important of these is that the evolution of the belief is an equilibrium object and thus if the seller would allocate the product differently (e.g. non-monotonically) the virtual surplus would also change. This also suggests that the negative virtual surplus types represent the dynamic cost of information that stems from the belief process and the buyers' ability to wait. The seller has to pay this cost in order to produce information.

As a final point, the apparatus we use in characterizing the commitment solution can readily be used to characterize the constrained efficient solution as well.

We only need to replace the virtual surplus with the total marginal surplus.[14] All the results, such as Lemma 6, then go through with similar arguments: the constrained efficient solution is such that the monopolist sells at full rate until the total marginal surplus equals zero and the value of future sales (now defined over the marginal total surplus) is negative.

5.2 Comparison

5.2.1 Expected sales

We next compare stationary equilibria to the commitment solution and analyze how the lack of commitment power affects the monopolist's behavior. We first have the following result regarding entry and exit (proof in Appendix C):

Proposition 4. *Suppose $\delta_G > 0$. Then there exists an initial belief such that the commitment monopolist enters and the equilibrium monopolist does not enter. Furthermore, there exists $\bar{\delta}_B > \delta_G$ such that for all $\delta_B < \bar{\delta}_B$ there exists an initial belief such that even though the equilibrium monopolist enters, the expected sales are greater in the commitment solution.*

Proposition 4 gives conditions under which the commitment monopolist sells more in expectation than the equilibrium monopolist. The result follows because the commitment monopolist internalizes the value of information better than the equilibrium monopolist whose value is limited by Coasian dynamics. Thus, for low enough initial beliefs, the commitment monopolist has stronger incentives to sell. Whenever the commitment monopolist sells more, she also delivers more buyer surplus because the commitment monopolist sells at full rate. The commitment solution then also generates more information.

Besides the expected amount of sales, the commitment solution may be more efficient in terms of faster sales because the commitment monpolist always sells at full rate (Proposition 3). An interesting observation is that while the result regarding the amount of sales, Proposition 4, assumes that learning must happen at least partially through good news, stationary equilibria feature delays only if

[14] The total marginal surplus is $\theta + \delta_G(V^G(\theta) + W^G(\theta)))x(\theta) - c$, where the after good news values are the same as in the stationary equilibrium.

learning happens through bad news (Theorem 2). Thus, the learning technology matters for the inefficiencies of the stationary equilibria.

Intuitively, the reason the *type* of news matters is because good news and bad news provide decision relevant information at different points in the game. The arrival of good news matters only at the exit type because good news is decision relevant only if the monopolist is willing to exit. In contrast, the arrival of bad news is decision relevant only if the monopolist is willing to sell or for types greater than the exit type. Because good news and bad news provide information at different margins, their arrival rates affect equilibrium behavior differently.

5.2.2 Prices

The previous subsection gives conditions when the commitment monopolist's behavior is more efficient in terms of larger or faster sales. When the commitment monopolist sells more, it immediately follows that prices are also lower because there are no delays in the commitment solution. A lower price in the commitment solution automatically means a higher consumer surplus for the same reason.

To understand how delays affect the initial price, consider the same pure bad news example with infinite capacity as in Section 4.4 but in addition let the buyers' types be uniformly distributed on $[0, 1]$. All stationary equilibria then feature delays and if the equilibrium monopolist enters, she must sell more than the commitment monopolist.[15] The comparison of prices reduces to the comparison of buyers' values in the two solution concepts so that if the buyers' value in the commitment solution is larger, the price must be lower. For our purposes, it is enough to concentrate on the equilibrium in which the buyers get all the surplus: the price equals the marginal cost and the buyers' value equals their valuation minus the marginal cost, $\theta x(\theta) - c$.[16] In our example, a buyer's value in the commitment solution equals $W^{\infty}(\theta; \theta, x_0) = x(\theta)(\theta - \theta_N)$ because purchasing times are zero for

[15] This follows because with pure bad news learning the marginal joint payoff is greater than the virtual surplus for every type.

[16] It is straightforward to check that this satisfies all the necessary conditions in Lemma 3 when $\rho \to \infty$.

types who buy. Thus, the initial price is lower in the commitment solution when

$$x_0(\overline{\theta} - \theta_N) > \overline{\theta}x_0 - c \iff x_0\theta_N < c.$$

The commitment exit type is such that $x(\theta_N)(2\theta_N - 1) = c$. Thus, the commitment monopolist sets a lower initial price whenever the initial belief is low enough so that $x_0 < (2 - (x_0 + (1 - x_0)e^{-\delta_B(1-F(\theta_N))})c$. Given the uniform distribution, the equilibrium monopolist enters if $x_0 > c$. Therefore, in this example, there is always an initial belief such that the equilibrium monopolist sells more but the initial prices are lower in the commitment solution.

The preceding example makes it clear that delays can cause equilibrium prices to be higher than those in the commitment solution. When this happens, the equilibrium monopolist is somewhat counter-intuitively setting higher initial prices and earning lower expected profits. The explanation for this is that the commitment monopolist is setting lower prices precisely to maximize expected profits: she sells fast in the beginning in order to produce information. Slowing down sales is sub-optimal, but it is the only way the equilibrium monopolist can reduce the buyers' value.

The interesting observation here is that Coasian dynamics, the equilibrium monopolist's temptation to speed up sales, can lead to higher prices. Under suitable conditions the monopolist's and buyer's preferences over monopolist's commitment power are thus fully aligned. Monopolist's commitment power can benefit all market participants in a market where buyers learn from each other.

6 Conclusions

The ultimate test for any product is releasing it to the market. Many new products come with monopoly power in the form of copyright and patent protections. This paper studies behavior in a market in which both a monopolistic seller and potential buyers are learning how valuable a new product is. The value, or quality, is slowly revealed as buyers purchase the product. A central part of the model is that buyers are timing their purchases which means future prices affect purchasing decisions today. We show that in a new market, where endogenous learning about quality is still present, stationary equilibria may well be inefficient in the sense of too small and too slow sales because the monopolist internalizes too little of the available surplus. Because the commitment monopolist is able to capture a larger share of the surplus, we show that commitment power can be beneficial for both the monopolist and the buyers. This contradicts the conventional wisdom from the Coase conjecture that sequential rationality leads to full efficiency while commitment power causes a loss in welfare.

One interpretation for our results is that the Coase conjecture has two separate implications in the complete information environment. First, it implies that the seller's profits equal zero because she speeds up sales until her value is zero. Second, it implies that the allocation corresponds to the first best. When endogenous learning is present the first statement no longer implies the second. This is because the seller's action today depends on her continuation value so it is not enough to leave her indifferent between selling and exiting unlike with complete information.

References

Ausubel, Lawrence and Raymond Deneckere, "Reputation in Bargaining and Durable Goods Monopoly," *Econometrica*, 1989, *57* (3), 511–531.

Board, Simon, "Selling options," *Journal of Economic Theory*, 2007, *136* (1), 324 – 340.

_ , "Durable-Goods Monopoly with Varying Demand," *The Review of Economic Studies*, 2008, *75* (2), 391–413.

Bonatti, Alessandro, "Menu pricing and learning," *American Economic Journal: Microeconomics*, 2011, *3* (3), 124–163.

Bose, Subir, Gerhard Orosel, Marco Ottaviani, and Lise Vesterlund, "Dynamic monopoly pricing and herding," *The RAND Journal of Economics*, 2006, *37* (4), 910–928.

Bulow, Jeremy, "Durable-Goods Monopolists," *Journal of Political Economy*, 1982, *90* (2), 314–332.

Che, Yeon-Koo and Johannes Hörner, "Recommender Systems as Mechanisms for Social Learning," *The Quarterly Journal of Economics*, 2017, pp. 871–925.

Deneckere, Raymond and Meng-Yu Liang, "Bargaining with interdependent values," *Econometrica*, 2006, *74* (5), 1309–1364.

Frick, Mira and Yuhta Ishii, "Innovation Adoption by Forward-Looking Social Learners," *Working Paper*, 2016.

Fuchs, William and Andrzej Skrzypacz, "Bargaining with arrival of new traders," *American Economic Review*, 2010, *100* (3), 802–36.

_ and _ , "Bridging the gap: Bargaining with interdependent values," *Journal of Economic Theory*, 2013, *148* (3), 1226–1236.

Fudenberg, Drew, David K. Levine, and Jean Tirole, "Infinite-horizon models of bargaining with one-sided incomplete information," *in Game-theoretic models of bargaining, ed. by Alvin E. Roth, Cambridge University Press, West Nyack, NY*, 1985, pp. 73–98.

Gul, Faruk, Hugo Sonnenschein, and Robert Wilson, "Foundations of dynamic monopoly and the coase conjecture," *Journal of Economic Theory*, 1986, *39* (1), 155–190.

Keller, Godfrey and Sven Rady, "Breakdowns," *Theoretical Economics*, 2015, *10* (1), 175–202.

