Compatibility, Interoperability, and Market Power in Upgrade Markets

James J. Anton
Duke University

Gary Biglaiser
UNC

April 2009

Abstract

We examine the market power of a seller who repeatedly offers upgraded versions of a product. In the case of pure monopoly, the seller also controls compatibility across versions. In the case of an entrant who offers an upgrade, the incumbent seller also controls subsequent interoperability across versions. We argue that control of compatibility and interoperability does not allow an incumbent seller to charge a price premium relative to when such control is absent and, consequently, neither is a necessary source of market power. (JEL Codes D42, L15, O32)

We thank the Fuqua Business Associates Fund and Microsoft for support, and Jacques Lawarree for helpful comments. All views expressed are solely those of the authors.

Introduction

Markets for upgrade goods have been the subject of intense scrutiny in recent years, the debate being fueled, in no small part, by the controversies surrounding the antitrust actions against Microsoft in the U.S. and Europe. Many commentators and economists have expressed concern regarding the market power of an incumbent firm who controls the underlying platform, whether for software or some other upgrade good. Specifically, the incumbent can offer upgrades and exploit either a lack of compatibility between product versions or a lack of interoperability with applications from other firms to gain market power.

Software is often the good that first comes to mind when thinking about upgrades. Computer operating systems from Apple and Microsoft have gone through many generations and upgrades in the decades since their introduction and a recurrent issue has been whether existing application programs will function with upgraded operating systems. Apple has emphasized the ability (and hence interoperability) of its Leopard operating system to support PC programs that are based on the Microsoft Windows operating system. Another prominent example comes from the market for enterprise resource planning software. The hostile takeover of PeopleSoft by Oracle triggered strong regulatory and customer concerns regarding upgrades and support of existing PeopleSoft products and interoperability between the two firms' products. By contrast, such concerns were much more muted when PeopleSoft acquired J.D. Edwards, a case where expectations about upgrades and interoperability were more positive.

Where interoperability refers to products from different firms, compatibility typically refers to the ability of products from a single firm, often different upgrade versions of a program, to function without conflicts. There are many familiar examples of applications programs, such as word processors, spreadsheets, and media players, where compatibility has been an issue. For example, the Adobe Acrobat Reader program will notify the user that some features are unavailable if the file was created for a later version of the program.

While compatibility and interoperability are familiar issues for upgrades in software markets, many capital goods markets also have a clear upgrade dimension. B-52 bombers are an interesting example -- the basic airframe has been in use since the 1950's but many of the critical systems such as avionics, engines and weaponry have been repeatedly upgraded and the planes are expected to remain in use for several decades to come. The evolving nature of upgrades in camera markets illustrates how compatibility and interoperability issues can persist in the face of major technology changes. Camera makers have long been accused of making it difficult for their camera bodies to work with lenses from a competitor. From the 1930s through the late 1980s this took the form of patenting the attachment mechanism. Since that time, technology has shifted and the camera-lens interface now involves digital communication protocols. The camera body makers, however, have again been criticized regarding interoperability, this time for not fully disclosing the necessary information. With respect to incompatibility, razors and blades across product generations are a classic example.

A common economic feature of software and these other products is that the buyer is purchasing a durable asset and then consuming a flow of services. While there are a number of examples of successful upgrade introductions, where buyer adoption gains momentum as later buyers are pushed to follow initial adopters of the upgrade, it is also possible to identify prominent failures. This leads us ask whether the control of compatibility/ interoperability (C/I) implies not only an increased likelihood of adoption of an upgrade, but also the ability to charge higher prices. If so, then control is necessarily a valuable option for the incumbent and it implies an ability to earn higher profits.

Imagine a market in which products are always C/I. Consider the outcome in such a market and suppose that the seller offers upgrades at a price p and that all buyers accept the offer. Thus, buyers upgrade on a regular basis and p divides the surplus between the seller and buyers. Now, suppose that the seller can control whether versions are C/I with each other. Is this a valuable tool for the seller, either because products will be made incompatible in practice or because the mere threat of doing so supports higher prices? If so, then the market outcome without C/I control must be vulnerable to the seller's ability to control C/I. In particular, the seller should be able to successfully increase the price from p. Without C/I control, buyers must reject any price above p, and this is what keeps the seller from raising prices. Thus, we need to assess whether buyers would necessarily accept a price above p when the seller has C/I control. If so, then we have a clear measure of increased market power: whatever the price was without C/I, a seller with C/I control will always be able to charge a higher price without sacrificing adoption.

