

Quality, Upgrades, and Equilibrium in a Dynamic Monopoly Model

James
Anton
Duke University

Gary
Biglaiser
University of North
Carolina Chapel Hill

May, 2009

ERID Working Paper Number 37

This paper can be downloaded without charge from
The Social Science Research Network Electronic Paper Collection:
<http://ssrn.com/abstract=1589748>

Economic Research Initiatives at Duke

WORKING PAPERS SERIES



Quality, Upgrades, and Equilibrium in a Dynamic Monopoly Model

James J. Anton
Duke University

Gary Biglaiser¹
University of North Carolina, Chapel Hill

May, 2009

¹Anton: james.anton@duke.edu; Biglaiser: gbiglais@email.unc.edu. We are grateful to the Fuqua Business Associates Fund and Microsoft for financial support. We thank Harvard Business School, the Portuguese Competition Authority, and UCSD for their hospitality where some of this research was conducted. We also thank Leslie Marx, Larry Samuelson, Jean Tirole, many colleagues at conferences and seminars, and, especially, Joel Sobel for many helpful conversations. The views in this work are solely are own.

Abstract

We examine an infinite horizon model of quality growth in a durable goods monopoly market. The monopolist generates new quality improvements over time and can sell any available qualities, in any desired bundles, at each point in time. Consumers are identical and for a quality improvement to have value the buyer must possess previous qualities: goods are upgrades. We show that subgame perfect equilibrium payoffs for the seller range from capturing the full social surplus all the way down to capturing only the current flow value of each good and that each of these payoffs is realized in a Markov perfect equilibrium that follows the socially efficient allocation path. This is true for all discount factors. We also show that inefficient equilibria exist for rates of innovation above a threshold.

1 Introduction

We examine the commercialization of quality innovations in a dynamic model of an upgrade good. Prominent examples are provided by technology markets, such as those for software, where cycles of upgrades to existing products have become the norm. Innovation is an ongoing process in these settings and buyers face a sequence of purchasing decisions. Thus, rather than timing a single purchase and then exiting the market, buyers have an ongoing incentive to return to the market to ‘upgrade’ and move to a higher quality level. Buyer expectations are pivotal for these decisions and, given the recurrent aspect of upgrading, bundling by the seller emerges as a critical aspect of the upgrade offers.

What determines the equilibrium division of the surplus generated by innovation in an upgrade market? Will equilibrium allocations be efficient? The answers hinge on the nature of a credible threat for buyers. An offer that is rejected today will likely be followed by a subsequent offer that includes not only the rejected quality increment but also additional ones (and bundling of quality increments is a necessary consideration). Because no equilibrium has perpetual rejection, it is possible that eventual buyer acquiescence causes the threat either to collapse entirely or be very sensitive to the extent of payoff discounting. In this regard, the distinguishing feature of an upgrade market with ongoing innovation is that the available joint surplus rises after a round of no trade. Surplus growth allows for expansive buyer expectations that support a buyer’s credible threat and, as we demonstrate, the level of this threat is endogenously determined.

A buyer always has the option of refusing an upgrade offer and staying with the existing version. Accepting moves the buyer to a higher quality level but a proper evaluation of the price must recognize that this is an interim position that will be superseded by subsequent upgrading. An additional concern is the behavior of other buyers and whether one is moving ahead of the market or lagging behind, as the seller may offer different bundles to attract buyers who differ in their current quality levels. In equilibrium, the price(s) for the commercialization of upgrades reflects the ability of the seller to tempt a buyer to purchase when others are not, the cost to a buyer of falling behind the market, and the resulting structure of credible buyer threats to refuse seller offers. Because a refusal leads directly to situations where bundling arises, credible threats and bundling are necessarily intertwined in equilibrium.

We examine an infinite horizon model of an upgrade market that features a very simple core economic structure. Innovation is exogenous but ongoing and in each period it is feasible for the seller to offer an additional quality increment. Buyers are homogeneous and have a fixed valuation per unit of quality; this corresponds to a horizontal demand curve in a static market setting. Building on the recent literature, we capture the notion these are ‘upgrade’ goods by assuming that the sequence of quality increments satisfies a downward complementarity property: for an additional quality increment to be valuable, a buyer must also hold all previous quality increments. The structure of offers by the seller is unconstrained with respect to bundling options in that any combination of quality increments (a single bundle or a set of bundles) may be offered in each period. Bundling is thus endogenous to the equilibrium. To keep the focus on the basic process of commercializing increases in quality, we abstract from complications in the valuation structure such as network effects, compatibility issues, and adjustment costs with adoption.

We construct Markov perfect equilibria for this dynamic game and show that every subgame perfect equilibrium payoff can be achieved by some Markov perfect equilibrium. Two classes of equilibria are identified. Efficient equilibria have buyers acquiring a new upgrade each period. In contrast, and despite the complete information setting, inefficient equilibria do exist and exhibit

cyclical delay in that multiple quality increments go unsold until they are bundled together for sale. Equilibrium payoffs span a significant economic range. At one extreme, there is an efficient equilibrium in which the seller captures all of the social surplus and each quality increment sells immediately at a price equal to the full present discounted value of the flow of benefits to a buyer (extraction of all buyer surplus). At the other extreme, there is an efficient equilibrium in which each quality increment sells only for the one period flow value, leaving a buyer with the entire residual surplus. Significantly, this range of equilibrium payoffs does not reflect a folk theorem: all equilibria are Markov perfect, including the support off the equilibrium path, and the result applies for any discount factor between zero and one. This highlights the endogenous nature of buyer expectations and, more subtly, the structure of a credible threat for buyers to refuse a seller offer when innovation is ongoing in an upgrade market.

To understand the structure of a credible threat in an upgrade market, it is helpful to begin with the case when the seller can only offer a single quality good. With homogeneous buyers, the seller is then in a strong position to make an offer that inevitably tempts buyers to purchase now relative to their expectations for how surplus will be shared in the future. This “speed-up” argument, for which an elegant version was developed by Fudenberg, Levine, and Tirole (1985) for a sequential offer game, is quite powerful and it undermines the credibility of a buyer’s threat to reject offers with high prices. In short, (homogeneous) buyers have no credible threat and in the unique equilibrium the seller captures all of the social surplus.

In contrast, we find that a range of credible threats can be sustained in equilibrium in an upgrade market. This is due to the interactions of three features of our model.

First, quality growth implies that the available social surplus rises in the event of delay. Thus, if buyers reject a current offer from the seller, it must be based on an expectation of how the larger future surplus will be divided. In contrast, the surplus does not change when the seller only has a single version of the good to offer and expectations are always anchored to the division of a fixed economic “pie.” Second, the horizon is infinite and it is feasible for the seller to offer a new upgrade quality unit every period. The third feature is that there are many buyers. Quality growth plus an infinite horizon guarantee that a buyer will always have an incentive to return to the market for possible future purchases. In other words, a buyer will always have a stake in the market. Quality growth together with multiple buyers creates the possibility that in equilibrium buyers have an incentive to coordinate their purchases (recall the model has no network effects or switching costs). Furthermore, an individual buyer expects that other buyers will also return to the market. This last point is critical for understanding how much or, perhaps, how little a buyer is willing to pay for any given quality upgrade. Purchasing an upgrade today implies a current flow benefit and, while the price reflects how current surplus is shared with the seller, we must consider how the purchase relates to buyer expectations regarding the division of surplus from future quality growth. We demonstrate that removing any of these three features (quality growth, infinite horizon, and multiple buyers) will result in the seller capturing the full social surplus. Thus, it is the interaction of these three features in an upgrade market that support an endogenous level for the credible threat by buyers and, hence, multiple equilibria.

The primary intuition for the credible buyer threat is as follows. Suppose that buyers expect to receive a positive share of the surplus on future quality improvements. Further, imagine that the seller offers a price above the candidate equilibrium for today’s upgrade. Is it credible for buyers to refuse the offer? Consider the willingness to pay of an individual buyer when other buyers are expected to refuse the offer. When others refuse, we have delay and the next period will have the larger surplus due to quality growth as the market position involves buyers who lack the previous

upgrade. When the typical buyer's share of this surplus is significant, a solitary individual buyer who purchased the high priced upgrade in the last period will wish to purchase again; despite the fact that this may require the buyer to purchase a bundle that includes quality increments already held, the assumed positive buyer share of future surplus makes it attractive to acquire the new upgrade and keep up with the market. But, then the initial upgrade purchase of a buyer who gets ahead of the market reduces to a one-period flow of value since such a buyer expects to acquire this upgrade next period. As a result, willingness to pay is limited to the one-period flow value of the upgrade. This is a credible threat for buyers to reject prices above the candidate equilibrium price. Moreover, we can apply the logic behind this threat - a willingness to pay that has been pushed to the one period flow value - at any stage of the game and for any given discount factor.

We construct Markov perfect equilibria by utilizing this one-step credible threat to support equilibria. By way of motivation, consider the upgrade market for operating systems. Relatively few users of Windows XP have upgraded to Windows Vista and Microsoft has announced that it will be expediting the introduction of Windows 7. Thus, rather than buying Vista at a later date, many XP users will actually jump to the new state of the art, Windows 7. This pattern suggests the following simple Markovian behavior. Buyers reject an attempted price increase by the seller (a deviation) based on the expectation that the seller will return to the market next period with a better version. Thus, our continuation strategies follow a 'cash-in' support at any state (distribution of buyers across quality levels) that is off of the equilibrium path: the seller offers buyers a bundle(s) that is sufficiently attractive to move the market up to the current state of the art. Given this seller response, rejecting unexpectedly high prices emerge as an equilibrium form of implicit coordination across buyers.

Significantly, the economic structure of the equilibrium payoff set is invariant with respect to the discount factor. That is, for any given discount factor, the seller's payoff always ranges from a maximum of all surplus to a minimum of only the one period flow value of each quality unit. In parallel, buyer payoffs always range from a minimum of zero to a maximum of the residual (after the initial one period flow) present discounted value of each quality unit. In our model, the discount factor reflects the rate of innovation as well as the interest rate. Thus, we find that whether innovation is rapid or slow, seller payoffs range from the flow value to the present discounted value on each unit of quality. Also significant is that strict Markovian behavior is sufficient to support these payoffs. In particular, the cash-in support returns players to the equilibrium path at a price(s) that depends only on the available unsold quality units and there is no need to tailor the support to specific past events. Finally, what is not important is a buyer's fear of falling behind the market. Indeed, we show that relatively high buyer payoffs are perfectly consistent with a zero continuation payoff for any buyer who falls behind the market. What is crucial is the seller's ability to tempt an individual buyer to jump ahead of the market. As we noted, the only offer that a buyer is guaranteed to accept is one that pays for itself in the current period (flow value). Because the seller cannot tempt an individual buyer to pay more than flow value to move ahead of other buyers, we are able to construct a cash-in support for any given discount factor.

Quality that improves over time is recurrent and significant in durable goods markets; see Waldman (2003). In addition to the earlier mentioned software markets, upgrades to cellular networks often allow vendors to offer, for an added charge, new or improved services such as web browsing, e-mail access and text messaging. Less obviously, B-52 bombers produced in the 1950s are still in use today and are expected to be in use in 2040, but the plane has been upgraded repeatedly in terms of electronics, weaponry, and other features; non-defense goods that are regularly upgraded include airports (terminals and runways) and oil refineries, among others. On the policy side,

upgrade markets have been at the center of recent antitrust cases including the Microsoft cases in the U.S. (bundling of operating system and internet browser) and in Europe (bundling with the media player).

Commercialization of innovation is achieved in a variety of ways. Many sellers offer versions of their product with a vertical quality dimension, such as Adobe who offers its Acrobat program in three versions, Standard, Pro, and Pro Extended. It is also common for sellers to condition offers on a buyer's purchasing history. In the Adobe case, there is an upgrade price for existing customers and a "full price" for new buyers: Acrobat 9 Pro Extended is 229 dollars for a current user and 699 dollars for a new user. Many software innovations are available initially as add-ons or separate programs. For example, the capability to e-mail a document, convert to PDF format, or make presentation slides used to be routinely provided by programs separate from the word processor program, such as Microsoft's Word and Mackitchen's Scientific Word. Now, these features are internal to the word processing programs. In all these cases, we see that bundling of quality units is featured prominently.

The extent of bundling by the seller is endogenously determined in our analysis. That is, we allow the seller to offer any feasible collection of quality units. For instance, we do not constrain the seller to make only offers that include all lower level quality units, although the seller is free to choose such an offer. We find that, in equilibrium, it is sufficient to consider only upgrade offers in which a contiguous set of quality units is bundled. This is true with respect to offers on the equilibrium path as well as supporting offers. Equivalently, we show how to interpret the upgrade offers in terms of a full bundle with pricing contingent on a buyer's current product holding. In Section 7, we discuss the implications of several variations of the upgrade and offer structure.

There is a relatively small literature on upgrade models, with most of the work involving a finite horizon. Waldman (1996) examines a two period model, focusing on the incentive to invest in quality growth and R&D time inconsistency. Fudenberg and Tirole (1998) examine a two-period model where consumers are heterogeneous and the period two (new) good renders the period one (old) good obsolete for a buyer. They focus on how the information structure of the monopolist impacts the pricing strategy for the upgrade product, whether the lower quality is sold in period two, and whether the firm may actually buy back the good it sold in period one. Ellison and Fudenberg (2000) analyze a series of static and two period models. These models feature network externalities and a cost to consumers of upgrading the good, and retain the upgrade structure in which the new good makes the old good obsolete for a buyer. They address the issue of whether there is excessive upgrading by the monopolist in a dynamic model and how heterogeneous preferences and network externalities interact. In our model, consumers are identical and there are no direct network externalities or adoption costs. In the finite horizon version of our model, the monopolist captures all the surplus. Thus, a key feature of our model is that the time horizon is infinite and every decision is made with respect to the prospect of future upgrades.

Fishman and Rob (2000) do examine an infinite horizon upgrade model. They focus on private versus social innovation incentives and analyze a rational expectations equilibrium. In contrast, we focus on pricing, taking innovation as exogenously given, and provide a game-theoretic analysis. We treat the seller choice of which bundles to offer as endogenous while the seller in Fishman and Rob is assumed to offer only a single bundle consisting of all prior quality levels. We return to this and other aspects of bundling in Section 7.

In the limiting case of a single upgrade, our model corresponds to a standard durable goods monopoly (Coase (1972)). Precisely because we have identical consumers, the standard Coasian

incentive to cut price over time and move down the demand curve is not present. Papers that examine whether the Coase conjecture holds for a given quality of a good and a fixed set of buyers include Stokey (1981), Bulow (1982), Gul, Sonnenschein, and Wilson (1986). Ausubel and Deneckere (1989), Fehr and Kuhn (1995) and Sobel (1991) provide folk theorems for the (single quality) durable goods model. Sobel (1991) analyzes a model where consumers only want a single unit of the good, but there is entry of new consumers over time. Methodologically, our paper is closest to Sobel, since both feature a market that never closes due to new demand (entry of new consumers and, in our case, quality growth). Our model of a dynamic upgrade monopoly market differs from this literature by including quality and surplus growth, buyers who never exit the market, and seller bundling options.

In contrast to folk theorem results, it is not necessary that the parties be sufficiently patient to achieve the full range of payoffs. In our case, it is the next upgrade that supports equilibrium coordination and creates a credible threat for buyers. As upgrades become increasingly frequent the upper bound for the buyers' share of the surplus converges to 1. At the other extreme, as the upgrades become increasingly infrequent (in the limit a single upgrade at time 0), the buyer's share of the surplus converges to 0. This result, where the seller captures all the surplus, is the same as that found by Fudenberg, Levine, and Tirole (1985).

In Section 2, we present the model. Benchmarks are generated in Section 3 to help differentiate our work from the literature and to understand the implications of the model assumptions. We provide basic results in Section 4, where we show that, in equilibrium, whenever a period has a sale, consumers always move to the current state of the art and purchase all feasible qualities that they do not possess. In Section 5, we examine efficient equilibria in which the monopolist sells each upgrade in the first period that it is available. We show that, for any positive discount factor, the monopolist's payoff can range from getting all the surplus to receiving only the single period flow value of each upgrade. In Section 6, we show that equilibria can be inefficient in that the sale of upgrades is delayed (and bundled). For inefficient equilibria, there are also necessary incentive conditions for delay. We show that there is a critical threshold for the discount factor such that a longer delay requires a higher discount factor. We discuss the upgrade structure of our model in Section 7 and consider directions for future research in Section 8. All proofs are in the Appendix.

2 The Model

We begin with a description of the basic elements of the game and then turn to strategies and payoffs. We allow for complete freedom with respect to bundling units and the formal framework, while involved, is essential for the construction of supporting continuation equilibria. This framework is employed primarily in the proofs in the Appendix rather than in the text and some readers may wish to skip this subsection. Finally, we define and discuss Markov perfect equilibrium.

2.1 Basic Elements

We examine an infinite horizon, discrete time model. Let $\tau = 1, 2, \dots$ index periods. There is a continuum of identical buyers with a measure of 1 represented by the unit interval and a single seller. A new perfectly durable good, unit τ , becomes available in each period τ . All seller costs are 0. Within each period τ , feasible offers for the seller consist of any collection of subsets of $\{1, 2, \dots, \tau\}$ and associated prices. For example, the seller can offer the bundle of all feasible qualities $\{1, 2, \dots, \tau\}$

for a price p , so that the new unit is made available only as part of a larger bundle. Alternatively, the seller can offer a collection of individual unit bundles, $\{1\}$ at price p_1 , quality $\{2\}$ at a price p_2 , and so on; a buyer would then have the option to purchase every feasible quality or any subset of the available unit bundles. Of course, the seller can also withhold some qualities or even make no offer.

Consider the feasible offer set for the seller in period τ . Let $\mathcal{P}_\tau \equiv \mathcal{P}(\{1, 2, \dots, \tau\})$ denote the power set for the first τ integers. Any set $z \in \mathcal{P}_\tau$ is called a *bundle*. An *offer* is a collection of bundles and associated (non-negative) prices, $(z, p_z)_{z \in Z}$ for some $Z \in \mathcal{P}(\mathcal{P}_\tau)$. Define the *offer set* Ω_τ by

$$\Omega_\tau \equiv \left\{ \omega \in \mathcal{P}(\mathcal{P}_\tau \times R_+) \mid (i) (\emptyset, 0) \in \omega, (ii) \text{ if } (z, p) \in \omega \text{ and } (z, p') \in \omega, \text{ then } p = p' \right\}.$$

By (i), we are including the null bundle in every offer by the seller. This is for two reasons: first, the seller can make no offer by choosing only the null bundle and, second, it streamlines the buyer choice formalism, as a buyer chooses to make no purchase by selecting the null bundle. By (ii), every offered bundle has a unique price. Clearly, if two prices were offered for the same bundle, no buyer would want to choose the higher price.¹

Given a seller offer, the buyers respond simultaneously with each buyer choosing which bundle(s) to accept in period τ . Thus, an acceptance choice by a buyer is an element of the set $\mathcal{P}(\mathcal{P}_\tau)$.

Any bundle that consists only of a set of contiguous qualities is defined as an *upgrade*. For example, an upgrade to the “state of the art” from a status quo of 0 is the bundle $\{1, \dots, \tau\}$; we also refer to this as a version. A partial upgrade is a bundle $\{\sigma, \dots, \sigma + k\}$, where $1 \leq \sigma \leq \sigma + k \leq \tau$. We will show that, in equilibrium, a seller need only make upgrade offers.

A buyer receives a flow utility of vq in period τ when contiguous units $1, \dots, q$ but not unit $q + 1$ are held by the buyer. We thus are imposing the condition that a buyer must have all lower quality units for quality q to have value. This “downward complementarity” assumption is the upgrade payoff structure in our model. For example, if a buyer holds quality units 1 and 3 but not 2 in a given period, then the flow utility is v . Thus, quality unit τ is always a complementary good with respect to prior quality units (the downward direction) while the prior quality units do not require unit τ in order to provide value to a buyer.

Players are all risk neutral and have a common discount factor $\delta < 1$. Because a new unit of quality becomes available in each period, the discount factor reflects the rate of innovation as well as the rate of time preference for the players. Thus, a large δ can be interpreted in terms of a rapid rate of innovation while a small δ means that innovations are infrequent. Also, for later interpretation, we can employ the familiar relationship of $\delta = e^{-r\Delta}$, where Δ is the length of a time period (the innovation rate) and r is the interest rate, together with an appropriate measure of flow utility, to assess limiting behavior as $\Delta \rightarrow 0$.

Consider the surplus for a buyer over the infinite horizon. In each period, a buyer holds some subset of the feasible qualities. For any $z \in \mathcal{P}_\tau$, define $M : \mathcal{P}_\tau \rightarrow \{0, 1, \dots, \tau\}$ by finding the unique

¹Formally (see below) this requires that buyers act as price takers and that the overall market does not respond to the actions of an individual buyer. Also, note that we do not impose any arbitrage structure across bundles. For example, if the seller offers a bundle with only good 1, and a bundle with only good 2, then there is no restriction on the price of a bundle that includes both goods 1 and 2. Rather, buyer choices determine which of these bundles will be purchased. Also, since buyers are identical, there are no possible gains for buyers from the possibility of resale.

$m \in \{0, \dots, \tau\}$ such that $m' \in z \forall m' \leq m$ and $m + 1 \notin z$, and set $M(z) = m$. Clearly, $M(z)$ is the maximal contiguous quality held by a buyer and $M(z)$ exists for any bundle z .

Consider an arbitrary sequence of holdings z_τ and payments p_τ for each $\tau \geq 1$, and let $q_\tau = M(z_\tau)$. From any period τ_0 , the net surplus of a buyer is the present discounted value from quality flows net of payments, as given by

$$\sum_{\tau=\tau_0}^{\infty} \delta^{\tau-\tau_0} (vq_\tau - p_\tau).$$

Similarly, the seller's surplus from any period τ_0 on is the present discounted value of revenues, r_τ for each τ , from sales to buyers, as given by

$$\sum_{\tau=\tau_0}^{\infty} \delta^{\tau-\tau_0} r_\tau.$$

Consider efficient allocations, where joint surplus is maximized. First, note that the payments and revenues are transfers that do not affect total surplus. Thus, for any path of quality holdings and payments, the sum of surplus for any given buyer and the seller is

$$\sum_{\tau=\tau_0}^{\infty} \delta^{\tau-\tau_0} vq_\tau.$$

Thus, the realized joint surplus is fully determined by the quality path. Since $q_\tau \leq \tau$ for any feasible path and $q_0 \equiv 0$, the joint surplus is maximized when each buyer holds the maximal quality, $q_\tau = \tau$. The maximal surplus from date τ is then given by

$$\begin{aligned} S_\tau &= v\tau + \delta v(\tau + 1) + \delta^2 v(\tau + 2) + \dots \\ &= \frac{v(\tau - 1)}{1 - \delta} + v \sum_{k=1}^{\infty} k\delta^{k-1} = \frac{v(\tau - 1)}{1 - \delta} + \frac{v}{(1 - \delta)^2}. \end{aligned}$$

We always have $S_\tau > \delta S_{\tau+1}$, as a delay always involves lost surplus and hence inefficiency. However, because each unit generates surplus, we also have $S_\tau < S_{\tau+1}$. Note that $S_1 = \frac{v}{(1-\delta)^2}$ is the maximal joint surplus at the start of the game; it is the surplus when buyers acquire one new unit in each period, where each new unit has a present discounted value of $\frac{v}{1-\delta}$.