_ , _ , and Martin Cripps, "Strategic Experimentation with Exponential Bandits," *Econometrica*, 2005, *73* (1), 39–68.

Kremer, Ilan, Yishay Mansour, and Motty Perry, "Implementing the "wisdom of the crowd"," *Journal of Political Economy*, 2014, *122* (5), 988–1012.

Laiho, Tuomas, Pauli Murto, and Julia Salmi, "Gradual Learning from Incremental Actions," *Working paper*, 2020.

Lomys, Niccolo, "Learning while Trading: Experimentation and Coasean Dynamics," *Working paper*, 2018.

Mcafee, R. Preston and Thomas Wiseman, "Capacity Choice Counters the Coase Conjecture," *The Review of Economic Studies*, 2008, *75* (1), 317–331.

Milgrom, Paul and Ilya Segal, "Envelope theorems for arbitrary choice sets," *Econometrica*, 2002, *70* (2), 583–601.

Ortner, Juan, "Durable goods monopoly with stochastic costs," *Theoretical Economics*, 2017, *12* (2), 817–861.

Simonoff, Jeffrey S and Lan Ma, "An empirical study of factors relating to the success of Broadway shows," *The Journal of Business*, 2003, *76* (1), 135–150.

Sobel, Joel, "Durable Goods Monopoly with Entry of New Consumers," *Econometrica*, 1991, *59* (5), 1455–1485.

Stokey, Nancy, "Intertemporal price discrimination," *The Quarterly Journal of Economics*, 1979, pp. 355–371.

Stokey, Nancy L, "Rational expectations and durable goods pricing," *The Bell Journal of Economics*, 1981, pp. 112–128.

von der Fehr, Nils-Henrik and Kai-Uwe Kühn, "Coase versus Pacman: Who Eats Whom in the Durable-Goods Monopoly?," *Journal of Political Economy*, 1995, *103* (4), 785–812.

Weng, Xi, "Dynamic pricing in the presence of individual learning," *Journal of Economic Theory*, 2015, *155*, 262 – 299.

Appendices

A Buyers' problem

Envelope theorem. Because purchasing times are optimal, we can use the envelope theorem to simplify the buyers' value (Milgrom and Segal (2002)). Let good news arrive when θ_k is the current highest type. The value of a buyer with type θ equals

$$W^G_\Delta(\theta;\theta_t,\theta_k) = \int_{\underline{\theta}}^{\theta} e^{-r(\tau^G(s;\theta_k)-t)} ds, \tag{10}$$

where θ_t is the current highest type at the time when the value is evaluated (at t). The lowest type never buys so his value, $W^G_\Delta(\underline{\theta};\theta_t,\theta_k)$, is zero. We can similarly use the envelope theorem to write the buyer's value absent news:

$$\begin{aligned} W_\Delta(\theta;\theta_t,x_0) = &\int_{\underline{\theta}}^{\theta} e^{-r(\tau(s)-t)}\beta(\theta_t,s)x(s)ds \\ &+ \int_{\underline{\theta}}^{\theta} \sum_{k=t+\Delta}^{\tau(s)} e^{-r(\tau^G(s;\theta_k)-t)} \Pr(\sigma_G = k|\sigma > t)ds. \end{aligned} \tag{11}$$

The envelope theorem representation together with (1) pins down the price at time t as a function of the purchasing times (τ^G, τ).

Continuous time limit. After good news, the only thing that depends on the period length is that purchasing times come from a discrete grid: $\tau^G(\theta;k) \in \{k, k+\Delta, k+2\Delta, \ldots\}$. Thus in the limit as $\Delta \to 0$ (10) becomes $W^G(\theta;\theta_t,\theta_k) = \int_{\underline{\theta}}^{\theta} e^{-r(\tau^G(s;\theta_k)-k)}ds$, where $\tau^G(\theta;k) \in [k,\infty)$. Then, as $\Delta \to 0$, the buyer's value absent news (11) becomes:

$$\begin{aligned} W(\theta;\theta_t,x_0) = &\int_{\underline{\theta}}^{\theta} e^{-r(\tau(s)-t)}\beta(\theta_t,s)x(s)ds \\ &+ \int_{\underline{\theta}}^{\theta} \int_t^{\tau(s)} e^{-r(\tau^G(s;\theta_k)-k)}\beta(\theta_t,\theta_k)\delta_G\lambda_k x_k dk ds, \end{aligned}$$

where $\beta(\theta_t,\theta_k)\delta_G\lambda_k x_k$ is the limit of $\Pr(\sigma_G = k|\sigma > t)$ (follows from the definition in Section 3). Notice also that $\beta(\theta_t,s)x(s) = x_t e^{-\delta_G(F(\theta_t)-F(s))}$. The law of motion for the current highest type, $dt = -f(\theta_t)/\lambda_t d\theta_t$, then allows us to rewrite the buyer's value as purely an integral over the types:

$$
\begin{aligned}
W(\theta;\theta_t,x_0) =& x(\theta_t)\int_{\underline{\theta}}^{\theta} e^{-r(\tau(s)-\tau(\theta_t))-\delta_G(F(\theta_t)-F(s))}ds \\
&+ x(\theta_t)\int_{\underline{\theta}}^{\theta}\int_{s}^{\theta_t} e^{-r(\tau^G(s;z)-\tau(\theta_t))-\delta_G(F(\theta_t)-F(z))}\delta_G f(z)dzds.
\end{aligned}
$$

As the likelihood, $q_t = x_t/(1-x_t)$, is a simpler process to handle, we make frequent use of it in the equations. Using the likelihood, the buyer's value becomes:

$$
\begin{aligned}
W(\theta;\theta_t,x_0) =& (1-x_t)\int_{\underline{\theta}}^{\theta} e^{-r(\tau(s)-t)-\delta_B(F(\theta_t)-F(s)))}q_{\tau(s)}ds \\
&+ (1-x_t)\int_{\underline{\theta}}^{\theta}\int_{t}^{\tau(s)} e^{-r(k-t)-\delta_B(F(\theta_t)-F(k))}\lambda_k\delta_G q_k e^{-r(\tau^G(s;k)-k)}dkds.
\end{aligned} \tag{12}
$$

B Stationary equilibria

B.1 Proof of the Coase conjecture

We provide here a general proof for the Coase conjecture which covers situations beyond our model such as different forms of uncertainty and mixed strategies.

In the appendix, we drop the subscript Δ whenever possible. We also use $W^G(\theta) := W^G(\theta;\theta,\theta_k)$ and $W(\theta) := W(\theta;\theta,x_0)$ to denote the current highest type's value after good news and absent news respectively.

Preliminaries

Uncertainty. We begin by introducing a mapping $\xi : [\underline{\theta},\overline{\theta}] \to [0,1]$ which maps the current highest type to beliefs. Together with a sales strategy $\theta_+ : [0,1] \times [\underline{\theta},\overline{\theta}] \to [\underline{\theta},\overline{\theta}]$ induced by Λ, ξ defines a path of realized beliefs and types so that the next highest type is given by $\theta_+(\xi(\theta),\theta)$, which we can then use to write the belief in the next period as $\xi(\theta_+(\xi(\theta),\theta))$ and so on.[17]

Let Ξ be the family of functions from $[\underline{\theta},\overline{\theta}]$ to $[0,1]$. We can then define a probability distribution functions $H_1 : \Xi \to [0,1]$, $H_2 : \triangle\Theta \to [0,1]$ and $H : \Xi \times \triangle\Theta \to [0,1]$. The function H assigns probability weights on all possible (and impossible) paths in the state space. Any equilibrium behavior can

[17]The proof easily extends to mixed strategies, when the uncertainty is over $\xi^* := (\xi,\theta_+)$.

then be characterized by a mapping H and any realized equilibrium path can be characterized by a mapping (ξ, θ_+) and an initial belief x_0.

The players do not observe ξ but can make inferences from the realized path. Because players are using Markovian strategies and the learning process is Markov, we can limit to Markovian randomization. We write $H_{(x,\theta)}$ for the conditional probability of a given state (x, θ). The current state is a sufficient statistic of the entire past and, therefore, specifies the most informative filtration on $\Xi \times \triangle\Theta$.

Notice that the proofs in this section would remain identical if instead of the exponential bandit used some other endogenous learning process that only depends on the current action. Any process where *the distribution* of net surplus, consumer's valuation minus the production cost, depends only on the cumulative past sales can be characterized as a randomization over ξ. This includes a learning model where each new sale generates a normally distributed signal (Brownian learning) and a model where the marginal cost of production depends, deterministically or stochastically, on past sales.

