Economic Fundamentals of the Digital Economy:

The Economics of Markets with Network Effects

Constructing a demand curve for a product with network effects (1)

Suppose that there are two types of consumers, denoted *L* and *H*, and type *H* consumers value the good more.

When a total of *Q* consumers use the product,

- $V_H(Q)$ denotes the value of the product for a type *H* consumer
- $V_L(Q)$ denotes the value of the product for a type *L* consumer

As the product has network effects, $V_H(Q)$ and $V_L(Q)$ are both increasing with *Q*

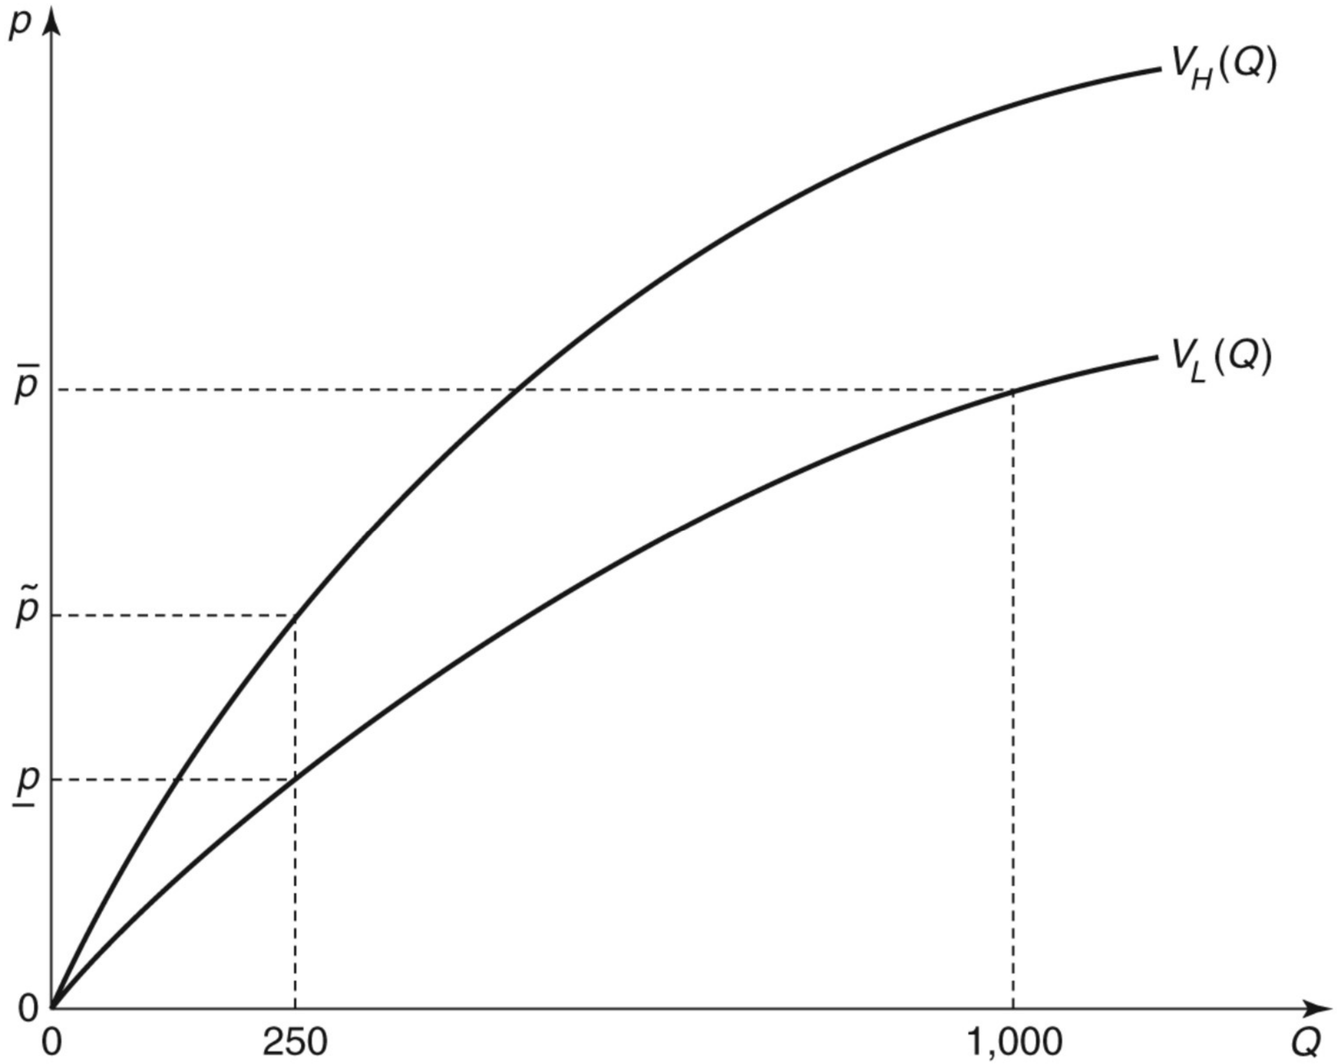

p
$V_H(Q)$
$V_L(Q)$
$\bar{p}$
$\tilde{p}$
$\underline{p}$
0
0
250
1,000
Q

Constructing a demand curve for a product with network effects (2)

- Assume a population of 1,000 consumers, of which
  - 250 are type *H* consumers
  - 750 are type *L* consumers
- Demand can take three possible values —0, 250, and 1,000— which correspond to no one buying, only type *H* consumers buying, and everyone buying, respectively (note that if a type *L* buys, then a type *H* buys, too)
- Demand is zero if price is sufficiently high. In particular:
  - If price exceeds $V_L(1{,}000)$, then no type *L* consumer will buy, as their value is below the price even when network effects are maximal
  - Furthermore, if only type *H* consumers buy, then a type *H* consumer is only willing to pay $V_H(250)$, which is less than $V_L(1{,}000)$ Thus, if price exceeds $\bar{p} = V_L(1{,}000)$, then demand is zero.

