Quality, Upgrades and Equilibrium in a Dynamic Monopoly Market

James J. Anton and Gary Biglaiser*

March 1, 2011 Revised

Abstract

We examine an infinite horizon model of quality growth for a durable goods monopoly. Quality improvements may be sold in any desired bundles. Consumers are identical and for a quality improvement to have value the buyer must possess previous qualities: goods are upgrades. Subgame perfect seller payoffs range from capturing the full social surplus down to only the initial flow value of each good. For any discount factor, each of these payoffs is realized in a Markov perfect equilibrium that follows the socially efficient path. However, inefficient delay equilibria, with bundling, exist for innovation rates above a threshold.

JEL: C72, C73, D42, L15

Keywords: Upgrades, durable goods, monopoly, market power, coordination

We examine the commercialization process - pricing and adoption - of an upgrade good in a dynamic monopoly market. Prominent examples are provided by technology markets, such as those for software, where cycles of upgrades to existing products have become the norm.[1] Ongoing innovation implies that buyers face a sequence of purchasing decisions.

---

*Anton: Fuqua School of Business, Duke University, Durham NC 27705, james.anton@duke.edu; Biglaiser: Department of Economics, University of North Carolina, Chapel Hill NC 27599, 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 the associate editor and referees, as well as Jacques Cremer, Leslie Marx, Larry Samuelson, Jean Tirole, Mike Waldman, Dennis Yao, many colleagues at conferences and seminars, and, especially, Joel Sobel for many helpful conversations. The views in this work are solely are own.

[1] Quality that improves over time is a significant feature of many durable goods markets, as emphasized by Waldman (2003). In addition to software, upgrades to cellular networks often allow vendors to offer,

Thus, rather than timing a single purchase and then exiting the market, buyers have an incentive to return to the market and 'upgrade' to a higher quality. 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.

The Microsoft antitrust cases highlight a fundamental question regarding prices in an upgrade market. Fudenberg and Tirole (2000) observed that the expert witnesses all appeared to agree that Microsoft was pricing the Windows operating system well below the static monopoly price. There was, however, wide disagreement as to why. Prominent arguments included network formation with low prices spurring adoption, limit pricing where a low price deters rivals, and leverage to gain sales in markets for application programs. Implicit in all of these arguments is the presumption that prices would be higher in the absence of these forces. There is, however, no model of dynamic monopoly that provides a basis for this claim. We provide a game theoretic analysis of dynamic monopoly pricing for an upgrade good and establish that, in equilibrium, high prices are not a necessary outcome. Significantly, low prices, as measured by a seller who captures a small share of the social surplus, emerge in equilibrium.

Upgrade markets, by definition, regularly confront buyers with the choice of adopting a new higher-quality version or remaining with their current version. Microsoft's recent introduction of Vista was an adoption failure as buyers overwhelmingly chose to stay with their existing XP version, echoing a previous episode with Windows Millennium in 2000. Microsoft moved quickly to introduce a new version. Windows 7 was launched in late October 2009 to a much more favorable buyer response. As early as May 2009, Microsoft CEO Steve Ballmer acknowledged that "If people want to wait [for Windows 7], they certainly can." This simple observation, which implicitly takes the failure of Vista as a given, leads to a more subtle set of questions regarding the full range of individual and collective adoption incentives.

Consider the initial offer of Vista. An individual buyer has the option of remaining with XP. If most other buyers had purchased Vista then we can expect a concern about "falling behind" the market to be pivotal for an individual buyer's willingness to pay. Given that others did not, an individual buyer who stayed with XP is in the position Ballmer described. By purchasing Vista a buyer would "jump ahead" of the market. This buyer would then be confronted with the choice of purchasing Windows 7 to "keep

---

for an added charge, new or improved services such as web browsing, e-mail access and text messaging. Many capital goods are regularly upgraded, including airports (terminals and runways) and oil refineries, among others.

up" with market, assuming that Windows 7 is widely adopted. How does this recurrent interplay of individual and collective decisions with respect to incentives to "fall behind" or "jump ahead" of the market work to determine prices and adoption in an upgrade market? We argue that the ability of the seller to tempt an individual buyer to "jump ahead" is the critical factor and that this incentive forms the basis of a credible threat for buyers to reject an upgrade offer. Moreover, low prices can emerge in equilibrium even when buyers have a very strong incentive not to "fall behind" the market.

We examine an infinite horizon model of an upgrade market with 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 setting. Building on the recent literature, we model 'upgrade' goods by assuming that the sequence of quality increments satisfies a downward complementarity property: an additional quality increment is valuable only if a buyer holds all previous quality increments. The seller is unconstrained with respect to bundling options and any combination of quality increments (a single bundle or a set of bundles) may be offered in each period. Bundling is thus endogenous.

This basic structure is arguably a very attractive setting for the seller - homogenous buyers with a fixed valuation per quality unit, unrestricted bundling, and a costless exogenous flow of upgrade innovations. It is natural to expect a "perfect" monopoly outcome in which the seller captures all of the social surplus, and this is exactly what happens in several benchmark cases. For example, when the seller can only offer a single quality good and buyers are homogeneous, the seller can make an offer that induces buyers to purchase now. 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 buyers to reject offers with high prices.[2] In sharp contrast, we find that surplus growth due to rising quality in an upgrade market provides buyers with an option to return to the market for future purchases and that this option leads each buyer in a group to reject an offer that a single buyer would not.

---

[2] The standard incentive (Coase (1972)) to cut price over time and move down the demand curve is not present when buyers are identical. Papers on the Coase conjecture with a single good and a set of heterogeneous buyers include Stokey (1981), Bulow (1982), and Gul, Sonnenschein, and Wilson (1986). Ausubel and Deneckere (1989), Fehr and Kuhn (1995) and Sobel (1991) provide folk theorems for the durable goods model. Bond and Samuelson (1984) examine a rational expectations equilibrium with depreciation and replacement sales. Methodologically, our paper is closest to Sobel (1991) where there is entry of new consumers over time. In both cases, the market never closes, due to new demand in the case of Sobel and to quality growth in our case.

The primary intuition 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 "jumps ahead" of the market reduces to a one-period flow of value. As a result, willingness to pay is limited to the one-period flow value of the upgrade. We can apply this result at any stage of the game and for any given discount factor. The credible threat to reject a seller offer, given that other buyers also reject, leads to an implicit form of coordination among buyers and, in turn, to multiple equilibria.

We construct Markov perfect equilibria for this dynamic game and show that every subgame perfect equilibrium payoff is achieved by some Markov perfect equilibrium. Two classes of equilibria are identified: efficient and generational. Efficient equilibria have buyers acquiring a new upgrade each period and payoffs span a significant economic range. At one extreme, the seller captures all of the surplus and each quality increment sells immediately for the full present discounted value. At the other extreme, each quality increment sells only for the one period flow value, leaving a buyer with the entire residual surplus. These equilibrium payoff results contrast sharply with a traditional folk theorem in two important respects. First, a high discount factor is not necessary; the result applies for any discount factor (innovation rate) between zero and one. Second, severe punishments are not required; after any deviation, the market returns to the equilibrium path after one period.

In contrast, and despite the complete information setting, inefficient equilibria do exist. These "generational" equilibria exhibit cyclical delay in which multiple quality increments go unsold until they are bundled together for sale and, necessarily, the market returns to the "state of the art" with a new generation. The cycle length reflects a second type of equilibrium coordination in an upgrade market. Importantly, relative to the set of efficient equilibria, we find that generational equilibria compress the range of payoffs. Intuitively, delay requires that deviations to make early trades are unattractive, and this implies that the seller and the buyers share the surplus more equally.