2.2 Strategies, Payoffs, and Equilibrium

A strategy for the seller is a *sequence* of offers, $\mathcal{O} = (\mathcal{O}_\tau)$. Each offer is a map from the *history* of play up through period $\tau - 1$ into the offer set Ω_τ . A history is the sequence of previous offers by the seller and acceptances by the buyers. Letting \mathcal{H}_τ denote the space of all histories up through period $\tau - 1$, we have

$$\mathcal{O}_\tau : \mathcal{H}_\tau \rightarrow \Omega_\tau.$$

That is, if $h_\tau \in \mathcal{H}_\tau$ is the observed history of play, then the seller's strategy specifies the offer $\omega_\tau = \mathcal{O}_\tau(h_\tau)$.

A buyer strategy profile is a *sequence* of acceptance decisions, $\mathcal{A} = (\mathcal{A}_\tau)$. Given a history h_τ and a seller offer ω_τ , each buyer $x \in [0, 1]$ needs to choose which bundles in ω_τ to accept. Thus, we have acceptance strategies for each buyer

$$\mathcal{A}_\tau^x : \mathcal{H}_\tau \times \Omega_\tau \rightarrow \mathcal{P}(\mathcal{P}_\tau).$$

Hence, for observed history $h_\tau \in \mathcal{H}_\tau$ and in response to a seller offer of $\omega_\tau \in \Omega_\tau$, buyer x chooses to accept the set of bundles $\mathcal{A}_\tau^x(h_\tau, \omega_\tau) \subseteq \mathcal{P}(\mathcal{P}_\tau)$. Of course, any accepted bundle, $z \in \mathcal{A}_\tau^x(h_\tau, \omega_\tau)$, must have been offered by the seller, $(z, p) \in \mathcal{O}_\tau(h_\tau)$ for some p . This is a feasibility restriction. Note that a buyer is free to accept one or more of the bundles (i.e., any subset) included in an offer ω_τ . For example, by “accepting” only the null bundle, a buyer makes no purchase in period τ . Finally, we use \mathcal{A}_τ for the strategy profile across buyers.

We need to specify the history space \mathcal{H}_τ . First, define $\Omega^\tau \equiv \Omega_1 \times \Omega_2 \times \dots \times \Omega_\tau$; this product space contains each feasible sequence of previous offers. Second, we need to calculate acceptance sets from buyer bundle purchases and this entails a measurability assumption on buyer strategies.

Let \mathcal{F}_τ denote the set of Borel measurable functions for $[0, 1] \rightarrow \mathcal{P}(\mathcal{P}_\tau)$. By definition, $f_\tau : [0, 1] \rightarrow \mathcal{P}(\mathcal{P}_\tau)$ is Borel measurable (that is, $f_\tau \in \mathcal{F}_\tau$) if for any $z \in \mathcal{P}_\tau$ we have $\mathcal{X}_\tau(z) \in \mathcal{B}$ (the Borel sets of $[0, 1]$), where $\mathcal{X}_\tau(z) = \{x \in [0, 1] \mid z \in f_\tau(x)\}$. Thus, the set of buyers who chose bundle z is a Borel set and we can calculate market share and revenues by using standard Lebesgue measure. Define the product space $\mathcal{F}^\tau \equiv \mathcal{F}_1 \times \mathcal{F}_2 \times \dots \times \mathcal{F}_\tau$.

Then the history space is specified by $H_1 = \emptyset$ and for $\tau > 1$,

$$\mathcal{H}_\tau = \Omega^{\tau-1} \times \mathcal{F}^{\tau-1}.$$

Note that the bundles and prices offered by the seller are recorded in $\Omega^{\tau-1}$ while the bundles accepted by each buyer are recorded in $\mathcal{F}^{\tau-1}$. Thus, we know the price a buyer paid for a bundle from the history.² We assume that for each $h_\tau \in \mathcal{H}_\tau$, and $\omega_\tau \in \Omega_\tau$, we have $\mathcal{A}_\tau \in \mathcal{F}_\tau$, i.e. $\mathcal{A}_\tau^x(h_\tau, \omega_\tau)$ is a Borel measurable function on $x \in [0, 1]$.

Turning to the calculation of player payoffs, we begin with the buyers. First, for each $h_{\tau+1}$, calculate the units acquired by buyer x in each period $k = 1, \dots, \tau$. These units are given by $Z_k(x) = \{i \in \{1, \dots, k\} \mid i \in z \text{ for some } z \in \mathcal{A}_k^x(h_k, \omega_k)\}$, the bundles accepted by buyer x . Thus, the set of units that buyer x has accumulated through the end of period τ is given by

$$Z^\tau(x) \equiv \bigcup_{k=1}^{\tau} Z_k(x) \subseteq \mathcal{P}_\tau.$$

Recalling that $M(z)$ is the maximal contiguous quality for any subset z of $\{1, \dots, \tau\}$, we see that the maximal contiguous quality unit held by buyer x is given by $m_\tau(x) \equiv M(Z^\tau(x))$.

Next, the total expenditure of buyer x in period τ is given by $p_\tau(x) \equiv \sum_{z \in \mathcal{A}_\tau^x(h_\tau, \omega_\tau)} p_z$, which is the sum of the payments for each bundle that the buyer accepted. Thus, the payoff to buyer x from strategy \mathcal{A}^x when other buyers follow $\mathcal{A}^{\sim x}$ and the seller follows \mathcal{O} is the present discounted value of surplus from the maximal unit held less expenditures in each period³:

$$U(\mathcal{O}, \mathcal{A}^x, \mathcal{A}^{\sim x}) = \sum_{\tau=1}^{\infty} \delta^{\tau-1} [vm_\tau(x) - p_\tau(x)].$$

²An equivalent, but less convenient, formulation would be to assign an index number to each element in the finite set $\mathcal{P}(\mathcal{P}_\tau)$ and then define measurability in the standard way for a real valued function.

³The infinite sum is always well defined, since (i) the sequence of maximal holdings m_τ is non-decreasing in τ , (ii) $m_\tau \leq \tau$, and (iii) $\sum_{\tau=1}^{\infty} \delta^{\tau-1} \tau = 1/(1-\delta)^2$.

We now compute the seller payoff. Given a history and an offer by the seller, $\mathcal{X}_\tau(z)$ as defined above is the set of buyers for whom $z \in \mathcal{A}_\tau^x(h_\tau, \omega_\tau)$. The Lebesgue measure of such buyers is $\alpha_\tau(z) \equiv \int_{\mathcal{X}_\tau(z)} dx$. Thus, the revenue of the seller in period τ is⁴

$$r_\tau = \sum_{z \in \mathcal{P}_\tau} \alpha_\tau(z) p_z.$$

The seller payoff under strategies $(\mathcal{O}, \mathcal{A})$ is then

$$\Pi(\mathcal{O}, \mathcal{A}) = \sum_{\tau=1}^{\infty} \delta^{\tau-1} r_\tau$$

The definitions for Nash and subgame perfect equilibrium are standard. The strategies $(\mathcal{O}, \mathcal{A})$ form a Nash equilibrium if

$$\begin{aligned} \Pi(\mathcal{O}, \mathcal{A}) &\geq \Pi(\widehat{\mathcal{O}}, \mathcal{A}) \quad \text{for all } \widehat{\mathcal{O}}, \\ U(\mathcal{O}, \mathcal{A}^x, \mathcal{A}^{\sim x}) &\geq U(\mathcal{O}, \widehat{\mathcal{A}}^x, \mathcal{A}^{\sim x}) \quad \text{for all } \widehat{\mathcal{A}}^x. \end{aligned}$$

A subgame perfect equilibrium requires that $(\mathcal{O}, \mathcal{A})$ form a Nash equilibrium at any given h_τ , where the seller makes an offer, and at any given h_τ and ω_τ , where the buyers respond to the offer.

We employ a continuum, the unit interval, for the set of buyers in order to capture the idea that individual buyers are insignificant with respect to market outcomes. Thus, we follow Gul, Sonnenschein, and Wilson (1986), Ausubel and Deneckere (1989), and Sobel (1991), among others, and restrict attention to equilibria that satisfy a zero-measure property: for any two histories that differ only with respect to the actions of a set of buyers of measure zero, the strategies of the seller and all other buyers are the same across the two histories. Thus, no individual buyer expects that their own acceptance/rejection decision will have any impact on subsequent play, such as affecting the set of bundles that will be available for purchase in the future. Hence, buyers act as price takers.

2.3 Markov Perfect Equilibrium

In this paper, we examine Markov perfect equilibria (MPE) as defined by Maskin and Tirole (2001), with the natural modification for a continuum of agents. By definition, Markov strategies depend only on the payoff relevant aspects of a history of the game. In our model, the seller's flow payoff depends only on revenues and each buyer's flow payoff depends only on the maximal contiguous unit held and the payments in a period. Thus, past prices and the timing of buyer acquisitions do not influence current period payoffs. For each bundle z , the set of buyers who hold z , $Q_\tau(z)$, is found directly from $Z^\tau(x)$, the units accumulated by buyer x through τ . Hence, the allocation of buyers across quality units, as given by $(Q_\tau(z))_{z \in \mathcal{P}_\tau}$, identifies all payoff relevant information as of the beginning of period $\tau + 1$.⁵

⁴Note that $\alpha_\tau(z)$ must be equal to zero if the seller did not offer the bundle z or if no buyer purchased z .

⁵This form of Markovian behavior would allow non-contiguous holdings to affect strategies. While a non-contiguous unit does not affect current buyer flow payoffs, it would change a future flow payoff if missing intermediate units were acquired.

A simpler form of Markovian behavior is to focus on the distribution of maximal contiguous units across buyers and the gap relative to the current period τ , which indexes the seller's feasible units. This notion not only has more economic appeal, but also, as we will see, allows us to generate all subgame perfect equilibrium seller payoffs. We need to define a state of the game. Consider any history that leads to period τ in which all buyers enter the period with the same maximal quality level Q (units 1 through Q). We define this to be state (τ, Q) .⁶ We then define Markovian behavior by the condition that players' strategies depend only on the size of the gap $\tau - Q$. This means, for instance, that if the seller offers an upgrade of σ units at a price p in state $(\tau, 0)$, then an upgrade from Q to $Q + \sigma$ at the same price p must be offered in state (τ', Q) , provided that the gaps coincide, $\tau' - Q = \tau$. Furthermore, except for a translation of the index number on quality units, buyers' accept/reject decisions are the same in states $(\tau, 0)$ and (τ', Q) .⁷ This implies that the seller's profits and buyers' utilities satisfy

$$\pi_\tau \equiv \pi(\tau, 0) = \pi(\tau', Q)$$

and

$$u(\tau', Q) = \frac{vQ}{1 - \delta} + u(\tau, 0) \quad \text{and} \quad u_\tau \equiv u(\tau, 0)$$

for $\tau' - Q = \tau$.

The main economic rationale for considering Markovian behavior in the quality gap is that feasible payoffs in our upgrade model have a simple stationary structure. Consider state $(\tau + 1, Q)$, where all buyers hold Q units at the end of period τ , and compare it to state $(\tau + 1 - Q, 0)$, where all buyers hold 0 units at the end of period $\tau - Q$. Then any subgame perfect equilibrium for the game that begins in state $(\tau + 1 - Q, 0)$ is also a subgame perfect equilibrium of the game that begins in state $(\tau + 1, Q)$; we only need relabel the indexes of the quality units in offer and accept strategies. The seller earns the same payoff in both situations and each buyer's payoff is translated by $v(\tau - Q)/(1 - \delta)$.

With buyers acting as price takers, in any subgame all buyers with the same quality holdings must receive the same payoff. Thus, all buyers earn the same equilibrium payoff from the start of the game. This allows us to focus on equilibria in which buyers follow symmetric strategies, because, as we will show, every possible buyer payoff can be generated by such strategies. Furthermore, we show that every equilibrium payoff can be implemented by an upgrade offer structure, where at each state (τ, Q) , the seller either delays by making no offer or offers one upgrade level $Q' \in \{Q + 1, \dots, \tau\}$ and an associated price. Thus, on the equilibrium path, all buyers have the same quality holdings.⁸

We will provide an explicit construction of the strategies and show that the equilibrium behavior of buyers and sellers necessarily follows a simple cyclical structure when strategies only depend on

⁶When the buyers are distributed across maximal holdings then the state is given by $(\tau, (Q_\tau^m)_{m=0, \dots, \tau-1})$, where Q_τ^m is the set of buyers with maximal contiguous quality m .

⁷More generally, when buyers are distributed as $(Q_\tau^m)_{m=0, 1, \dots, \tau}$, then the translated state is given by $(\tau + 1 - \underline{m}_\tau, (Q_\tau^m)_{m=\underline{m}_\tau, \dots, \tau})$ where \underline{m}_τ is the smallest index of Q_τ^m with a non-zero measure.

⁸States where buyers have asymmetric holdings are off-the-equilibrium path as are histories where the seller makes multiple upgrade offers. Note that mixing by buyers in response to a seller offer would lead to asymmetric holdings. This is often required for continuation equilibria in the durable goods literature. In our case, because buyers never exit the market, we are able to construct pure strategy continuation equilibria in all states.

the quality gap. Furthermore, the definition of Markov perfect equilibrium is flexible enough to allow for both efficient and inefficient equilibria. Henceforth, we use equilibrium to refer to a buyer symmetric Markov perfect equilibrium in the quality gap.

3 Benchmarks for the Quality Growth Model

We begin our analysis by identifying the equilibrium outcomes for several simplified versions of our model. These benchmarks help to illuminate the roles of quality growth, the infinite horizon, and multiple buyers.

3.1 Finite Horizon $T > 1$

An infinite horizon with quality growth implies that buyers will always seek to acquire higher quality units. We consider, then, a finite horizon model where the prospect of acquiring higher quality units is truncated, in order to highlight how the equilibrium outcomes depend on the continued presence of buyers in the market. We need to specify how buyers value their quality holdings after the final period. Let $w \in \left[0, \frac{v}{1-\delta}\right]$ denote the “scrap value” for each (contiguous) quality unit that a buyer holds after the final period T , where a good can have no value after period T up to a flow value of v forever.

Consider the final period. Suppose the state is (T, q_{T-1}) , where $q_{T-1} \leq T - 1$ is the quality held by buyers at the start of period T .⁹ Then there exists a unique outcome (subgame perfect) in which the seller offers an upgrade from q_{T-1} to T units, and prices the upgrade at an extraction level (subgame perfection is being employed to rule out non-credible threats in which buyers do not accept positive surplus offers). All buyers will accept the offer. Thus, $u_T = (v + \delta w)q_{T-1}$ and the buyers are held to their status quo utility as of the start of period T .

Now consider period $T - 1$ and suppose the state is $(T - 1, q_{T-2})$, where $q_{T-2} \leq T - 2$. Since buyers know that they will not receive any incremental surplus in period T , they will necessarily accept any offer that provides a positive utility increment in period $T - 1$. The seller clearly prefers to sell a unit in period $T - 1$ rather than period T . Thus, there exists a unique outcome (subgame perfect) in which the seller offers an upgrade of $(T - 1 - q_{T-2})$ units at the extraction price. Working backwards to period 1, the seller always offers an upgrade to the current state of the art at the extraction price. This outcome does not depend on whether we have a single buyer or a continuum. It also prevails if the quality units are independent goods (no upgrade payoff structure). To summarize, the absence of future transactions implies that the seller captures all of the social surplus.

⁹If buyers are distributed across holdings from 0 to $T - 1$, then the seller offers an upgrade to T units to each buyer segment that fully extracts the surplus of each segment. A buyer selects the intended offer, since the offer to a lower segment is more expensive and the offer to a higher segment lacks necessary contiguous units relative to the buyer’s current position. Alternatively, the seller could offer a version $\{1, \dots, \tau\}$ contingent on a buyer’s holdings.

3.2 Infinite Horizon, Single Buyer

Now, suppose we have a single buyer instead of a continuum. Under the Markovian hypothesis, individual buyers in a continuum have no effect on the state. By contrast, with a single buyer, the state necessarily depends on the buyer's purchasing decision. We find that any Markov perfect equilibrium (MPE) necessarily has the properties that the seller will follow the efficient path, selling the new unit in each period, and price each unit at extraction, $\frac{v}{1-\delta}$, so that all surplus accrues as profits to the seller. Let us start with a simple example to see why sales must occur without delay. Suppose that there is delay and two units are sold in period 2 at price p . MPE then implies

$$\pi_1 = \delta p + \delta^2 \pi_1$$

and

$$u_1 = \delta(2v - p) + \delta^2 \left(\frac{2v}{1-\delta} + u_1 \right).$$

We can now apply a modified version of the familiar argument of Fudenberg, Levine, and Tirole (1985) to obtain a profitable speed up deviation by the seller. Suppose the seller offered one unit at a price \hat{p} in period 1. If the buyer accepts (note that by doing so the single buyer changes the continuation state), then the seller earns $\hat{\pi} = \hat{p} + \delta\pi_1$. The buyer accepts provided that $\hat{u} = v - \hat{p} + \frac{\delta v}{1-\delta} + \delta u_1 > u_1$. Thus, the deviation is profitable for the seller, $\hat{\pi} > \pi_1$, and acceptable to the buyer, $\hat{u} > u_1$, provided

$$\frac{v}{1-\delta} - (1-\delta)u_1 > \hat{p} > (1-\delta)\pi_1,$$

as follows from the above expressions for u_1 and π_1 . Such a \hat{p} exists if and only if

$$S_1 = \frac{v}{(1-\delta)^2} > u_1 + \pi_1.$$

This always holds since S_1 is the maximal surplus and $u_1 + \pi_1$ is necessarily smaller due to the assumed delay of a sale.

Thus, the seller can profitably speed up the candidate equilibrium. Intuitively, the buyer and seller can share the larger surplus of S_1 by selling a unit in period 1 and it is simple to find a mutually beneficial price for that transaction. More generally, we always have $S_\tau > \delta S_{\tau+1}$, and the extra surplus allows us to apply a similar speed up argument to any state $(\tau + 1, q)$ with a sale that is preceded by a delay. Hence, with a single buyer, the equilibrium path from the start of the game follows the efficient path with a sale every period and the continuation path from any state must involve an immediate upgrade to the state of the art.

We now argue that this implies extraction of the buyer. For each state $(\tau, 0)$ we know that the continuation is an upgrade offer to the state of the art at price p_τ for payoffs of $\pi_\tau = p_\tau + \delta\pi_1$ and $u_\tau = \frac{v\tau}{1-\delta} - p_\tau + \delta u_1$. Then,

$$\pi_\tau + u_\tau = \frac{v\tau}{1-\delta} + \delta(\pi_1 + u_1)$$

is the joint payoff. We must have $u_\tau = \delta u_{\tau+1}$: if $u_\tau < \delta u_{\tau+1}$, the buyer would reject p_τ , since the $\tau + 1$ offer is more attractive; if $u_\tau > \delta u_{\tau+1}$, then the seller could raise the price and the buyer would still accept. This implies that $u_1 = \delta^{\tau-1} u_\tau$. Substituting for u_τ in the equation for the joint payoff and simplifying, we find

$$\pi_\tau = \frac{v\tau}{1-\delta} + \delta(\pi_1 + u_1) - \frac{u_1}{\delta^{\tau-1}} = \frac{v\tau}{1-\delta} + \delta \frac{v}{(1-\delta)^2} - \frac{\delta u_1}{\delta^\tau}.$$

Suppose u_1 is positive. Then, despite the growth in quality, as τ goes to infinity the required exponential growth in the buyer’s utility will eventually push the seller’s profit below zero. Obviously, this cannot happen in equilibrium. Thus, the buyer lacks a credible threat and is necessarily extracted.

The above argument does not extend to a continuum of buyers: an individual buyer cannot change the state, either by delaying or accepting the seller’s offer. For example, in state $(\tau - 1, q)$ if a single buyer accepts an offer to move to the state of the art, but no other buyer accepts, then in the next period the state is (τ, q) . The seller can only earn a profit by making an offer that targets a positive mass of buyers with quality q .

3.3 Infinite Horizon, No Growth, Continuum of Buyers.

With no growth, the model reduces to the case of a single good: the seller has one unit to offer to buyers. Thus, when all buyers are identical we essentially have a special case of the problem studied by Fudenberg, Levine, and Tirole (1985), who allow for buyer valuation heterogeneity. Using simpler versions of the arguments employed above, we then find that there is never delay and buyers are always extracted in any MPE when there is no quality growth.

These benchmarks demonstrate the robustness of the seller’s market power with respect to the time horizon, the number of buyers, and quality growth. Thus, taken individually, none of these three factors can change the equilibrium outcome. We defer to Section 7 a discussion of a fourth factor, the upgrade structure, including the independent goods benchmark, as this is best done after our results are in place. We now turn to our model where there is an infinite horizon, growth in quality, and a set of buyers who never leave the market, to show how the necessity of extraction breaks down and a credible threat for buyers emerges.

4 Preliminary results

We begin with basic results on the necessary structure of equilibria. These serve as building blocks for the main analysis. First, we show that by pricing at a very low level the seller can always induce buyers to make a purchase.

Lemma 1 (*Flow Dominance*) *Consider any history such that, at the start of period τ , all buyers hold the first Q quality units and no buyer holds unit $Q + 1$, where $\tau > Q$. Suppose the seller makes an upgrade offer for units $\{Q + 1, \dots, \tau\}$ at price p , where $p < v(\tau - Q)$. Then, in any continuation equilibrium, every buyer accepts the upgrade offer.*

The intuition for “flow dominance” is simple. The upgrade from Q to τ is priced sufficiently low that that it pays for itself in the current period, since $v\tau - p > vQ$. Moreover, even if all other buyers were to reject the offer, an individual buyer who accepts is always weakly better off in the future. This follows from (1) the upgrade payoff structure, since an accepting buyer has a flow surplus of at least $v\tau$ in future periods, and (2) all buyers have the same opportunities for purchasing from the seller, so an accepting buyer always has the option of making the same choices in the future as other buyers. Essentially, a buyer who holds all of the first τ units in period $\tau + 1$ is never at a disadvantage relative to any other buyer.

It then follows directly that the seller must have a positive payoff both at the start of the game and at any point in the future. This is due to quality growth and flow dominance. At any point in time, the seller always has the option of offering a bundle that includes the new quality unit at a (flow dominant) upgrade price.

Lemma 2 *In any equilibrium, the payoff of the seller is at least $v/(1 - \delta)$. For any history in which all buyers hold quality units $\{1, \dots, Q\}$ and no buyer holds unit $Q + 1$ at the start of period τ , the continuation payoff of the seller is at least $v(\tau - Q) + \delta \frac{v}{1-\delta}$.*

It is important to note that the above results are very basic and, as the proofs demonstrate, they do not depend on Markovian behavior or symmetric buyer strategies. Rather, these two lemmas rely only on buyers acting as price takers. The lower bound on the seller payoff provides a reference point and we will show that in a buyer symmetric MPE, every payoff ranging from the full surplus, $S_1 = \frac{v}{(1-\delta)^2}$, all the way down to the flow dominance lower bound, $v/(1 - \delta)$, can be supported.