Prices in discrete time. Current price depends on the state, (θ, x), and on the last type the monopolist serves this period, θ_+. The discrete time equilibrium price is hence a function of three variables: $P(\theta_+; \theta, x)$. In any stationary equilibrium, the price must be such that the lowest type is indifferent between buying and waiting: $P(\theta_+; \theta, x) = x\theta_+ - e^{-r\Delta}\mathbb{E}_\xi[W(\theta_+)|\theta, x]$. This implies the following recursive pricing behavior:

$$P(\theta_+; \theta, x) = (1 - e^{-r\Delta})x\theta_+ + e^{-r\Delta}\mathbb{E}_\xi[P(\theta_{++}; \theta_+, \xi(\theta_+))|\theta, x], \tag{13}$$

where $\theta_{++} = \xi_2(\xi_1(\theta_+), \theta_+)$ is the marginal buyer and $\xi_1(\theta_+)$ the belief in the following period. Notice that the belief is a martingale and therefore the expected quality tomorrow equals the expected quality today.

Proof idea. The Coase conjecture is usually proved by looking at deviations which jump over one equilibrium state and arguing that if the monopolist is unwilling to do so her value must be low (see e.g. Gul, Sonnenschein and Wilson (1986) and Fuchs and Skrzypacz (2010)). Thus, either the monopolist's value is low or there must not be delays in sales. We run into two problems with this argument in our model. First, once learning is introduced sales at full rate may lead to a negative payoff for the monopolist. Hence, we do not have a strictly positive lower bound for the monopolist's payoff. Second, jumps over equilibrium states may violate the capacity constraint. Instead, we need to show that the monopolist would prefer speeding up sales to off path states. To do that, we need to prove additional continuity results between the equilibrium and off-the-equilibrium states.

Lemma 7. *There exists $B > 0$ such that for all $\Delta > 0$ and for all stationary equilibria,*

$$
\begin{aligned}
&|P(\theta_+;\theta,\xi_1(\theta)) - P(\theta'_+;\theta',\xi_1(\theta'))| \\
&\leq B(|F(\theta) - F(\theta')| + |F(\theta_+) - F(\theta'_+)| + |\xi_1(\theta) - \xi_1(\theta')| + \Delta + A)
\end{aligned}
$$

where A is such that $|\tau(\theta_+) - \tau(\theta'_+)| \leq A$ almost surely.

Proof. Let $\theta_+ > \theta'_+$ without loss of generality. We add and subtract a term $P(\theta'_+;\theta,\xi(\theta))$ to split the problem into two parts: $|P(\theta_+;\theta,\xi(\theta)) - P(\theta'_+;\theta',\xi(\theta'))|$ equals

$$
\begin{aligned}
&|P(\theta_+;\theta,\xi(\theta)) - P(\theta'_+;\theta,\xi(\theta)) + P(\theta'_+;\theta,\xi(\theta)) - P(\theta'_+;\theta',\xi(\theta'))| \\
&\leq |P(\theta_+;\theta,\xi(\theta)) - P(\theta'_+;\theta,\xi(\theta))| + |P(\theta'_+;\theta,\xi(\theta)) - P(\theta'_+;\theta',\xi(\theta'))|,
\end{aligned}
$$

where the inequality follows from the triangle inequality. The first term to the right of the inequality keeps the initial state constant but varies the amount of sales. The second term to the right of the inequality keeps the continuation state constant but varies the initial state. Part (i): We start by finding an upper bound for the second term. The effect of the initial state is purely mechanical: the initial state affects the relative probabilities of different continuation scenarios but does not change them. The following upper bound thus holds for the second term:

$$
\begin{aligned}
&|P(\theta'_+;\theta,\xi(\theta)) - P(\theta'_+;\theta',\xi(\theta'))| \\
&= \Big|\mathbb{E}\Big[(1-e^{-r\Delta})\xi(\theta)\theta'_+ + e^{-r\Delta}P(\theta'_{++};\theta'_+,\xi(\theta'_+))\Big|\theta,\xi(\theta)\Big] \\
&- \mathbb{E}\Big[(1-e^{-r\Delta})\xi(\theta')\theta'_+ + e^{-r\Delta}P(\theta'_{++};\theta'_+,\xi(\theta'_+))\Big|\theta',\xi(\theta')\Big]\Big| \\
&\leq |\xi(\theta)-\xi(\theta')|(1-e^{-r\Delta})\theta'_+ + e^{-r\Delta}\Big|\mathbb{E}\Big[P(\theta'_{++};\theta'_+,\xi(\theta'_+))\Big|\theta,\xi(\theta)\Big] \\
&- \mathbb{E}\Big[P(\theta'_{++};\theta'_+,\xi(\theta'_+))\Big|\theta',\xi(\theta')\Big]\Big|.
\end{aligned}
$$

We can further evaluate the second term. Let the current time be t, then:

$$\begin{aligned}
&\left|\mathbb{E}\Big[P(\theta'_{++};\theta'_{+},\xi(\theta'_{+}))\Big|\theta,\xi(\theta)\Big] - \mathbb{E}\Big[P(\theta'_{++};\theta'_{+},\xi(\theta'_{+}))\Big|\theta',\xi(\theta')\Big]\right| \\
&= \Big|(Pr(\sigma > t+\Delta|\theta,\xi(\theta)) - Pr(\sigma > t+\Delta|\theta',\xi(\theta')))\mathbb{E}\Big[P(\theta'_{++};\theta'_{+},x(\theta'_{+}))\Big] \\
&+ (Pr(\sigma_G = t+\Delta|\theta,\xi(\theta)) - Pr(\sigma_G = t+\Delta|\theta',\xi(\theta')))\mathbb{E}\Big[P(\theta'_{++};\theta'_{+},1)\Big]\Big| \\
&\leq |Pr(\sigma > t+\Delta|\theta,\xi(\theta)) - Pr(\sigma > t+\Delta|\theta',\xi(\theta'))|\mathbb{E}\Big[P(\theta'_{++};\theta'_{+},x(\theta'_{+}))\Big] \\
&+ |Pr(\sigma_G = t+\Delta|\theta,\xi(\theta)) - Pr(\sigma_G = t+\Delta|\theta',\xi(\theta'))|\mathbb{E}\Big[P(\theta'_{++};\theta'_{+},1)\Big].
\end{aligned}$$

We have that

$$\begin{aligned}
&|Pr(\sigma > t+\Delta|\theta,\xi(\theta)) - Pr(\sigma > t+\Delta|\theta',\xi(\theta'))| \\
&\leq B_1(|\xi(\theta)-\xi(\theta')| + |F(\theta)-F(\theta')|), \\
&|Pr(\sigma_G = t+\Delta|\theta,\xi(\theta)) - Pr(\sigma_G = t+\Delta|\theta',\xi(\theta'))| \\
&\leq B_2(|\xi(\theta)-\xi(\theta')| + |F(\theta)-F(\theta')|),
\end{aligned}$$

for some $B_1, B_2 > 0$. Therefore, there must exist B' such that $|P(\theta'_{+};\theta,\xi(\theta)) - P(\theta'_{+};\theta',\xi(\theta'))|$ is bounded above by $B'(|\xi(\theta)-\xi(\theta')| + |F(\theta)-F(\theta')|)$.

Part (ii): we find an upper bound for $|P(\theta_{+};\theta,\xi(\theta)) - P(\theta'_{+};\theta,\xi(\theta))|$. In this difference the processes end up in different states and hence expectations about future prices may be potentially different. Let $(\theta_{+} = \theta_0, \theta_1, \dots)$ and $(\theta'_{+} = \theta'_0, \theta'_1, \dots)$ be some realized equilibrium paths (where news may or may not arrives) consistent with the initial state $(\theta,\xi(\theta))$. We use a shorthand notation $(P_0, P_1, \dots)$ and $(P'_0, P'_1, \dots)$ for the corresponding realized price paths.

The idea is to use the buyers' incentive compatibility to get an upper bound for $P_0 - P_m$ and $P'_0 - P'_{m'}$. We then use the monopolist's sequential optimality to show that $P_m - P'_{m'}$ must be small for some m and m'. Combining the two gives an upper bound for the initial price difference, $P_0 - P'_0$.

First, we use buyers' incentive compatibility (use (13) repeatedly) so that we can write the initial price difference as

$$\begin{aligned}
P_0 - P'_0 = \mathbb{E}\Big[&(1-e^{-r\Delta})\sum_{i=0}^{m-1} e^{-r\Delta i}x_i\theta_i + e^{-r\Delta m}P_m \\
&- (1-e^{-r\Delta})\sum_{i=0}^{m'-1} e^{-r\Delta i}x'_i\theta'_i + e^{-r\Delta m'}P'_{m'})|\theta,\xi(\theta)\Big] \qquad (14)
\end{aligned}$$

where we use a shorthand θ_i for $\theta_+(\xi(\theta_{i-1}), \theta_{i-1})$ and x_i for $\xi(\theta_i)$.