The seller is free to offer any feasible collection of quality units. On the equilibrium path, it is sufficient to consider only upgrade offers with a contiguous set of quality units. Equivalently, we show how to interpret these upgrade offers in terms of a full bundle (new version of the product) with pricing contingent on a buyer's current product holding, much as the owner of an existing product faces an upgrade price to acquire a new version. Furthermore, we find no role for the commitment period (time between seller offers), in contrast to the literature on the Coase Conjecture. What matters for equilibrium outcomes is the frequency at which quality improves: allowing the seller to make offers more frequently has no impact with homogenous buyers.

There is a relatively small literature on upgrade models, with most of the work involving a finite horizon. Waldman (1996) and Nahm (2004) each examine 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; Hoppe and Lee (2003) extend this model to allow entry. Ellison and Fudenberg (2000) analyze a series of static and two period models that feature network externalities and a cost to consumers of upgrading the good. In the finite horizon version of our model, the monopolist captures all the surplus because the basis for a credible buyer threat is undermined by the terminal period. Fishman and Rob (2000) examine an infinite horizon upgrade model, focusing on innovation incentives, and analyze a rational expectations equilibrium in which the seller is assumed to offer only a single bundle consisting of all prior quality levels. We focus on pricing and adoption, taking innovation as exogenously given, and provide a game-theoretic analysis in which the seller choice of which bundles (and prices) to offer is endogenous.

In the next section, we present the model. In Section 2 we examine efficient equilibria and in Section 3 we examine generational equilibria. We discuss the upgrade structure of our model in Section 4 and consider directions for future research in Section 5. Proofs are in the Appendix; all omitted proofs are in Anton and Biglaiser (2010b).

1 The Model

We first describe the basic elements of the game. We next turn to strategies and payoffs, and then define and discuss Markov perfect equilibrium. We present the formal theoretic framework for bundling (strategies and equilibrium) in Appendix A.

1.1 Basic Elements

We examine an infinite horizon, discrete time model. Let $\tau = 1, 2, ...$ 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 $\{1, 2, ..., \tau\}$ and associated prices. For example, the seller can offer the bundle of all feasible qualities $\{1, 2, ..., \tau\}$ 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 could 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. Given a seller offer, the buyers respond simultaneously with each buyer choosing which bundle(s) to accept in period τ.[3]

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 $\{1, ..., \tau\}$; we also refer to this as a version. A partial upgrade is a bundle $\{\sigma, ..., \sigma + k\}$, where $1 \leq \sigma \leq \sigma + k \leq \tau$. We will show that, in equilibrium, a seller need only make upgrade offers.

A buyer holding contiguous units $1, ..., q$ but not $q + 1$ has a flow utility of vq in a period. Thus, 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.

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, we can interpret a

---

[3] 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 in equilibrium for buyers from the possibility of resale.

large (small) δ in terms of rapid (slow) rate of innovation and assess limiting behavior.

Consider the payoff for a buyer. In each period, a buyer holds some subset of the feasible qualities. Let q_τ denote the maximal contiguous quality held by a buyer after any purchase in period τ. That is, a buyer holds units 1 up through q_τ but does not hold unit $q_\tau + 1$. From any point in the game, the payoff of the buyer is the present discounted value from quality flows net of payments. From the start of the game this is given by

$$\sum_{\tau=1}^{\infty} \delta^{\tau-1}(vq_\tau - p_\tau),$$

where p_τ is the payment made by the buyer in period τ. Similarly, the payoff of the seller from any point onward is the present discounted value of revenues, r_τ, from sales to buyers. From the start of the game, this is given by

$$\sum_{\tau=1}^{\infty} \delta^{\tau-1} r_\tau.$$

Consider efficient allocations. Payments and revenues are transfers that do not affect total surplus. Thus, for any path of quality holdings, the sum of surplus for any given buyer and the seller from any period τ_0 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 surplus in an efficient allocation from the start of the game is

$$S_1 = v + \delta 2v + \delta^2 3v + ... = \frac{v}{(1-\delta)^2}.$$

Intuitively, S_1 is the surplus created when buyers acquire one new unit in each period, where each new unit has a present discounted value of $\frac{v}{1-\delta}$. Starting from any period τ, the maximal available surplus is

$$S_\tau = v\tau + \delta v(\tau+1) + \delta^2 v(\tau+2) + ... = \frac{v(\tau-1)}{1-\delta} + S_1.$$

Intuitively, the difference between S_τ and $S_{\tau+1}$ is the flow value of τ units in period τ.

Thus, we always have $S_\tau > \delta S_{\tau+1}$, as delay necessarily involves lost surplus and hence inefficiency. However, because each unit generates surplus, we also have $S_\tau < S_{\tau+1}$.

1.2 Markov Perfect Equilibrium

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.

The simplest 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 allows us to generate all subgame perfect equilibrium seller payoffs. To proceed, consider any history in which all buyers enter period τ with the same maximal quality level Q (units 1 through Q). We define this to be state (τ, Q) and refer to $\tau - Q$ as the quality gap. Markovian behavior is defined by the condition that players' strategies depend only on the size of the quality gap. Thus, 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). 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 $\tau' - Q = \tau$.[4] Thus, buyer payoffs are always the sum of the PDV of current holdings, $vQ/(1-\delta)$, and the incremental utility, u_τ.

This definition of Markovian behavior implies that the same number of units are included in an upgrade bundle whenever the quality gaps coincide. In particular, the

---

[4] Feasible payoffs in our model have a simple stationary structure. Any subgame perfect equilibrium from state $(\tau + 1 - Q, 0)$ is also subgame perfect from state $(\tau + 1, Q)$ once we relabel units and translate the buyer's payoff.

seller is not required to offer the full bundle of all feasible units; we discuss contingent pricing and versions in section 4. Henceforth, we use equilibrium to refer to a pure strategy buyer symmetric Markov perfect equilibrium in the quality gap.[5]

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 (past seller offers and buyer acceptances) 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. As a result, buyers act as price takers: no buyer expects that their own decision will have any impact on subsequent play, such as affecting future seller offers.

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 $vk/(1-\delta)$.[6] We can easily allow for higher buyer catch-up continuation values as long as they do not exceed the equilibrium payoff. For the analysis, however, it is helpful to follow the above specification of a zero increment in utility for a buyer who falls behind the market. This will highlight the critical role played by the incentive for a buyer to jump ahead of the market.

2 Efficient Equilibria

We begin with a basic result on the necessary structure with respect to all equilibrium payoffs: the seller can always induce buyers to make a purchase.

---

[5] States 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 continuatiuon equilibria in all states. The Appendix provides a detailed analysis of continuation equilibria for any distribution of buyer holdings and for any bundles offered by the seller.

[6] One can interpret this as (i) the seller (optimally) ignores individual buyers (measure zero) who differ from the market path and the missing units necessary to catch up are never offered, or (ii) the seller offers the necessary units but sets the upgrade price so as to extract all of the continuation surplus. Because the seller's strategy does not depend on a deviation by a set of measure zero buyers, the seller must either completely refrain from making "catch-up" offers, or always make such offers. Either of (i) or (ii) is sufficient for the equilibrium construction.

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* $\{Q+1, ..., \tau\}$ *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.

Lemma 1 implies that in state $(\tau, 0)$ the continuation payoff of the seller is at least $v\tau + \delta v/(1-\delta)$. First, the seller can offer τ units at flow value. Second, in the future each new unit can be sold at v. Lemma 1 and this flow dominance payoff bound are basic results. They apply to any subgame perfect payoff and do not depend on Markovian behavior or symmetric buyer strategies: they only rely on buyers acting as price takers.

In an efficient equilibrium, a good is sold in each period when it first becomes available. At the start of the game, the seller offers the first unit at price p_1 and all buyers accept. In the second period, the state is then $(2, 1)$ and the quality gap is again 1. Under Markovian behavior, the seller offers the second unit at price p_1 and all buyers accept, leading to state $(3, 2)$, and so on. Thus, the firm earns a profit of $\pi_1 = p_1/(1-\delta)$ and buyers receive a payoff of $u_1 = \frac{1}{1-\delta}\left[\frac{v}{1-\delta} - p_1\right]$. In an efficient equilibrium, the firm and buyers divide the maximal social surplus, $S_1 = v/(1-\delta)^2 = \pi_1 + u_1$.