A simple consequence of a positive seller payoff in any continuation is that the quality gap never grows without bound. That is, all new quality units are eventually sold within some fixed number of periods.

Lemma 3 *In any equilibrium, for any state (τ, Q) , the continuation path has a bounded quality gap.*

Now, we show that equilibria must have a simple cyclical structure. To see this, we introduce the notion of a $t - cycle$ equilibrium. In a $t - cycle$ equilibrium a sale occurs every t periods, and t units are sold in each sale period. Thus, states $(1, 0)$ through $(t - 1, 0)$ are delay states with no sales, and state $(t, 0)$ has a sale of units 1 through t . Hence, once a sale occurs in state $(t, 0)$, the quality gap falls back to 1 at the start of the next period and, effectively, the state returns to $(1, 0)$. Thus, the equilibrium path will cycle and the buyers will increase their quality holdings every t periods, each time buying an upgrade to go from t to $2t$ and so on. Note that, as a special case, this includes the possibility that $t = 1$, where the current quality unit is sold to buyers in every period (the efficient path).

Proposition 1 *Every equilibrium follows a $t - cycle$ equilibrium path: the buyers purchase quality units $\{1, \dots, t\}$ from the seller in state $(t, 0)$, all payments to the seller occur in state $(t, 0)$, and the maximal buyer quality is zero until period t .*

What makes this argument work is flow dominance and the fact that the seller can profitably deviate by speeding up a cycle that does not have buyers moving to the state of the art in $(t, 0)$. Thus, if the sale to buyers only involves $\tau < t$ units, the seller can feasibly offer these units in state $(t - 1, 0)$. By pricing these units at $\hat{p} = v\tau + \delta p - \varepsilon$, where p is the price for τ units in state $(t, 0)$, a seller improves his payoff if all the buyers accept since

$$\begin{aligned} \hat{p} + \delta\pi(t, \tau) &> \delta [p + \delta\pi(t + 1, \tau)] \Leftrightarrow \\ (v\tau + \delta p - \varepsilon) + \delta^2\pi(t + 1, \tau) &> \delta p + \delta^2\pi(t + 1, \tau) \Leftrightarrow \\ v\tau &> \varepsilon \end{aligned}$$

where we have substituted for \hat{p} and the fact that (t, τ) is a delay state.

The candidate equilibrium cannot have buyers rejecting this offer. If other buyers reject, an individual will always find it optimal to purchase the deviation offer (for small $\varepsilon > 0$). By accepting, an individual buyer receives $\delta u(t, 0) + \varepsilon$. To see this, note that the deviating buyer does not change the state, so τ units will be offered next period. Since the buyer already has these units, the purchase in period t can be skipped and the buyer will have the same holdings as all other buyers as of $t + 1$. Thus, her payoff is improved relative to waiting whenever $\varepsilon > 0$. Hence, all buyers rejecting the offer is not an equilibrium continuation. But, as we showed above, when all buyers accept the offer the seller can profit by making the deviation offer. Thus, an equilibrium with sales of τ less than t cannot be supported.

By contrast, the speed up argument does not apply to a $t - cycle$ equilibrium, where the seller offers an upgrade for all units $\{1, \dots, t\}$, for two reasons. The first is feasibility. The seller does not have t units to sell in period $t - 1$. Second, an individual buyer who accepts the deviation offer in $t - 1$ is not in an analogous position. By acquiring τ units when no other buyers accept, an individual buyer can no longer safely skip all purchases in state $(t, 0)$, since other buyers will be acquiring units 1 through t . For example, if the seller only offers the bundle of units 1 through t , then the deviating buyer will either have to buy the same bundle as the other buyers and pay for the τ units that were previously purchased or continue on with holding only the first τ units at the end of period t .

To summarize, a seller must either sell units as soon as they are feasible, thus following the efficient path, or delay to a maximal set of units periodically, inducing an inefficient path.

Finally, to streamline the equilibrium analysis, we specify strategies such that an individual buyer who deviates by not following other buyers in a purchase that increases the maximal buyer quality will obtain no future additional surplus. Thus, if an individual buyer has the first k units of the good, when all other buyers also have additional contiguous units, then this buyer's continuation payoff is $\frac{vk}{1-\delta}$. One can interpret this in two ways. First, the seller (optimally) ignores individual buyers (measure zero) who differ from the market path. Thus, the missing units necessary for the buyer to benefit from further purchases will never be offered. Alternatively, the seller can always make the necessary units available, thus allowing the individual buyer to achieve parity with other buyers, but set the price for the needed units at an appropriate upgrade price, so as to extract all the continuation surplus.

The critical feature is that the seller's offer does not depend on a deviation by a set of measure zero buyers. This means that the seller's offer strategy specifies that he either completely refrains from making "catch-up" offers, or always makes such offers. As will be clear from the equilibrium construction, we could also allow for higher buyer catch-up continuation values as long as they do not exceed the equilibrium payoff. It will also be clear, from the range of payoffs that are supported in equilibrium, that the exact details are inessential for the equilibrium construction. For the analysis, however, it is helpful to adopt the specific upgrade construction mentioned earlier. Thus, in a $t - cycle$ equilibrium the seller either makes no offer or offers one upgrade (multiple upgrades if buyers are asymmetrically distributed) consisting of units from the status quo buyer holding(s) to the state of the art.

5 Efficient Equilibria

In an efficient equilibrium, a good is sold in each period when it first becomes available. By equilibrium hypothesis, this occurs at price p_1 in each period. Thus, the firm's profits and buyers'

utilities are $\pi_1 = \frac{p_1}{1-\delta}$ and $u_1 = \frac{1}{1-\delta} \left[\frac{v}{1-\delta} - p_1 \right]$, respectively. In an efficient equilibrium, the firm and the consumers divide the maximal social surplus: $S_1 = \frac{v}{(1-\delta)^2} = \pi_1 + u_1$.

We will show that any buyer utility level $u_1 \in [0, \delta S_1]$ can be supported as an equilibrium payoff for any δ . That is, the seller may be limited to only the flow payoff of v per period, which has a present discounted value of $\frac{v}{1-\delta}$. Thus, the seller may only receive the minimum possible payoff (flow dominance and Lemma 1). To show this result, we first find support constraints for all off-the-equilibrium path states $(\tau, 0)$ where the quality gap is above 1. These constraints, which must be satisfied in equilibrium, reduce to “cash-in” constraints in which the seller offers an upgrade to move the buyers up to the state of the art for every quality gap. Next, we generate a set of supporting utilities, the continuation payoffs for buyers in the event that there is no current transaction, that will be used to implement the set of equilibrium payoffs. Essentially, these payoffs constitute a credible threat that induces the seller to make the equilibrium offer.

A period in our model corresponds to the length of time before the next unit of quality can feasibly be offered by the seller. As noted above, frequent innovation corresponds to a large discount factor while relatively infrequent upgrades correspond to a small discount factor. The main analysis focuses on $\delta > 1/2$. We conclude the section with the simpler case of infrequent innovation, $\delta \leq 1/2$.

5.1 Buyer and Seller Support Constraints

To derive the equilibrium payoffs, we must make sure that players cannot do better by deviating.¹⁰ To know that a deviation is not profitable, we must specify the continuation payoffs from state $(2, 0)$ and other off-equilibrium states. By the Markovian structure of equilibrium, we can determine the continuation payoff for any state once we specify the continuation payoffs for all states of the form $(\tau, 0)$. We construct continuation payoffs so that in state $(\tau, 0)$, the seller offers τ units at a price p_τ and this is accepted by all buyers. Thus, the next state is $(\tau + 1, \tau)$, where the quality gap has returned to 1, and the players are back on the (incremental) equilibrium path from $(1, 0)$. The payoffs with a cash-in support at $(\tau, 0)$ are $\pi_\tau = p_\tau + \delta\pi_1$ for the seller and $u_\tau = v\tau - p_\tau + \delta u(\tau + 1, \tau) = \frac{v\tau}{1-\delta} - p_\tau + \delta u_1$ for the buyers. Note that since this is the efficient path from $(\tau, 0)$, we have $S_\tau = \frac{v\tau}{1-\delta} + \delta S_1 = \pi_\tau + u_\tau$.

For a continuation equilibrium to follow this “cash-in” support, we must specify the accompanying buyer and seller strategies. The seller has three ways of deviating. The first option is to make no offer, a “delay,” which necessarily leads to state $(\tau + 1, 0)$ and buyers make no decision. The second option is to offer an upgrade of less than τ units, a “partial cash-in.” The final option is to offer an upgrade of τ units at a price different from p_τ . It must be optimal for the seller to follow the strategy of offering τ units at the price p_τ in state $(\tau, 0)$. For buyer strategies in state $(\tau, 0)$ we specify a simple cut-off rule: a buyer accepts the seller offer of price p for σ units in state $(\tau, 0)$ if and only if $p \leq p(\sigma, \tau)$. Thus, we must find both the cash-in price p_τ for all $\tau \geq 1$ and cut-off rules $p(\sigma, \tau)$ for all $\sigma \leq \tau$, where $\tau \geq 1$.

First, we derive the buyer cut-off strategies. It must be optimal for an individual buyer to accept any offer $p \leq p(\sigma, \tau)$, given that all other buyers are accepting the offer (symmetric strategies). When all other buyers accept, an individual buyer earns $v\sigma - p + \delta u(\tau + 1, \sigma)$ by accepting, while

¹⁰We apply the one-stage-deviation principle to verify the proposed strategies constitute an equilibrium; our model conforms to the necessary requirement of “continuity at infinity,” since the limit of $\tau\delta^\tau$ is 0 as τ goes to infinity (see Fudenberg and Tirole (1991) pp. 108-110).

rejecting (when others all accept) yields 0 by construction. Thus, it is an equilibrium for all buyers to accept p for σ units in state $(\tau, 0)$ if $v\sigma - p + \delta u(\tau + 1, \sigma) \geq 0$ or, equivalently, $\frac{v\sigma}{1-\delta} + \delta u(\tau + 1 - \sigma, 0) \geq p$. Further, it must be that an offer of $p > p(\sigma, \tau)$ is rejected by all buyers. Rejecting when all other buyers reject yields a payoff of $\delta u(\tau + 1, 0)$. Accepting an offer when all other buyers reject yields a flow of $v\sigma - p$ today plus the option of purchasing the continuation offer for $\tau + 1$ units in the next period. Thus, an individual buyer optimally rejects when others reject if $\delta u(\tau + 1, 0) > v\sigma - p + \delta \max \left\{ \frac{v\sigma}{1-\delta}, u(\tau + 1, 0) \right\}$.

To understand this rejection condition, we define $g(\sigma, u) \equiv v\sigma + \delta \max \left\{ \frac{v\sigma}{1-\delta}, u \right\} - \delta u$ as the “net surplus” value of the option for a deviating buyer who makes a purchase when other buyers do not. When the other buyers purchase in period $\tau + 1$, this buyer has two options. If $u > \frac{v\sigma}{1-\delta}$, it is optimal to purchase when the other buyers do and, thus, the deviating buyer is willing to pay at most the flow value of the units, $v\sigma$, in period τ . Otherwise, the buyer will not make the purchase in $\tau + 1$ and is willing to pay up to $\frac{v\sigma}{1-\delta} - \delta u$. Thus, recalling $u_\tau \equiv u(\tau, 0)$, the cut-off strategy must satisfy

$$g(\sigma, u_{\tau+1}) \leq p(\sigma, \tau) \leq \frac{v\sigma}{1-\delta} + \delta u_{\tau+1-\sigma} \quad (1)$$

for all $0 < \sigma \leq \tau$ and all $\tau \geq 1$. The cut-off strategies apply to full ($\sigma = \tau$) and partial ($\sigma < \tau$) cash-in offers. Since $g(\sigma, u)$ is less than or equal to $\frac{v\sigma}{1-\delta}$, cut-off strategies for the buyers exist for any non-negative utility sequence. The upper bound in (1) says that prices must be sufficiently low that rejection is not optimal for an individual buyer, while the lower bound says that it is always optimal to accept offers below this level. Note that $g(\sigma, u_{\tau+1})$ is at least as large as $v\sigma$; flow dominance implies that a buyer is always willing to pay at least $v\sigma$.

We call the difference, $\frac{v\sigma}{1-\delta} + \delta u_{\tau+1-\sigma} - g(\sigma, u_{\tau+1})$, the buyers’ price wedge. It demonstrates that a buyer’s willingness to pay depends on what other buyers are expected to do. It is important to note that there are no network externalities in our model, which is a standard reason for why buyers make their purchasing decisions based on expectations of other buyers’ choices. Instead, the linkage of decisions in our model arises from quality growth and the resulting incentive for a buyer to return to the market for another upgrade together with the fact that what the seller will offer in the future depends on “the state of the market” and, thus, an individual buyer will be affected by his position relative to the market when making future purchasing decisions.

Given these buyer responses, the seller must find it optimal to offer τ units at price p_τ in state $(\tau, 0)$. Beginning with partial cash-ins, note that $p(\sigma, \tau)$ is the optimal price choice for any such offer and it generates a payoff of $p(\sigma, \tau) + \delta \pi(\tau + 1, \sigma)$. This implies that for an equilibrium

$$\pi_\tau - \delta \pi_{\tau+1-\sigma} \geq p(\sigma, \tau) \quad (2)$$

for $\sigma = 1, \dots, \tau - 1$.

The other two deviations are delay and offering τ units at a price different than p_τ . Delay, $\sigma = 0$, is not optimal if $\pi_\tau \geq \delta \pi_{\tau+1}$. Defining $p(0, \tau) \equiv 0$, (2) applies. Finally, consider a cash-in offer of τ units. Buyers will accept any price below $p(\tau, \tau)$, so we must have $p_\tau = p(\tau, \tau)$ or else the seller could successfully offer a price above p_τ . In other words, buyers must reject any price above p_τ for τ units. Note that (2) holds with equality by construction of the equilibrium continuation.

We can now combine the buyer and seller support conditions, (1) and (2), to identify when there exist supporting prices $p(\sigma, \tau)$ such that the cash-in outcome constitutes a continuation equilibrium. Note that the conditions also apply at $\tau = 1$ for $\sigma = 0$ (delay) and $\sigma = 1$ (the equilibrium path).

Combining the seller profit expression (2) with the buyer lower bound on prices, the following condition must be satisfied:

$$\pi_\tau - \delta\pi_{\tau+1-\sigma} \geq p(\sigma, \tau) \geq g(\sigma, u_{\tau+1}).$$

Recalling that $S_\tau = \pi_\tau + u_\tau$, we see that the profit difference exceeds the net surplus value, $g(\sigma, u_{\tau+1})$, if and only if

$$S_\tau - \delta S_{\tau+1-\sigma} \geq u_\tau - \delta u_{\tau+1-\sigma} + g(\sigma, u_{\tau+1}). \quad (3)$$

Note that the surplus difference on the left hand side is an exogenous sequence that is increasing in τ . So, as τ grows and more units are “on the table,” a larger set of payoff utilities can be supported. Given a sequence of utilities that satisfies (3), we can clearly construct the supporting prices $p(\sigma, \tau)$ for conditions (1) and (2). At $\sigma = \tau$, the utility u_τ determines the cash-in price p_τ via $u_\tau = \frac{v\tau}{1-\delta} - p_\tau + \delta u_1$ and conditions (1) and (2) can be verified directly. At $\sigma < \tau$, the supporting price $p(\sigma, \tau) = g(\sigma, u_{\tau+1})$ always satisfies conditions (1) and (2); typically this is the smallest price in an interval of such prices.

When (3) holds, the optimal upgrade offer for the seller is to offer τ units for the price p_τ . Each buyer then finds it optimal to accept the upgrade offer, given that all other buyers also accept. We then have

Lemma 4 *Suppose the sequence of buyer utilities u_τ satisfies (3) for all (σ, τ) where $0 \leq \sigma \leq \tau$ and $\tau \geq 1$. Then there exists an efficient equilibrium with supporting prices $p(\sigma, \tau)$.*

The proof that a sequence of utilities satisfying (3) is sufficient for the existence of an equilibrium outcome with payoff u_1 is by construction. Taking a given u_1 , the rest of the utility sequence is specified in the next section. For this sequence, we must show that it is not profitable for the seller to deviate by offering multiple upgrade options (as well as options with non-contiguous units); note that (3) only rules out seller deviations involving a single upgrade offer. This requires that we specify buyer strategies in response to any such offer from the seller. In addition, we must specify strategies for continuation equilibria in the event that buyer holdings are distributed asymmetrically across units in $\{0, \dots, \tau - 1\}$ in any period τ , even though such states do not arise on the equilibrium path. We do this in Appendix B.

We now turn to the task of finding buyer utilities that support equilibria.

5.2 Support Utilities for Frequent Innovation

For each payoff $u_1 \in [0, \delta S_1]$, we will construct an associated supporting path of u_2, u_3, \dots such that the seller finds it optimal to make an acceptable offer to achieve a cash-in outcome in every state. First, we take up the case of frequent innovation, $\delta > 1/2$.

To gain intuition for how to support the set of u_1 payoffs, we begin with two special cases of support utilities. First, we assume that the buyers’ support utilities are constant, $u_1 = u_2 = \dots = u$. This means that, in the event of delay, all of the gains from the growing surplus will accrue to the seller. What we will show is that this gives the seller an incentive to delay whenever the buyers’ utility exceeds $\frac{v}{1-\delta}$, which is less than δS_1 for $\delta > 1/2$. To see this, consider a delay in period 1: by the support condition (3) at $\tau = 1$ and $\sigma = 0$, we have

$$S_1 - \delta S_2 \geq (1 - \delta)u \quad \Leftrightarrow \quad \frac{v}{1 - \delta} \geq u.$$

The second inequality is violated if $\frac{v}{1-\delta} < u$. Since the seller is the residual claimant of surplus, the loss from delay is just v . On the other hand, the gain from delay is the saving in utility given to buyers of $(1-\delta)u$. If $u > \frac{v}{1-\delta}$, then the seller will prefer to delay and earn $\delta(S_2 - u)$ rather than $S_1 - u$ from selling today. Intuitively, the discounted share of a larger residual will always dominate once u becomes large enough. This is a direct consequence of the growth in surplus that is generated when quality improves over time.

What this example shows is that buyer utility in the event of delay must be increasing to support a higher u_1 equilibrium payoff. Now suppose that the buyers' utilities are always increasing, such that the seller is always indifferent between delay and a cash-in, i.e., the support condition holds with equality at $\sigma = 0$ for all τ . The efficiency gains from the early cash-in are $v\tau$; this is the one period flow value of τ units and is precisely $S_\tau - \delta S_{\tau+1}$. Thus, by (3),

$$u_{\tau+1} = \frac{u_\tau - v\tau}{\delta}. \quad (4)$$

Since the difference between u_τ and $\delta u_{\tau+1}$ is exactly $v\tau$, the difference between π_τ and $\delta\pi_{\tau+1}$ must be 0. The payoff to the seller rises over time, but this is exactly offset by the discount factor δ . The support condition (3) for a cash-in, $\sigma = \tau$, is

$$S_\tau - \delta S_1 \geq u_\tau - \delta u_1 + g(\tau, u_{\tau+1}).$$

For τ sufficiently large, we have $1 > \delta + \delta^\tau$, and it is straightforward to show that $g(\tau, u_{\tau+1}) = \frac{v\tau}{1-\delta} - \delta u_{\tau+1}$. That is, $\frac{v\tau}{1-\delta} > \delta u_{\tau+1}$ for τ sufficiently large. Then the support condition for a cash-in requires that

$$\delta u_1 \geq v\tau$$

which is violated for τ sufficiently large.

What is happening here is the following. Relative to the future, buyers always receive an extra $v\tau$ today, since $u_\tau = \delta u_{\tau+1} + v\tau$. For τ sufficiently large, the utility growth cannot be maintained. Eventually, a flow-dominance offer at price $v\tau$, which is always available to the seller, allows the seller to charge a higher price.

Thus, we seek a support that is increasing, but not increasing either too fast or long, in order to generate the set of payoffs $u_1 \in [0, \delta S_1]$ for any $\delta > 1/2$. To do so, we combine aspects of the two special cases we just examined. From state $(1, 0)$ to $(T, 0)$, the support will make the seller indifferent between a cash-in and delaying until the next period; for states with a larger quality gap the support will keep buyer's utility constant at u_T . Thus, for $T \geq 2$ we define a T -stage support utility sequence by

$$u_\tau = v\tau + \delta u_{\tau+1} \text{ for } \tau = 1, \dots, T-1, \quad (5)$$

and

$$u_\tau = u_T \text{ for } \tau \geq T.$$

Clearly, $u_{\tau+1}$ is increasing in u_τ in the initial stage, and all u_τ are increasing in u_1 . Also, for any given u_1 the sequence (u_2, \dots, u_T) is determined. For the special case of $T = 1$, we specify a constant sequence $u_\tau = u_1$ for all τ . For another perspective, starting from u_T and working backwards to u_1 , it is clear that u_1 includes all of the interim surplus generated by quality growth over the next T periods: $u_1 = v + \delta 2v + \dots + \delta^{T-1} u_T$. As we will see, the higher u_1 or the larger δ , the longer will be the period of time that the support utilities must be strictly increasing in order to implement the higher payoff.

A direct consequence of a T -stage support is that we only need to satisfy the support constraints, equations (3), over the range $\tau = 1, \dots, T$; see Lemma A1 in Appendix B. This is because, when (3) holds at $\tau = T$, then it necessarily holds at all larger τ whenever the buyer utility remains constant. Thus, an advantage of a T -stage support is that we only have to check a finite set of conditions. This is for two reasons. The first is that if the support works in period τ for all sales σ that are large enough to put the state back in the range where the value of buyer utility is rising, then the support works the next period $\tau + 1$ for any sales σ' that also induce a state where the value of buyer utility is rising. Second, if sales to buyers are small enough that the continuation state has a constant buyer utility, then the support holds if it holds at $\sigma = 0$ (delay). Thus, the support condition is relatively straightforward when the buyer's utilities are constant.

We now turn to finding the appropriate length for the T -stage support.

5.3 Frequent Innovations: Existence

We use the following algorithm to determine the T -stage support utilities that are needed to support a given utility level $u_1 \in [0, \delta S_1]$:

- Pick a utility level u_1 between 0 and δS_1 .
- If $u_1 \leq (1 - \delta) S_1$, then set $u_\tau = u_1$ for all $\tau > 1$.
- If $u_1 \in [(1 - \delta) S_1, \delta S_1]$, set $u_2 = \left(\frac{u_1 - v}{\delta}\right)$.
- If $u_1 < (1 - \delta^2) S_1$, set $u_\tau = u_2$ for all $\tau > 2$. If not, set $u_3 = \left(\frac{u_2 - 2v}{\delta}\right)$.
- Keep following this logic until T such that $u_1 \leq (1 - \delta^T) S_1$.