We argue that this is not the case and that C/I control implies no such ability to command higher prices. Stated a bit differently, this means that the original price of p will remain a stable market outcome even when the seller has C/I control. Thus, if the seller were to raise the price and exploit a C/I option, then buyers must be willing to refuse the offer. A key aspect of the credibility of this refusal is the ability of the seller to tempt an individual buyer to purchase when others are not. It is on this dimension that C/I control fails to deliver a wedge that alters an individual buyer's evaluation of an offer and, consequently, the original market outcome is robust. Thus, neither incompatibility nor non-interoperability is a necessary source of increased market power for a monopoly seller of upgrades.

Paradoxically, market power is not driven by an individual buyer's fear of falling behind the market. As we explain below, it is perfectly consistent for individual buyers to have a very strong incentive to keep up with the market and for the seller to lack market power. In fact, we find that a lack of C/I only reinforces the incentive to keep up with the market. Instead, the source of market power lies with altering the incentive to adopt an upgrade even when others are not.

The next section describes our basic assumptions concerning upgrade markets. We then turn to the argument for compatibility and buyer willingness to pay for an upgrade. Finally, we extend the argument to settings with interoperability and demonstrate that C/I control does not ensure an increase in market power even when an incumbent monopolist of a platform good can offer a perfect substitute for a competitor's add-on good.

Basic Assumptions and Upgrade Market Structure: Compatibility

We examine a dynamic market in which a monopolist generates upgraded versions of a good in every period (e.g., annual upgrades). In response to each seller offer, buyers decide whether to remain with an old version of the good or to buy a new version (upgrade). A canonical example is that of software programs such as word processing, digital photo processing, and spreadsheet analysis. In all these cases a buyer can stay with their current version or upgrade when the seller offers a new version. Software programs that do not fit our framework are those with limited lifetimes, such as tax preparation software, where the previous year's version is rendered obsolete by tax code changes.

We begin by considering the seller's commercialization options for quality upgrades. At each point in time, the seller can offer versions of the good that vary in quality, from low to high, up to the current state of the art. As time progresses, new innovations become embodied in higher quality versions. Moreover, the seller can offer any given version at a price that depends on the buyer's current version of the good. This structure is common for software programs and other upgrade goods. For example, many software packages including Microsoft Office, Adobe Acrobat, and Scientific Word are offered on this basis. As a case in point, Adobe offers its Acrobat program in three versions, Standard, Pro, and Pro Extended; each version then has an upgrade price for existing customers and a "full price" for new buyers. Current users can buy Acrobat 9 Pro Extended for 229 dollars while a new user must pay 699 dollars.

Our multi-version, multi-price framework for a seller's offering is quite flexible.[1] It does, however, reflect a choice of the seller and is not a primitive aspect of the upgrade good. In this regard, we note that there are at least two underlying structures that lead directly to the multi-version, multi-price framework.

The first case is the simpler of the two. Suppose that each version is 'unbreakable' in the sense that it incorporates all prior innovations. Thus, the versions are nested with regard to quality. This allows, for instance, that one version is a lower capability program while another, by incorporating additional innovations, is the current state of the art product. Then, as long as the seller can condition a purchase of any version on the buyer's purchasing history, we arrive at a multi-version, multi-price framework.

---

[1] If the seller must sell all upgrades in a package and cannot price according to prior purchases, then the seller's market power is necessarily limited (see Fishman and Rob (2000)).

For the second case, imagine that each innovation is embodied in a distinct good (breakable upgrades) but that they form a cumulative technological sequence in which the latest good will not work unless all prior goods are in place (downward complementarity). [2] The seller can then implement a multi-version, multi-price offering and effectively charge different prices to buyers with different current holdings by relying on downward complementarity: no buyer who lacks previous units will buy an advanced good without also purchasing the missing units. A buyer with the missing units, however, will be willing to purchase the advanced good by itself in a number of circumstances.