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. As this is the minimum possible payoff for the seller (flow dominance and Lemma 1), every subgame perfect payoff is achieved in some Markov perfect equilibrium. Then, to highlight the fundamental economic forces at work, we present the extremal equilibria. We then develop sufficient conditions and characterize efficient equilibria.

2.1 Buyer and Seller Support Constraints

Equilibrium requires that no player can benefit from a deviation.[7] To rule out deviations, we must specify the continuation payoffs from state $(2,0)$ and other off-equilibrium states; the continuation payoff for any state is pinned down (Markov) once we specify those for states of the form $(\tau,0)$. Suppose 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. Then, continuation payoffs with this "cash-in" support at $(\tau,0)$ are $\pi_\tau = p_\tau + \delta\pi_1$ for the seller and $u_\tau = \frac{v\tau}{1-\delta} - p_\tau + \delta u_1$ for the buyers. Note that this is the efficient path from $(\tau,0)$, and 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 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 $(\tau+1,0)$. 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 but 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 $p \leq p(\sigma,\tau)$. On the acceptance side, it must be optimal for an individual buyer to accept any offer $p \leq p(\sigma,\tau)$, given that all other buyers are accepting (symmetric strategies) and the quality gap moves to $\tau+1-\sigma$. Rejecting when others all accept yields 0 by construction, as the buyer falls behind the market. Accepting yields a current flow of $v\sigma - p$ plus a future value of $\delta u(\tau+1,\sigma)$; recall that $u(\tau+1,\sigma)$ is the sum of the PDV of σ units and the increment $u_{\tau+1-\sigma}$. Thus, it is optimal for all buyers to accept p for σ units in state $(\tau,0)$ if $\frac{v\sigma}{1-\delta} + \delta u_{\tau+1-\sigma} \geq p$. This reflects the incentive of a buyer to "keep up" with the market.

The rejection side of the cut-off rule reflects the incentive not to "jump ahead" of the market, and an offer of $p > p(\sigma,\tau)$ must be rejected by all buyers. Rejecting when others reject yields a payoff of $\delta u_{\tau+1}$. More subtly, accepting when all other buyers reject yields a flow of $v\sigma - p$ plus the option of purchasing the cash-in offer for $\tau+1$ units next period.

---

[7] We 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).

Thus, an individual buyer optimally rejects when others reject if

$$p > v\sigma + \delta \max\left\{\frac{v\sigma}{1-\delta}, u_{\tau+1}\right\} - \delta u_{\tau+1} \equiv g(\sigma, u_{\tau+1}).$$

To understand this condition, consider the position of a deviating buyer in the next period. When the other buyers purchase the cash-in offer in period $\tau + 1$, this buyer has two options. If $u_{\tau+1} > \frac{v\sigma}{1-\delta}$, it will be optimal to purchase when these other buyers do. Thus, the deviating buyer is initially 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 $\frac{v\sigma}{1-\delta} - \delta u_{\tau+1}$.

Thus, combining the acceptance and rejection sides of the cut-off strategy we have

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

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 is bounded above by $\frac{v\sigma}{1-\delta}$, cut-off strategies exist for any non-negative utility sequence. The upper bound in (1) says that prices must be sufficiently low that falling behind the market is not optimal for an individual buyer, while the lower bound says that it is optimal to jump ahead and accept offers below this level. The lower bound on g of $v\sigma$ reflects flow dominance.

The distinct upper and lower bounds on the cut-off rule show that an individual buyer's willingness to pay depends on the actions of other buyers. It is important to recall 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 arises from i) quality growth and the resulting incentive for a buyer to return to the market for another upgrade, ii) the seller's offer in the future depends on "the state of the market" and, iii) an individual buyer will be affected by his position relative to the market when making future purchasing decisions.

Given these buyer cut-off strategies, 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

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

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

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. Buyers must reject any price above p_τ for τ units. Note that (2) holds with equality at $\sigma = \tau$ by construction of the equilibrium continuation. Similarly, note that the buyer condition (1) also applies in the event of delay ($\sigma = 0$) and on the equilibrium path ($\tau = \sigma = 1$).

2.2 Equilibria with extreme payoffs

To provide insight for the general characterization of all subgame perfect equilibrium payoffs, it is helpful to begin by considering the economic structure of two extremal equilibria with respect to payoffs. In the first equilibrium, the seller captures all of the social surplus while the buyers receive a payoff of zero. In the second, the seller payoff is held to the flow value level, identified via Lemma 1 as the minimum subgame perfect payoff, while the buyers capture the full discounted surplus of each unit after the initial period. While both equilibria follow the efficient path, the contrast in prices and off-equilibrium behavior serves to identify the critical roles played by an infinite horizon, growth, and a continuum of buyers.

In the equilibrium where the seller captures the full surplus, each new upgrade unit is sold at the price $p_1 = \frac{v}{1-\delta}$. Thus, the seller earns $\pi_1 = S_1$ leaving the buyer payoff at $u_1 = 0$. This is supported by cash-in prices of $p_\tau = \frac{v\tau}{1-\delta}$ so that in the off-equilibrium situation where more than one unit is available, either because of a delay by the seller or an offer refusal by buyers, the seller only offers a bundle of τ units at a price equal to their full discounted value. Given the expectation of no surplus in the future, each buyer can do no better than to accept the offer of p_1 for each upgrade unit.

These turn out to be the only equilibrium payoffs for three benchmark cases. With an infinite horizon, buyers always seek to acquire higher quality units. Consider a finite horizon. Once the final period arrives, buyers have no prospect of acquiring upgrades in the future. With no surplus in the future, subgame perfection then implies that each buyer will necessarily accept any offer that provides a positive payoff. The unique equilibrium outcome in the final period, given any current units held by the buyers, is that the seller offers an upgrade to the state of art at a price equal to the full value of the upgrade to the buyers. Backward induction then implies that this holds for all prior periods since buyers never expect a positive payoff in the future. Hence, the equilibrium with a finite horizon follows the efficient path and the seller captures the full social surplus.[8]

For the benchmark of no growth, the seller has one unit to offer and trade may take place at any time over an infinite horizon. Since there is no heterogeneity among buyers, we have a special case of the standard durable goods model (gap case). The result of Fudenberg, Levine, and Tirole (1985) implies that there is never delay and the full surplus is always extracted from buyers in any subgame perfect equilibrium.

Suppose that there is only one buyer in an infinite horizon model with growth. A variation on the "speed-up" argument of Fudenberg, Levine, and Tirole implies that the seller can always profitably tempt the buyer to purchase all available units immediately.[9] The purchasing decision of a single buyer necessarily changes the state, in contrast to the case with a continuum of buyers. When responding to a current offer, a refusal is optimal only if the expected continuation payoff, δu_{t+1}, exceeds the payoff from accepting the offer, u_t. Then a positive buyer payoff, $u_1 > 0$, requires a supporting utility path that rises indefinitely at an exponential rate. Because this exceeds the available surplus in finite time, a credible threat for refusing a price increase, relative to u_1, unravels. The only equilibrium has $u_1 = 0$.

Thus, all three benchmarks suggest strong monopoly power for the seller. We find, however, that there are equilibria where the monopoly power is quite limited. With an infinite horizon and growth, there are always mutually beneficial trades in the future. With a continuum of buyers, an individual buyer's purchasing decision must be assessed relative to the behavior of other buyers because subsequent seller offers are based on the aggregate state of the market.