We know that there exists a T for the final step of the algorithm, $u_1 \leq (1 - \delta^T) S_1$, since for T sufficiently large we have $\delta S_1 \leq (1 - \delta^T) S_1$. The discount factor δ determines how large T must be in order to cover the entire range of buyer payoffs $[0, \delta S_1]$. To see this, we need to define a set of critical δ cutoffs: let δ_τ be the root of $\delta^\tau + \delta = 1$ for $\delta \in (0, 1)$. The cut-off sequence δ_τ is strictly increasing in τ , from $\delta_1 = 1/2$ to $\lim_{\tau \rightarrow \infty} \delta_\tau = 1$, and satisfies $\delta^{\tau-1} < 1 - \delta < \delta^\tau$ for $\delta \in (\delta_{\tau-1}, \delta_\tau)$. Figure 1 then illustrates the relationship between u_1 , δ , and T . For example, when $1/2 < \delta < \delta_2$, we use a 1-stage support for the area A in Figure 1, where $u_1 < (1 - \delta) S_1$, and then a 2-stage support for larger u_1 in the area B. Because the δS_1 curve is below the $(1 - \delta^2) S_1$ surplus curve, we have covered all possible buyer payoffs for this range of δ . In the next range, $\delta \in (\delta_2, \delta_3)$, after areas C and D, we must also use a 3-stage support to cover the highest buyer payoffs (in the area E). As δ continues to rise, we use the critical δ cutoffs to identify the appropriate (maximal) length for the T -stage support. Referring back to Figure 1, for each δ range we rise vertically with a 1-stage, then a 2-stage, and so on up to the appropriate T -stage to support each payoff interval on the u_1 axis. It is also helpful to note that for each u_1 in a given payoff interval, the region between the surplus curves in Figure 1 shows that a given length T -stage support will support the u_1 payoff across all discount factors. Thus, a 1-stage support covers the areas A, C and F which includes all payoffs below the $(1 - \delta) S_1$ surplus curve. Since $(1 - \delta) S_1 = \frac{v}{1 - \delta}$, this is exactly what we found with our first example for support utilities (constant utility). Similarly, with $T = 2$ we cover all payoffs in the areas B, D and G, the region between the surplus curves $(1 - \delta^2) S_1$ and $(1 - \delta) S_1$.

Let us consider the incentive structure for the seller and the buyers that arises with a T -stage support constructed according to the algorithm. The algorithm identifies a set of intervals that




Figure 1: Boundary for T-Stage Support ($v = 1$)

covers the entire range of buyer payoffs, 0 to δS_1 . Thus, for each feasible buyer payoff, u_1 , there is a unique T such that u_1 lies in the interval from $(1 - \delta^{T-1})S_1$ to $(1 - \delta^T)S_1$. The right end point is determined by incentives on the seller side of the market, while the left end point is determined by incentives on the buyer side.

For the seller, the initial phase of indifference between cash-in and delay is followed by a strict preference to cash-in once the quality gap reaches T . For this second phase, $\tau \geq T$, we have $\pi_\tau > \delta\pi_{\tau+1}$ as long as $u_1 < (1 - \delta^T)S_1$. Any u_1 above this upper bound (given T) would provide the seller with a strict incentive to delay selling in period T and the support would unravel. To understand better the economic link between buyer and seller payoffs in periods 1 and T , consider the division of surplus. In period 1, a buyer utility of $u_1 = (1 - \delta^T)S_1$ together with the efficient outcome imply that the seller receives $\pi_1 = \delta^T S_1$ and therefore $\pi_T = \delta S_1$, by seller indifference. In period T , the cash-in supporting outcome has the parties divide the surplus of S_T which consists of the value of T units plus the discounted efficient surplus of S_1 . Thus, the buyer payoff must be $u_T = \frac{vT}{1-\delta}$, since the seller's share is δS_1 . This reflects the fact that the T -stage support is constructed so that u_1 capitalizes the surplus growth over T periods (see Lemma A2 in Appendix B for the algebra with the partial sums).

The economic link between the seller's delay incentive and the buyer payoff u_T operates as follows. Since the T -stage support holds the buyer payoff constant for any period $\tau \geq T$, the seller can give the buyers u_T now or wait and give them u_T next period. Hence, the cost of selling today is $(1 - \delta)u_T$. The benefit of selling today is the efficiency gain $v\tau = S_\tau - \delta S_{\tau+1}$, since the seller is now the residual claimant for surplus growth. So, selling today is more profitable than

waiting to cash-in tomorrow when

$$v\tau \geq (1 - \delta)u_T. \tag{6}$$

Otherwise, waiting is better. By choosing T so that u_T is less than $\frac{vT}{1-\delta}$ whenever u_1 is below $(1 - \delta^T)S_1$, the seller has the necessary incentive to avoid delay along the support path. Of course, longer initial support phases will also be sufficient for the seller incentive. Thus, the upper bound $(1 - \delta^T)S_1$ on the equilibrium payoff u_1 is tied directly to when the seller will first have a strict preference to cash-in. In other words, T is the first time the flow dominance payoff exceeds the savings from delay.

The left end point, the lower bound on u_1 of $(1 - \delta^{T-1})S_1$, is determined by buyer incentives. Suppose that the seller offered τ units at a price above the supporting price, p_τ . Given that all other buyers reject this offer, an individual buyer who deviates must assess the option value of having τ units in later periods. In a $T - stage$ support, the seller is expected to make a cash-in offer of $\tau + 1$ units next period and a deviating buyer will have to decide whether to accept or reject this offer. A rejection means means that the deviating buyer will lack unit $\tau + 1$. Due to the upgrade structure, a buyer who falls behind the market will not be able to make any future purchases that yield positive surplus. Thus, accepting the $\tau + 1$ offer is best if $u_{\tau+1}$ exceeds the status quo position of $\frac{v\tau}{1-\delta}$. Clearly, when $\tau \geq T$, a deviating buyer rejects the offer in favor of the status quo holding because $u_{\tau+1} = u_T$ is now below the value $\frac{v\tau}{1-\delta}$. This follows directly because $\frac{vT}{1-\delta}$ exceeds u_T when u_1 is below $(1 - \delta^T)S_1$, the same property that drove the seller's incentive.

A $T - stage$ support has more subtle buyer incentives in the initial phase, where $\tau < T$, since both $u_{\tau+1}$ and $\frac{v\tau}{1-\delta}$ vary with τ . Because the $T - stage$ support capitalizes the surplus growth up through period T , by taking u_1 above the left end point, $(1 - \delta^{T-1})S_1$, we generate utilities such that (i) $u_\tau > \frac{v\tau}{1-\delta}$ for $\tau = 1, \dots, T - 1$ and (ii) $\frac{v(T-1)}{1-\delta} < u_T$ (see Lemma A3 in Appendix B).¹¹ From the buyer price wedge expression (1), a deviating buyer in period $\tau \leq T - 1$, expects to purchase again in period $\tau + 1$. Thus, the most the deviating buyer would be willing to pay in τ is just the flow value of the units, $v\tau$. In this way, the $T - stage$ support provides buyers with a credible threat when presented with a higher than expected price from the seller.

The key underlying forces behind the credible buyer threat are surplus growth and implicit buyer coordination. With respect to surplus growth, in our finite horizon and no quality growth benchmarks, there is always a successful speed up deviation by the seller for any candidate equilibrium with a positive buyer payoff. By contrast, we have equilibria with positive buyer payoffs because surplus growth allows us to construct the $T - stage$ support that maintains incentives by making buyers and then the seller the residual claimant of the changes in surplus generated by new quality units. With respect to coordination, in our single buyer benchmark case a buyer has a zero payoff. With multiple buyers, no single buyer can influence the state of the market and an individual buyer's decision focuses on their position relative to the market, which matters because of the upgrade payoff structure (we return to this point in the conclusion where we consider variations on the upgrade structure). Because quality grows, buyers are forced to return to the market to acquire new units, and we find that implicit coordination is supported when a buyer's willingness to pay depends on whether they hold more or less than other buyers. For example, when other buyers are expected to refuse the current offer, a buyer who purchases today and also expects to purchase tomorrow is then only willing to pay the flow value for the units today.

¹¹When u_1 is at an endpoint of one the intervals in the algorithm, (i) and (ii) typically involve a weak inequality.

Implicit coordination does not, of course, guarantee that buyers pay only the flow value for units in equilibrium. The price p_τ is determined by u_τ , and this is the highest price that a seller can set before the buyers reject. This price must be at least $v\tau$, as implied by flow dominance. The price p_τ can be strictly above $v\tau$, but a seller cannot raise the price any further without losing all current sales.

Lemma A4 in Appendix B demonstrates that we only need to check the support conditions with respect to cash-in outcomes, $\sigma = \tau \leq t$. The cash-in constraints are the most difficult ones to satisfy since partial cash-ins, $\sigma < t$, will only delay the time when the seller can receive π_1 . To see when the support conditions (3) will break down for a given T – *Stage* support, recall that $u_1 = (1 - \delta^T)S_1$ implies that $u_T = \frac{vT}{1-\delta}$. If we try to push u_1 any higher without changing T , the seller will strictly prefer to delay a sale all the way until state $(T + 1, 0)$ instead of selling in state $(1, 0)$.

We can now prove our main result for this section.

Proposition 2 *Let $\delta > 1/2$. Then every $u_1 \in [0, \delta S_1]$ can be supported in equilibrium. Specifically, for each u_1 , there exists a unique T such that $\delta + \delta^{T-1} \geq 1$ and u_1 lies between $(1 - \delta^{T-1})S_1$ and $(1 - \delta^T)S_1$ and the T – stage support satisfies condition (3) for all (σ, τ) where $0 \leq \sigma \leq \tau$ and $\tau \geq 1$.*

Significantly, the minimum possible equilibrium payoff for the seller is the flow dominance lower bound from Lemma 1. Thus, there is an equilibrium in which the seller’s market power is reduced to the static flow value of each unit with all of the future surplus from a unit accruing to buyers.

In this section we examined the situation where innovation is relatively frequent, $\delta \geq 1/2$. Consider the range of equilibrium payoffs as innovation occurs at an increasingly rapid rate. Thus, we let the time between innovations, Δ (period length) converge to 0. Recall that $\delta = e^{-r\Delta}$. Taking care to adjust the flow value of buyer surplus, v , for the rate of innovation via $v = \int_0^\Delta \lambda e^{-r\tau} d\tau = \lambda(1 - e^{-r\Delta})/r$, where λ is the instantaneous flow value of quality to buyers, limiting outcomes can be calculated directly. First, we observe that the flow dominance lower bound on the seller payoff of $v/(1 - \delta)$ (Lemma 2) converges to λ/r . Thus, in the limit flow dominance reduces to the seller collecting a payment of λ at each instant. This is, however, a vanishingly small fraction of the total surplus, since $v/(1 - \delta)^2$ grows without bound. Given this, it follows directly that in an efficient equilibrium the maximum buyer share of the surplus converges to 1. Thus, we have.

Corollary 1 *In the limit, as upgrades become increasingly frequent ($\Delta \rightarrow 0$), the seller’s minimum share of the surplus goes to zero and the buyers’ maximum share goes to one.*

Intuitively, as innovation becomes increasingly frequent, the flow value becomes a smaller and smaller portion of the total surplus. This reflects a necessary limitation on the extent of the seller’s market power. Despite the fact that the surplus grows without bound, flow dominance only ensures a finite profit for the seller. Thus there is no guarantee of market power, as measured by profit as a share of the total surplus, when innovations arrive very frequently.

5.4 Infrequent Upgrades: $\delta \leq 1/2$

We now take up the case when innovations are infrequent, $\delta \leq 1/2$. Flow dominance will play a stronger role, since the flow value of a unit of quality, v , is now larger than the future discounted value of a unit, $\frac{\delta v}{1-\delta}$. When $\delta \leq 1/2$, a buyer necessarily values one unit today more than two

units tomorrow, other things equal. Stated a bit differently, given a prospect of receiving two units tomorrow free of charge, there is always a positive price that a buyer would be willing to pay to acquire, instead, one unit today. This immediately implies that the seller will necessarily be able to induce a ‘speed up’ in a number of situations where it is not possible to do so when the discount factor is larger than $1/2$. As we show in the next section, it rules out the possibility of delay in equilibrium when $\delta \leq 1/2$.

The increased power of flow dominance, the driving force behind the speed-up argument, does not, however, imply a necessary increase in market power for the seller. We are still able to support efficient equilibria in which buyers receive any payoff $u_1 \in [0, \delta S_1]$ and, as a result, the range of equilibrium payoffs for seller continues to include the flow dominance lower bound. The speed-up argument does change the qualitative nature of the equilibrium support. In particular, no equilibrium can be supported with a constant buyer continuation payoff (i.e., $u_1 = u_2 = \dots = \bar{u}$); the relatively low continuation value implies that buyers will not reject a deviation offer of a price above p_1 (for the candidate u_1). Instead, equilibrium requires a relatively high continuation payoff for buyers so that the seller cannot successfully increase price in period one.

The following proposition formalizes the above argument.

Proposition 3 *Suppose $\delta \leq 1/2$. Then any $u_1 \in [0, \delta S_1]$ is supported as a buyer payoff in an efficient equilibrium with a continuation utility of $u_\tau = \bar{u} \equiv \frac{1-\delta}{\delta} u_1$, for all $\tau \geq 2$.*

Consider the limiting case as innovation becomes less and less frequent, $\Delta \rightarrow \infty$. This provides a useful reference point for market power. As Δ rises, the value of the future sequence of innovation upgrades declines (innovation occurs less frequently) and the current flow value of v comes to dominate the future surplus of δS_1 . In the limit, as $\Delta \rightarrow \infty$, we find that $v \rightarrow \lambda/r$ and we have a unique equilibrium outcome in which $\pi_1 = p_1 = v$ and $u_1 = 0$. Formally, this is now equivalent to a static model and the seller is able to extract all surplus from the buyers. Thus, the flow dominance lower bound on the seller’s payoff can be viewed in terms of limiting the seller’s market power to that of a static monopolist who derives no added value from the upgrade market.

6 Delay and Inefficient Equilibria

We now show that equilibria do not necessarily have to be efficient. By Proposition 1, every equilibrium is a t -cycle equilibrium. Consider, therefore, $t \geq 2$ so that equilibria exhibit (inefficient) delay. Then, on the equilibrium path, we have no sales in periods 1 through $t-1$ and $(1, 0)$ through $(t-1, 0)$ are “delay states.” In period t , when the quality gap has reached t , there is a sale of t units at a price p_t . Thus, the continuation state after the sale is then $(t+1, t)$ which, by Markovian behavior, reduces to state $(1, 0)$. This returns the quality gap to 1 and the cycle begins again.

Payoffs in a t -cycle equilibrium are then $\pi_t = \frac{p_t}{1-\delta^t}$ for the seller, as the revenue flow of p_t is received once every t periods, and $u_t = \frac{1}{1-\delta^t} \left[\frac{vt}{1-\delta} - p_t \right]$ for the buyers, as a purchase of t units at a price p_t is made once every t periods. Due to delay, the realized joint surplus in a t -cycle equilibrium is less than the maximal surplus S_1 . Delay also implies that the continuation path from any state in a t -cycle equilibrium always has a smaller realized surplus than that of the efficient path. Letting Ψ_t denote the realized joint surplus in equilibrium at the time of a sale, the seller

and buyer payoffs imply the equilibrium relationship

$$\Psi_t \equiv \pi_t + u_t = \frac{vt}{(1-\delta)(1-\delta^t)}.$$

Note that $\Psi_\tau = \delta^{t-\tau}\Psi_t$ for $\tau = 1, \dots, t-1$ in the delay states. Thus, $\Psi_1 = \delta^{t-1}\Psi_t$, $\Psi_2 = \delta^{t-2}\Psi_t$, and so on as we approach the sale date of t . The same discounting pattern holds for seller profits and buyer utilities.

We first show that delay cannot occur in equilibrium when the discount factor is below $1/2$. We then consider seller and buyer incentives in the delay states and derive approach conditions that must be satisfied in an equilibrium. Next, as we did with efficient equilibria, we construct a ‘‘cash-in’’ continuation support for (off-equilibrium) states in which the quality gap exceeds t . We then derive conditions for the sale at date t and, finally, show existence.

6.1 Delay Equilibria and Upgrade Frequency

An equilibrium with delay has no sales until period $t > 1$. What prevents the seller from profitably deviating to make a ‘speed-up’ offer? Suppose the seller were to offer a bundle of $t-1$ units for a price of \hat{p} in state $(t-1, 0)$. If all other buyers reject such an offer, then a deviating buyer would accept if

$$v(t-1) - \hat{p} + \delta \max \left\{ \frac{v(t-1)}{1-\delta}, u_t \right\} > \delta u_t.$$

Intuitively, if δu_t is relatively small when δ is small then this will hold and the offer of \hat{p} would necessarily be accepted by all buyers in any symmetric equilibrium. Would the offer then be profitable for the seller relative to delaying and selling in the next period? When $\delta \leq 1/2$, this is necessarily the case.

Proposition 4 *Suppose $\delta \leq 1/2$. Then there does not exist an equilibrium with delay.*

When $\delta \leq 1/2$ the seller and the buyers both value current flows more heavily than future ones. Intuitively, upgrade innovations are sufficiently infrequent that a mutually beneficial speed-up deviation to avoid delay is possible. An individual buyer with $t-1$ units on hand would not purchase the t bundle in state $(t, 0)$ and this makes the seller’s speed-up offer for $t-1$ units in state $(t-1, 0)$ attractive to an individual buyer. When all buyers accept, the current revenue dominates the payoff from waiting to sell in the next period.

The longer the delay, the easier it is to find a speed-up deviation for a given δ . That is, as the delay t rises, a speed-up deviation destroys a delay equilibrium for a range of δ that exceeds $1/2$. This suggests that for a given delay, we will need a sufficiently high discount factor to support the equilibrium: as feasible upgrades become more frequent, it is possible that they are bundled on the equilibrium path.

6.2 Delay and Approach Conditions

To have delay on the equilibrium path starting at state $(1, 0)$, it must be that the seller can find no offer that is acceptable to buyers and also profitable relative to waiting to sell in state $(t, 0)$ when the quality gap has reached t units. Thus, we must derive approach conditions for the equilibrium. For the buyers, we must specify cut-off rules of when to accept offers for σ units in period $\tau < t$.

Given these cut-off prices, $p(\sigma, \tau)$, the seller must find it optimal not to make any offers until the state is $(t, 0)$. Note that it is never credible for buyers to reject all seller offers due to flow dominance.

The cut-off rules for buyers in periods $\tau < t$ are to reject any (upgrade) offer for $\sigma \leq \tau$ units at a price greater than $p(\sigma, \tau)$, where $p(\sigma, \tau)$ satisfies

$$\frac{v\sigma(1 - \delta^{t-\tau})}{(1 - \delta)} + \delta^{t-\tau} \max \left[\frac{v\sigma}{(1 - \delta)}, u_t \right] - \delta^{t-\tau} u_t \leq p(\sigma, \tau) \leq \frac{v\sigma}{(1 - \delta)} + \delta^{t-(\tau-\sigma)} u_t. \quad (7)$$

Note that we have used the continuation properties $\delta^{t-\tau} u_t = u(\tau + 1, 0)$ and $\delta^{t-(\tau-\sigma)} u_t = u(\tau + 1 - \sigma, 0)$ as the first sale on the equilibrium path occurs in state $(t, 0)$. The left-hand-side of (7) provides the lower bound on the cut-off price; otherwise, an individual buyer would be better off accepting when other buyers reject. This bound reflects the difference in gross surplus for an individual buyer between buying and rejecting, since other buyers are expected to reject and, hence, the continuation state would be $(\tau + 1, 0)$. The first term is the buyer's interim flow payoff, from τ until t , generated by σ units while the second terms correspond to the option of buying (or not) with the other buyers once the state reaches $(t, 0)$. The right-hand-side provides an upper bound on the cut-off price; if it failed, an individual buyer would be better off rejecting when others accept. Given that all other buyers are expected to buy the package and the state will be $(t + 1, \sigma)$, the bound reflects the payoff difference for an individual buyer between buying and not buying the offer for σ units. Clearly, a set of prices exists that satisfies (7) for $\sigma = 1, \dots, \tau$ and $\tau = 1, \dots, t - 1$. As with efficient equilibria, this is due to the implicit coordination on cut-off prices among buyers.

Given the cut-off prices, delay must be optimal for the seller. Thus, in state $(\tau, 0)$ for $\tau < t$, the seller prefers the equilibrium path payoff of $\pi_\tau = \delta^{t-\tau} \pi_t$ to selling σ units in period τ at a price of $p = p(\sigma, \tau)$ and receiving a payoff of $p(\sigma, \tau) + \delta^{t-(\tau-\sigma)} \pi_t$. We then have

$$\delta^{t-\tau} (1 - \delta^\sigma) \pi_t \geq p(\sigma, \tau) \quad (8)$$

for $\sigma = 1, \dots, \tau$ and $\tau = 1, \dots, t - 1$. The next lemma greatly simplifies the analysis of these buyer and seller approach conditions.

Lemma 5 *If the buyer and seller approach conditions, (7) and (8), hold for $\sigma = \tau$, at each $\tau = 1, \dots, t - 1$, then the conditions hold for all feasible pairs (σ, τ) .*

As a result, we need only find $t - 1$ distinct prices, $p(1, 1), \dots, p(t - 1, t - 1)$ and it is sufficient to deter the seller from selling the maximum feasible number of units, “cashing-in,” in each delay period. Intuitively, if it is not profitable to sell τ units in period τ , the first time it is possible to do so, it will not be profitable to sell τ units in a later delay state. For example, if the seller does not offer one unit in state $(1, 0)$, then there will be no temptation to sell one unit at a later date when the additional unsold units will create a longer delay in the continuation state.

Delay implies that buyers necessarily receive a positive payoff in equilibrium.

Lemma 6 *In any inefficient $t - 1$ cycle equilibrium it is necessary that $u_t > 0$.*

Thus, the seller cannot extract the full surplus of Ψ_t and buyers must receive positive utility. The intuition is quite simple. The approach condition for buyers at $u_t = 0$ dictates that buyers will accept any offer that provides positive utility. Then, the seller can successfully make a deviating

offer that extracts buyers in periods before t , thus speeding up sales and increasing profits. For example, in the case when the cycle length is two periods, then conditions (7) and (8) imply that u_2 must exceed $\frac{v}{\delta(1+\delta)}$. At a smaller utility level, the seller would be able to offer 1 unit in period 1 at a price that attracts buyers and provides a profitable speed-up deviation.

In view of the need to provide a positive payoff for buyers, how can we ensure that the delay incentives are satisfied? Combining the buyer and seller approach conditions, (7) and (8), we see that supporting prices exist if and only if

$$\frac{v\tau(1 - \delta^{t-\tau})}{(1 - \delta)} + \delta^{t-\tau} \max \left[\frac{v\tau}{(1 - \delta)}, u_t \right] - \delta^{t-\tau} u_t \leq p(\tau, \tau) \leq \delta^{t-\tau} (1 - \delta^\tau) \pi_t \quad (9)$$

for $\tau = 1, \dots, t - 1$. Recalling $\pi_t + u_t = \Psi_t$, (9) reduces to

$$(\delta^{t-\tau} - \delta^t)(\Psi_t - u_t) \geq \frac{v\tau(1 - \delta^{t-\tau})}{(1 - \delta)} + \delta^{t-\tau} \max \left[\frac{v\tau}{(1 - \delta)}, u_t \right] - \delta^{t-\tau} u_t \quad (10)$$

for $\tau = 1, \dots, t - 1$. Thus, satisfying the approach conditions for a $t - cycle$ (delay) equilibrium reduces to finding a buyer utility for the sale date, u_t , that satisfies (10) for a given t and δ . Note that buyer incentives can change significantly during the approach when u_t lies between $\frac{v}{1-\delta}$ and $\frac{v(t-1)}{1-\delta}$. Initially, when τ is small, a deviating buyer who acquired τ units would be willing to purchase again at the sale date t . However, when τ is closer to t , such a buyer would not purchase again at t since the incremental surplus from $t - \tau$ units is insufficient.