Software often has elements of each of these scenarios. Succeeding generations of an operating system illustrate the first case. The second case, which is a bit more subtle, arises whenever an innovation can be implemented in conjunction with an existing program. Many familiar examples illustrate this point. For a number of years users of word processors, such as MSWord 2003 and Scientific Word 3.0, used Adobe Acrobat to make versions of their documents in PDF format files (note that the issue of interoperability across competing products arises here). Currently, MSWord 2007 and Scientific Word 5 generate PDF format files internally. Thus, while innovations can be bundled into a larger program or remain a separate program, each case can be accommodated by the seller's commercialization options in a multi-version, multi-price framework.[3]

Now, we turn to assumptions on buyer preferences. In order to concentrate on the basic economic structure of a dynamic upgrade market, it is instructive to focus on settings where the added complexity of price discrimination across buyer types is absent.[4] Thus, we assume that all buyers have identical preferences, as a result buyers would only choose different seller offers if their current versions differed. In a static market, with only a single quality of a good, this is the familiar idea of a horizontal demand curve. A profit-maximizing monopoly seller then has a very simple solution to the pricing problem -- set price equal to the willingness to pay, which is the same for all buyers. Since all buyers purchase the good the outcome is efficient and social surplus is maximized. However, since the entire surplus accrues as profit, the market power of the seller is maximal in this simple static setting. Moreover, as shown by Fudenberg, Levine and Tirole (1985) if there is a single durable good of unchanging quality, the dynamic outcome is

[2] This is a standard assumption in the economics literature on upgrade goods see, for example, Fudenberg and Tirole (1998), Ellison and Fudenberg (2000), and Fishman and Rob (2000).

[3] Other software examples include the use of graphics programs and spreadsheets in conjunction with other programs. For examples outside of software, note that the fore mentioned B-52 example fits the second case, while televisions with respect to screen size fit the first case.

[4] As discussed above, offering high, medium, and low quality versions corresponds to a textbook form of second-degree price discrimination. It is also common for sellers of software to offer commercial, academic, and student versions often requiring identity verification of the buyer, much like the textbook form of third-degree price discrimination.

very simple. There is no delay and all buyers purchase the good immediately, paying a price equal to their willingness to pay, the full social surplus.

It is tempting to conclude that with homogeneous buyers this extreme form of market power for the seller will necessarily carry over to a dynamic upgrade market. Thus, one might conjecture that the seller will offer upgrades at prices that capture the entire social surplus. But the economic structure of an upgrade market is fundamentally different from a single durable good environment -- social surplus increases as upgrades become available and this transforms the buyer-seller relationship from a “zero-sum” situation to one where a larger surplus must be divided in the event of delay.

Given identical buyers, let us consider the flow utility (value gross of price) that a typical buyer receives in a period based on their holdings and the holdings of other buyers. We divide the flow utility into intrinsic stand alone and network components. The stand alone value can exhibit compatibility effects and these can operate in both the forward and backward directions. At this point, it is standard practice in the software industry for programs to be backward compatible – the files one created in the old version will work with the new version; we follow others, including Ellison and Fudenberg (2000), and make this backward compatibility assumption. To address the issue of incompatibility, we will assume some lack of forward compatibility – files created with the new version have conflicts with the old version.

Network externalities, generated by direct or indirect effects arising from the added value to an individual when other buyers also use a product, amplify the impact of incompatibility and interoperability on the flow utility to a buyer. Thus, a buyer who is "ahead" of other buyers suffers no direct network loss from being ahead (backward compatibility) as for instance when files from other users are needed. By contrast, a buyer who is "behind" the market, may receive a lower value from his holdings, a compatibility loss, as when a file from an advanced buyer is unusable. Thus, the possibility of forward incompatibility with newer versions of the product combines with the network effect to lower the consumer's utility from interactions with buyers who have a more advanced version of the good.

An example may help to frame the issues regarding market power, incompatiblity, and the willingness to pay for the buyers. Suppose that a monopoly seller of a software product offers annual upgrades when they first become available. To keep things simple, let the typical buyer value the incremental surplus over each of the coming years from each upgrade at 1 dollar. Further, set all seller costs and buyer adoption costs (except the product's price) to zero. For a tangible example, we can imagine a word processing program that adds a new feature each year such as spell checking, enhanced printing capabilities for internal pdf generation, web access, and so on.