Next we bound the difference for an arbitrary realized path of news by using the monopolist's sequential rationality. We do that in three steps: 1) We argue that the effect of the sales that happen before the path (θ_i) reaches θ' can be ignored and re-index the path starting from the last period before reaching θ'; 2) We show that the difference in payoffs between the two paths, $\sum_{i=0}^{m-1} e^{-r\Delta i} x_i \theta_i - \sum_{i=0}^{m-1} e^{-r\Delta i} x_i' \theta_i'$, is bounded by θ_0 as long as $\theta_i' \geq \theta_{i+1}$ for all $i < \hat{m}$; 3) then we show the result separately for the cases when there exists/ does not exist i such that $\theta_i' < \theta_{i+1}$.

Step 1): Let the path $(\theta_0, \theta_1, \dots)$ reach θ' at θ_n, i.e. $\theta_i > \theta'$ for all $i < n$ and $\theta_n \geq \theta'$. Path (θ_i) reaches θ' in finite time because $\tau(\theta_0') - \tau(\theta_0) \leq A$. The difference inside the expectation in (14) is then at most

$$\begin{aligned} &(1 - e^{-rA}) \max_{i \in \{0,\dots,n-1\}} (x_i \theta_i) \\ &+ (1 - e^{-r\Delta}) \sum_{i=n}^{m} e^{-r\Delta i} (x_i \theta_i - x_{i-n}' \theta_{i-n}') + e^{-r\Delta m} (P_m - P_{m-n}'). \end{aligned} \tag{15}$$

The first term is at most equal to $rA\overline{\theta}$. We then re-index the sum and the the price difference $P_m - P_{m-n}'$ so that $\hat{\theta}_i = \theta_{i+n}$, $\hat{x}_i = x_{i+n}$ and $\hat{m} = m - n$.

Step 2): Let $\hat{m}$ be such that $\theta_i' \geq \hat{\theta}_{i+1}$ for all $i < \hat{m}$. This gives $\hat{x}_i \hat{\theta}_i - x_i' \theta_i' = \hat{x}_i (\hat{\theta}_i - \theta_i') + (\hat{x}_i - x_i') \theta_i' \leq \hat{x}_i (\hat{\theta}_i - \hat{\theta}_{i+1}) + (\hat{x}_i - x_i') \hat{\theta}_{i+1}$. This allows us to find an upper bound for the remaining two terms in (15):

$$(1 - e^{-r\Delta}) \sum_{i=0}^{\hat{m}-1} e^{-r\Delta i} \Big(\hat{x}_i (\hat{\theta}_i - \hat{\theta}_{i+1}) + (\hat{x}_i - x_i') \hat{\theta}_{i+1} \Big) + e^{-r\Delta \hat{m}} (\hat{P}_{\hat{m}} - P_{\hat{m}}'). \tag{16}$$

We can then use that $\sum_{i=0}^{\hat{m}-1} \hat{x}_i (\hat{\theta}_i - \hat{\theta}_{i+1}) \leq \hat{\theta}_0 - \hat{\theta}_{\hat{m}}$. To evaluate $\sum_{i=0}^{\hat{m}-1} (\hat{x}_i - x_i') \hat{\theta}_{i+1}$, notice that $\hat{x}_i - x_i > 1 - e^{-(\delta_G + \delta_B)\Delta\rho}$ at most once when $|F(\hat{\theta}_i) - F(\theta_i')| \leq \rho\Delta$. This gives $\sum_{i=0}^{\hat{m}-1} (\hat{x}_i - x_i') \hat{\theta}_{i+1} \leq \theta_0 (1 - e^{-(\delta_G + \delta_B)\Delta\rho}) + \theta_0$.

Step 3): To find an upper bound for the last term of (16), $e^{-r\Delta \hat{m}} (\hat{P}_{\hat{m}} - P_{\hat{m}}')$, assume first that $\theta_i' \geq \hat{\theta}_{i+1}$ for all i. We can choose an arbitrary large $\hat{m}$ in (16). Because we have assumed $\theta_i' \geq \hat{\theta}_{i+1}$ for all i, the lowest θ' cannot be lower than the lowest $\hat{\theta}$. Hence, $\hat{P}_{\hat{m}} - P_{\hat{m}}'$ has an upper bound close to exit proportional to $\theta_{\hat{m}}' - \hat{\theta}_{\hat{m}}$: $\hat{P}_{\hat{m}} - P_{\hat{m}}' \leq D(\theta_{\hat{m}}' - \hat{\theta}_{\hat{m}}) < D\rho\Delta$ where D is some constant.

We can now focus on the case where there exists $\hat{m}$ such that $\theta_{\hat{m}}' < \hat{\theta}_{\hat{m}+1}$. Take the smallest such $\hat{m}$. Then we have $F(\theta_{\hat{m}-1}') - F(\hat{\theta}_{\hat{m}+1}) \leq \rho\Delta$ and $F(\hat{\theta}_{\hat{m}}) - F(\theta_{\hat{m}}') \leq$

$\rho\Delta$ and hence it is feasible to sell up to $\theta'_{\hat{m}}$ from $\hat{\theta}_{\hat{m}}$ and up to $\hat{\theta}_{\hat{m}+1}$ from $\theta'_{\hat{m}-1}$. The seller's sequential optimality requires:

$$
\begin{aligned}
&(P(\hat{\theta}_{\hat{m}+1};\theta_{\hat{m}},\hat{x}_{\hat{m}})-c)(F(\hat{\theta}_{\hat{m}})-F(\hat{\theta}_{\hat{m}+1}))+e^{-r\Delta}\mathbb{E}(V(\hat{\theta}_{\hat{m}+1},x)|\theta_{\hat{m}},\hat{x}_{\hat{m}})\geq\\
&(P(\theta'_{\hat{m}};\hat{\theta}_{\hat{m}},\hat{x}_{\hat{m}})-c)(F(\theta_{\hat{m}})-F(\theta'_{\hat{m}}))+e^{-r\Delta}\mathbb{E}(V(\theta'_{\hat{m}},x)|\hat{\theta}_{\hat{m}},\hat{x}_{\hat{m}})\\
&(P(\theta'_{\hat{m}};\theta'_{\hat{m}-1},x'_{\hat{m}-1})-c)(F(\theta'_{\hat{m}-1})-F(\theta'_{\hat{m}}))+e^{-r\Delta}\mathbb{E}(V(\theta'_{\hat{m}},x)|\theta'_{\hat{m}-1},x'_{\hat{m}-1})\geq\\
&(P(\hat{\theta}_{\hat{m}+1};\theta'_{\hat{m}-1},x'_{\hat{m}-1})-c)(F(\theta'_{\hat{m}-1})-F(\hat{\theta}_{\hat{m}+1}))+e^{-r\Delta}\mathbb{E}(V(\hat{\theta}_{\hat{m}+1},x)|\theta'_{\hat{m}-1},x'_{\hat{m}-1}).
\end{aligned}
$$

Combining these two optimality conditions then gives

$$
\begin{aligned}
&(P(\hat{\theta}_{\hat{m}+1};\hat{\theta}_{\hat{m}},\hat{x}_{\hat{m}})-c)(F(\hat{\theta}_{\hat{m}})-F(\hat{\theta}_{\hat{m}+1}))-(P(\theta'_{\hat{m}};\hat{\theta}_{\hat{m}},\hat{x}_{\hat{m}})-c)(F(\hat{\theta}_{\hat{m}})-F(\theta'_{\hat{m}}))\\
&+(P(\theta'_{\hat{m}};\theta'_{\hat{m}-1},x'_{\hat{m}-1})-c)(F(\theta'_{\hat{m}-1})-F(\theta'_{\hat{m}}))-(P(\hat{\theta}_{\hat{m}+1};\theta'_{\hat{m}-1},x'_{\hat{m}-1})-c)\\
&\times(F(\theta'_{\hat{m}-1})-F(\hat{\theta}_{\hat{m}+1}))+e^{-r\Delta}\Big(\mathbb{E}(V(\hat{\theta}_{\hat{m}+1},x)|\hat{\theta}_{\hat{m}},\hat{x}_{\hat{m}})-\mathbb{E}(V(\hat{\theta}_{\hat{m}+1},x)|\theta'_{\hat{m}-1},x'_{\hat{m}-1})\Big)\\
&+e^{-r\Delta}\Big(\mathbb{E}(V(\theta'_{\hat{m}},x)|\theta'_{\hat{m}-1},x'_{\hat{m}-1})-\mathbb{E}(V(\theta'_{\hat{m}},x)|\hat{\theta}_{\hat{m}},\hat{x}_{\hat{m}})\Big)\geq 0. \qquad (17)
\end{aligned}
$$

Notice that we can use the same logic as in Part (i) to show that the initial state has a bounded effect on the prices and expected continuation values: exists a constant B_3 such that $B_3\Delta(F(\hat{\theta}_{\hat{m}})-F(\theta'_{\hat{m}-1}))$ is an upper bound for the combined the effect of the different initial states. We then ignore the differences in initial states and get $(\hat{P}_{\hat{m}+1}-c)(F(\hat{\theta}_{\hat{m}})-F(\hat{\theta}_{\hat{m}+1}))-(P'_{\hat{m}}-c)(F(\hat{\theta}_{\hat{m}})-F(\theta'_{\hat{m}}))+(P'_{\hat{m}}-c)(F(\theta'_{\hat{m}-1})-F(\theta'_{\hat{m}}))-(\hat{P}_{\hat{m}+1}-c)(F(\theta'_{\hat{m}-1})-F(\hat{\theta}_{\hat{m}+1}))=(P'_{\hat{m+1}}-\hat{P}_{\hat{m}})(F(\hat{\theta}_{\hat{m}})-F(\theta'_{\hat{m}-1})))$.