This change in buyer deviation incentives highlights the fact that the approach condition can have a binding constraint at an interior τ . This is important for maintaining the seller's incentives during the approach. Intuitively, when u_t is too large relative to δ^t the seller is unwilling to wait until t for a sale and (10) will fail at $\tau = 1$ as a profitable deviation offer will exist: a deviating buyer would purchase again and a cut-off price has a significant component in the form of the interim surplus. On the other hand, when u_t is too small relative to δ^t the seller will find a profitable deviation at τ closer to t . While a complete characterization of (10) is quite involved, it turns out that many of the complications only arise at relatively low discount factors.¹² Thus, we are able to develop a sufficient condition on δ^t such that if δ^t is above a threshold then the approach conditions (10) are satisfied for an interval of utility levels. From Proposition 4, we know that delay equilibria require relatively high discount factors. While the sufficient condition excludes some equilibria that arise when δ is at the low end of this range, it has the advantage of allowing us to study behavior for all larger δ without having to deal explicitly with the subtle pattern of binding incentives at interior delay states.

The sufficient condition is generated by a slope condition on the terms involving discounted payoffs in $t - \tau$, the length of time until the sale, and this rate of change is proportional to $\ln \delta$. We define $a(d) \equiv -\ln \left[-\frac{d \ln(d)}{1-d} \right]$, and note that there exists a unique root $d^* \in (0, 1)$ for the equation

¹²This is because the conditions at $\tau = 1$ and $\tau = t - 1$ are not sufficient, in general, and we have to consider when a binding constraint arises at an interior integer as well as the endpoints. The cases of $t = 2$ and $t = 3$ can be solved explicitly and reveal that the complications with integer constraints arise primarily at the lowest discount factors for which (10) can be satisfied.




Figure 2: Approach Conditions ($v = 1$ and $t = 2$)

$d = a(d)$. We then introduce an upper and lower bound on buyer utility:

$$\begin{aligned}\underline{u}^A &= \Psi_t + \frac{v}{\delta^t \ln \delta} \left(\frac{1 - a(\delta^t)}{1 - \delta} \right) \\ \bar{u}^A &= \Psi_t + \frac{v}{\delta^t \ln \delta} \left(\frac{1 - \delta^t}{1 - \delta} \right).\end{aligned}$$

This leads to the following result.

Lemma 7 *If $\delta^t > d^*$, then there exist bounds \underline{u}^A and \bar{u}^A such that the approach conditions (10) are satisfied for any $u_t \in (\underline{u}^A, \bar{u}^A)$ in a t -cycle equilibrium.*

Numerically, the equation root d^* is about .439. For example, if $t = 2$, then δ must be at least $\sqrt{.439} = 0.663$. One can interpret the t -cycle as having two stages, the approach and the sale date, where the discount factor between stages is δ^t . Hence, the longer delay in equilibrium, the higher must be δ so that the seller will not find a profitable deviation. Specifically, δ must exceed $d(t) \equiv \sqrt[t]{d^*}$, which is clearly increasing in t .

Lemma 7 provides a lower and upper bound on the payoff to buyers. The bounds on utility, \underline{u}^A and \bar{u}^A , depend on δ and t and they are derived in Appendix C. At d^* , $\underline{u}^A = \bar{u}^A$, and for all t and δ pairs where $\delta^t > d^*$, we have $\underline{u}^A < \bar{u}^A$. This pattern is shown for $t = 2$ in Figure 2.

6.3 Existence and Payoffs for Delay Equilibria

In full analogy to the analysis of efficient equilibria, the cash-in (off equilibrium) support conditions for inefficient t -cycle equilibria are given by buyer cut-off rules that satisfy (1) and seller profits that satisfy (2). The only difference relative to the efficient case is that we must replace the

efficient surplus S_τ with the realized surplus, Ψ_τ . Furthermore, the set of support utilities for off the equilibrium path states greater than t need to be modified for delay equilibria. For expositional purposes, we provide the details in Appendix C. We can now turn directly to the existence of delay equilibria. Sufficient conditions for approach states, $(\tau, 0)$ where $\tau < t$, are provided in Lemma 7 and summarized by the bounds, $\underline{u}^A < u_t < \bar{u}^A$, on the equilibrium payoff. We can combine the approach conditions with the cash-in support conditions to establish the existence of delay equilibria.

Proposition 5 *Let $t \geq 2$ and suppose that $\delta^t > d^*$. Then every $u_t \in (\underline{u}^A, \bar{u}^A)$ can be supported in a $t - cycle$ inefficient equilibrium.*

Thus, the binding constraints on what buyer utilities can be supported are generated solely by the approach incentives. Intuitively, as δ rises we can employ longer support lengths to support higher buyer payoffs.

The relationship between delay equilibria and market power is best understood in terms of the rate of innovation. Note that the length of time between sales in a $t - cycle$ equilibrium is given by $D \equiv \Delta t$. As we let Δ decline, any given $t - cycle$ will continue to exist, since $\delta^t = e^{-r\Delta t} > d^*$ will continue to hold, but the delay length D will go to zero. However, longer length cycles can be supported as Δ declines. As with efficient equilibria, we assess seller market power as the ratio of seller profits relative to realized equilibrium surplus; we also assess the efficiency loss from delay, namely $\delta^{t-1}\Psi_t/S_1$.

The utility bounds relative to realized equilibrium surplus, $\bar{s}_b \equiv \bar{u}^A/\Psi_t$ and $\underline{s}_b \equiv \underline{u}^A/\Psi_t$, index the range of equilibrium payoffs. Simplifying, we find an invariance property in that the bounds depend only on the length of delay, D , and not on t and Δ individually. The limiting cases are when delay vanishes, $D \rightarrow 0$ (for example when $\Delta \rightarrow 0$ for a given $t - cycle$) and when delay is maximized, $D \rightarrow D_{\max} \equiv (-\ln d^*)/r$. Straightforward limit calculations yield the following result.

Corollary 2 *As equilibrium delay vanishes, the bounds on buyer utility converge to those for efficient equilibria, $\bar{s}_b \rightarrow 1$ and $\underline{s}_b \rightarrow 0$ as $D \rightarrow 0$. As equilibrium delay approaches D_{\max} , the bounds on buyer utility converge to each other at approximately .13 of the realized surplus. Further, the maximum efficiency loss from delay is approximately .35 and occurs as $D \rightarrow D_{\max}$ and $\lambda \rightarrow 0$.*

In the limit where the buyer utility is 13% of the realized surplus and thus the seller receives 87% of this surplus, the 35% efficiency loss at the maximum delay implies that the seller would prefer 57% of the efficient surplus to 87% of the smaller delay surplus.

7 Discussion of Upgrade Structure

In this section, we consider how our endogenous bundling framework corresponds to observed practice in a variety of upgrade markets and discuss our results in relation to assumptions on bundling and the upgrade literature including price contingencies, bundling restrictions, and product complementarities.

7.1 Bundling in Practice

In practice, the upgrade process varies greatly with respect to how buyers move to higher quality levels. In our model, the seller is free to offer an unrestricted set of bundles. In equilibrium, offers

take the form of upgrades, defined as a bundle of the form $\{\sigma + 1, \dots, \sigma + k\}$. Thus, a buyer with units $\{1, \dots, \sigma\}$ purchases the upgrade offer and the new units are used in conjunction with the current ones. Effectively, the upgrade is an add-on. Alternatively, the seller in our model can offer bundles that constitute new generations (versions), bundles of the form $\{1, \dots, \sigma, \dots, \sigma + k\}$. As we will show, there is a very close relationship between upgrade offers and generations once downward complementarity and pricing contingent on a buyer's current holdings are taken into account.

A number of familiar examples help to illustrate the nature of upgrade (add-on) and generation (version) offers. For example, products such as media players from Windows and Real.com, or browsers such as Internet Explorer and Mozilla Firefox, are add-on programs in the sense that they can be installed on an existing computer. On the other hand, quality improvements for a given software program often involve replacing an older version with a newer version. However, as the anti-trust debate illustrates, the distinction between upgrades and generations is not necessarily driven by technological considerations. In this regard, recall that one of the United States's anti-trust cases against Microsoft centered on bundling of the internet browser and the operating system, while in Europe concern focused on the bundling of the media player with the operating system. Other cases demonstrate how the capabilities that were provided by a separate program, such as creating a PDF with Acrobat or presentation slides with Beamer have become part of word processing programs. Thus, the form of offered bundles, whether determined by firms or policy makers, can be viewed as an endogenous outcome.

In cases where a specific good has multiple quality dimensions, upgrade add-ons and generations can co-exist. Consider, for example, television sets or computer monitors. On one dimension, screen size, higher quality requires a larger television or monitor. On another dimension, picture quality, however, an existing analog television can be upgraded by adding a digital tuner for high-definition programming and existing monitors can be upgraded by installing a better video card. In order to accommodate quality increases via either replacement or add-on technologies, a bundling framework for seller offers is necessary.

Contract contingencies, especially with respect to a buyer's current holdings, are a frequently observed component of seller offers. MacKichan, for example, offers the technical word processor Scientific Word 5.5 in a number of versions differentiated by features and each version has an upgrade price for prior users (serial number required) as well as a (higher) price for new users. Airlines typically offer seat upgrades, club memberships, and other amenities at prices that vary with frequent flier status, a result of past purchases.

A generation version with a price contingent on a buyer's current version is very close and often will be equivalent to an upgrade offer. Consider a buyer who holds units $\{1, \dots, \sigma\}$ and two offers. One is an upgrade bundle $\{\sigma + 1, \dots, \sigma + k\}$ for price p . The other is a new version $\{1, \dots, \sigma, \dots, \sigma + k\}$ at price p that is only available to buyers who hold units $\{1, \dots, \sigma\}$. The direct value to the buyer is the same with either bundle. Now, consider the same offers, but suppose that the buyer does not hold any units. By downward complementarity, the upgrade bundle has no direct value (non-contiguous units) while the buyer does not qualify for the other offer. Further, observe that if the pricing contingency is stated as a minimum requirement then buyers who hold at least the minimum will place a common value on the two bundles. Such a minimum requirement is common in practice. Microsoft allows any 2000-2007 Office program or suite to qualify a buyer for Office Professional at the upgrade (discount) price.

7.2 Upgrades, versions, and customer contingent pricing

What, if any, are the differences between an upgrade bundle and a generation bundle where the price is contingent on a buyer's holdings? From the buyer's point of view, there is no difference. In our analysis, every seller offer in every state, on and off the equilibrium path, involves only upgrade bundles. Suppose that we replaced every such upgrade, $\{\sigma + 1, \dots, \sigma + k\}$, with a generation bundle, $\{1, \dots, \sigma, \dots, \sigma + k\}$, identically priced, but where only buyers who hold units $\{1, \dots, \sigma\}$ qualify for purchase. Then, as just noted, every buyer will make exactly the same purchasing decisions as with the upgrade bundles. Thus, our equilibria will be robust to allowing price contingencies if we can demonstrate that it is still optimal for the seller to make the same offers (either in the upgrade form or the appropriate generation form with a holding contingency).

This is easy to see in a finite horizon setting and the unique outcome (efficient path and all surplus accrues to the seller) will continue to prevail. With our downward complementarity structure, the ability to offer partial upgrade bundles (lower quality units are not included) already allows the seller to induce buyers with different holdings to accept different bundles.¹³ Thus, whether we regard the offers as upgrade bundles or generations with price contingencies, as in Fudenberg and Tirole (1998) or Ellison and Fudenberg (2000), does not matter.¹⁴

In our primary case with an infinite horizon, we need to consider the possibility that, by conditioning offers on current holdings, a seller may be able to curtail the credible threat of buyers to reject offers with high prices and thus eliminate equilibria with low seller payoffs. In this regard, our equilibria are robust. First, recall that buyers all have the same holdings on any equilibrium path and a contractual contingency in this regard has no force. With respect to the support for the equilibrium path, the same observation applies to the cash-in support for states of the form $(\tau, 0)$. Finally, for states in which buyers have asymmetric holdings, we constructed continuation equilibria with offers in which the seller captured the available surplus and, as a result, they are robust to the addition.

At a more intuitive level, recall that the credible threat to reject seller offers with high prices is based on the expectation of a sufficiently high future surplus (when all buyers reject the offer). Consider, for example, the support condition (3) for an efficient equilibrium and the impact of allowing contract contingencies on buyer holdings. An individual buyer who fails to purchase when others do will fall behind the equilibrium path, but such a buyer is already extracted in our analysis. On the other hand, a buyer who purchases when others do not will be ahead of the market and, in our analysis, such a buyer does have a strict preference for purchasing the subsequent equilibrium support cash-in offer from the seller. The seller could then employ a contractual contingency to isolate such a buyer and eliminate the (valuable) option to accept an offer designed for buyers with fewer units and "rejoin" the equilibrium path. But if the contingency is used to make an offer that extracts the deviating buyer then this will only serve to reduce further the deviation payoff to

¹³The proof is very similar to the finite case. The only potential difference is that the seller can use a contingency to prevent a buyer who is ahead of the market from purchasing bundles that are intended for buyers with lower holdings. However, since no buyer earns an incremental surplus in any continuation, the contingency has no impact.

¹⁴In both Fudenberg and Tirole and Fudenberg and Ellison the second period offers can distinguish between buyers who purchased in period 1 and those who did not. These models allow for buyer heterogeneity and, thus, the contract contingency is much more salient for influencing buyers with different valuations than in our model where buyers are homogenous.

purchasing when others do not (refer to the max condition for $g(\tau, \tau)$ in (3)). Thus, the support condition (3) continues to be satisfied by our T -stage utility path and buyers retain a credible threat based on future surplus to reject offers with high prices.

In short, the ability to identify customers based on past purchases does not affect the equilibrium structure. In terms of information structures, it is worth noting that our original offer set corresponds to the semi-anonymous regime of Fudenberg and Tirole (1998), where harsher terms cannot be imposed on buyers who hold more units than others.

This establishes that our results do not depend on the precise form of the seller’s offers. Bundles can be presented as upgrades or as generations with a contingency on past purchases. In this regard, our formulation of the offer space is simpler as it avoids the added complications of contingencies. We note, however, that limiting the offer space to generation version bundles while adding price contingencies will also result in the same set of equilibria.

7.3 Unbreakable version offers

Suppose that, as an exogenous condition, all bundles must be versions and no contingencies on purchase are allowed. That is, any offer for the current quality increment, unit τ , is necessarily also an offer for units $\{1, \dots, \tau\}$ and similarly for any lower quality level. This might reflect a necessary property of the production technology, where a quality increment cannot be ‘broken out’ for separate sale. This offer structure is examined in Waldman (1996) for a two-period model and in Fishman and Rob (2000) for an infinite horizon model, both of whom focus on innovation incentives.

In our setting, where buyers have identical preferences, this “unbreakable” upgrade structure necessarily limits the market power of the seller. The reason is that a buyer always has the option of passing on a current offer and waiting to purchase a later offer. As long as the seller eventually offers a higher quality, the cost of waiting (relative to purchasing when other buyers do) is the lost flow value. Any subsequent higher-quality offer that attracts prior buyers will necessarily provide the buyer who delays with a strictly larger surplus than that received by prior buyers. As a result, there is no equilibrium in which a monopoly seller is able to capture the full surplus from buyers in a market where upgrades are unbreakable.

In the simplest case of a two-period unbreakable goods version of our model, it is straightforward to derive the equilibrium outcomes. For large discount factors, the seller delays until period 2 and then offers units $\{1, 2\}$ at the extraction price for the remaining surplus. For low discount factors, however, unit 1 is sold in period 1 at price v and then period 2 has extraction pricing for units $\{1, 2\}$. In both cases, the seller is unable to capture the full surplus, in contrast to the finite horizon equilibrium with upgrade offers or generation offers with price contingencies.

In the case of an infinite horizon, non-contingent version in the presence of ongoing quality growth necessarily imposes a more severe limit on the seller. Fishman and Rob (2000) point out that the option to wait implies that the seller can charge no more than the flow value.¹⁵ In their

¹⁵Full extraction of total surplus by the seller requires implementing the efficient path and, in the unbreakable version of our model, the seller would be limited to prices that reflect only the flow value and not the present discounted value to buyers from quality increments. Formally, consider the efficient path and let $p(\tau, \tau - 1)$ be an equilibrium price for version τ when all buyers hold $\{1, \dots, \tau - 1\}$. By rejecting an offer and resuming purchases next period, an individual buyer obtains $v(\tau - 1) + \delta u(\tau + 1, \tau)$. Purchasing today

rational expectations equilibrium, they find that the low price leads to a rate of innovation that is inefficiently low. In our analysis, where bundling is endogenous, the option to wait (fall behind the market) is inconsequential. We find that the basis of a credible threat for buyers resides with the extent to which the seller can tempt a buyer to move ahead of others.

7.4 Independent goods

The case of ‘independent’ goods involves a buyer who receives a flow utility of v from a good independently of whether the buyer holds any other units. With a complete absence of complementarity across quality levels, units $1, 2, \dots$ are effectively independent goods. Then, as one might expect, the speed-up logic of Fudenberg, Levine, and Tirole (1985) implies that the seller regains the ability to extract buyers because the buyers lack a credible threat for refusing a seller offer. Consider, for example, payoffs in a Markov equilibrium for the efficient path and suppose the equilibrium has a cash-in support at price p_τ in state $(\tau, 0)$. If buyers are to reject a price $p > p_\tau$ for τ units in state $(\tau, 0)$, then we must have

$$\begin{aligned} (v\tau - p) + \delta v\tau + \delta^2 \left(\frac{v\tau}{1 - \delta} + u_1 \right) &\leq \delta u_{\tau+1} = \frac{v(\tau + 1)}{1 - \delta} - p_{\tau+1} + \delta u_1 \Rightarrow \\ u_\tau - \delta(1 - \delta)u_1 &\leq \delta u_{\tau+1} \end{aligned}$$

since an individual buyer can accept the offer at price p (while other buyers reject), skip the cash-in offer in the continuation state of $(\tau + 1, 0)$, and then resume purchasing in state $(\tau + 2, \tau + 1)$. Because the goods are independent, there is no payoff consequence due to complementarity from the deviating buyer’s lack of unit $\tau + 1$. But this immediately implies the recursion

$$\left(\delta + \frac{1 - \delta}{\delta^\tau} \right) u_1 \leq u_{\tau+1}$$

and, as with our benchmark case of a single buyer, the threat to refuse the price increase is not credible since buyer utility would have to grow faster (at rate $1/\delta^\tau$) than joint surplus. Thus, with independent goods, we necessarily have $u_1 = 0$. Intuitively, complementarity is essential for building a credible threat to refuse price increases as a deviating individual buyer faces the extra cost of having to acquire the missing unit to avoid falling behind the market. This is why, in our upgrade structure, it matters to an individual buyer whether or not other buyers are expected to purchase the seller’s offer.

8 Directions for Future Work

We have assumed homogeneous buyers in our model. This was done primarily to focus on the structure of credible threats for buyers in a dynamic upgrade model in what one would expect to be the ideal situation for a seller to capture the full (efficient) joint surplus. Allowing for buyer heterogeneity is an important direction for subsequent work. In practice, it is common for sellers in upgrade markets to offer simultaneously different versions or quality levels of their products. This is typically taken to be a form of price discrimination. As noted before, several papers examine a finite horizon model, but there has been very little theoretical work on infinite horizon models in which

yields $v\tau - p(\tau, \tau - 1) + \delta u(\tau + 1, \tau)$. Combining, we must have $p(\tau, \tau - 1) \leq v$.

buyers are always in the market and quality improvements are ongoing. We are currently exploring this problem in our model by allowing for buyer heterogeneity with high or low valuation buyers who are privately informed of their type. This allows for an endogenous determination of pricing and whether the buyer segments remain distinct over time or whether the seller chooses to price over a cycle that periodically brings high and low types together at a common quality level (a generational cycle). An interesting feature of equilibrium price discrimination in this dynamic context is that incentive constraints can bind in both directions (with low-value buyer types choosing to mimic high-value buyer purchases as well the standard downward incentive constraint).

We also assumed an exogenous rate for the increase in quality. Of course, a model that addresses the question of how rewards for a given quality innovation are determined is a necessary step toward an endogenous determination of quality change. We are currently studying innovation and pricing incentives in a model where innovations can be generated not only by an incumbent but also by potential entrants. In this setting, property rights for innovation in relation to imitation incentives are crucial for buyer decisions regarding adopting the products of an incumbent or an entrant and, in turn, for assessing public policy choices and welfare in upgrade markets.

9 Appendix A- Preliminary Results.

In this Appendix we prove Lemmas 1, 2, and 3 as well as Proposition 1.

Proof of Lemma 1. Depending on the history, buyers may also hold a subset, possibly null, of units $\{Q + 2, \dots, \tau - 1\}$. Without unit $Q + 1$, a buyer who rejects the upgrade offer will receive a flow payoff of vQ and have the same quality holdings in period $\tau + 1$. A buyer who accepts will receive a flow payoff of $v\tau$ in period τ and hold $\{1, \dots, \tau\}$ next period. We will show that accepting yields a strictly higher payoff than rejecting, for any strategy choices of other buyers and the seller following the upgrade offer.

Obviously, accepting yields a higher flow payoff in period τ since $v\tau - p > vQ$. Consider the continuation payoff. An accepting buyer (i) begins with more units than a non-accepting buyer and (ii) has the option of mimicking the strategy of any non-accepting buyer (any seller offer may be purchased by any buyer). By the upgrade payoff structure, the continuation payoff of an accepting buyer is therefore at least as large as that of a non-accepting buyer. This is because the mimicking option implies the same payments, but the accepting buyer will never hold fewer units.

It is now clear that every buyer will choose to accept the upgrade offer. Given any strategy choices of other buyers and the seller following the upgrade offer, a buyer who accepts always has a weakly larger payoff from $\tau + 1$ onward, with respect to the continuation sequence of offers implied by the strategies, and a strictly larger flow payoff in period τ . Note that, by the zero measure property, the continuation sequence of offers does not depend on the choice of a specific individual buyer to accept or reject the upgrade offer. ■

Proof of Lemma 2. By Lemma 1, at the start of the game the seller can offer unit 1 for a price of $p_1 < v$ and every buyer will accept. Also, by Lemma 1, in period 2 when all buyers hold unit 1 the seller can offer unit 2 for a price of $p_2 < v$ and every buyer will accept. By induction, in any period τ and for any history in which all buyers hold units $\{1, \dots, \tau - 1\}$ we can apply Lemma 1 to see that the seller can sell unit τ for a price $p_\tau < v$. Each price can be arbitrarily close to v , so letting $v - \epsilon = p_\tau$ for all τ , the seller's payoff from the start of the game must be at least $(v - \epsilon)(1 + \delta + \delta^2 + \dots) = (v - \epsilon)/(1 - \delta)$. As this must hold for any $\epsilon > 0$, we are done.