The efficient outcome would be for all buyers to upgrade each year. The joint seller and buyer surplus generated is then the full present discounted value (PDV) of the upgrade good. A seller with perfect or complete market power would be able to capture the full present value of each upgrade. For instance, with a flow value of 1 dollar, the full present value is approximately 10 if the interest rate is 10%. At the other extreme, a price of 1 provides a lower bound on the market power of the seller: it is clearly a dominant strategy for an individual buyer (negligible in the overall market) to purchase at any price up to the flow value independent of any future considerations, since the upgrade will pay for itself in one period.

What market price, from the low of the one period flow value to the high of the full PDV, will the monopoly seller be able to charge for each upgrade? In other words, what constitutes the appropriate notion of willingness to pay on the part of buyers? Willingness to pay is central to the notion of a credible threat to refuse an offer from the seller. That is, if no individual buyer is willing to refuse an upgrade offer at a price below the full PDV, then the market price must be the full PDV: the seller will have perfect market power and be able to capture the full surplus generated by each upgrade. On the other hand, if buyers do have a credible threat of rejecting some offers at prices below the full PDV, then the seller will not have complete market power. In fact we will argue that buyers do have a credible threat in an upgrade market and that, in turn, the market power of the seller may be limited to charging a price as low as the one period flow value of each upgrade. In general, any price between the one period flow value and the full PDV can be supported as an equilibrium price. Furthermore, since making goods incompatible will only lower a buyer's willingness to move ahead of other buyers in the market when network effects are present, it will not improve a seller's profitability.

The Argument for Willingness to Pay

How could it be that control of compatibility does not guarantee the ability to successfully raise prices? A natural intuition is that compatibility issues lead an individual buyer to be fearful of falling behind the market and this is the key element to increasing willingness to pay and hence high prices. A more subtle intuition is that it is the opposite position in which an individual buyer moves ahead of other buyers that is the key element in the market power argument. In both cases, however, we need to understand buyer expectations regarding the consequences of a buyer falling behind or moving ahead of the market. We argue that it is the nature of the second position, in particular the ability of the seller to tempt a buyer to move ahead of the market, that explains why platform control does not imply increased market power.

We consider market adoption paths that satisfy two properties. First, we assume that there is no persistent inefficiency in any continuation play. This property has implications both on and off the market equilibrium path. We focus on equilibrium outcomes that follow the efficient outcome path (see above) for the adoption of upgrades -- buyers always end up moving to the "state of the art" version. However, a critical question for market power involves expectations for what would happen if the market were to veer off of this path. For example, suppose the seller offers an unexpectedly high price and some or possibly all of the buyers refuse and hence fail to acquire the latest upgrade. Then, in the continuation play following any such episode, no persistent inefficiency means that for any distribution of buyer holdings of upgrades, the off-the-equilibrium continuation outcome has all buyers moving immediately up to the state of the art (this continuation support can be justified as in Anton and Biglaiser (2008)).[5]

The second property specifies what happens to an individual buyer who either falls behind or jumps ahead of other buyers and, thus, becomes "out of sync" with the rest of the market. If a buyer falls behind the market, then we specify that the buyer receives no increments in surplus from any market purchase in the future. Such a buyer might make purchases, but the price extracts the value of catching up with the market. In practice, for example, Adobe Acrobat owners can buy Adobe Acrobat 9 Pro Extended for 229 dollars while "new buyers" pay the new higher price of 699 dollars. Observe that this property can only enhance a seller's market power - a buyer is fearful that if he does not buy when others do then he will lose out on any future net benefits from participating in the market, and this makes the incentive to keep up with the market as strong as possible. We note that this refers to an individual buyer's choice relative to overall market and not the movement of the market itself.

If, instead, a buyer is ahead of the market, then a seller's offer that returns the mass of buyers to the efficient adoption path will also attract a buyer who is ahead of the market. That is, we assume that there is sufficient future surplus that this buyer will want to keep up with the rest of the market and avoid letting the market pass them by. Since a buyer who falls behind the market obtains no further surplus, this property is conservative with respect to the required future surplus for a buyer who jumped ahead of others to keep up with the market.[6]

We are now ready to proceed with the main argument regarding willingness to pay, compatibility, and market power.