As $\hat{\theta}_{\hat{m}}<\theta'_{\hat{m}-1}$, this then implies that the inequality (17) can only hold if $P_{\hat{m}+1}-P'_{\hat{m}}\leq B_3\Delta$. Hence, (16) is bounded in this case, too.

A lower bound for the difference $P(\theta_+;\theta,\xi(\theta))-P(\theta'_+;\theta,\xi(\theta))$ can be found similarly and the analysis is hence omitted. The claim then follows from the existence of these upper and lower bounds. □

An important part of the proof is that intervals $[\theta'_i,\hat{\theta}_i]$ and $[\theta'_{i+1},\hat{\theta}_{i+1}]$ do not overlap. Because of this property, $\sum_i(\hat{\theta}_i-\theta'_i)\leq\theta_0$.

Lemma 7 guarantees that the expected price does not jump if both the types and the expected purchasing times are close to each other. The price may jump however if there are periods when the monopolist does not sell with a positive probability. We have the following result for this case.

Lemma 8. *For all $\epsilon > 0$ and $A > 0$, there exists $\bar{\Delta}$ such that for all $\Delta < \bar{\Delta}$ and any equilibrium (θ, x) at least one of the following holds for all stationary equilibria:*

- $\tau_\Delta(\theta') - \tau_\Delta(\theta) \leq A$ *a.s. for all θ' such that $F(\theta) - F(\theta') \in (0, \rho\Delta)$. Or*
- $V_\Delta(\theta, x) \leq \epsilon$.

Proof. Assume instead that $F(\theta) - F(\theta') \leq \rho\Delta$ and that $\tau^\xi_\Delta(\theta') - \tau^\xi_\Delta(\theta) > A$ for some ξ. It is a feasible deviation to sell up to θ' immediately. The monopolist's sequential rationality then requires that:

$$\sum_{i=1}^{m-1} e^{-r\Delta(i-1)}(F(\theta_{i-1} - F(\theta_i))(P(\theta_i; \theta_{i-1}, x_{i-1}) - c) + e^{-r\Delta(m-1)}\mathbb{E}[V_\Delta(\theta_m)|(x, \theta)]$$
$$\geq (F(\theta) - F(\theta'))(P(\theta'; \theta, x) - c) + e^{-r\Delta}\mathbb{E}[V_\Delta(\theta_m)|(x, \theta)],$$

where m is such that the equilibrium path from θ reaches θ' in $m + 1$ periods.

By rewriting the necessary condition, we get

$$\sum_{i=1}^{m-1} e^{-r\Delta(i-1)}(F(\theta_{i-1} - F(\theta_i))(P(\theta_i; \theta_{i-1}, x_{i-1}) - c)$$
$$- (F(\theta) - F(\theta'))(P(\theta'; \theta, x) - c) \geq e^{-r\Delta}(1 - e^{-r\Delta(m-2)})\mathbb{E}[V_\Delta(\theta_m)|(x, \theta)].$$

The left hand side is at most $(F(\theta - F(\theta'))(\theta - c) \leq \rho\Delta(\theta - c)$. The right-hand side is at least $e^{-r\Delta}(1 - e^{-r(A-2\Delta)})\mathbb{E}[V_\Delta(\theta_m)]$. As $\Delta \to 0$, the left-hand side goes to zero but the right-hand side converges to $(1 - e^{-rA})V(\theta)$ almost surely. Thus, the monopolist's value must go to zero. □

Proof of Theorem 1

Proof. Assume instead that for some θ' such that $F(\theta) - F(\theta') \in (0, K\Delta)$, there exists a $\Xi' \subset \Xi$ such that $H_{(x,\theta)}(\Xi') > 0$ and

$$\tau^\xi_\Delta(\theta') - \tau^\xi_\Delta(\theta) > \frac{F(\theta) - F(\theta')}{\rho} + \Delta \tag{18}$$

for all $\xi \in \Xi'$. We show that then $V_\Delta(\theta, x)$ approaches zero as $\Delta \to 0$.

We prove the result in three steps. We first derive a necessary condition that has to hold in any equilibrium and which gives an upper bound for the monopolist's continuation value. In the second step, we use Lemma 7 to rewrite the upper bound. In the third step, we show that the bound goes to zero as $\Delta \to 0$.

Step 1: necessary condition

Let θ' be the highest type that satisfies (18). Any realized equilibrium path can be written as $(\theta = \theta_0, \theta_1, \dots, \theta_m = \theta', \theta_{m+1}, \dots)$ for some $m \geq 2$. If there are periods at which the monopolist does not sell with positive probability on the equilibrium path, the claim follows immediately from Lemma 8. Therefore, we focus on the case where there are sales in every period, i.e. $\theta_i < \theta_{i-1}$ almost surely.

We compare a proposed equilibrium sales path to one that specifies full rate sales. Suppose that this alternative sales path follows the equilibrium for $\xi \notin \Xi'$ but specifies full rate sales $\xi \in \Xi'$ until reaching type θ_m. We can write the full rate sales path as $(\theta = \tilde{\theta}_0, \tilde{\theta}_1, \dots, \tilde{\theta}_{m-2}, \theta_m, \theta_{m+1}, \dots)$ where $F(\tilde{\theta}_{i-1}) - F(\tilde{\theta}_i) = \rho\Delta$ for all $i \leq m-2$ and $\tilde{\theta}_{m-1} = \theta_m = \theta'$. The alternative path satisfies $\theta_{i+1} < \tilde{\theta}_i \leq \theta_i$ for $1 \in \{1, \dots, m-2\}$. In other words, the alternative path uses the full capacity to reach θ' one period earlier than the equilibrium path and follows the equilibrium otherwise. The alternative path differs from the proposed equilibrium path only when $\xi \in \Xi'$. Hence, both paths yield the same value for $\xi \notin \Xi'$.

The proposed equilibrium yields $V_\Delta(\theta) = \mathbb{E}\Big[\sum_{i=1}^m e^{-r\Delta(i-1)}(F(\theta_{i-1}) - F(\theta_i))$ $\times(P(\theta_i;\theta_{i-1},x_{i-1}) - c) + e^{-r\Delta m}V_\Delta(\theta_m)\Big|(\theta,x)\Big]$ and the alternative full rate path yields $\tilde{V}_\Delta(\theta) = \mathbb{E}\Big[\sum_{i=1}^{m-1} e^{-r\Delta(i-1)}(F(\tilde{\theta}_{i-1}) - F(\tilde{\theta}_i))(\tilde{P}_i - c) + e^{-r\Delta(m-1)}V_\Delta(\theta_m)\Big|(\theta,x)\Big]$. The equilibrium path must give at least a weakly higher value than the alternative path:

$$\frac{V_\Delta(\theta) - \tilde{V}_\Delta(\theta)}{\Delta} \geq 0. \tag{19}$$

We evaluate this payoff difference element-wise separately for each realized path $\xi \in \Xi'$:

$$\begin{aligned} &\frac{1}{\Delta}\sum_{i=1}^{m} e^{-r\Delta(i-1)}(F(\theta_{i-1}) - F(\theta_i))(P_i - c) + \frac{1}{\Delta}e^{-r\Delta m}V_\Delta(\theta_m) \\ &-\frac{1}{\Delta}\sum_{i=1}^{m-1} e^{-r\Delta(i-1)}(F(\tilde{\theta}_{i-1}) - F(\tilde{\theta}_i))(\tilde{P}_i - c) - \frac{1}{\Delta}e^{-r\Delta(m-1)}V_\Delta(\theta_m) \end{aligned} \tag{20}$$

After simplifying and re-indexing, (20) becomes

$$\frac{1}{\Delta}\sum_{i=1}^{m-1} e^{-r\Delta(i-1)}(F(\theta_i)-F(\tilde{\theta}_i))(e^{-r\Delta}P_{i+1}-P_i) \tag{21}$$
$$+\frac{1}{\Delta}\sum_{i=1}^{m-1} e^{-r\Delta(i-1)}(F(\tilde{\theta}_{i-1})-F(\tilde{\theta}_i))(P_i-\tilde{P}_i)$$
$$+\frac{1}{\Delta}(1-e^{-r\Delta})c\sum_{i=1}^{m-1} e^{-r\Delta(i-1)}(F(\theta_i)-F(\tilde{\theta}_i))-\frac{1}{\Delta}e^{-r\Delta(m-1)}(1-e^{-r\Delta})V_\Delta(\theta_m).$$

We cannot take the limit $\Delta \to 0$ directly from (21): the equilibrium path and m are specific to each Δ. Hence in Step 2, we derive an upper bound, which does not depend on the details of the equilibrium.