In the equilibrium where the seller payoff is reduced to the flow value of the upgrades, each upgrade unit is sold at the price $p_1 = v$. Thus, the seller earns $\pi_1 = \frac{v}{1-\delta}$, which is equal to the share $(1 - \delta)$ of the total surplus, S_1. The supporting path in the off-equilibrium states where multiple units are available always has a positive payoff for buyers, in contrast to the maximal seller payoff equilibrium. Our supporting path, which also applies to a broad range of payoffs, involves two distinct phases. Whenever the state involves a sufficiently large number of units, τ above a threshold T, we specify that buyer utility is held at a fixed value u_T that does not vary with the number of units τ. This makes the seller the residual claimant of the increase in surplus in the event of a delay or offer refusal. Due to discounting, the seller then has a strict incentive to cash-in today rather than delay. The efficiency gain from selling today is exactly $v\tau$, the flow value to

---

[8] Further detail on this and the other benchmarks are provided in the working paper version.

[9] In Section 3, we show that the speed-up argument implies all equilibria are cyclical for our model.

buyers from having the units one period earlier. The cost of sharing u_T of the available surplus one period earlier is $(1-\delta)u_T$. As a result, this phase has $vT \geq (1-\delta)u_T$.

The buyer incentives highlight the role of a continuum. A single buyer would always accept a slightly higher price today. In contrast, a buyer credible threat exists with a continuum. A single buyer who accepts a price increase when other buyers refuse will find that the state moves from $(\tau, 0)$ to $(\tau+1, 0)$. The seller will then offer the bundle of units 1 to $\tau+1$ at the price $p_{\tau+1}$. By purchasing this bundle, the deviating buyer would rejoin other buyers on the equilibrium path for a payoff of u_T. However, because $\frac{v\tau}{1-\delta}$ is above u_T, the deviating buyer will choose not to buy this bundle.

Of course, the seller is unconstrained with respect to multiple offers. There are several alternative ways to support the equilibrium in this regard. First, and most simply, the seller makes only the $p_{\tau+1}$ offer. Second, the seller could also offer unit $\tau+1$ by itself. This would create a second way for a deviating buyer (who holds τ units) to keep up with the other buyers. A wide range of prices for unit $\tau+1$ will support an individual buyer's decision not to deviate and accept the price increase in period τ; all that is needed is that the price is high enough that the payoff to a deviating buyer falls below δu_T.[10]

For the first phase of the support, instead of having the seller be the residual claimant of growth, we assign all of the growth in surplus to the buyers. To see how this works start with u_T and set $u_{T-1} = vT + \delta u_T$. This pushes the seller to indifference with respect to delay. As we repeat this step and work back to u_1 we ultimately assign all of the surplus growth from period 1 through T to the buyers. By choosing T large enough, we are able to lift u_1 up to δS_1. During this phase, we have $v\tau < (1-\delta)u_{\tau+1}$. This changes the structure of the buyer's credible threat. Now, when confronted with a price increase, a buyer who buys when others do not will find it optimal to keep up with the other buyers and buy with them in the next period. Thus, a deviating buyer would never pay more than the flow value $v\tau$ to purchase when others do not. This willingness to pay is necessarily lower than p_τ, the cash-in price that a buyer is willing to pay when others purchase. Thus, a seller who tries to increase price will make no sales and will not gain.

---

[10] With a single deviating buyer, there are no revenue implications for the seller, since all other buyers lack the first τ units and would never purchase unit $\tau+1$ by itself. Because a single deviating buyer is a set of measure 0, the seller either refrains from making this offer or always makes the offer. In off-equilibrium states where masses of buyers hold different units, the seller will find it striclty optimal to make multiple offers (see Appendix).

2.3 Sufficient Conditions

We now combine the buyer and seller support conditions, (1) and (2), to identify when there exist supporting prices $p(\sigma, \tau)$ such that an efficient equilibrium is supported by the cash-in outcome. Combining the seller profit from (2) with the buyer lower bound on prices from (1), we must have

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

Recalling that $S_\tau = \pi_\tau + u_\tau$, we see that the above condition is equivalent to

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

for $0 \leq \sigma \leq \tau$ and $\tau \geq 1$. 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.[11] Given a sequence of utilities that satisfies (3), we can clearly construct the supporting prices $p(\sigma, \tau)$ for conditions (1) and (2). 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 2 *Suppose the sequence of buyer utilities u_τ satisfies (3). 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 $\{0, ..., \tau - 1\}$ in any period τ, even though such states do not arise on the equilibrium path. In all cases, the support returns to the equilibrium path after 1 period. See Appendix B for the support construction.

---

[11] In the benchmark case of a single good, where a buyer exits the market after a purchase, (3) reduces to $0 \geq u_1$ and all surplus accrues as profit.

2.4 Frequent Innovations: Existence and Payoffs

For each payoff $u_1 \in [0, \delta S_1]$, we will construct an associated supporting path of $u_2, u_3, ...$ such that the seller finds it optimal to make an acceptable offer to achieve a cash-in outcome in every state. 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. First, we take up the main case of frequent innovation, $\delta > 1/2$, in which future units are relatively more valuable. We conclude the section with the simpler case of infrequent innovation, $\delta \leq 1/2$.

The sequence of utilities that we construct to satisfy the sufficient condition (3) has two phases. When the quality gap is T or smaller, states $(1, 0)$ to $(T, 0)$, the support makes the seller indifferent between a cash-in and delay; for states with a larger quality gap the support keeps buyer utility constant at u_T. Thus, for $T \geq 2$ we define a $T - stage$ support utility sequence by

$$u_\tau = \begin{cases} v\tau + \delta u_{\tau+1} & \text{for} \quad \tau = 1, ..., T-1 \\ u_T & \text{for} \quad \tau \geq T. \end{cases} \tag{4}$$

For any given u_1 and length T, the sequence $(u_2, ..., u_T)$ is determined. As we will see, a higher u_1 or a larger δ will require a larger length T.

A direct consequence of a *T-stage* support is that we only to need to satisfy the support constraints, equations (3), over the range $\tau = 1, ..., T$; 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 and the seller is the residual claimant of surplus growth. Thus, an advantage of a $T - stage$ support is that we only have to check a finite set of conditions. We now turn to finding the appropriate length for the $T - stage$ support.

The following algorithm determines the length of the $T - stage$ support for any given utility level $u_1 \in [0, \delta S_1]$ and discount factor above 1/2:

- 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 $u_1 < \left(1 - \delta^2\right) S_1$, set $u_\tau = u_2$ for all $\tau > 2$. If not, set $u_3 = \left(\frac{u_2 - 2v}{\delta}\right)$.

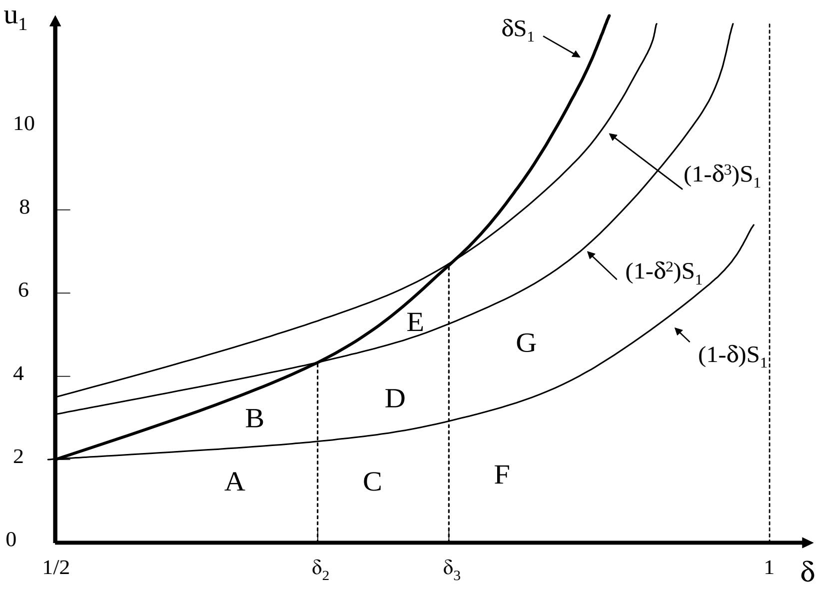

Figure 1: Boundary for T-Stage Support ($v = 1$)

- Keep following this logic until T such that $u_1 \leq \left(1 - \delta^T\right) S_1$.