Suppose, as required by efficiency, that all buyers enter the current period with the prior state of the art (they have all previous upgrades). At this point, an efficient outcome implies that the seller is supposed to offer the latest upgrade for some price p and all buyers are supposed to accept the offer. Imagine that p falls short of the full present discounted value of the upgrade. Suppose that, instead of p, the seller offers a higher price for the upgrade. What should a buyer do? Buyers who lack a credible threat will necessarily acquiesce to the higher price and the seller

[5] Anton and Biglaiser (2008) also consider the possibility of delay. In the event of delay, efficiency is necessarily compromised and with it market power of the seller, since delay reduces the surplus available.

[6] We are implicitly assuming that a buyer who is ahead of the market is not discriminated against based on holding a more advanced version and can purchase the same package the seller offers to other buyers. For example, any 2000-2007 Office program or suite qualifies a buyer to purchase Office Professional at the upgrade price. If, instead, the seller did price discriminate against a buyer who is ahead of the market, then that buyer would obtain even less future surplus and this would only strengthen our willingness to pay argument.

will inevitably capture the full surplus from buyers. We claim, however, that the buyers do have a credible threat to refuse the higher price.

To understand the idea of a buyer's credible threat, suppose all other buyers will reject the offer and consider whether an individual buyer will reject the offer even if it is only for a price slightly above p. An individual buyer would then find himself in the position of deciding whether to accept, moving ahead of other buyers, or reject. In the following period, the continuation outcome for the situation where the market has fallen behind the state of the art, a temporary inefficiency, is that all buyers acquire the latest upgrade along with any missing previous upgrades. So, a buyer who moves "ahead" of the market by making a purchase when others do not will in turn be confronted with the decision about whether to keep up with the market when other buyers move back up to the state of the art (return to the efficient path). If he does not purchase, then he will fall behind the market and as noted above this will preclude any future surplus increments. On the other hand, if the buyer keeps up with the market and makes the purchase, he will resume a path that is identical to that of other buyers. Provided that the utility increment from keeping up with the market is sufficient, as we hypothesized above, then the buyer will make the purchase.

Thus, and this is crux of the argument, the most that the buyer would ever be willing to pay to move ahead of the market is exactly the single period flow value from the upgrade. Intuitively, not purchasing generates an expectation that the missing upgrade will be acquired next period. Purchasing thus provides only a single period of additional flow value, since the buyer expects to rejoin the efficient path next period along with the buyers who refused the offer.

To complete the argument, return to the initial attempt by the seller to increase the price above p. If the initial price p is above the single period flow value, then the buyer should reject the offer. As we just argued, the buyers do have a credible threat to reject a seller's offer when the seller attempts to raise his price above the equilibrium price. Given that other buyers are expected to reject the price increase, an individual buyer will also reject and the seller will be unable to tempt a buyer into moving ahead of the market.

How does compatibility affect the willingness to pay of buyers and the seller’s market power? Suppose the product becomes less (forward) compatible. A buyer who falls behind the market will experience a diminishing utility level as the market leaves them further behind and will thus pay a higher price to reestablish parity with the market. Thus, the cost of falling behind is larger. The effect on a buyer who moves ahead of the market is more subtle. Initially, there is no effect, because the product is backward compatible. However, when the market subsequently returns to the efficient path the prospect of not keeping up is now less attractive. Thus, paradoxically, making the product less compatible will actually reduce the need for as large a future surplus increment to justify keeping up with the market. As a result, the seller is unable to exploit his control of compatibility to affect an increase in willingness to pay and hence market power.

A similar conclusion applies to the extent of network externalities and willingness to pay. The main impact of stronger network effects is on a buyer who falls behind the market and such a buyer will pay a correspondingly larger amount to catch up with the market. For a buyer who moves ahead of the market, however, there is no direct effect. As with compatibility, the only effect is to reduce the amount of future surplus such a buyer will subsequently need to justify keeping up with the market. Thus, our argument regarding willingness to pay when others are not purchasing is robust to the extent of network effects.

Since we chose an arbitrary price when discussing a buyer's willingness to pay when other buyers refuse to pay, any price that is at least the single period flow value but less than the full PDV can be an equilibrium price. Distinguishing a buyer’s willingness to pay relative to the behavior of other buyers is crucial in this argument. If others are purchasing at a price p, then the individual is also willing to pay p in order to keep up with the market; falling behind the market, which generates no future surplus increment, is the relevant concern when others are buying. It is when others are expected not to purchase that the individual willingness to pay drops to the flow value of an upgrade. Moving ahead of the market and then confronting the subsequent decision to go to the state of the art with the market takes over as the relevant concern when others are not buying.