Step 2: upper bound

The continuity of prices (Lemma 7) yields $\sum_{i=1}^{m-1} e^{-r\Delta(i-1)}(F(\theta_i)-F(\tilde{\theta}_i))(e^{-r\Delta}P_{i+1}-P_i) \leq (F(\theta_1)-F(\theta_{m-1}))\max_{i\in\{1,\ldots,m-1\}}(P_{i+1}-P_i) \leq (F(\theta)-F(\theta'))D\Delta$, for some constant D. To see how Lemma 7 implies the last inequality, observe that $\tau(\theta_{i+1})-\tau(\theta_i)+\Delta$ by the definition of θ_i and θ_{i+1}.

Similarly, $(1-e^{-r\Delta})c\sum_{i=1}^{m-1} e^{-r\Delta(i-1)}(F(\theta_i)-F(\tilde{\theta}_i)) \leq (1-e^{-r\Delta})c(F(\theta)-F(\theta'))$, $\sum_{i=1}^{m-1} e^{-r\Delta(i-1)}(F(\tilde{\theta}_{i-1})-F(\tilde{\theta}_i))(P_i-\tilde{P}_i) \leq (F(\theta)-F(\theta'))D\Delta$.

Combining these gives that the payoff difference (21) is at most

$$\frac{1}{\Delta}(F(\theta)-F(\theta'))(2D\Delta+(1-e^{-r\Delta})c)-\frac{1}{\Delta}e^{-r\Delta(m-1)}(1-e^{-r\Delta})V_\Delta(\theta_m)$$
$$\frac{K}{\rho}\Delta(2D+rc)-e^{-r\Delta(m-1)}\frac{1-e^{-r\Delta}}{\Delta}V_\Delta(\theta_m). \tag{22}$$

Step 3: continuous time limit

Let $\Delta \to 0$. The upper bound (21) then goes to $-rV(\theta\downarrow_\xi) \leq 0$ where $\theta\downarrow_\xi$ is the right limit of θ along the path specified by ξ.

So far we have only considered a fixed ξ while the necessary condition (19) only holds in expectation. When we take the expectation over all possible ξ, the necessary condition implies that the expectation over (22) must be non-negative. For small Δ, there is an upper bound for (22), say $\epsilon(\Delta)$, which goes to zero as $\Delta \to 0$ for all ξ. Therefore, in order to (19) to hold we must have

$$(F(\theta)-F(\theta'))\Big(2B+\frac{1-e^{-r\Delta}}{\Delta}c\Big)+\epsilon(\Delta) \geq e^{-r\Delta(m-1)}\frac{1-e^{-r\Delta}}{\Delta}V_\Delta(\theta_m).$$

for almost all $\xi \in \Xi'$. This then implies that each $V_\Delta(\theta\downarrow_\xi)$ goes to zero as $\Delta \to 0$, which further implies that $V_\Delta(\theta)$ must be very small for small Δ. □

B.2 Continuous time

Proof of Lemma 3

Proof. *(i)* Follows immediately by taking the continuous time limit of Theorem 1.

(ii) Assume instead that $\lim_{s\to\theta^+}\pi(s) < 0$ and $V(\theta) = 0$. Then, there exists $\epsilon > 0$ such that the monopolist's value must be negative for some interval $(\theta, \theta+\epsilon)$. When $\Delta \to 0$, this implies that the monopolist's value would be negative somewhere on the equilibrium path when she sells to types above θ with a positive probability. Lemma 2 then rules out that the monopolist does not sell to types just above θ because m is continuous by V^G being continuous.

(iii) The result follows immediately from the non-negativity of the value function. Because $m(\theta)$ is an upper bound for the monopolist's marginal payoff, the same reasoning as above shows that the monopolist's value would be negative for some types just above θ if she sold with a positive probability.

(iv) The right-continuity of V follows directly from the integral representation (5) because τ is independent of the initial state. The integral representation also gives left-continuity at the continuity points of τ. Part *(i)* implies that τ is continuous at θ whenever $V(\theta) > 0$.

We are left to rule out $V(\theta) = 0$ and $\lim_{s\to\theta(-)} V(s) > 0$. Sequential rationality implies that $V_\Delta(\theta) \geq (F(\theta) - F(\theta'))(P(\theta';\theta) - c) + e^{-r\Delta}\beta(\theta',\theta)V_\Delta(\theta')$ $+e^{-r\Delta}x(\theta;x_0)(1-e^{-\delta_G(F(\theta)-F(\theta'))})V_\Delta^G(\theta')$ for all θ' such that $F(\theta)-F(\theta') \in [0,\rho\Delta]$. By rearranging and taking the limit as $\Delta \to 0$ this then yields $V(\theta : x_0) - \lim_{s\to\theta(-)} V(s) \leq 0$, contradicting $V(\theta) = 0$ and $\lim_{s\to\theta^-} V(s) > 0$. □

Proof of Proposition 1

Proof. This follows directly from the first condition in Lemma 3 once we show that $m(\theta)$ has a unique point at which it is zero and that the monopolist's flow payoff is strictly above zero everywhere above that point (and negative below). After good news, $m(\theta) = \theta - c$ and hence the first condition is trivially satisfied. For the second condition, from the envelope theorem we have that the buyer's value equals $W^G(\theta_t, \theta) = \int_{\underline{\theta}}^{\theta} e^{-r(\tau^G(s;k)-\tau(\theta_t))}ds$. This is strictly smaller than $m(\theta) = \theta - c$ for all $\theta > c$ because the capacity constraint implies $e^{-r(\tau^G(s;k)-\tau(\theta_t))} < 1$. Therefore, if

$m(\theta) > 0$, the monopolist's flow payoff $\theta - c - W^G(\theta, \theta)$ is strictly positive. $\square$

Proposition 1 implies that the monopolist's value after good news arrives at θ can be written as $V^G(\theta) = \int_c^\theta e^{-r(\tau^\rho(s) - \tau(\theta))}(s - \frac{F(\theta)-F(s)}{f(s)} - c)f(s)ds$ where we have used Lemma 5 to write the payoff with the help of the virtual surplus.

Proof of Proposition 2

Proof. We prove the result in two part by first arguing that Lemma 3 rules out exit before θ_E and then arguing that it rules out exit after θ_E.

Assume instead that the monopolist sells until $\theta > \theta_E$ and then exits. We have two possibilities: either $m(\theta) \neq 0$ or $\lim_{s \to \theta} V(s) > 0$. The latter case would contradict the first necessary condition in Lemma 3. The former case consists of two parts: 1) $m(\theta) > 0$ and 2) $m(\theta) < 0$. The part 1) violates the second necessary condition. For part 2), first notice that if monopolist exits at θ , $V(\theta) = 0$. By the continuity of m (and by m being an upper bound for the monopolist's marginal payoff), V must be (weakly) decreasing in buyer's type in some neighborhood of θ. Therefore, there exists some $\theta' > \theta$ such that $V(\theta') = 0$ and $m(\theta') < 0$ and by the third necessary condition the monopolist must exit at that type already, contradicting sales until θ.

Finally, notice that the third necessary condition implies $\lambda(\theta) = 0$ ruling out sales to any $\theta < \theta_E$ with a positive probability in the no news path. $\square$

Proof of Lemma 4

Proof. To guarantee that the monopolist enters, assume $m(\overline{\theta}; x_0) > 0$. The equilibrium must satisfy the necessary conditions of Lemma 3 and therefore the exit type is in $\hat{\Theta}_E(x_0)$. Suppose then that for every $\hat{\theta}_E \in \hat{\Theta}_E(x_0)$ there exists $\theta' > \hat{\theta}_E$ such that $\Pi^\rho(\hat{\theta}_E; \theta', x(\theta')) < 0$ and suppose that there exists an equilibrium without delays: $\lambda(\theta) = \rho$ whenever the monopolist sells. This leads to a contradiction as, by assumption, there exists a θ' such that the monopolist has a profitable deviation to exit at θ' and thus the claim follows. $\square$