We know the algorithm always has a final step, since $\delta S_1 \leq (1 - \delta^T) S_1$ holds for T sufficiently large. The discount factor δ determines how T must be set 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 $\delta \in (0, 1)$. 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 $\left(1 - \delta^2\right) S_1$ 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. Thus, we have

Proposition 1 *Let $\delta > 1/2$. Then, for every $u_1 \in [0, \delta S_1]$, there exists an efficient equilibrium with buyer payoff u_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.

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. The discount factor is given by $\delta = e^{-r\Delta}$, where r is the interest rate. Adjusting 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 $v/(1-\delta)$ 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 total surplus, since $v/(1-\delta)^2$ grows without bound. It follows directly that in an efficient equilibrium the maximum buyer share of the surplus converges to 1.

Corollary 1 *In the limit, as upgrades become increasingly frequent ($\Delta \to 0$), 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. Stated a bit differently, a market may allow a seller to earn a high level of absolute profit while capturing only a small share of the total surplus.[12]

---

[12] In our model, δ reflects the time between innovations as well the commitment period of the seller (time between offers). This is without loss of generality. We can allow the seller to make offers more frequently, holding the innovation rate constant, and follow the same principles for constructing the support utilities. In contrast to Coasian settings, these offers will be off-the-equilibrium path. It is straightforward to check that an efficient equilibrium with a $T - stage$ support is robust to allowing the seller to make an interim offer of any feasible bundle at any time between innovations. Specifically, in the constant utility part of the $T - Stage$ support we can specify a cash-in at the same utility level for the interim offer. When utility is rising it is simplest to specify a delay outcome for the support. Where as allowing interim offers in the Coasian setting with heterogenous buyers and a single good will speed up sales, we find that the original equilibrium path where the arrival of quality increments coincides with sales does not change with the commitment period of the seller.

2.5 Infrequent Upgrades: Existence and Payoffs

When innovations are infrequent, $\delta \leq 1/2$, flow dominance plays a stronger role, since the flow value of a unit of quality, v, is now larger than the future discounted value of a unit, $\delta v/(1-\delta)$. Now, a buyer necessarily values one unit today more than two units tomorrow, other things equal. The increased power of flow dominance does not, however, imply a necessary increase in seller market power. We are still able to support efficient equilibria in which buyers receive any payoff $u_1 \in [0, \delta S_1]$, as described next.

Proposition 2 *Suppose $\delta \leq 1/2$. Then every $u_1 \in [0, \delta S_1]$ can be supported in an efficient equilibrium with a continuation utility of $u_\tau = \bar{u} \equiv u_1(1-\delta)/\delta$, for all $\tau \geq 2$.*

Consider the limiting case as innovation becomes less and less frequent, $\Delta \to \infty$. 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 $\Delta \to \infty$, we find that $v \to \lambda/r$ 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.

3 Delay and Generational Equilibria

We now consider inefficient equilibria. First, we show that equilibria must have a simple cyclical structure and, second, that innovation needs to be sufficiently frequent for delay to occur. We then consider seller and buyer incentives in the delay states, derive approach conditions, and show existence.

3.1 Cyclical Equilibria

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 and the state returns to $(1,0)$.

Proposition 3 *Every equilibrium follows a $t-$cycle equilibrium path: the buyers purchase quality units $\{1, ..., t\}$ 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 $(t, 0)$, the seller payoff rises if all accept since

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

and, upon substituting for $\hat{p}$ and noting that (t, τ) is a delay state, this reduces to $v\tau > \varepsilon$.

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 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. We refer to a $t - cycle$ equilibrium with $t \geq 2$ as a generational equilibrium, since buyers upgrade to the quality frontier.

Payoffs in a $t-cycle$ equilibrium are then $\pi_t = p_t/(1-\delta^t)$ 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. These are short run efficiency losses since each cycle resolution ends with buyers holding all feasible units as of the sale date. From the above seller and buyer payoffs, joint surplus is

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

Thus, we will be able assess the relative efficiency loss by comparing the joint surplus Ψ_t to the maximal surplus S_1. Note that $\Psi_\tau = \delta^{t-\tau}\Psi_t$ for $\tau = 1, ..., t-1$ in the delay states; thus, Ψ_τ as well as seller and buyer payoffs are increasing in τ.

3.2 Delay Equilibria and Upgrade Frequency

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

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

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, current revenue dominates the payoff from waiting to sell 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.

3.3 Delay and Approach Conditions

Equilibrium delay requires that the seller can find no offer that is acceptable to buyers and also profitable relative to waiting to sell in state $(t, 0)$. Thus, we must derive approach conditions for the equilibrium. In analogy with the support conditions for efficient equilibria, buyer strategies follow cut-off rules of when to accept offers and in conjunction the seller does not find it profitable to make a sale before the quality gap reaches t units.

The problem of satisfying these buyer and seller delay incentives reduces to finding a buyer payoff u_t that satisfies

$$(1-\delta^{\tau})(\Psi_t - u_t) \geq \frac{v\tau(\delta^{\tau-t}-1)}{(1-\delta)} + \max\left[\frac{v\tau}{(1-\delta)}, u_t\right] - u_t \tag{5}$$

for $\tau = 1, ..., t-1$. As developed in Appendix C, when the buyer payoff satisfies (5), there exist cut-off prices such that delay is supported. The analysis of (5) is involved because the buyer and seller delay incentives can change significantly as the quality gap rises. To begin, note that (5) implies a lower as well as an upper bound on u_t. As u_t approaches 0, (5) necessarily fails (at all τ). In this case, the incentive for an individual buyer to jump ahead is very strong when u_t is small and the seller can profitably attract buyers as early as $\tau = 1$. If the buyers get too much of the available surplus (u_t approaches Ψ_t), then (5) necessarily fails (at all τ). Intuitively, recalling that $\pi_t = \Psi_t - u_t$, this is because the seller can exploit flow dominance to sell early and profitably attract buyers by offering τ units at a price of $v\tau$.

An extreme payoff on either the buyer or the seller side thus allows the seller to profitably induce a speed up. As a consequence, both buyer and seller payoffs are compressed relative to the range of payoffs for efficient equilibria.

Lemma 3 *The set of payoffs for all generational equilibria is a strict subset of the set of payoffs for efficient equilibria.*

In order to be willing to wait until period t, a buyer must decline seller deviation offers. This requires a strictly positive payoff for buyers. By contrast, there is an efficient equilibrium in which buyers receive no surplus. Similarly, the seller must be willing to delay and this includes foregoing the option to sell units prematurely at flow value. Thus, the seller necessarily earns more than the flow dominance lower bound. The upper bounds then follow from the lower bound on the other side of the market. Thus, while delay generates less surplus, both sides of the market necessarily receive a larger payoff relative to the minimum subgame perfect equilibrium payoff.

Satisfying the approach conditions for a $t-$*cycle* (delay) equilibrium reduces to finding a buyer utility for the sale date, u_t, that satisfies (5) 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. In this case, a sufficiently large quality increase will accumulate by the sale date that an (early) deviating buyer would choose to keep up with the market. However, when τ is closer to t, a deviating buyer would not purchase again at t since the incremental surplus from $t-\tau$ units is insufficient and this buyer would choose to fall behind the market.