For the continuation result, simply apply Lemma 1 with $p_\tau = (v - \epsilon)(\tau - Q)$ in period τ and then use the same argument as above in period $\tau + 1$. ■

Proof of Lemma 3. It is sufficient to prove the result for states of the form $(\tau, 0)$, since any state of the form (σ, Q) has the same quality increments and payments as $(\sigma - Q, 0)$. Consider state $(\tau, 0)$ and a continuation path $(\sigma, q_{\sigma-1})$ for $\sigma \geq \tau + 1$. By Lemma 2, we know the seller's payoff in state $(\tau, 0)$ is positive. This implies that $q_{\sigma-1} > 0$ for some σ . Otherwise, we have $q_{\sigma-1} = 0$ for all σ and buyers never acquire unit 1. Hence, buyer payments to the seller must be zero. As a result, the seller's payoff would be zero, which is not possible. Thus, $q_{\sigma-1} > 0$ for some σ . Relabel so that σ denotes the first such period, and we have $q_{\sigma-1} > 0$ and $q_{\sigma'-1} = 0$ for $\sigma' < \sigma$. Thus, the quality gap rises from τ in state $(\tau, 0)$ to $\sigma - 1$ in state $(\sigma - 1, 0)$ and then goes to a gap of $\sigma - q_{\sigma-1}$, which is less than or equal to the previous gap of $\sigma - 1$, in state $(\sigma, q_{\sigma-1})$. By equilibrium, the continuation path from $(\sigma, q_{\sigma-1})$ has the same quality increments as state $(\sigma - q_{\sigma-1}, 0)$. Thus, the quality gap will thereafter cycle repeatedly from size $\sigma - q_{\sigma-1}$ up to $\sigma - 1$ and the continuation path from $(\tau, 0)$ has a bounded quality gap. ■

Proof of Proposition 1. Starting from state $(1, 0)$, we know from Lemma 3 that the quality gap is bounded and, therefore, that there is a first date, say t , at which a sale involving unit 1 takes place. If $t = 1$, we are done as equilibrium (MPE) implies we have a 1-cycle equilibrium. So, consider $t > 1$. By construction, the maximal quality held by buyers before period t is zero, so the state is $(t, 0)$, and $q_t > 0$ results from sales in period t . A potential complication with state $(t, 0)$ that does not occur with $(1, 0)$ is that $(t, 0)$ corresponds to histories in which buyers acquired no quality units as well as histories in which they acquired some subset of $\{2, \dots, t - 1\}$. By definition, however, Markovian behavior requires that the seller offer in $(t, 0)$ and buyer acceptance choice(s) are the same across these histories since strategies only depend on the state $(t, 0)$.

Suppose that the sale at date t does not result in $q_t = t$ or, in other words, buyers do not acquire the full feasible set of units $\{1, 2, \dots, t\}$. This implies that, for some τ where $1 \leq \tau < t$, buyers hold units $\{1, \dots, \tau\}$ but not unit $\tau + 1$ as they enter period $t + 1$. Also, let p denote the total payment made by a buyer to the seller for all bundles purchased in state $(t, 0)$. Finally, note that whether or not any of the units in $\{\tau + 2, \dots, t\}$ are held by buyers before period t or acquired in t , the state in period $t + 1$ will be $(t + 1, \tau)$.

By construction, the equilibrium buyer continuation payoff from state $(t, 0)$ is given by

$$u(t, 0) = v\tau - p + \delta u(t + 1, \tau)$$

as the quality flow utility is $v\tau$ and the payment is p in $(t, 0)$, and next period's state is $(t + 1, \tau)$. We will show that the seller has a profitable deviation in period $t - 1$, namely, offering the feasible (as $\tau < t$) upgrade bundle to quality τ (units $1, \dots, \tau$) for some price \hat{p} .

Before proceeding with the main argument, we need to develop two properties of buyer payoffs. First, the equilibrium path will follow a cycle, since state $(t + 1, \tau)$ has the same quality gap as state $(t + 1 - \tau, 0)$. Thus, the maximal buyer quality remains at τ until period $t + \tau$, when the state reaches $(t + \tau, \tau)$, at which time the maximal buyer quality rises to 2τ and the cycle begins again. Equilibrium also implies that a buyer only needs to make purchases in the states where maximal quality rises in order to achieve the equilibrium buyer payoff. As noted above, the history of play only matters to the extent that it impacts maximal buyer quality. Thus, the bundle(s) offered by the seller in any state of the form $(t + k\tau, k\tau)$, where $k = 1, 2, \dots$, must, at a minimum, always include the next τ units of quality. In particular, this is true for the history where buyers hold exactly the first $k\tau$ quality units (and no other units), since the maximal quality for this history is $k\tau$. Thus,

an individual buyer never needs to hold more than these units in order to be able to reach the next equilibrium path level of maximal quality via purchases in state $(t + k\tau, k\tau)$. Furthermore, such a buyer can always choose from the same offered bundle(s) and price as any other buyer. It follows directly that the continuation payoff of a buyer only depends on holding the current maximal quality and it is independent of whether the buyer holds higher but non-contiguous units. This is the first property of buyer payoffs that we will need.

The second property is that, in equilibrium, the seller only receives revenues in states of the form $(t + k\tau, k\tau)$, where $k = 1, 2, \dots$. In $(t + k\tau, k\tau)$, in equilibrium, the seller offer must include units $\{k\tau + 1, \dots, k\tau + \tau\}$ and all buyers must acquire these units. Thus, no buyer ever pays a positive price for any bundle in states $(1, 0)$ through $(\tau, 0)$, since only units in $\{2, \dots, \tau\}$ can be offered by the seller in equilibrium and these units will necessarily be acquired via the upgrade bundle in state $(t, 0)$ when buyers also acquire unit 1. The same logic then applies for the next τ units, and so on.

We now proceed with the main deviation argument. To keep things simple, let us first consider the case where the history for state $(t, 0)$ has buyers holding no quality units. For the seller deviation in period $t - 1$, choose the price \hat{p} for the bundle of units $\{1, \dots, \tau\}$ so that

$$\hat{u} \equiv v\tau - \hat{p} + \delta v\tau + \delta^2 u(t + 1, \tau) = \delta u(t, 0) + \epsilon,$$

for a small $\epsilon > 0$. Combining this with the earlier expression for $u(t, 0)$, we find that $\hat{p} = v\tau + \delta p - \epsilon$. We claim that in any equilibrium continuation after this offer, all buyers will accept. By symmetry of strategies, in response to this offer in state $(t - 1, 0)$, all buyers must either accept or reject. Suppose the buyer strategy calls for a rejection and consider the decision of an individual buyer. Because no other buyer accepts, the continuation state will be $(t, 0)$. By accepting, the payoff for an individual buyer is $\delta u(t, 0) + \epsilon$. To see this, note first that the individual buyer receives a flow of $v\tau - \hat{p} = -\delta p + \epsilon$ in period $t - 1$. In period t by (optimally) making no purchase in period t , the buyer receives a flow of $v\tau$.

Finally, consider the continuation state $(t + 1, \tau)$. A complication is that, in addition to the first τ units, the outcome in state $(t, 0)$ may involve buyers acquiring units in $\{\tau + 2, \dots, t\}$. By making no purchases in period t , the deviating buyer will lack these units in the future while other buyers possess them. But, this is of no consequence in equilibrium: a buyer holding exactly τ units obtains the same continuation payoff of $u(t + 1, \tau)$, by our first property of buyer payoffs.

Now, adding the terms in periods $t - 1$, t , and $t + 1$ for a deviating buyer, we arrive at \hat{u} as in the above equation. Thus, accepting the seller's deviation offer in period $t - 1$ for τ units results in a higher payoff than rejecting and waiting whenever $\epsilon > 0$. Thus, all buyers rejecting the offer is never an equilibrium continuation. In a symmetric equilibrium, it must be that all buyers accept the offer in $t - 1$.

Now, to see that the deviation is profitable for the seller, note that the payoff to the deviation offer in period $t - 1$ (where all buyers accept) is

$$\hat{\pi} = \hat{p} + \delta\pi(t, \tau) = v\tau + \delta p - \epsilon + \delta^2\pi(t + 1, \tau) = v\tau - \epsilon + \delta\pi(t, 0) > \delta\pi(t, 0),$$

where we have used the definition of \hat{p} and the equilibrium hypothesis for $(t, 0)$, which implies $\pi(t - 1, 0) = \delta\pi(t, 0)$ and $\pi(t, 0) = p + \delta\pi(t + 1, \tau)$. Thus, we cannot have $\tau < t$ in equilibrium.

Finally, we must verify that the same deviation will work for the seller when the history for state $(t, 0)$ has buyers holding quality units (but not unit 1). By equilibrium, seller's payoffs $\pi(t, \tau)$ and $\pi(t + 1, \tau)$ are independent of these holdings. The only remaining possible complication is that

the deviation offer in $t - 1$ sacrifices revenues that would otherwise have been received by the seller from an offer of units in $\{2, \dots, t - 1\}$. Equilibrium, however, rules out any such revenues for the seller as we showed above. ■

10 Appendix B - Efficient Equilibria.

In this Appendix we first prove Lemma 4, which provides sufficient conditions for equilibrium existence. We then prove several lemmas dealing with properties of the $T - stage$ support. This is followed by proofs of Propositions 2 and 3 for existence of equilibrium with frequent and infrequent innovation, respectively.

(i) Sufficient Conditions for Equilibrium

Proof of Lemma 4. To establish existence of an equilibrium we need to show that (i) our candidate upgrade offer is optimal for the seller with respect to the offer set Ω_τ , and (ii) construct a continuation equilibrium for states where buyers are distributed asymmetrically providing, for any such state, both an optimal offer for the seller and optimal buyer responses for any given seller offer.

Now, define an upgrade offer \mathcal{B} as any offer in Ω_τ , with the property that if $(z, p_z) \in \mathcal{B}$, then $M(z) = \sup\{i \mid i \in z\}$. By construction, every bundle in \mathcal{B} is an upgrade bundle, since the maximal contiguous quality in z coincides with the largest quality unit in z . Thus, we can denote any $(z, p_z) \in \mathcal{B}$ by (b, p_b) , where $b \equiv \sup\{i \mid i \in z\}$ is the upgrade level for z and $p_b \equiv p_z$ is the price. Note that an upgrade offer need not include all of the feasible upgrade bundle levels $1, \dots, \tau$. A buyer will optimally choose at most one bundle in \mathcal{B} .

The restriction to upgrade offers can be shown to be without loss of generality. This is because we construct continuation equilibria in which only upgrade offers are made by the seller. Thus, even if a period τ offer includes non-upgrade bundles, all players expect that every possible period $\tau + 1$ continuation state will involve only upgrade offers. Furthermore, in every possible period $\tau + 1$ continuation state, every buyer will move to a quality holding of at least τ in the continuation equilibrium outcome. Consequently, a buyer will value bundles in an offer only to the extent that the bundle allows the buyer to move to a higher quality level in period τ . If a buyer's purchases in τ result in the acquisition of non-contiguous quality levels, these non-contiguous units have no current or future payoff effect due to the structure of the continuation equilibria.

More generally, if buyers are asymmetrically distributed across maximal quality levels and if buyers hold non-contiguous quality units (above a buyer's maximal level), we can still work with upgrade offers without loss of generality. This is because any offer in conjunction with a buyer's current maximal quality and non-contiguous holdings can always be reduced to an implied set of payments for achievable (higher) maximal quality levels. With the continuation equilibria noted above, any resulting non-contiguous quality units will have no payoff impact.

To demonstrate (i), it is sufficient to show that the seller cannot profitably deviate in state $(\tau, 0)$ to some other upgrade offer (with multiple upgrade bundles). First, define a buyer preference relation, \succeq_B , for any two upgrade bundles by

$$(b, p_b) \succeq_B (i, p_i) \iff vb - p_b + \delta u(\tau + 1, b) \geq vi - p_i + \delta u(\tau + 1, b \mid i),$$

where $u(\tau + 1, b \mid i)$ is equal to $\frac{vi}{1-\delta}$ if $i < b$ and equal to $\max\left\{\frac{vi}{1-\delta}, u(\tau + 1, b)\right\}$ if $i \geq b$. Note that the \succeq_B relation reflects implicit coordination in that an individual buyer has no incentive to

choose (i, p_i) if all other buyers choose (b, p_b) . Now, define an upgrade $(b, p_b) \in \mathcal{B}$ to be a buyer continuation equilibrium (*BCE*) in state $(\tau, 0)$ for offer \mathcal{B} if $(b, p_b) \succeq_B (i, p_i) \forall (i, p_i) \in \mathcal{B}$. We must show that for any offer \mathcal{B} there exists a *BCE* such that the seller cannot gain in state $(\tau, 0)$ by deviating to offer \mathcal{B} instead of making the cash-in offer of p_τ . The proof is lengthy so we only provide a sketch. First, one shows that any two upgrade offers are comparable under \succeq_B . Next, one can show that

$$\arg \max_{(b, p_b) \in \mathcal{B}} \left\{ \frac{vb}{1-\delta} - p_b + \delta u_{\tau+1-b} \right\}$$

is a *BCE* (existence is trivial as \mathcal{B} has a finite number of bundles; if it is not unique then select the arg max with largest upgrade level). Essentially, this follows because the argmax is the highest possible coordinated payoff for buyers and because utility differences across a T -stage support satisfy the bound $u_\sigma - u_{\sigma'} \leq \frac{v(\sigma-\sigma')}{1-\delta}$.

We then have two cases for the offer \mathcal{B} . If every $(b, p_b) \in \mathcal{B}$ satisfies $p_b \geq G_b \equiv g(b, u_{\tau+1})$, then it is easy to show that $(0, 0)$, where all buyers refuse to purchase, is a *BCE* for \mathcal{B} . This is equivalent to a delay outcome and we know from (3) that the seller prefers to make the cash-in offer p_τ . If $p_b < G_b$ for some $(b, p_b) \in \mathcal{B}$ then we first find the arg max specified above for the subset of all such bundles in \mathcal{B} , call it (b^*, p_{b^*}) . We then show that $(b^*, p_{b^*}) \succeq_B (i, p_i) \forall (i, p_i)$ where $p_i \geq G_i$. Then, (b^*, p_{b^*}) is a *BCE* for \mathcal{B} and the resulting seller payoff of $p_{b^*} + \delta u_{\tau+1-b^*}$ is, by (3), not profitable relative to p_τ . This completes (i).

Consider (ii) and any state in which buyers are asymmetrically distributed across quality levels. Since all buyers with the same maximal contiguous quality level are treated identically in the continuation it is sufficient to keep track only of market shares. Thus, let us denote such a state by (τ, α) where $\alpha = (\alpha_0, \dots, \alpha_{\tau-1})$ specifies for each $\sigma = 0, \dots, \tau-1$ the fraction $\alpha_\sigma \in [0, 1]$ of buyers entering period τ with maximal quality level of σ . By hypothesis, $1 > \alpha_0 > 0$ and $\sum_{\sigma=0}^{\tau-1} \alpha_\sigma = 1$. We specify a continuation equilibrium for (τ, α) as follows. The seller makes an upgrade offer $\{(b_\sigma, p_\sigma)\}_{\sigma=0, \dots, \tau-1}$ where each (b_σ, p_σ) is an upgrade from σ to τ , that is the bundle $\{\sigma+1, \dots, \tau\}$, for price

$$p_\sigma = \frac{v(\tau-\sigma)}{1-\delta} + \delta u_1.$$

For buyer strategies, we specify that a buyer with σ units chooses to accept (b_σ, p_σ) . It is straightforward to verify that, when all other buyers follow this strategy, it is optimal for an individual buyer with σ to do so as well. Since these upgrade offers leave each buyer with a payoff of $v\sigma/(1-\delta)$, the payoff to the seller is equal to the continuation surplus of S_τ less these individual-rationality payoffs aggregated across buyers according to the distribution α . By feasibility, this is an upper bound on the seller's payoff in any continuation.

Finally, to complete the argument that the above upgrade offer is an optimal choice for the seller, we need to specify a *BCE* if the seller makes some other upgrade offer. Allowing for partial upgrades, denote such an offer by $\mathcal{B} = \{(b, \sigma; p_{b,\sigma}) \mid \sigma \leq b \leq \tau, 0 \leq p_{b,\sigma}\}$ in (τ, α) , where each $(b, \sigma; p_{b,\sigma})$ denotes an upgrade bundle for units $\{\sigma+1, \dots, b\}$ at price p_b . Since the offer \mathcal{B} has upgrades that begin at different levels and a buyer is free to purchase multiple bundles, we construct from \mathcal{B} for each possible buyer quality level $\sigma = 0, \dots, \tau-1$, the set \mathcal{B}_σ of all upgrade bundles that move σ to a higher quality level; note \mathcal{B}_σ might contain only the refusal option for some σ . We then have each buyer with quality level σ choose to accept the (largest index)

$$\arg \max_{\mathcal{B}_\sigma} \left\{ \frac{vb}{1-\delta} - p_{b,\sigma'} \right\}.$$

Then, these choices can be shown to form a *BCE* for \mathcal{B} . The proof is trivial if the buyer choices result in a nondegenerate distribution across quality levels since the buyers are then held to their individual-rationality payoffs in the continuation state. If not, then all buyers move to some common quality level, say \hat{b} , and we must use the continuation payoff $u_{\tau+1-\hat{b}}$ from the T -stage support. ■

(ii) *Properties of the T – stage Support*

Lemma A1 *Consider a T -stage support, where $u_\tau = u_T \equiv \bar{u}$ for $\tau \geq T$. If the support condition (3) holds at $\tau = T$ for $\sigma = 0, \dots, T$, then (3) holds at $\tau > T$ for $\sigma = 0, \dots, \tau$.*

Proof. We begin by establishing two claims. These are sufficient to prove the lemma for $T = 1$. We then establish a third claim and prove the lemma for $T \geq 2$.

First, we show that if (3) holds at $(0, \tau)$, then it holds at $(0, \tau + 1)$. By (3) at $(0, \tau)$ we have

$$v\tau \geq u_\tau - \delta u_{\tau+1} + g(0, u_{\tau+1}) = (1 - \delta)\bar{u}$$

and this directly implies that (3) holds at $(0, \tau + 1)$, since

$$v(\tau + 1) \geq u_{\tau+1} - \delta u_{\tau+2} + g(0, u_{\tau+2}) = (1 - \delta)\bar{u}.$$

The second claim is that if (3) holds at $(0, \tau)$, then it holds at (σ, τ) for $\sigma = 1, \dots, \tau + 1 - T$. Since $\sigma \leq \tau + 1 - T$ if and only if $T \leq \tau + 1 - \sigma$, the quality gap in the continuation state, $(\tau + 1, \sigma)$, is at least T and therefore $u_{\tau+1-\sigma} = \bar{u}$. The support condition at (σ, τ) is then

$$\begin{aligned} S_\tau - \delta S_{\tau+1-\sigma} &\geq u_\tau - \delta u_{\tau+1-\sigma} + g(\sigma, u_{\tau+1}) \Leftrightarrow \\ v\tau + \frac{\delta v\sigma}{1 - \delta} &\geq \bar{u}(1 - \delta) + g(\sigma, \bar{u}). \end{aligned}$$

At $\sigma = 0$, we have $g(0, \bar{u}) = 0$. Thus, the $(0, \tau)$ condition is $v\tau \geq \bar{u}(1 - \delta)$. Now consider $\sigma = 1, \dots, \tau + 1 - T$. We have two cases. First, if $\bar{u} \leq \frac{v\sigma}{1 - \delta}$, then support at σ becomes

$$\begin{aligned} v\tau + \frac{\delta v\sigma}{1 - \delta} &\geq \bar{u}(1 - \delta) + \frac{v\sigma}{1 - \delta} - \delta\bar{u} \Leftrightarrow \\ v(\tau - \sigma) &\geq (1 - 2\delta)\bar{u}, \end{aligned}$$

which always holds for $\tau \geq \sigma$ and $\delta \geq 1/2$. Second, if $\bar{u} > \frac{v\sigma}{1 - \delta}$, then $\max\{\frac{v\sigma}{1 - \delta}, \bar{u}\} = \bar{u}$ and support at σ becomes

$$\begin{aligned} v\tau + \frac{\delta v\sigma}{1 - \delta} &\geq \bar{u}(1 - \delta) + v\sigma \Leftrightarrow \\ v\tau + v\sigma \left(\frac{\delta}{1 - \delta} - 1 \right) &\geq \bar{u}(1 - \delta). \end{aligned}$$

Since $v\tau \geq \bar{u}(1 - \delta)$, by support at $\sigma = 0$, and $\frac{\delta}{1 - \delta} \geq 1$ for $\delta \geq 1/2$, we have established the second claim.

We can now prove Lemma A1 for $T = 1$. By hypothesis, (3) holds for $\sigma = 0$ and 1 at $T = 1$. Consider $\tau > 1$. By the first claim, (3) holds for $\sigma = 0$ at τ . By the second claim, (3) then holds for $\sigma = 1, \dots, \tau + 1 - T$. Since $\tau + 1 - T = \tau$ in this case, we are done.

From now on take $T \geq 2$. The third claim is that if (3) holds at τ for $\sigma = (\tau + 2 - T)$ to $\sigma = \tau$, then it holds at $\tau + 1$ for $\sigma = (\tau + 2 - T) + 1$ to $\sigma = \tau + 1$. Note that the quality gap

in the continuation state $(\tau + 1, \sigma)$ is less than T exactly when $\sigma \geq \tau + 1 - T$. Hence, we have $u_{\tau+1-\sigma} \in \{u_1, \dots, u_{T-1}\}$, and we are on the rising utility part of the T -stage support. Also, note that $u_{\tau+1} = \bar{u}$.

Condition (3) holds at (σ, τ) if and only if

$$\begin{aligned} S_\tau - \delta S_{\tau+1-\sigma} &\geq u_\tau - \delta u_{\tau+1-\sigma} + g(\sigma, u_{\tau+1}) \Leftrightarrow \\ v\tau + \frac{\delta v\sigma}{1-\delta} &\geq \bar{u} - \delta u_{\tau+1-\sigma} + g(\sigma, \bar{u}), \end{aligned}$$

and at $(\sigma + 1, \tau + 1)$ if and only if

$$v(\tau + 1) + \frac{\delta v(\sigma + 1)}{1-\delta} \geq \bar{u} - \delta u_{\tau+1-\sigma} + g(\sigma + 1, \bar{u}).$$

For the (σ, τ) condition to imply the $(\sigma + 1, \tau + 1)$ condition, it is sufficient to show that

$$\frac{v}{1-\delta} \geq g(\sigma + 1, \bar{u}) - g(\sigma, \bar{u})$$

There are three cases that we need to check.