This argument has direct implications for the important question of how innovation, measured here by the frequency of upgrade offers, affects the market power of the seller. Clearly, without the prospect of a subsequent upgrade, a buyer’s concern about moving ahead of the market vanishes. As a result, the willingness to pay logic breaks down and there is no longer a credible threat for buyers to refuse any offer that leaves them with a positive surplus, no matter how small that surplus might be. Thus, the total absence of subsequent innovation restores full monopoly power to the seller.[7] By contrast, any ongoing innovation, independent of whether it is rapid or slow, makes the flow value of an upgrade pivotal and supports the willingness to pay argument.

---

[7] This limiting case of an upgrade market where innovation eventually terminates corresponds to the economic structure in a standard durable goods model, in which there is only a single good. With identical buyers, the seller faces a static horizontal demand curve and the “speed-up” logic of Fudenberg, Levine, Tirole (1985) will apply. First, without the need to purchase more than once, a buyer can always be tempted into paying a premium to purchase today: For any expectation of purchasing at a price p in the future, a buyer will pay slightly more than p to acquire the good today due to the added flow value. Then, given that all buyers purchase immediately, buyers have no credible threat for refusing an offer that leaves them any positive surplus. As a result, the only equilibrium has the seller charging a price equal to the full surplus and the market price is at the theoretical maximum. With an upgrade good, buyers know that they will be in the market repeatedly and the option to purchase again is necessary to support a credible threat to reject a seller's offer.

Interoperability

Interoperability involves three additional elements relative to a pure compatibility setting. First, we need to allow for a competing firm. Competitors often enter a market by offering a substitute good to compete directly with the incumbent firm's product. However, in order to focus on control of interoperability by the incumbent, it is more appropriate to consider innovation by the competing firm. Thus, the second element is that the entrant introduces a new product that complements the existing platform of the incumbent. For example, Adobe Acrobat works in conjunction with documents generated by existing products to produce PDF versions. More generally, any software application program must work in conjunction with the underlying operating system. The third element is that the incumbent has control over interoperability.

We incorporate these three elements with a very simple modification of the monopoly and compatibility framework. Figure 1, where I stands for incumbent and E for entrant, provides a summary of the market timeline. Initially, the incumbent introduces and upgrades the platform. To keep things simple, the entrant only has a single opportunity to offer the complementary good. After this, the incumbent returns to the prior path of regular upgrade offers, but with two additions. First, the incumbent can decide whether to offer a competing version of the entrant's product. Second, the incumbent decides whether to make subsequent upgrades interoperable with the entrant's product. Except for C/I issues, the entrant's product and the incumbent's version are perfect substitutes.

Figure 1

Past Years	T	T+1	Future Years
I innovates and upgrades	E innovates and offers product	I innovates, imitates E, and chooses operability	I innovates and chooses operability

As with compatibility, we focus on the question of market power as a necessary consequence of control of interoperability. Thus, imagine that in the absence of control the incumbent sells upgrades at a price p and that the entrant offers their product at some price q. For this starting point we suppose that all products are fully interoperable. Note that the incumbent's imitation of the entrant's product commands no premium- it is bundled with the incumbent's current innovation in T+1 and together they sell for p. This is because of full interoperability. In this scenario we further suppose that the incumbent's price p is low, near the flow value of the good.

Now, we confer on the incumbent full control over interoperability. Does this allow the incumbent to charge a price above the old price of p? Suppose that after the entrant has introduced her product (in period T), the incumbent then offers (in period T+1) his upgraded platform, which includes the new innovation and the imitation version, and this is non-interoperable with the entrant's product. Thus, when buyers purchase the incumbent's upgrade in T+1 it must be that the entrant prices its product at flow value. This is because interoperability will not persist and the product only has value for one period. The real question, however, is how much buyers will pay for the incumbent's upgrade in period T+1?

We argue that buyers need pay no more than p, the original price. That is, control of interoperability simply destroys the entrant's ability to command a price premium without creating a necessary ability to force buyers to pay the incumbent anymore than before. The logic follows the same basic outline as with compatibility and the critical element is buyer expectations. Specifically, if the incumbent tries to charge a price above p, then all buyers rationally refuse the offer based on the expectation that the incumbent will offer a subsequent upgrade that does induce buyers to move to the new state of the art. At that point, the upgrade will include two rounds of innovation, the imitation version of the entrant's product, and will be non-interoperable.