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 (5) 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 (5) is quite involved, it turns out that many of the complications only arise at relatively low discount factors.[13] In Appendix C, we develop a sufficient condition on δ^t such that if δ^t is above a threshold, d^*, then the approach conditions (5) are satisfied for an interval, $\left(\underline{u}^A, \overline{u}^A\right)$, of utility levels. We then have

Lemma 4 *If $\delta^t > d^*$, then there exist bounds $\underline{u}^A$ and $\overline{u}^A$ such that the approach conditions (5) are satisfied for any $u_t \in \left(\underline{u}^A, \overline{u}^A\right)$ in a $t-cycle$ equilibrium.*

Numerically, the threshold, 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.

Lemma 4 provides a lower and upper bound on the buyers' payoffs. The utility bounds, $\underline{u}^A$ and $\overline{u}^A$, depend on δ and t and are derived in Appendix C. At d^*, $\underline{u}^A = \overline{u}^A$, and for all t and δ pairs where $\delta^t > d^*$, we have $\underline{u}^A < \overline{u}^A$. See Figure 2.

3.4 Existence and Payoffs for Delay Equilibria

We must also consider the cash-in (off equilibrium) support conditions for inefficient $t - cycle$ equilibria in addition to the approach conditions. The buyer cut-off rules and

---

[13] This is because, in general, when condition (5) holds at $\tau = 1$ and at $\tau = t-1$ it does not necessarily follow that (5) holds at $1 < \tau < t-1$. This can be seen by explicitly solving the cases of $t = 2$ and $t = 3$.

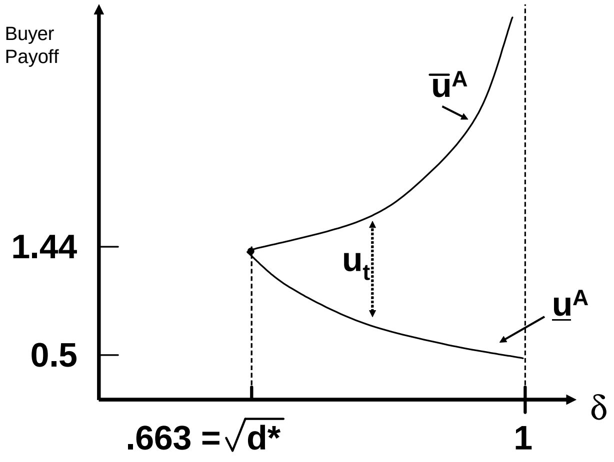

Figure 2: Approach Conditions ($v = 1$ and $t = 2$)

seller profits must satisfy the analog of (3). The main 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 where the quality gap exceeds t need to be modified for delay equilibria. We provide the details in Appendix C and now turn directly to the existence of delay equilibria. Sufficient conditions for approach states, $(\tau, 0)$ where $\tau < t$, are provided in Lemma 4 and summarized by the bounds, $\underline{u}^A < u_t < \overline{u}^A$, 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 $u_t \in (\underline{u}^A, \overline{u}^A)$ 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 the division of surplus is best understood in terms of the rate of innovation, Δ. The length of time between sales in a $t-cycle$ equilibrium is $D \equiv \Delta t$. As we let Δ decline, so that innovation is more rapid, any given $t-cycle$ will continue to exist, since $\delta^t = e^{-r\Delta t} > d^*$ still holds, 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, $\overline{s}_b \equiv \overline{u}^A/\Psi_t$ 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,* $\overline{s}_b \rightarrow 1$ *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.

4 Discussion of Upgrade Structure

We now consider how our endogenous bundling framework corresponds to observed practice in a variety of upgrade markets. We also discuss our results in relation to assumptions on bundling and the upgrade literature.

4.1 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 $\{\sigma + 1, ..., \sigma + k\}$. Thus, a buyer with units $\{1, ..., \sigma\}$ 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 versions, bundles of the form $\{1, ..., \sigma, ..., \sigma + k\}$. As we will show, there is a very close relationship between upgrade offers and versions once downward complementarity and pricing contingent on a buyer's current holdings are taken into account.

Contract contingencies, especially with respect to a buyer's current holdings, are frequently observed in 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. Airliners typically offer seat upgrades, club memberships, and other amenities at prices that vary with frequent flier status, a result of past purchases.

A new 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 $\{1, ..., \sigma\}$ and two offers. One is an upgrade bundle $\{\sigma+1, ..., \sigma+k\}$ for price p. The other is a new version $\{1, ..., \sigma, ..., \sigma+k\}$ at price p that is only available to buyers who hold units $\{1, ..., \sigma\}$. 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.

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 version form with a holding contingency). This is easy to see in a finite horizon setting and the same 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.

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, when buyers have asymmetric holdings, we constructed continuation equilibria in which the seller captured the available surplus and, as a result, this support is 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 jump 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 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.

Thus, whether we regard the offers as upgrade bundles or versions with price contingencies, as in Fudenberg and Tirole (1998) or Ellison and Fudenberg (2000), does not affect the equilibrium structure.[14] In terms of information structures, 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.[15] Bundles can be presented as upgrades or as versions with a contingency on past purchases. Our formulation of the offer space is simpler as it avoids the 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.

4.2 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

---

[14] In both Fudenberg and Tirole and Ellison and Fudenberg the second period offers can distinguish between buyers who purchased in period 1 and those who did not. These models also feature buyer heterogeneity.

[15] In our working paper we also address network effects, compatibility issues, and adoption costs. We argue that our equilibrium results are robust to these forces as they all reinforce the incentive for a buyer to keep up with the market while reducing the payoff of a buyer who jumps ahead of the market. See the related policy discussion in Anton and Biglaiser (2010a).

τ, is necessarily also an offer for units $\{1, ..., \tau\}$ 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 captures the full surplus when upgrades are unbreakable.

In the simplest case of our model with two-periods and unbreakable goods, 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 with non-contingent pricing, 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.[16] In their rational expectations equilibrium, 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. The basis of a credible threat for buyers resides with the extent to which the seller can tempt a buyer to jump ahead of others.

---

[16] Full 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 $\{1, ..., \tau-1\}$. By rejecting an offer and resuming purchases next period, an individual buyer obtains $v\,(\tau-1) + \delta u(\tau+1, \tau)$. Purchasing today yields $v\tau - p(\tau, \tau-1) + \delta u(\tau+1, \tau)$. Combining, we must have $p(\tau, \tau-1) \leq v$.

4.3 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 $1, 2, ...$ 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 due to the lack a credible threat for refusing a seller offer.[17] The essential difference is that a buyer can accept a current offer when others do not, skip the subsequent cash-in offer from the seller, and then resume purchasing. Because the goods are independent, there is no payoff consequence due to complementarity from any missing units. Thus, with independent goods, we necessarily have $u_1 = 0$. Intuitively, complementarity is essential for a credible threat to refuse price increases as a deviating individual buyer faces the extra cost of having to acquire the missing unit. This is why, in our upgrade structure, it matters to an individual buyer whether or not others are expected to purchase the seller's offer.

5 Directions for Future Work

Buyers are homogeneous in our model. This was assumed 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 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

---

[17] 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 $(\tau, 0)$. In order to reject a price increase we must have $(\delta + (1-\delta)/\delta^\tau)\, u_1 \leq u_{\tau+1}$ and this implies $u_1 = 0$.

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.

References

[1] Anton, James J., and Gary Biglaiser. 2010a. “Compatibility, Interoperability, and Market Power in Upgrade Markets.” *Economics of Innovation and New Technology*, 19(4): 373-85.

[2] Anton, James J., and Gary Biglaiser. 2010b. “Quality, Upgrades and Equilibrium in a Dynamic Monopoly Market.” Working paper.

[3] Ausubel, Lawrence M., and Raymond J. Deneckere. 1989. “Reputation in Bargaining and Durable Goods Monopoly.” *Econometrica,* 57(3): 511-31.