Proof of Theorem 2

Proof. To prove the claim we show the the monopolist's marginal payoff, $\pi(\theta) = (\theta + \delta_G V^G(\theta) - W(\theta)/x(\theta))x(\theta) - c$, is negative for all types greater than the exit type if she sells at full rate as $\rho \to \infty$ and thus there has to be delays in any stationary equilibrium. We start by simplifying the buyers' value (6). Suppose that the monopolist sells at full rate so that all purchasing times up to the exit type θ_E are zero in the limit. The buyers' value is then

$$\frac{W^\infty(\theta)}{x(\theta)} = \int_{\theta_E}^{\theta} e^{-\delta_G(F(\theta)-F(s))} ds + \int_c^\theta \int_s^\theta e^{-\delta_G(F(\theta)-F(z))}\delta_G f(z) 1_{z\geq\theta_E} dz ds,$$

where the indicator function $1_{z\geq\theta_E}$ keeps track of sales absent news. Using Fubini's theorem for the double integral yields $\int_c^\theta e^{-\delta_G(F(\theta)-F(z))}\delta_G f(z)(z-c)1_{z\geq\theta_E}dz$. We can then apply integration by parts to simplify this further to $(\theta - c) - (\theta_E - c)e^{-\delta_G(F(\theta)-F(\theta_E))} - \int_{\theta_E}^\theta e^{-\delta_G(F(\theta)-F(z))}dz$. Plugging this into the expression for the buyers' value gives $W^\infty(\theta)/x(\theta) = (\theta_E - c) - (\theta_E - c)e^{-\delta_G(F(\theta)-F(\theta_E))}$ Thus, our marginal payoff in the limit becomes $\pi^\infty(\theta) = e^{-\delta_G(F(\theta)-F(\theta_E))}(\theta_E - c)x(\theta) - (1 - x(\theta))c$. We are then left to argue that it is negative when $\delta_B > 0$. Using $\theta_E = c/x(\theta_E)$ and simplifying gives that $\pi^\infty(\theta)$ is negative if and only if $e^{-\delta_G(F(\theta)-F(\theta_E))}(1-x(\theta_E))x(\theta) - (1-x(\theta))x(\theta_E) < 0$. Dividing with $(1-x(\theta_E)(1-x(\theta))$ and writing the likelihood as $q(\theta_E) = q(\theta)e^{-(\delta_G-\delta_B)(F(\theta)-F(\theta_E))}$ then gives

$$e^{-\delta_G(F(\theta)-F(\theta_E))} - e^{-(\delta_G-\delta_B)(F(\theta)-F(\theta_E))} < 0 \iff 1 - e^{\delta_B(F(\theta)-F(\theta_E))} < 0.$$

This is true if and only if $\delta_B > 0$. □

Full rate payoff approaches zero from below at the exit type. Here we derive the rate of sales when the monopolist's marginal payoff approaches zero from below at the exit type. The relevant conditions are $\pi = 0$ and $\pi_\theta = 0$ as $\theta \to \theta_E$ so that the monopolist is indifferent between selling and exiting. The derivative of the marginal payoff is

$$\pi_\theta(\theta) = (1 + \delta_G V_\theta^G)x + (1 + \delta_G V^G)x' - c - W_\theta, \tag{23}$$

where we have dropped all dependencies for notational simplicity. To evaluate the derivative further we take the derivative of the buyers' value. Recall that we can write the value $W(\theta;\theta_t,x_0)$ as

$$x(\theta_t)\int_{\underline{\theta}}^{\theta} e^{r\tau(\theta_t))-\delta_G F(\theta_t)}(e^{-r\tau(s)+\delta_G F(s)} + \int_s^{\theta_t} e^{-r\tau^G(s;k)+\delta_G F(z)}\delta_G f(z)dz)ds.$$

The chain rule then gives (both the type $\theta = \theta_t$ and state θ_t change): $\frac{dW}{d\theta} =$

$\frac{\partial W}{\partial \theta}\frac{\partial \theta}{\partial \theta_t} + \frac{\partial W}{\partial \theta_t}$. On the equilibrium path ($\theta = \theta_t$) this equals

$$W_\theta = x(\theta_t) - \Big(r\frac{f(\theta_t)}{\lambda} + \delta_G f(\theta_t)x(\theta_t) + \delta_B f(\theta_t)(1 - x(\theta_t)) \Big) W(\theta_t; \theta_t, x_0)$$
$$+ x(\theta_t)\delta_G f(\theta_t) W^G(\theta; \theta_t, \theta_t).$$

The derivative of the monopolist's value after good news is $V^G_\theta(\theta_t) = (\theta_t - W^G(\theta_t; \theta_t) - c)f(\theta_t)$. Given W_θ and V^G_θ, (23) becomes

$$\pi_\theta(\theta) = \delta_G x(\theta_t - W^G - c)f + (\theta_t + \delta_G V^G)(\delta_G - \delta_B)x(1 - x)$$
$$+ \Big(r\frac{f}{\lambda} + \delta_G f x + \delta_B f(1 - x) \Big) W - x\delta_G f W^G.$$

The above must equal zero when the monopolist sells at interior rate. Solving for λ yields $\lambda(\theta) = rW(\theta)/D(\theta)$ where

$$D(\theta) = x\delta_G f W^G(\theta) + (\theta_t + \delta_G V^G(\theta))(\delta_B - \delta_G)x(\theta)(1 - x(\theta))$$
$$- \delta_G x(\theta)(\theta_t - W^G(\theta) - c)f(\theta) - (\delta_G f(\theta)x(\theta) + \delta_B f(1 - x(\theta)))W(\theta).$$

We see from this that $\lambda(\theta) \to 0$ as $\theta \to \theta_E$ and that $\lambda(\theta) > 0$ for $\theta > \theta_E$. All types above θ_E purchase absent news but purchasing times go to infinity as $\theta \to \theta_E$.

C Commitment solution and comparison

Proof of Lemma 5

Proof. Once we substitute in W^G from Appendix A, we can write the monopolist's payoff from good news at time k as

$$\Pi^G(\mathcal{T}; \theta_k) = \int_{\underline{\theta}}^{\theta_k} e^{-r(\tau^G(m;k) - \tau(\theta_k)))} \Big(m - c - \int_{\underline{\theta}}^{m} e^{-r(\tau^G(s;k) - \tau^G(m;k))} ds \Big) f(m)dm,$$

where θ_k is the time at which good news arrived. Using Fubini's theorem to change the order of integration for the double integral gives

$$\Pi^G(\mathcal{T}; \theta_k) = \int_{\underline{\theta}}^{\theta_k} e^{-r(\tau^G(m;k) - k)} \Big(m - c - \frac{F(\theta_k) - F(m)}{f(m)} \Big) f(m)dm.$$

For the payoff before news, we then plug in the buyers' value (12) to simplify the monopolist's payoff (recall $q = x/(1-x)$):

$$\begin{aligned}\frac{\Pi(\mathcal{T};\theta_t,x_0)}{1-x_t} =& \int_{\underline{\theta}}^{\theta_t} e^{-(r\tau(\theta)-t)-\delta_B(F(\theta_t)-F(\theta))}\Big(\theta q(\theta)-(1+q(\theta))c\Big)f(\theta)d\theta \\ &+ \int_{\underline{\theta}}^{\theta_t} e^{-r(\tau(\theta)-t)-\delta_B(F(\theta_t)-F(\theta))}\delta_G q(\theta)\Pi^G(\mathcal{T};\theta)f(\theta)d\theta \\ &- \int_{\underline{\theta}}^{\theta_t}\int_{\underline{\theta}}^{\theta} e^{-r(\tau(s)-t)-\delta_B(F(\theta_t)-F(s))}q(s)f(\theta)dsd\theta \\ &- \int_{\underline{\theta}}^{\theta_t}\int_{\underline{\theta}}^{\theta}\int_{s}^{\theta} e^{-r(\tau^G(s;k)-t)-\delta_B(F(\theta_t)-F(k))}\delta_G q(k)f(k)f(\theta)dkdsd\theta.\end{aligned}$$

We use Fubini's theorem to simplify the triple integral:

$$\begin{aligned}&\int_{\underline{\theta}}^{\theta_t}\int_{\underline{\theta}}^{\theta}\int_{s}^{\theta} e^{-r(\tau^G(s;k)-t)-\delta_B(F(\theta_t)-F(k))}\delta_G q(k)f(k)f(\theta)dkdsd\theta \\ =&\int_{\underline{\theta}}^{\theta_t}\int_{\underline{\theta}}^{\theta}\int_{s}^{\theta} e^{-r(\tau^G(s;k)-t)-\delta_G(F(\theta_t)-F(k))}\delta_G q_0 f(k)f(\theta)dkdsd\theta \\ =&\int_{\underline{\theta}}^{\theta_t}\int_{\underline{\theta}}^{k}\int_{k}^{\theta_t} e^{-r(\tau^G(s;k)-t)-\delta_G(F(\theta_t)-F(k))}\delta_G q_0 f(k)f(\theta)d\theta dsdk \\ =&\int_{\underline{\theta}}^{\theta_t}\int_{\underline{\theta}}^{k} e^{-r(\tau^G(s;k)-t)-\delta_G(F(\theta_t)-F(k))}\delta_G q_0\frac{F(\theta_t)-F(k)}{f(s)}f(s)f(k)dsdk \\ =&\int_{\underline{\theta}}^{\theta_t} e^{-r(\tau(k)-t)-\delta_B(F(\theta_t)-F(k))}\delta_G q(k)\int_{\underline{\theta}}^{k} e^{-r(\tau^G(s;k)-\tau(k))}\frac{F(\theta_t)-F(k)}{f(s)}f(s)f(k)dsdk.\end{aligned}$$

We can simplify the double integral similarly and get:

$$\begin{aligned}\Pi(\mathcal{T};\theta_t,x_0) =&(1-x_t)\int_{\underline{\theta}}^{\theta_t} e^{-r(\tau(\theta)-t)-\delta_B(F(\theta_t)-F(\theta))} \\ &\times\left(\left(\theta-\frac{F(\theta_t)-F(\theta)}{f(\theta)}+\delta_G\Phi_G(\mathcal{T};\theta)\right)q(\theta)-(1+q(\theta))c\right)f(\theta)d\theta,\end{aligned}$$

where we have combined the payoff after good news and the triple integral to get: $\Phi_G(\mathcal{T};\theta) := \int_{\underline{\theta}}^{\theta} e^{-r(\tau^G(s;k)-\tau(\theta))}\left(s-c-\frac{F(\theta_t)-F(s)}{f(s)}\right)f(s)ds$.