1) If $\bar{u} \leq \frac{v\sigma}{1-\delta}$, then $g(\sigma + 1, \bar{u}) = \frac{v(\sigma+1)}{1-\delta} - \delta\bar{u}$ and $g(\sigma, \bar{u}) = \frac{v\sigma}{1-\delta} - \delta\bar{u}$, hence the condition holds with equality.

2) If $\frac{v\sigma}{1-\delta} < \bar{u} \leq \frac{v(\sigma+1)}{1-\delta}$, then $g(\sigma + 1, \bar{u}) - g(\sigma, \bar{u}) = \frac{v(\sigma+1)}{1-\delta} - \delta\bar{u} - v\sigma$, which is less than $\frac{v}{1-\delta}$ if and only if $\frac{v\sigma}{1-\delta} < \bar{u}$.

3) If $\frac{v(\sigma+1)}{1-\delta} < \bar{u}$, then $g(\sigma + 1, \bar{u}) - g(\sigma, \bar{u}) = v(\sigma + 1) - v\sigma < \frac{v}{1-\delta}$.

Thus, we have established the third claim.

To complete the proof, consider $\tau = T + 1$ and suppose (3) holds at T for $\sigma = 0, \dots, T$. We know from the third claim that since, (3) holds at T for $\sigma = 2, \dots, T$, it must hold at $T + 1$ for $\sigma = 3, \dots, T + 1$. From the first claim, the support holds at $(0, T + 1)$, since it holds at $(0, T)$. From the second claim, we know the support at $(0, T + 1)$ is sufficient for support at $(\sigma, T + 1)$ for $\sigma = 1, \dots, \tau + 1 - T$. But, $\tau + 1 - T = 2$ in this case and thus (3) holds for all $(\sigma, T + 1)$. The same logic applies between any τ and $\tau + 1$. ■

Lemma A2 Consider a T -stage support. Then for any $\tau \leq T$, we have

$$u_\tau = \frac{1}{\delta^{\tau-1}} \left[u_1 - \frac{v}{1-\delta} \left(\frac{1-\delta^{\tau-1}}{1-\delta} - (\tau-1)\delta^{\tau-1} \right) \right].$$

Proof. It is clearly valid at $\tau = 1$. To verify for $\tau + 1$, assume it holds at τ and we then have

$$\begin{aligned} u_{\tau+1} &= \frac{1}{\delta} [u_\tau - v\tau] = \frac{1}{\delta^\tau} \left[u_1 - \frac{v}{1-\delta} \left(\frac{1-\delta^{\tau-1}}{1-\delta} - (\tau-1)\delta^{\tau-1} \right) \right] - \frac{v\tau}{\delta} \\ &= \frac{1}{\delta^\tau} \left[u_1 - \frac{v}{1-\delta} \left(\frac{1-\delta^{\tau-1}}{1-\delta} - (\tau-1)\delta^{\tau-1} + \tau\delta^{\tau-1}(1-\delta) \right) \right] \\ &= \frac{1}{\delta^\tau} \left[u_1 - \frac{v}{1-\delta} \left(\frac{1-\delta^\tau}{1-\delta} - \tau\delta^\tau \right) \right] \end{aligned} \tag{11}$$

and the formula is valid for $\tau + 1$. ■

Lemma A3 *The T -stage support sequence $(\underline{u}_1, \dots, \underline{u}_T)$ defined by $\underline{u}_1 = (1 - \delta^{T-1})S_1$ satisfies (i) $\underline{u}_\tau = (1 - \delta^{T-\tau})S_1 + \frac{v(\tau-1)}{1-\delta}$, (ii) $\underline{u}_\tau \geq \frac{v\tau}{1-\delta}$ if and only if $\delta \geq \delta^{T-\tau}$ and (iii) $\underline{u}_{T-1} \equiv \frac{v(T-1)}{1-\delta}$ and $\underline{u}_{T-2} \equiv \frac{v(T-1)}{1-\delta} - v$.*

Proof. We have $\underline{u}_1 \equiv \frac{v}{1-\delta} \left[\frac{1-\delta^{T-1}}{1-\delta} \right]$ and $\underline{u}_\tau = v\tau + \delta\underline{u}_{\tau+1}$ for $\tau = 1, \dots, T$. We first show (i). This holds for \underline{u}_1 by construction. Assume that it holds for \underline{u}_τ and consider $\underline{u}_{\tau+1}$. Then

$$\begin{aligned} \underline{u}_{\tau+1} &= \frac{1}{\delta} [\underline{u}_\tau - v\tau] = \frac{1}{\delta} \left[\frac{v}{1-\delta} \left(\frac{1-\delta^{T-\tau}}{1-\delta} + (\tau-1) \right) - v\tau \right] \\ &= \frac{v}{\delta(1-\delta)^2} [\delta + \delta\tau - \delta^{T-\tau} - \delta^2\tau] = \frac{v}{1-\delta} \left[\frac{1-\delta^{T-(\tau+1)}}{1-\delta} + \tau \right]. \end{aligned}$$

Next, we show (ii): $\underline{u}_\tau \geq \frac{v\tau}{1-\delta}$ if and only if $\delta \geq \delta^{T-\tau}$. Using (i) for \underline{u}_τ , we have

$$\begin{aligned} \frac{v}{1-\delta} \left[\frac{1-\delta^{T-\tau}}{1-\delta} + (\tau-1) \right] &\geq \frac{v\tau}{1-\delta} \Leftrightarrow \\ \frac{1-\delta^{T-\tau}}{1-\delta} + (\tau-1) &\geq \tau \Leftrightarrow \delta \geq \delta^{T-\tau}. \end{aligned}$$

For (iii), apply the previous formula for u_τ in terms of u_1 at $\tau = T-1$ and simplify, and then use (5) to find \underline{u}_{T-2} in terms of \underline{u}_{T-1} and simplify. ■

Lemma A4 *Consider a T -stage support with (i) $u_\tau \geq \frac{v\tau}{1-\delta}$ for $\tau = 1, \dots, T-1$ and (ii) $\frac{v(T-1)}{1-\delta} < u_T < \frac{vT}{1-\delta}$. If the support condition (3) holds at (τ, τ) for $\tau = 1, \dots, T$, then the T -stage support satisfies (3) for all (σ, τ) , where $0 \leq \sigma \leq \tau$ and $\tau \geq 1$.*

Proof. Note that we choose T so that (i) and (ii) of Lemma A4 are satisfied. By Lemma A1, it is sufficient to show (3) holds for $\tau \leq T$. Consider $\sigma = 0$ and $\tau = 1, \dots, T$. Since $g(0, u_\tau) = 0$, we see that (3) reduces to $v\tau \geq u_\tau - \delta u_{\tau+1}$ for $\tau < T$. This holds with equality by construction of a T -stage support. At $\tau = T$, (3) reduces to $vT \geq u_T(1-\delta)$ and this holds by (i) of the lemma.

Now consider $1 \leq \sigma \leq \tau$ with $\tau < T$. We claim (3) at (σ, τ) implies (3) at $(\sigma, \tau+1)$. The condition for (σ, τ) is

$$\begin{aligned} S_\tau - \delta S_{\tau+1-\sigma} &\geq u_\tau - \delta u_{\tau+1-\sigma} + g(\sigma, u_{\tau+1}) \Leftrightarrow \\ v(\tau - \sigma) + \frac{\delta v\sigma}{1-\delta} &\geq u_\tau - \delta u_{\tau+1-\sigma}, \end{aligned}$$

since, by (i) and (ii), $u_{\tau+1} > \frac{v\sigma}{1-\delta}$ for any $\sigma \leq \tau < T$. At $(\sigma, \tau+1)$, we need

$$\begin{aligned} S_{\tau+1} - \delta S_{\tau+2-\sigma} &\geq u_{\tau+1} - \delta u_{\tau+2-\sigma} + g(\sigma, u_{\tau+2}) \Leftrightarrow \\ v(\tau - \sigma + 1) + \frac{\delta v\sigma}{1-\delta} &\geq u_{\tau+1} - \delta u_{\tau+2-\sigma}, \end{aligned}$$

since $u_{\tau+2} > \frac{v\sigma}{1-\delta}$ for any $\sigma \leq \tau < T$. By the T -stage support, the $(\sigma, \tau+1)$ condition becomes

$$\begin{aligned} v(\tau - \sigma + 1) + \frac{\delta v\sigma}{1-\delta} &\geq \frac{1}{\delta}(u_\tau - v\tau) - u_{\tau+1-\sigma} + v(\tau+1-\sigma) \Leftrightarrow \\ v\tau + \frac{\delta^2 v\sigma}{1-\delta} &\geq u_\tau - \delta u_{\tau+1-\sigma}, \end{aligned}$$

Thus, we need only show that $v\tau + \frac{\delta^2 v\sigma}{1-\delta} \geq v(\tau - \sigma) + \frac{\delta v\sigma}{1-\delta}$. This clearly holds for any $\delta \in [0, 1]$ and any non-negative σ .

Consequently, (3) at $(1, 1)$ implies (3) at $(1, \tau)$ for $\tau = 2, \dots, T$, (3) at $(2, 2)$ implies (3) at $(2, \tau)$ for $\tau = 3, \dots, T$, and so on up through $(T-1, T-1)$. Thus, (3) at (τ, τ) for $\tau = 1, \dots, T$ is sufficient and the lemma is established. ■

(iii) *Existence of Equilibrium*

Proof of Proposition 2. First, note that conditions (i) and (ii) of Lemma A4 are valid when $(1 - \delta^{T-1})S_1 \leq u_1 \leq (1 - \delta^T)S_1$. This follows by applying Lemma A3 to the reference sequences, \underline{u}_τ , for T and for $T+1$. Then, by Lemma A4, it is sufficient to verify condition (3) at (τ, τ) for $\tau = 1, \dots, T$.

It is immediate that (3) at $\tau = 1$ requires $\delta S_1 \geq u_1$. Hence, we are done if $T = 1$. Now, consider $T \geq 2$ and note that the same observation implies that (3) holds at $(1, 1)$.

Now, consider (3) at (τ, τ) for $\tau \leq T-1$. Then $u_{\tau+1} > \frac{v\tau}{(1-\delta)}$ and we have $g(\tau, u_{\tau+1}) = v\tau$. Thus, the equilibrium support condition (3) becomes

$$\frac{\delta v\tau}{1-\delta} + \delta u_1 \geq u_\tau.$$

We claim that condition (τ, τ) implies condition $(\tau+1, \tau+1)$ for $\tau \leq T-2$. In other words, we claim that $\frac{\delta v\tau}{1-\delta} + \delta u_1 \geq u_\tau$ implies $\frac{\delta v(\tau+1)}{1-\delta} + \delta u_1 \geq u_{\tau+1}$. Recall that $u_{\tau+1} = \frac{1}{\delta}(u_\tau - v\tau)$. So, condition $(\tau+1, \tau+1)$ can be written as

$$\frac{\delta^2 v(\tau+1)}{1-\delta} + \delta^2 u_1 + v\tau \geq u_\tau.$$

Thus, it is sufficient to show that $\frac{\delta^2 v(\tau+1)}{1-\delta} + \delta^2 u_1 + v\tau > \frac{\delta v\tau}{1-\delta} + \delta u_1$. But, this holds if and only if $\delta S_1 + \frac{v\tau}{\delta} > u_1$, which is always the case for $\tau \geq 1$. Thus, (3) holds at $(1, 1)$ and this implies (3) holds at (τ, τ) for $\tau = 2, \dots, T-1$.

We are then left with the (T, T) condition, which reduces to

$$\delta u_1 \geq (1-\delta)u_T,$$

since $g(T, u_{T+1}) = \frac{vT}{1-\delta} - \delta u_T$ by Lemma A3. We know that the condition holds at $(T-1, T-1)$ and we have

$$\begin{aligned} \frac{\delta v(T-1)}{1-\delta} + \delta u_1 &\geq u_{T-1} = v(T-1) + \delta u_T \Leftrightarrow \\ \frac{1}{\delta} \left[\frac{\delta v(T-1)}{1-\delta} - (T-1)v + \delta u_1 \right] &\geq u_T. \end{aligned}$$

Thus, it is sufficient for (T, T) to show that

$$\frac{\delta}{1-\delta} u_1 \geq \frac{1}{\delta} \left[\frac{\delta v(T-1)}{1-\delta} - (T-1)v + \delta u_1 \right]. \quad (12)$$

Simplifying and noting that $\delta \geq 1/2$, condition (12) holds if and only if $\delta u_1 \geq v(T-1)$.

From $u_1 \geq (1 - \delta^{T-1})S_1$, it is sufficient to show that $\delta(1 - \delta^{T-1})S_1 \geq v(T-1)$. At $T = 2$ this reduces to $\delta \geq 1/2$. Now, we carry out an induction: assume it holds for T if $\delta + \delta^{T-1} > 1$ and

show that it holds for $T + 1$ if $\delta + \delta^T > 1$. So, we must show that $\delta(1 - \delta^T)S_1 \geq vT$ or, equivalently, that

$$\delta(1 + \dots + \delta^{T-1}) > T(1 - \delta).$$

The condition at T is $\delta(1 - \delta^{T-1})S_1 \geq v(T - 1)$, which holds if and only if

$$(1 - \delta) + \delta(1 + \dots + \delta^{T-2}) > T(1 - \delta).$$

But,

$$\delta(1 + \dots + \delta^{T-1}) > (1 - \delta) + \delta(1 + \dots + \delta^{T-2}) \Leftrightarrow \delta + \delta^T > 1,$$

which establishes the induction. Thus, the (T, T) condition holds, and we have therefore shown that the T -stage support satisfies (3) for all (σ, τ) , where $0 \leq \sigma \leq \tau$ and $\tau \geq 1$.

To see that every buyer payoff, $u_1 \in [0, \delta S_1]$ can be supported in this way, simply note that each $\delta \in [1/2, 1]$ lies in exactly one of the δ_τ cutoff sequence intervals. With $\delta \in [\delta_\tau, \delta_{\tau+1})$, we then see that every $u_1 \in [0, \delta S_1]$ lies in exactly one of the $[(1 - \delta^{T-1})S_1, (1 - \delta^T)S_1]$ intervals, where T ranges from 1 up to the index on the δ_τ root. ■

Proof of Proposition 3. By construction of an efficient equilibrium, for $u_1 \in [0, \delta S_1]$ we have a payoff of $\pi_1 = \frac{p_1}{1-\delta}$ for the seller and $u_1 = \frac{1}{1-\delta} \left[\frac{v}{1-\delta} - p_1 \right]$ for the buyers. It is sufficient for an equilibrium to show that the support condition

$$S_\tau - \delta S_{\tau+1-\sigma} \geq u_\tau - \delta u_{\tau+1-\sigma} + g(\sigma, u_{\tau+1})$$

holds for $\sigma = 0, 1, \dots, \tau$ at each $\tau \geq 1$. To verify that the support condition holds, we note three useful facts: (i) $\bar{u} \in \left[0, \frac{v}{1-\delta}\right]$, as follows from the range for u_1 , (ii) $g(0, \bar{u}) = 0$ and $g(\sigma, \bar{u}) = \frac{v\sigma}{1-\delta} - \delta\bar{u}$ for $\sigma \geq 1$, also by the range for u_1 (iii) $S_\tau - \delta S_{\tau+1-\sigma} = v\tau + \frac{\delta v\sigma}{1-\delta}$, by definition of surplus.

Begin with $\sigma = 0$. For $\tau = 1$, the support condition reduces to $v \geq \delta u_1$. But $\delta < 1/2$ implies $v > \delta^2 S_1$, so we are done as $u_1 \leq \delta S_1$. For $\tau > 1$, the support condition reduces to $v\tau \geq (1 - \delta)\bar{u}$. It is sufficient to show this holds at $\tau = 2$, which reduces to $\frac{2v}{1-\delta} \geq \bar{u}$, and we are done by the upper bound on \bar{u} .

Now consider $\sigma = 1, \dots, \tau - 1$ and $\tau > 1$. The condition reduces to $v(\tau - \sigma) \geq (1 - 2\delta)\bar{u}$. As $\bar{u} \leq \frac{v}{1-\delta}$, we need only show that at $\tau = 2$ we have $1 \geq \frac{1-2\delta}{1-\delta}$. This reduces to $\delta > 0$.

Finally, the case of $\sigma = \tau$ corresponds to the equilibrium path and continuation outcomes. The condition reduces, by construction, to

$$\pi_\tau = S_\tau - u_\tau \geq \delta(S_1 - u_1) + g(\tau, \bar{u}) = \delta\pi_1 + \frac{v\tau}{1-\delta} - \delta\bar{u}.$$

Since $\pi_1 = \frac{p_1}{1-\delta}$, $\pi_\tau = p_\tau + \delta\pi_1$, $\bar{u} = \frac{v\tau}{1-\delta} - p_\tau + \delta u_1$ by construction of the candidate equilibrium, this reduces to $\delta u_1 \geq (1 - \delta)\bar{u}$ and this holds with equality. ■

11 Appendix C - Inefficient Equilibria.

This Appendix has four parts. We begin with infrequent innovation and prove Proposition 4. Next, we prove results for the approach conditions. We then provide the analysis of the off equilibrium support. Finally, we prove Proposition 5 to show existence of equilibria for frequent innovation.

(i) *Infrequent Innovation* ($\delta \leq 1/2$).

Proof of Proposition 4. We first show that $\frac{v(t-1)}{1-\delta} > u_t$ is necessary in any delay equilibrium when $\delta \leq 1/2$. By equilibrium construction, a first sale in state $(t, 0)$ of t units implies $\pi_1 = \delta^{t-1}\pi_t$ and

$$\pi_t + u_t = \frac{vt}{(1-\delta)(1-\delta^t)}.$$

By flow dominance, we have $\pi_1 \geq \frac{v}{1-\delta}$ and, hence, $\pi_t \geq \frac{v}{\delta^{t-1}(1-\delta)}$. Combining with the joint payoff expression from above, we have

$$u_t \leq \frac{vt}{(1-\delta)(1-\delta^t)} - \frac{v}{\delta^{t-1}(1-\delta)}.$$

It is then sufficient to show that $\frac{v(t-1)}{1-\delta}$ exceeds the right-hand-side in the above expression. Simplifying the resulting inequality, this holds if

$$1 > \delta^{t-1} [1 + \delta + t\delta^t - \delta^t].$$

At $t = 2$ this reduces to $1 > \delta [1 + \delta + \delta^2]$. It holds at $\delta = 1/2$ since $1 > 7/8$ and, hence, for all smaller δ . By induction, it holds for all $t > 2$ if

$$\begin{aligned} \delta^{t-1} [1 + \delta + t\delta^t - \delta^t] &> \delta^t [1 + \delta + (t+1)\delta^{t+1} - \delta^{t+1}] \Leftrightarrow \\ 1 &> \delta^2 + \delta^t + t\delta^t(\delta^2 - 1). \end{aligned}$$

But $\delta^2 + \delta^t < 1$ for $\delta \leq 1/2$ and the second term is negative. Therefore, we have $\frac{v(t-1)}{1-\delta} > u_t$.

Now suppose the seller offers $t-1$ units for a price of $\hat{p} = \frac{v(t-1)}{1-\delta} - \delta u_t - \epsilon$ in state $(t-1, 0)$. It cannot be that all buyers reject this offer since

$$\begin{aligned} v(t-1) - \hat{p} + \delta \max \left\{ \frac{v(t-1)}{1-\delta}, u_t \right\} &> \delta u_t \Leftrightarrow \\ \frac{v(t-1)}{1-\delta} - \delta u_t &> \hat{p}. \end{aligned}$$

Thus, in a symmetric equilibrium, it must be that all buyers accept the offer and, hence, the continuation state is $(t, t-1)$. The payoff to the seller is then $\hat{\pi} = \hat{p} + \delta\pi_1$. This exceeds $\delta\pi_t$ if

$$\frac{v(t-1)}{1-\delta} + \delta\pi_1 > \delta(\pi_t + u_t).$$

Substituting for the equilibrium joint surplus and using $\pi_1 \geq \frac{v}{1-\delta}$, by flow dominance, it is sufficient to show

$$\begin{aligned} \frac{v(t-1)}{1-\delta} + \delta \frac{v}{1-\delta} &> \frac{\delta vt}{(1-\delta)(1-\delta^t)} \Leftrightarrow \\ t(1-\delta) &> (1-\delta)(1-\delta^t) + t\delta^t. \end{aligned} \tag{13}$$

At $t = 2$, (13) reduces to $1 > \delta(1 + \delta + \delta^2)$, which is valid for all $\delta \leq 1/2$. To show that an induction argument holds, note that (13) at $t+1$ is

$$\begin{aligned} (t+1)(1-\delta) &> (1-\delta)(1-\delta^{t+1}) + (t+1)\delta^{t+1} \Leftrightarrow \\ t(1-\delta) &> \delta^{t+2} + t\delta^{t+1}. \end{aligned}$$

So, to show that the induction argument is valid, and we need only show

$$\begin{aligned} (1 - \delta)(1 - \delta^t) + t\delta^t &> \delta^{t+2} + t\delta^{t+1} \Leftrightarrow \\ (1 - \delta - \delta^t) + \delta^{t+1}(1 - \delta) + t\delta^t(1 - \delta) &> 0. \end{aligned}$$

The first term is positive for $\delta \leq 1/2$ and the other two terms are clearly positive. Thus, the deviation is profitable for the seller. ■

(ii) *Approach Conditions.*

Proof of Lemma 5. Beginning with the seller approach conditions, (8), note that $\delta^{t-\tau}$ is strictly increasing in τ . Thus, we can set $p(\sigma, \tau) = p(\sigma, \sigma)$ and the condition holds for $\sigma < \tau$. Now, for the buyer approach conditions, (7), note that the right hand side is increasing in τ . We claim that the left hand side is decreasing in τ . When $\frac{v\sigma}{(1-\delta)} \geq u_t$, then the left hand side is $\frac{v\sigma}{(1-\delta)} - \delta^{t-\tau}u_t$, which is falling in τ . If $\frac{v\sigma}{(1-\delta)} < u_t$, then the left hand side is $\frac{v\sigma(1-\delta^{t-\tau})}{(1-\delta)}$, which is also strictly decreasing in τ . Again setting $p(\sigma, \tau) = p(\sigma, \sigma)$ for $\sigma < \tau$ is sufficient. ■

Proof of Lemma 6. Note that when $u_t = 0$, for any deviation offer by the seller in period τ for σ units, condition (7) immediately reduces to $p(\sigma, \tau) = \frac{v\sigma}{(1-\delta)}$. Thus, take $\sigma = \tau = t - 1$ and consider condition (8). Since $\pi_t = \Psi_t$, when $u_t = 0$, we have

$$\delta(1 - \delta^{t-1})\Psi_t \geq \frac{v(t-1)}{(1-\delta)} \Leftrightarrow \frac{1 - \delta^t}{1 - \delta} \geq t$$

after simplification. Clearly, this is always false for $t \geq 2$ and any δ . ■

Proof of Lemma 7. In (10), τ assumes integer values $1, \dots, t-1$. Let us replace τ with a continuous variable, x , that assumes values in the interval $[0, t]$. This greatly simplifies the derivation of the sufficiency condition. It is useful to define three functions:

$$A(x, u, \delta, t) \equiv (\delta^{t-x} - \delta^t) \left[\frac{vt}{(1-\delta)(1-\delta^t)} - u \right]$$

$$B(x, \delta, t) \equiv \frac{vx}{1-\delta}(1 - \delta^{t-x})$$

$$C(x, u, \delta, t) \equiv \frac{vx}{1-\delta} - \delta^{t-x}u,$$

where $u \equiv u_t$. For (δ, t) , consider $u \in [0, \frac{vt}{1-\delta}]$; we treat the case of $u > \frac{vt}{1-\delta}$ later in the proof. In terms of x , condition (10) becomes

$$\begin{aligned} A(x, u, \delta, t) &\geq B(x, \delta, t) \quad \text{for } 0 \leq x \leq \frac{(1-\delta)u}{v}, \\ A(x, u, \delta, t) &\geq C(x, u, \delta, t) \quad \text{for } \frac{(1-\delta)u}{v} < x \leq t. \end{aligned}$$

First, consider when $A(x, u, \delta, t) \geq B(x, \delta, t)$ for all x in the interval $\left[0, \frac{(1-\delta)u}{v}\right]$. Suppressing arguments, note that A is increasing and convex in x , and equals 0 when $x = 0$, while B is strictly concave in x and equals 0 when $x = 0$. Thus, if we have $\frac{\partial A}{\partial x} \geq \frac{\partial B}{\partial x}$ at $x = 0$, then $A \geq B$ must hold for all positive x . Calculating the partial derivatives, this yields the condition

$$\bar{u}^A \equiv \frac{v}{(1-\delta)} \left[\frac{t}{1-\delta^t} - \frac{1-\delta^t}{(-\ln \delta)\delta^t} \right] \geq u. \quad (14)$$




Figure 3: Payoff Lower Bound for Approach Conditions

Next, consider when $A(x, u, \delta, t) \geq C(x, u, \delta, t)$ for all x in the interval $[\frac{(1-\delta)u}{v}, t]$. This condition simplifies to

$$\frac{\delta^t(1-\delta)u}{v} \geq x - \frac{t(\delta^{t-x} - \delta^t)}{1-\delta^t} \equiv h(x, t, \delta). \quad (15)$$

It is easy to show $h(x, t, \delta)$ is strictly concave and equals 0 at $x = 0$ and $x = t$. Thus, $h(x, t, \delta)$ has a unique interior maximum at some $x^*(\delta, t)$ which is implicitly defined by $\delta^{x^*}(1-\delta^t) = -\delta^t \ln \delta^t$. Note that when condition (14) holds, the term in brackets must be positive; this implies that $\delta^t > \frac{\partial h}{\partial x}$ at $x = 0$. Refer to Figure 3.