[4] Bond, Eric W., and Larry Samuelson. 1984. “Durable Good Monopolies with Rational Expectations and Replacement Sales.” *RAND Journal of Economics*, 15(3): 336-45.

[5] Bulow, Jeremy I. 1982. “Durable-Goods Monopolists.” *Journal of Political Economy*, 90(2): 314-32.

[6] Coase, Ronald H. 1972. “Durability and Monopoly.” *Journal of Law and Economics,* 15(1): 143-49.

[7] Ellison, Glenn, and Drew Fudenberg. 2000. “The Neo-Luddite’s Lament: Excessive Upgrades in the Software Industry.” *RAND Journal of Economics*, 31(2): 253-72.

[8] Fehr, Nils-Henrik, and Kai-Uwe Kuhn. 1995. “Coase versus Pacman: Who Eats Whom in the Durable-Goods Monopoly?” *Journal of Political Economy*, 103(4): 785-812.

[9] Fishman, Arthur, and Rafael Rob. 2000. “Product Innovation by a Durable Goods Monopoly.” *RAND Journal of Economics*, 31(2): 237-52.

[10] Fudenberg, Drew, and Jean Tirole. 1991. *Game Theory.* Cambridge, MA: MIT Press.

[11] Fudenberg, Drew, and Jean Tirole. 1998. “Upgrades, Tradeins and Buybacks.” *RAND Journal of Economics*, 29(2): 235-58.

[12] Fudenberg, Drew, David Levine, and Jean Tirole. (1985), “Infinite Horizon Models of Bargaining with One-Sided Incomplete Information.” In *Game Theoretic Models of Bargaining,* ed. Alvin E. Roth, 73-98. Cambridge: Cambridge University Press.

[13] Gul, Faruk, Hugo Sonnenschein, and Robert Wilson. 1986. “Foundations of Dynamic Monopoly and the Coase Conjecture.” *Journal of Economic Theory,* 39(1): 155-90.

[14] Hoppe, Heidrun and In Ho Lee. 2003. “Entry Deterrence and Innovation in Durable-Goods Monopoly.” *European Economic Review*, 47(6): 1001-1036.

[15] Maskin, Eric, and Jean Tirole. 2001. “Markov Perfect Equilibrium I: Observable Actions.” *Journal of Economic Theory,* 100(2): 191-219.

[16] Nahm, Jae. 2004. “Durable-Goods Monopoly with Endogenous Innovation.” *Journal of Economics and Management Strategy,* 13(2): 303-19.

[17] Sobel, Joel. 1991. “Durable Goods Monopoly with Entry of New Consumers.” *Econometrica,* 57(5), 1455-85.

[18] Stokey, Nancy L. 1981. “Rational Expectations and Durable Goods Pricing.” *Bell Journal of Economics,* 12(1): 112-28.

[19] Waldman, Michael. 1996. “Planned Obsolescence and the R&D Decision.” *RAND Journal of Economics*, 27(3): 583-95.

[20] Waldman, Michael. 2003. “Durable Goods Theory for Real World Markets.” *Journal of Economic Perspectives*, 17(1): 131-54.

6 Appendix

6.1 Appendix A - Formal Structure

We present the formal structure of the model. First, we define bundles and offers. Next, we define strategies and histories, present payoffs and formally define equilibrium.

(i) The Bundle Offer Structure.

Consider the feasible offer set for the seller in period τ. Let $\mathcal{P}_\tau \equiv \mathcal{P}(\{1, 2, ..., \tau\})$ 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

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

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 choose the higher price. This implies that buyers act as price takers and that the market does not respond to the actions of an individual buyer. An acceptance by a buyer is an element of the set $\mathcal{P}(\mathcal{P}_\tau)$.

Consider the maximal contiguous quality. For any $z \in \mathcal{P}_\tau$, define $M : \mathcal{P}_\tau \to \{0, 1, ..., \tau\}$ by finding the unique $m \in \{0, ..., \tau\}$ 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. For an arbitrary sequence of holdings z_τ, define $q_\tau = M(z_\tau)$.

(ii) Strategies and Histories.

A (pure) 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 $\mathcal{H}_\tau$ denote the space of all histories up through period $\tau - 1$, we have

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

Given an observed history, $h_\tau \in \mathcal{H}_\tau$, the seller's strategy specifies an offer $\omega_\tau = \mathcal{O}_\tau(h_\tau)$.

A buyer (pure) 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 $\mathcal{A}_\tau$ for the strategy profile across buyers.

We need to specify the history space $\mathcal{H}_\tau$. First, define $\Omega^\tau \equiv \Omega_1 \times \Omega_2 \times ... \times \Omega_\tau$; 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 $\mathcal{F}_\tau$ 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 $\mathcal{F}^\tau \equiv \mathcal{F}_1 \times \mathcal{F}_2 \times ... \times \mathcal{F}_\tau$.

Then the history space is specified by $H_1 = \varnothing$ 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. 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 $x \in [0,1]$. An equivalent, but less convenient, formulation would be to assign an index to each element in the finite set $\mathcal{P}(\mathcal{P}_\tau)$ and define measurability in the standard way for a real valued function.

(iii) Payoffs and Equilibrium.

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 $k = 1, ..., \tau$. These units are given by $Z_k(x) = \{i \in \{1, ..., k\} \mid i \in z \text{ for some } z \in \mathcal{A}_k^x(h_k, \omega_k)\}$, 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 $\{1, ..., \tau\}$, 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 $p_{\tau}(x) \equiv \sum\limits_{z \in \mathcal{A}^x_{\tau}(h_{\tau}, \omega_{\tau})} p_z$, 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 period:

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

The 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) $\sum\limits_{\tau=1}^{\infty} \delta^{\tau-1} \tau = 1/\left(1-\delta\right)^2$.

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 $z \in \mathcal{A}^x_{\tau}(h_{\tau}, \omega_{\tau})$. The Lebesgue measure of such buyers is $\alpha_{\tau}(z) \equiv \int\limits_{\mathcal{X}_{\tau}(z)} dx$. Thus, the revenue of the seller in period τ is

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

$\alpha_{\tau}(z) = 0$ must hold if the seller did not offer bundle z or if no buyer purchased 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.

When buyers are distributed across maximal holdings, the state is given by $(\tau,(Q^m_\tau)_{m=0,\ldots,\tau-1})$, where Q^m_τ is the set of buyers with maximal contiguous quality m. More generally, when buyers are distributed as $(Q^m_\tau)_{m=0,1\ldots,\tau}$, then the translated state is given by $(\tau+1-\underline{m}_\tau,(Q^m_\tau)_{m=\underline{m}_\tau,\ldots,\tau})$ where $\underline{m}_\tau$ is the smallest index of Q^m_τ with a non-zero measure.

6.2 Appendix B - Efficient Equilibria

We first prove Lemma 2 and then state several lemmas dealing with properties of the $T-stage$ support. This is followed by the proof of Proposition 1.

(i) Sufficient Conditions for Equilibrium

Proof of Lemma 2. 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. An upgrade offer need not include all of the feasible upgrade bundle levels $1,\ldots,\tau$. 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

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

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 seller payoff of $p_{b^*} + \delta u_{\tau+1-b^*}$ is, by (3), not profitable relative to p_τ.