The position of an individual buyer who expects other buyers to refuse the incumbent's attempted price increase is remarkably similar to that in the compatibility scenario. By purchasing and moving ahead of the market, the individual buyer is immediately confronted with a decision to keep up with the market when the incumbent subsequently moves buyers to the new state of the art. We argue that the low price of p will induce the individual buyer to keep up with the market. This is exactly the scenario that supported low prices under compatibility, but the comparison is more involved due to the entrant's product and the interoperability issue.

To make the comparison, we first describe the market path following the incumbent's attempt to raise price from p in period T+1. At this point, buyers refuse the offer and continue to use the incumbent's previous platform and the entrant's product. The expectation for the future is that, given the failed attempt to raise prices, the incumbent will make an offer next period that does attract these buyers. As a result, the entrant's product will be rendered non-interoperable and buyers will be on the incumbent's platform using the incumbent's version of the entrant's product as well as the two rounds of incumbent upgrades.

Let us contrast this buyer experience with that of an individual buyer who accepts the incumbent's offer with an increased price in period T+1, even though all other buyers are expected to refuse the offer. By accepting, this individual buyer acquires the incumbent's version of the entrant's product as well as the incumbent's new T+1 upgrade. Relative to the market path, the only difference in value for such a buyer is the T+1 upgrade. Non-interoperability makes the entrant's product worthless, and the incumbent's version simply replaces this loss. Now, consider the decision of this buyer in the following period when the incumbent makes an offer that attracts all other buyers and returns the market to the state of the art. To be sure, this buyer is starting from a higher base level and has less to gain from keeping up with the market than the buyers who refused the initial price increase. On the one hand, falling behind the market guarantees no future surplus increments. On the other hand, the gain from keeping up with the market includes not only the incremental surplus from current upgrades but also that of future upgrade offers. When the gain large, as is necessarily the case when p is low, this is sufficient to induce an individual buyer who is ahead of the market to purchase along with other buyers. Thus, the only difference in value for an individual buyer who jumps ahead of the market by accepting the incumbent's price increase is that they acquire the incumbent's T+1 upgrade one period before the other buyers. The conclusion, just as with compatibility, is that no buyer will pay more than the flow value of the upgrade. Hence, the incumbent is unable to exploit control of interoperability to raise prices.

Another way to see the limitations of interoperability control is to eliminate the incumbent's ability to offer a version of the entrant's product. By choosing non-interoperability following the entrant's product offer, the incumbent can limit the entrant's price to flow value, but, lacking a version of its own, the incumbent cannot regenerate the surplus associated with the entrant's product. Thus, in this case, non-interoperability guarantees an inefficient outcome. As before, however, limiting the profits of the entrant does not translate into increased profits for the incumbent.

Conclusions

A common concern in upgrade markets is that incumbent control of compatibility and interoperability leads to high prices and market power for the seller. In this paper, we have examined the added value of such control in terms of the ability of an incumbent firm to raise prices relative to their levels when such control is lacking. We argued that market power, as measured by the ability to raise prices, is not linked to control. Buyers do have a credible threat to reject high prices when innovation is an ongoing process and this threat is robust both to a lack of compatibility across versions of an incumbent's product and to lack of interoperability with respect to competitor's product.

References

Anton, James J. and Gary Biglaiser (2008), "Quality, Upgrades, and (the Loss of) Market Power in a Dynamic Monopoly Model," working paper.

Ellison, G. and D. Fudenberg (2000), "The Neo-Luddite's Lament: Excessive Upgrades in the Software Industry," *RAND Journal of Economics*, 31, 253-272.

Fishman, A. and R. Rob (2000), "Product Innovation by a Durable-Good Monopoly," *Rand Journal of Economics*, 31, 237-252.

Fudenberg, D. and J. Tirole (1998), "Upgrades, Trade-ins and Buybacks," *RAND Journal of Economics*, 29, 235-258.

Fudenberg, D., D. Levine, and J. Tirole (1985), "Infinite Horizon Models of Bargaining with One-Sided Incomplete Information," in *Game Theoretic Models of Bargaining*, Ed. by A. Roth, Cambridge: Cambridge University Press.