Define \underline{u}^A by

$$\begin{aligned} \underline{u}^A &\equiv \frac{v}{\delta^t(1-\delta)} h(x^*, t, \delta) \\ &= \frac{v}{\delta^t(1-\delta)} \left[x^* - \frac{t\delta^{t-x^*} - t\delta^t}{1-\delta^t} \right] \\ &= \Psi_t + \frac{v(1-a(\delta^t))}{\delta^t(1-\delta) \ln \delta}, \end{aligned}$$

where the function $a(\delta^t)$ is as defined in the text. Then any $u \geq \underline{u}^A$ will necessarily satisfy (15). See Figure 3.

In order to find when (14) and (15) hold at a candidate u , we compare the bounds to obtain,

$$\bar{u}^A \geq \underline{u}^A \Leftrightarrow \delta^t \geq a(\delta^t),$$

It is easy to show that $a(d)$ is strictly decreasing for $d \in (0, 1)$ and, by L'Hospital's Rule, that $\lim_{d \rightarrow 1} a(d) = 0$. Then, the equation $d = a(d)$ has a unique root, d^* in $(0, 1)$. Thus, we have established that $\delta^t \geq d^*$ implies the interval $(\underline{u}^A, \bar{u}^A)$ is non-empty.

The last step of the proof is to consider the range of values for u that can be supported. A straightforward argument establishes that $\underline{u}^A < \frac{vt}{1-\delta}$ holds for $\delta^t \geq d^*$. The comparison of \bar{u}^A with $\frac{vt}{1-\delta}$ reveals that \bar{u}^A crosses $\frac{vt}{1-\delta}$ exactly one time, from below, at the root of the equation $-\ln \delta^t = (\delta^{-t} - 1)^2$; numerically, the root is .572, which exceeds d^* . Thus, for δ^t below this root, we have $\bar{u}^A < \frac{vt}{1-\delta}$ and any $u \in (\underline{u}^A, \bar{u}^A)$ satisfies the approach condition (10). For δ^t above this root, we have $\bar{u}^A > \frac{vt}{1-\delta} > \underline{u}^A$ and there are two cases. First, by the above analysis, any $u \in (\underline{u}^A, \frac{vt}{1-\delta})$ satisfies the approach condition (10). Second, for the case of $u \in [\frac{vt}{1-\delta}, \bar{u}^A)$, the approach condition (10) requires that $A(x, u, \delta, t) \geq B(x, \delta, t)$ for all x in the interval $[0, t]$; note that since u is large, the case of $\frac{(1-\delta)u}{v} < x \leq t$ never arises. Clearly, we have $A(x, u, \delta, t) \geq B(x, \delta, t)$ and utilities in this payoff range satisfy (10). ■

(iii) *Off Equilibrium Support (Cash-in offers)*

Now, we deal with off equilibrium states $(\tau, 0)$ where $\tau > t$. Since we are using cash-in supports, the payoffs are $\pi_\tau = p_\tau + \delta\pi_1$ for the seller and $u_\tau = v\tau - p_\tau + \delta u(\tau + 1, \tau) = \frac{v\tau}{1-\delta} - p_\tau + \delta u_1$ for buyers. Note that from $(\tau, 0)$ the surplus on the continuation path is

$$\Psi_\tau \equiv \frac{v\tau}{1-\delta} + \delta\Psi_1 = \pi_\tau + u_\tau \text{ for } \tau > t$$

and that cash-in states contrast with delay states, $\tau < t$, where we have $\Psi_\tau \equiv \delta^{t-\tau}\Psi_t$.

Recall from the text that we must satisfy (1) and (2). Thus, we seek a utility sequence that satisfies the analog of (3), as given by

$$\Psi_\tau - \delta\Psi_{\tau+1-\sigma} \geq u_\tau - \delta u_{\tau+1-\sigma} + g(\sigma, u_{\tau+1}) \quad (16)$$

for $\tau \geq t$ and $\sigma = 0, \dots, \tau$.

Analogous to our support utility sequence for efficient equilibria, we define a T -stage support sequence, where $T \geq t + 1$, for inefficient equilibria by

$$u_\tau = \delta u_{\tau+1} + \Psi_\tau - \delta\Psi_{\tau+1} \text{ for } \tau = t, \dots, T-1, \quad (17)$$

and

$$u_\tau = u_T \text{ for } \tau \geq T.$$

For the special case of $T = t$, we specify a constant sequence $u_\tau = u_t$ for all τ . As before, the seller is indifferent with respect to delay and the cash-in up to period T , and the seller prefers to cash-in in period T provided that the buyers' utilities, u_T , are not too large.

We can satisfy (16) by an appropriate choice of the support length T for a given δ and buyer payoff. Recalling Figure 1 for efficient equilibria, we find a similar structure for the relationship between u_t , δ , and the length of a T . For any given δ we can support successively higher buyer payoffs by increasing the support horizon, T . In particular as u_t rises, then T must also rise until we achieve $u_t \in [\Psi_t - \delta^{T-t}S_1, \Psi_t - \delta^{T-t+1}S_1]$. Thus, we demonstrate

Lemma A5 *Let $T \geq t$ and suppose $\delta + \delta^T \geq 1$. Let buyer utilities follow a T -stage support*

sequence of horizon T , where u_t satisfies (i) $0 \leq u_t \leq \Psi_t - \delta S_1$ for $T = t$ and (ii) $\Psi_t - \delta^{T-t} S_1 \leq u_t \leq \Psi_t - \delta^{T-t+1} S_1$ for $T > t$. Then (16) holds for every $\tau \geq t$ and $\sigma = 0, \dots, \tau$.

Proof. We only sketch the arguments, since they are similar to those of the efficient case. First, we note that if expression (16) holds when buyer utility first becomes constant at $\tau = T$ and for all $\sigma \leq T$, then it holds for all $\tau > T$ for all $\sigma \leq \tau$. This follows from arguments very similar to those used in Lemma A1. We must, however, replace S_τ with Ψ_τ . Also, we need to take into account a new range of large σ , $\sigma \geq T + 1 - t$, where the continuation state involves delay. Thus, it only remains to verify the support conditions for $\tau = t, \dots, T$.

Second, it is straightforward to show that every partial cash-in constraint, where $0 < \sigma < \tau$, is implied by the cash-in constraint $\sigma = \tau$ and the delay constraint $\sigma = 0$ for each $\tau = t, \dots, T - 1$. Recall that delay constraints for $\tau = t, \dots, T - 1$ hold by construction of the T -stage support. The delay constraint, (16) for $\tau = T$ and $\sigma = 0$, reduces to

$$\frac{\delta v}{1 - \delta} \leq (1 - \delta)(\Psi_T - u_T).$$

This holds if $u_T \leq \Psi_T - \delta S_1$, which is implied by the assumed range for u_t . Thus, only the cash-in constraints $\tau = t, \dots, T$ remain to be checked. Also, (16) is easy to verify for the special case of $t = T$. Henceforth, we take $T > t$.

We begin with the cash-in constraint at $\tau = T$. The condition for $\tau = T$ and $\sigma = T$ is

$$g(T, u_T) \leq (1 - \delta^T)(\Psi_T - u_T). \quad (18)$$

By the range for u_t and the T -stage construction, we find that u_T must be less than $\frac{vT}{1-\delta}$. Thus, (18) becomes

$$(1 - \delta - \delta^T)u_T \leq (1 - \delta^T)\Psi_t - \frac{vT}{1 - \delta}.$$

This holds since the l.h.s. is negative, by the lemma hypothesis, and, via calculations, the r.h.s. is positive. Thus, we have verified that condition (16) holds for $\tau = T$ and $\sigma = T$.

Verifying (16) is more subtle when $\tau < T$. We need the following two claims for the relationship between u_τ and $\frac{v(\tau-1)}{1-\delta}$ for $\tau = t, \dots, T$ (the proofs involve only routine algebra with the T -stage utility path and are hence omitted). The first claim is

Claim 1. (i) If $u_\tau < \frac{v(\tau-1)}{1-\delta}$, then $u_{\tau+1} < \frac{v\tau}{1-\delta}$ for all $\tau = t + 1, \dots, T$; (ii) if $u_{\tau+1} \geq \frac{v\tau}{1-\delta}$, then $u_\tau \geq \frac{v(\tau-1)}{1-\delta}$ for all $\tau = t, \dots, T$.

This claim provides a simple algorithm for verifying the interior support constraints. Recall that $u_T < \frac{vT}{1-\delta}$. If $u_T > \frac{v(T-1)}{1-\delta}$, then Claim 1 implies $u_\tau > \frac{v(\tau-1)}{1-\delta}$ for all $\tau \leq T - 1$. If $u_T < \frac{v(T-1)}{1-\delta}$, however, then we move to $\tau = T - 2$ and compare u_{T-1} with $\frac{v(T-2)}{1-\delta}$. If u_{T-1} is greater, then $u_\tau > \frac{v(\tau-1)}{1-\delta}$ for all $\tau \leq T - 2$. If not, continue down until either $u_\tau < \frac{v(\tau-1)}{1-\delta}$ or $\tau = t$. If at every τ , it is always the case that $u_\tau < \frac{v(\tau-1)}{1-\delta}$, then we are in the "small" utility case. Otherwise, there always exists a $\hat{\tau}$ such that if $\tau \geq \hat{\tau}$, then $u_{\tau+1} < \frac{v\tau}{1-\delta}$ and if $\tau < \hat{\tau}$, then $u_{\tau+1} > \frac{v\tau}{1-\delta}$. Since $u_T < \frac{vT}{1-\delta}$ there is always a $\hat{\tau} \geq t$ such that $u_\tau < \frac{v(\tau-1)}{1-\delta}$.

Claim 2. (i) Suppose $u_\tau < \frac{v(\tau-1)}{1-\delta}$, then if (16) holds at τ it holds at $\tau - 1$ for any $\tau \leq T$. (ii) Suppose $u_\tau > \frac{v(\tau-1)}{1-\delta}$, then if (16) holds at $\tau - 1$ it holds at τ for any $\tau \leq T$.

We now employ these results to complete the proof of Lemma A5. We need to show (16) holds at $\tau = t, \dots, T - 1$. There are two cases. First, suppose $u_{t+1} > \frac{vb}{1-\delta}$. Then it is straightforward to show that (16) holds at $\tau = t$. Now apply (ii) in Claim 1 at $\tau = t$ and (ii) in Claim 2 to see that (16) holds at $\tau = t + 1, \dots, T - 1$. The second case is $u_{t+1} \leq \frac{vb}{1-\delta}$. From before, we know (16) holds at $\tau = T$. We can now apply the (i) in Claim 1 at $\tau = T$ and (i) in Claim 2 to see that (16) holds at $\tau = T - 1, \dots, t$. This completes the proof. ■

(iv) *Equilibrium Existence*

We have the following lemma that provides sufficient conditions for equilibrium existence.

Lemma A6 *Suppose a sequence of buyer utilities u_τ satisfies (10) and (16) for some $t \geq 2$ and for all σ and τ , such that $0 \leq \sigma \leq \tau$ for any $\tau \geq 1$. Then there exists an inefficient $t - cycle$ equilibrium with supporting prices $p(\sigma, \tau)$.*

Proof. The approach conditions are satisfied by Lemma 7 for $u_t \in (\underline{u}^A, \bar{u}^A)$. With (16) satisfied as well, we know that no single-bundle deviation is profitable for the seller. As with the efficient case, we must now show that the seller cannot profitably deviate to an upgrade offer set in state $(\tau, 0)$. While the logic is similar to that for efficient equilibria, the proof is more complicated and we only provide a brief sketch. We must distinguish between approach states, where $\tau < t$, and support states where $\tau \geq t$. We define a *BCE* based on the same buyer preference relation \succeq_B but now $u(\tau + 1, b)$ and $u(\tau + 1, b | i)$ must account for whether there is a delay or a cash-in the continuation state. We can then show the same argmax construction yields a *BCE*. Next, note that any upgrade offer can be written in the form $\mathcal{B} = \mathcal{I} \cup \mathcal{J} \cup \mathcal{K}$ where an upgrade to b is in \mathcal{I} if $p_b \leq G_b$, in \mathcal{J} if $p_b \geq H_b$, and in \mathcal{K} if $G_b < p_b < H_b$. The threshold prices, following the buyer and seller decisions in the text at (16), are specified by

$$G_b = \begin{cases} \frac{vb}{1-\delta} (1 - \delta^{t-\tau}) + \delta^{t-\tau} \max \left\{ \frac{vb}{1-\delta}, u_t \right\} - \delta^{t-\tau} u_t & \text{if } \tau < t \\ vb + \delta \max \left\{ \frac{vb}{1-\delta}, u_{\tau+1} \right\} - \delta u_{\tau+1} & \text{if } \tau \geq t \end{cases}$$

and

$$H_b = \begin{cases} \delta^{t-\tau} (1 - \delta^b) (\psi_t - u_t) & \text{if } \tau < t \\ (\psi_\tau - u_\tau) - \delta (\psi_{\tau+1-b} - u_{\tau+1-b}) & \text{if } \tau \geq t \end{cases} .$$

The proof then proceeds by considering the various cases for the decomposition $\mathcal{B} = \mathcal{I} \cup \mathcal{J} \cup \mathcal{K}$. For instance, if \mathcal{J} is empty then we are done as the argmax property can be applied to $\mathcal{I} \cup \mathcal{K}$ and we can specify that buyers choose this *BCE*, which is priced below H_b . The most difficult case arises when none of the component sets in $\mathcal{I} \cup \mathcal{J} \cup \mathcal{K}$ is empty. We can show that there always exists a *BCE* priced below H_b in this case.

The next step is to specify continuation equilibrium strategies at asymmetric buyer states, (τ, α) ; see part (i) of Appendix B for the definition. These states are complicated by the fact that joint surplus is not maximized in an inefficient $t - cycle$ equilibrium. This contrasts with the efficient case, where the (τ, α) offers fully extracted buyers relative to equilibrium payoffs for the continuation state of $(\tau + 1, \tau)$, since feasibility will no longer imply this is an optimal offer for the seller. Instead, we specify a continuation equilibrium that leaves buyers asymmetric. Thus, this part of the support neither returns to the equilibrium path nor to a cash-in support state.

Let $\alpha_m \equiv \inf \{ \alpha_\sigma \mid \alpha_\sigma > 0, \}$ denote the smallest non-zero mass of buyers in α and if the inf is not unique let m be the largest index (e.g., if $\alpha_\sigma = 1/\tau$ for all σ then $\alpha_m = 1/\tau$ and $m = \tau - 1$).

We specify that the seller offers a collection of upgrades $\{(b_\sigma, p_\sigma)\}_{\sigma=0, \dots, \tau-1}$ with prices of

$$p_m = \frac{v(\tau - 1 - m)}{1 - \delta}, \quad \text{and} \quad p_\sigma = \frac{v(\tau - \sigma)}{1 - \delta} \quad \text{for } \sigma \neq m.$$

It is straightforward to verify that all buyers optimally accept and that the continuation state is of the form $(\tau + 1, \alpha')$ where $\alpha' = (0, \dots, 0, \alpha_m, 1 - \alpha_m)$. Thus, in period $\tau + 1$ the smallest mass group is at quality level $\tau - 1$ units, while all other buyers are at τ . We then apply the same logic to specify offers for $(\tau + 1, \alpha')$ and all resulting states so that the distribution of buyers is always concentrated in the same proportions $(\alpha_m, 1 - \alpha_m)$ at the two highest quality levels.

The argument that this is an optimal strategy for the seller is involved, so we only offer a brief sketch. For any arbitrary seller offer in (τ, α) and pattern of buyer choices, one can bound the seller payoff via individual rationality of the buyer choices, using the hypothesis that the seller makes the offer in the continuation state in period $\tau + 1$ that induces a concentration of buyers at the two highest quality levels (if the buyer choices lead to a degenerate distribution, then payoffs must instead follow the $t - cycle$ equilibrium path and cash-in support). Then, since the repeated concentration of buyers in proportions $(\alpha_m, 1 - \alpha_m)$ at the two highest quality levels allows the seller to extract all surplus generated on this path, the associated seller payoff can be shown to be at least as large as the bound derived from individual buyer rationality for any other seller offer. In this regard, note that every equilibrium with delay ($t \geq 2$) has a buyer quality level that is strictly below the mean for the $(\alpha_m, 1 - \alpha_m)$ concentration.

Finally, finding a *BCE* for an arbitrary upgrade offer \mathcal{B} in state (τ, α) follows the same basic approach as that for an efficient equilibrium. The only difference is that we must take care to account for any possible delay if the continuation state is degenerate. ■

Proof of Proposition 5. Define $B(t, T, \delta) \equiv \Psi_t - \delta^{T-t} S_1$, the upper bound of the payoffs that can be supported with a $T - stage$ support. We will show that $B(t, T, \delta) \geq \bar{u}^A(t, \delta)$, whenever δ satisfies the approach conditions, thus proving the proposition. Comparing, $B(t, T, \delta) \geq \bar{u}^A(t, \delta) \Leftrightarrow$

$$\frac{(1 - \delta^t)(1 - \delta)}{\delta^T} \geq -\ln \delta \tag{19}$$

It is clear that both sides of (19) are positive, falling in δ and equal to 0 at $\delta = 1$. The l.h.s. of (19) is rising in T and t . Differentiating both sides of (19) with respect to δ we find that the slope of the l.h.s. is greater/less than the slope of the r.h.s. if and only if

$$\delta^T + \delta^{t+1} \geq T(1 - \delta) + \delta - (T - t)\delta^t(1 - \delta) \tag{20}$$

In (20), the l.h.s. is positive and convex in δ and the r.h.s. is decreasing in δ . Thus, the l.h.s. and the r.h.s. of (19) cross a single time in δ . It easy to see that (19) is negative in a neighborhood of 1 and positive at $\delta = \sqrt[t]{d^*}$ for any $T \geq t$. So, there is a $b(t, T) \in (\sqrt[t]{d^*}, 1)$ such that (19) holds whenever $\delta \leq b(t, T)$. Clearly, $b(t, T)$ is increasing in T .

Now, we will show that $b(t, T) > \delta_T$ for all $t \geq 2$. At $b = b(t, T)$ we have

$$\frac{1 - b}{b^T} = \frac{-\ln b}{1 - b^t}$$

The r.h.s. is falling in b . Since for any b this is less than 1, $b(t, T) > \delta_T$, the result is proved. ■

References

- [1] Ausubel, L. and R. Deneckere (1989), "Reputation in Bargaining and Durable Goods Monopoly," *Econometrica*, 57, 511-531.
- [2] Bulow, J. (1982), "Durable Goods Monopolists," *Journal of Political Economy*, 90, 314-332.
- [3] Coase, R. (1972), "Durability and Monopoly," *Journal of Law and Economics*, 15, 143-149.
- [4] Ellison, G. and D. Fudenberg (2000), "The Neo-Luddite's Lament: Excessive Upgrades in the Software Industry," *RAND Journal of Economics*, 31, 253-272.
- [5] Fehr, N. and K. Kuhn (1995), "Coase versus Pacman: Who Eats Whom in the Durable-Goods Monopoly?" *Journal of Political Economy*, 103, 785-812.
- [6] Fishman, A. and R. Rob (2000), "Product Innovation by a Durable Goods Monopoly," *RAND Journal of Economics*, 31, 237-252.
- [7] Fudenberg, D. and J. Tirole (1991), *Game Theory*, Cambridge: MIT Press.
- [8] Fudenberg, D. and J. Tirole (1998), "Upgrades, Trade-ins and Buybacks," *RAND Journal of Economics*, 29, 235-258.
- [9] Fudenberg, D., D. Levine, and J. Tirole (1985), "Infinite Horizon Models of Bargaining with One-Sided Incomplete Information," in *Game Theoretic Models of Bargaining*, Ed. by A. Roth, Cambridge: Cambridge University Press.
- [10] Gul, F., H. Sonnenschein, and R. Wilson (1986), "Foundations of Dynamic Monopoly and the Coase Conjecture," *Journal of Economic Theory*, 39, 155-190.
- [11] Maskin, E. and J. Tirole (2001), "Markov Perfect Equilibrium I: Observable Actions," *Journal of Economic Theory*, 100, 191-219.
- [12] Sobel, J. (1991), "Durable Goods Monopoly with Entry of New Consumers," *Econometrica*, 57, 1455-1485.
- [13] Stokey, N. (1981), "Rational Expectations and Durable Goods Pricing," *Bell Journal of Economics*, 12, 112-128.
- [14] Waldman, M. (1996), "Planned Obsolescence and the R&D Decision," *RAND Journal of Economics*, 27, 583-595.
- [15] Waldman, M. (2003), "Durable Goods Theory for Real World Markets," *Journal of Economic Perspectives*, 17, 131-154.