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 $\alpha = (\alpha_0, ..., \alpha_{\tau-1})$ specifies for each $\sigma = 0, ..., \tau - 1$ 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 $\{(b_\sigma, p_\sigma)\}_{\sigma=0,...,\tau-1}$ where each (b_σ, p_σ) is an upgrade from σ to τ, that is the bundle $\{\sigma + 1, ..., \tau\}$, 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 bounds 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 $\{\sigma + 1, ..., b\}$ 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 $\sigma = 0, ..., \tau - 1$, the set $\mathcal{B}_\sigma$ of all upgrade bundles that move σ to a higher quality level; note $\mathcal{B}_\sigma$ 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 $u_\tau = u_T \equiv \overline{u}$ for $\tau \geq T$. If the support condition (3) holds at $\tau = T$ for $\sigma = 0, ..., T$, then (3) holds at $\tau > T$ for $\sigma = 0, ..., \tau$.*

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

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

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

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

(iii) Existence of Equilibrium

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

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 $u_{\tau+1} = \frac{1}{\delta}\left(u_\tau - v\tau\right)$. 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 $\tau = 2, ..., T-1$.

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

$$\frac{\delta}{1-\delta} u_1 \geq \frac{1}{\delta}\left[\frac{\delta v(T-1)}{1-\delta} - (T-1)v + \delta u_1\right]. \tag{6}$$

Simplifying and noting that $\delta \geq 1/2$, condition (6) 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

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

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

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

But,

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

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. Note that the cut-off sequence δ_τ is strictly increasing in τ, from $\delta_1 = 1/2$ to $\lim_{\tau\to\infty} \delta_\tau = 1$, and satisfies $\delta^{\tau-1} < 1 - \delta < \delta^\tau$ for $\delta \in (\delta_{\tau-1}, \delta_\tau)$. 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 $\left[\left(1-\delta^{T-1}\right) S_1, \left(1-\delta^T\right) S_1\right]$ intervals, where T ranges from 1 up to the index on the δ_τ root. ■

6.3 Appendix C - Inefficient Equilibria

First, we examine the approach conditions and then provide the analysis of the off equilibrium support. Finally, we prove Proposition 5.

(i) Approach Conditions.

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

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

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 $\sigma = 1, ..., \tau$ and $\tau = 1, ..., t-1$. 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 $p(\sigma, \tau) + \delta^{t-(\tau-\sigma)}\pi_t$. Hence,

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

for $\sigma = 1, ..., \tau$ and $\tau = 1, ..., t-1$. We then have

Lemma A5. If the buyer and seller approach conditions, (7) and (8), hold for $\sigma = \tau$, at each $\tau = 1, ..., t-1$, then the conditions hold for all feasible pairs (σ, τ).

As a result, we need only find $t-1$ distinct prices, $p(1,1), ..., p(t-1, t-1)$ 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.

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

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

for $\tau = 1,...,t-1$. Condition (5) in the text follows from (9) upon simplifying and substituting with $\pi_t + u_t = \Psi_t$.

Proof of Lemma 4. In (5), τ assumes integer values $1,...,t-1$. 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 (5) becomes

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

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

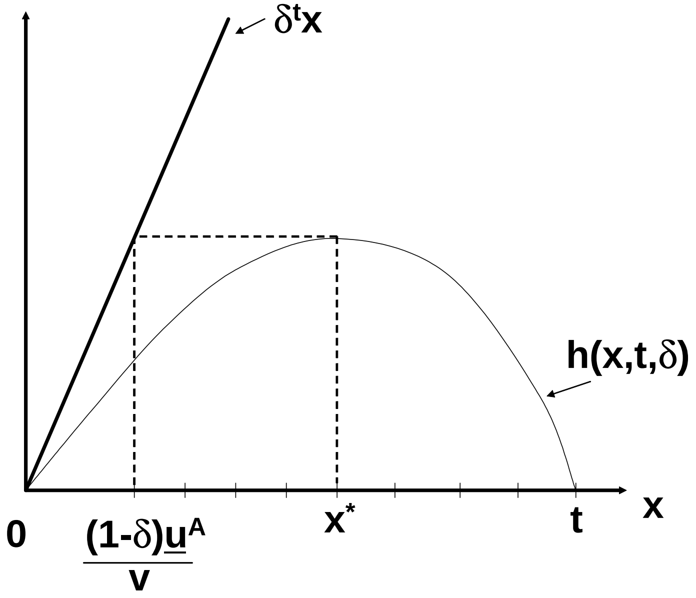

Figure 3: Payoff Lower Bound for Approach Conditions

the condition

$$\bar{u}^A \equiv \frac{v}{(1-\delta)}\left[\frac{t}{1-\delta^t} - \frac{1-\delta^t}{(-\ln\delta)\delta^t}\right] \geq u. \tag{10}$$

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

$$\frac{\delta^t(1-\delta)u}{v} \geq x - \frac{t(\delta^{t-x}-\delta^t)}{1-\delta^t} \equiv h(x,t,\delta). \tag{11}$$

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 (10) 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

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

where the function $a(\delta^t)$ is defined below. Then any $u \geq \underline{u}^A$ will necessarily satisfy (11). See Figure 3.

To find when (10) and (11) hold at a candidate u, we compare the bounds

$$\overline{u}^A \geq \underline{u}^A \quad \Leftrightarrow \quad \delta^t \geq a\left(\delta^t\right),$$

where we define $a(d) \equiv -\ln\left[-\frac{d\ln(d)}{1-d}\right]$ for $d \in (0,1)$. It is easy to show that $a\left(d\right)$ is strictly decreasing for $d \in (0,1)$ and, by L'Hospital's Rule, that $\lim_{d\to 1} a(d) = 0$. Then, the equation $d = a\left(d\right)$ has a unique root, d^* in $(0,1)$. Thus, we have established that $\delta^t \geq d^*$ implies the interval $\left(\underline{u}^A, \overline{u}^A\right)$ 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 $\overline{u}^A$ with $\frac{vt}{1-\delta}$ reveals that $\overline{u}^A$ 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 $\overline{u}^A < \frac{vt}{1-\delta}$ and any $u \in \left(\underline{u}^A, \overline{u}^A\right)$ satisfies the approach condition (5). For δ^t above this root, we have $\overline{u}^A > \frac{vt}{1-\delta} > \underline{u}^A$ and there are two cases. First, by the above analysis, any $u \in \left(\underline{u}^A, \frac{vt}{1-\delta}\right)$ satisfies the approach condition (5). Second, for the case of $u \in \left[\frac{vt}{1-\delta}, \overline{u}^A\right)$, the approach condition (5) 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. Then $A(x,u,\delta,t) \geq B(x,\delta,t)$ and utilities in this payoff range satisfy (5). ■

(ii) 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

$$\Psi_\tau - \delta\Psi_{\tau+1-\sigma} \geq u_\tau - \delta u_{\tau+1-\sigma} + g(\sigma, u_{\tau+1}) \tag{12}$$

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

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

$$u_\tau = \delta u_{\tau+1} + \Psi_\tau - \delta\Psi_{\tau+1} \text{ for } \tau = t, ..., T-1, \tag{13}$$

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 (12) 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,

Lemma A6 *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 (12) holds for every $\tau \geq t$ and $\sigma = 0, ..., \tau$.*

(iii) Equilibrium Existence

The following lemma provides sufficient conditions for equilibrium existence.

Lemma A7 *Suppose a sequence of buyer utilities u_τ satisfies (5) and (12) 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 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 $B(t, T, \delta) \geq \overline{u}^A(t, \delta)$, whenever δ satisfies the approach conditions, thus proving the proposition. Comparing,

$B(t,T,\delta) \geq \overline{u}^A(t,\delta) \Leftrightarrow$

$$\frac{(1-\delta^t)(1-\delta)}{\delta^T} \geq -\ln\delta \tag{14}$$

It is clear that both sides of (14) are positive, falling in δ and equal to 0 at $\delta = 1$. The l.h.s. of (14) is rising in T and t. Differentiating both sides of (14) 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

$$\delta^T + \delta^{t+1} \gtrless T(1-\delta) + \delta - (T-t)\delta^t(1-\delta) \tag{15}$$

In (15), 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 (14) cross a single time in δ. It easy to see that (14) 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 (14) 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, $b(t,T) > \delta_T$. ■