The expression in the lemma follows from taking $1+q(\theta)$ as a common factor, noting that $q(\theta)/(1+q(\theta)) = x(\theta)$ and simplifying the resulting expression. □

Proof of Lemma 6

Proof. We prove the claim by showing that selling at a lower than the full rate cannot be strictly optimal. Suppose an optimal selling scheme $\hat{\mathcal{T}}$ involves delays. This implies that $\hat{\tau}(s) - \frac{1-F(s)}{\rho} > \varepsilon$ for some $s > \theta_N$ and $\varepsilon > 0$.

Now define an alternative selling scheme $\tilde{\mathcal{T}}$ where $\tilde{\tau}(\theta) = \hat{\tau}(\theta)$ for $\theta > s$ and

$\tilde{\tau}(\theta) = \hat{\tau}(\theta) - \varepsilon$ for $\theta \leq s$. The alternative scheme yields:

$$\Pi(\tilde{\mathcal{T}}; \overline{\theta}, x_0) = (1 - x_0) \int_s^{\overline{\theta}} e^{-r\hat{\tau} - \delta_B(1-F)} \Big((\theta + \delta_G \Phi_G^* - \frac{1-F}{f}) q - (1+q)c \Big) f d\theta$$
$$+ (1 - x_0) \int_{\theta_N}^{s} e^{-r(\hat{\tau} - \varepsilon) - \delta_B(1-F)} \Big((\theta + \delta_G \Phi_G^* - \frac{1-F}{f}) q - (1+q)c \Big) f d\theta$$
$$= (1 - x_0) \int_{\theta_N}^{\overline{\theta}} e^{-r\hat{\tau} - \delta_B(1-F)} \Big((\theta + \delta_G \Phi_G^* - \frac{1-F}{f}) q - (1+q)c \Big) f d\theta$$
$$+ (e^{r\varepsilon} - 1)(1 - x_0) \int_{\theta_N}^{s} e^{-r\hat{\tau} - \delta_B(1-F)} \Big((\theta + \delta_G \Phi_G^* - \frac{1-F}{f}) q - (1+q)c \Big) f d\theta,$$

where we have dropped all dependencies on θ for notational convenience. The first term equals $\Pi(\hat{\mathcal{T}}; \overline{\theta}, x_0)$ and the second term is positive by the optimality of θ_N. This is because if the integral (virtual payoff from continuing to sell until θ_N) was negative for some s, the monopolist would be better off by stopping at s. As the multiplier $e^{r\varepsilon} - 1$ is positive, $\Pi(\tilde{\mathcal{T}}; \overline{\theta}, x_0) \geq \Pi(\hat{\mathcal{T}}; \overline{\theta}, x_0)$. Therefore, there is always an optimal scheme such that the monopolist sells at full rate. $\square$

Proof of Proposition 3

Proof. Part(i): from lemma 6 we know the monopolist must be selling at full rate. The necessity of the first order condition (virtual surplus equal to zero) follows from the standard arguments of optimality. The sufficiency of $\Psi(\theta_t; x_0) < 0$ then follows from the construction of $\Psi(\theta_t; x_0)$ as it gives the maximum payoff over the set of critical points, which by necessity must contain the optimal exit type. Hence, if $\Psi(\theta_t; x_0) < 0$, then the maximum the monopolist can gain by continuing to sell is negative and she must be willing to exit at θ_N.

Part(ii): The purchasing times after good news, τ^G, only show up in the virtual payoff from good news: $\Phi^G(\mathcal{T}, \theta; \overline{\theta}) = \int_{\underline{\theta}}^{\theta} e^{-r\tau^G(s;\theta)}(s - c - \frac{1-F(s)}{f(s)})ds$ (Lemma 5), which is monotonically increasing in θ due to the monotone hazard rate. Thus, there is a unique θ_G such that the integrand equals zero and the integral is maximized by setting minimal purchasing times for types $\theta \geq \theta_G$ and not selling to types $\theta_G > \theta$. $\square$

Proof of Proposition 4

Proof. We first prove the expected sales result without conditioning on entry. When $\delta_G > 0$, $\phi^*(\overline{\theta}; \overline{\theta}, x_0) > m(\overline{\theta}; x_0)$ because $\Phi_G(\overline{\theta}) > V^G(\overline{\theta})$. By continuity, there exists an initial belief x_0 such that the payoff from sales in the commitment solution equals zero, which means that the commitment monopolist is indifferent between entering and not entering. Because the payoff from sales is strictly less in any equilibrium, the equilibrium monopolist does not enter and thus the claim follows.

We then prove the expected sales result assuming the equilibrium monopolist enters. The latter is equivalent to the initial belief being such that $m(\overline{\theta}; x_0) > 0$. First, notice that the virtual surplus is greater than the marginal joint surplus, $\phi^*(\theta; \overline{\theta}, x_0) > m(\theta; x_0)$, for types greater than some $\theta_p \in (\theta_G, \overline{\theta})$. We have that the following upper bound holds for the equilibrium value:

$$V(\theta; x_0) < M(\theta; x_0) := \max_{\theta_L} \int_{\theta_L}^{\theta} e^{-r\tau^{\rho}(s)} \beta(\theta, s) m(s; x_0) f(s) ds.$$

The claim follows if we can show that we can find an initial belief such that both $m(\overline{\theta}; x_0) > 0$ and $M(\theta_p; x_0) < 0$ hold.

To see when $M(\theta_p; x_0) < 0$ let θ^* denote the optimal θ_L and observe that the following expression provides an upper bound for $M(\theta; x_0)$:

$$x(\theta) \int_{\theta^*}^{\theta} e^{-\delta_G(F(\theta)-F(s))}(\theta + \delta_G V^G(\theta) - c) f(s) ds - (1 - x(\theta)) \int_{\theta^*}^{\theta} e^{-\delta_B(F(\theta)-F(s))}) c ds.$$

We can evaluate the integral in the upper bound to get $x(\theta)((\theta + \delta_G V^G(\theta) - c)/\delta_G)(1 - e^{-\delta_G(F(\theta)-F(\theta^*))}) - (1 - x(\theta))(c/\delta_B)(1 - e^{-\delta_B(F(\theta)-F(\theta^*))})$. The upper bound implies that $M(\theta_p; x_0) < 0$ if

$$q_0 < e^{(\delta_G - \delta_B)(1-F(\theta))} \frac{\delta_G}{\delta_B} \frac{c}{\theta_p + \delta_G V^G(\theta_p) - c} \frac{(1 - e^{-\delta_B(F(\theta_p)-F(\theta^*))})}{(1 - e^{-\delta_G(F(\theta_p)-F(\theta^*))})},$$

where we have used the likelihood $q = x/(1-x)$. Similarly, the assumption $m(\overline{\theta}; x_0) > 0$ can be written with the likelihood as $(\overline{\theta} + \delta_G V^G(\overline{\theta}) - c) q_0 - c > 0$, which gives a lower bound for the initial likelihood $q_0 > c/(\overline{\theta} + \delta_G V^G(\overline{\theta}) - c)$.

Thus, we have a lower and an upper bound for the initial likelihood q_0. The upper bound is larger than the lower bound for any $\theta_p, \theta^* < \overline{\theta}$ when

$$\frac{(\theta_p + \delta_G V^G(\theta_p) - c}{(\overline{\theta} + \delta_G V^G(\overline{\theta}) - c)} < e^{(\delta_G - \delta_B)(1-F(\theta))} \frac{\delta_G}{\delta_B} \frac{(1 - e^{-\delta_B(F(\theta_p)-F(\theta^*))})}{(1 - e^{-\delta_G(F(\theta_p)-F(\theta^*))})}.$$

It is clear that this holds when $\delta_G \geq \delta_B$, because the left-hand side is less than one. Furthermore, there exists $\bar{\delta}_B > \delta_G$ such that the inequality is true for $\delta_B < \bar{\delta}_B$. Therefore, the claim in the second part of the statement follows. $\square$