Constructing a demand curve for a product with network effects (3)

What if price is less than $\bar{p}$? Three cases:

1. *Q = 0*
   - if each consumer expects $Q = 0$, then, since $V_H(0) = 0$ and $V_L(0) = 0$, they are not willing to buy at any positive price
   - Thus, there is a vertical segment at $Q = 0$ that indicates that demand can be zero regardless of price
2. *Q = 1000*
   - If all consumers expect $Q = 1000$, then type L (and therefore type H) consumers will buy as long as $p \leq \bar{p}$
3. *Q = 250*
   - For that to occur, price must satisfy $V_L(250) \leq p \leq V_H(250)$ so when $Q = 250$ is expected, type H consumers want to buy but type L consumers do not. This holds when $\underline{p} \leq p \leq \tilde{p}$

Summing up, when price exceeds $\bar{p}$, demand is zero. When price is between $\tilde{p}$ and $\bar{p}$, demand is 0 or 1,000. When price is between $\underline{p}$ and $\tilde{p}$ demand is 0, 250, or 1,000. And when price is less than $\underline{p}$, demand is once again 0 or 1,000

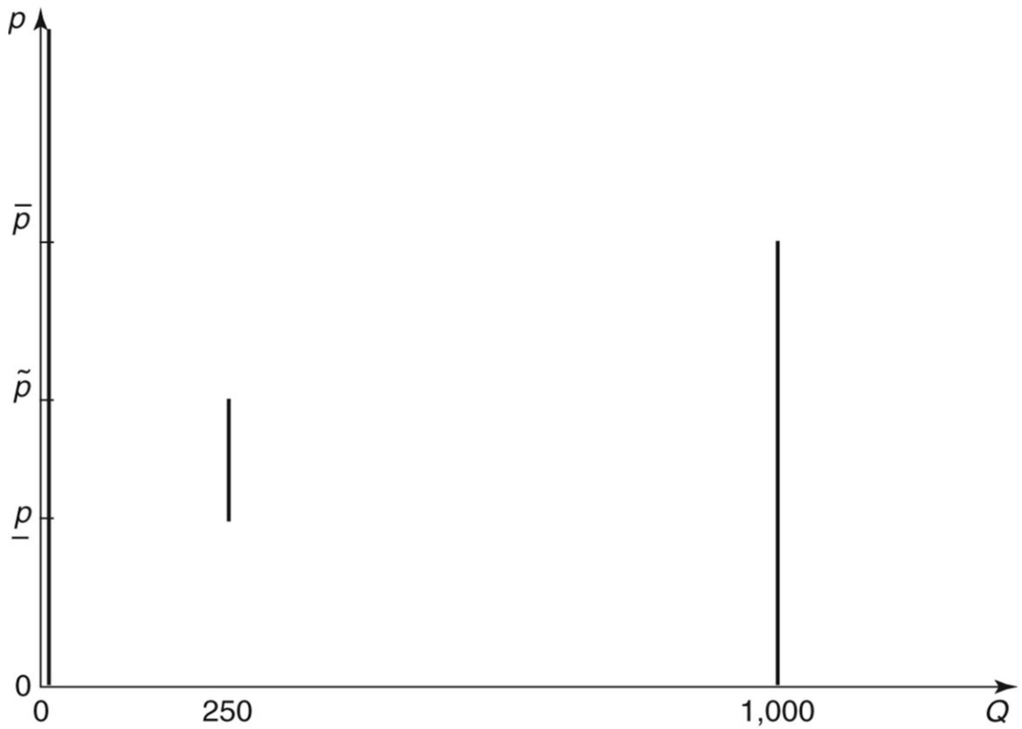

N.B.

demand need not be higher when price is lower. For example, if price lies between $\underline{p}$ and $\tilde{p}$ and consumers only expect type H consumers to buy, then demand is 250. If price is between $\tilde{p}$ and $\bar{p}$ and consumers expect everyone to buy, then demand is 1,000.

Raising price changes consumers' expectations (and their belief as to the magnitude of network effects), what leads to higher demand at a higher price

Fundamental properties of markets with network effects: Consumer expectation and critical mass

- Consumer expectations about the popularity of a product matter in determining demand, and a critical mass of consumer support can be instrumental in the success of a product with network effects
- Suppose:
  - a firm must get all 1,000 consumers to buy in order to achieve profitability;
  - consumers base their decisions on the size of the installed base, that is, the number of consumers who previously purchased and thereby are using the product.
- If the firm is able to induce the type H consumers to buy, so that initially $Q = 250$, it can induce the type L consumers to buy by setting price below $V_L(250)$
- However, if it cannot get that initial mass of consumers to buy, there is no price that will induce any consumer to buy
- At work is a mechanism known as *positive feedback*: the more consumers that buy, the easier it is to induce additional consumers to buy. The trick is getting the critical mass to jumpstart the positive feedback process

Fundamental properties of markets with network effects: it is natural that market dominance emerge

- Due to positive feedback, there can be *tipping*
  - The firm with a higher installed base will offer a more attractive product because of network effects and this will attract consumers at a higher rate than rival firms
  - Once a firm gets enough of a lead in its installed base, it goes on to dominate
- The initial firm in an industry can persist in being dominant (so called *sustained dominance*), even when superior products come along, because it has a large installed base to offset a rival's superiority
  - Suppose all 1,000 consumers are using the current technology at the price p, which means the net surplus to type H consumers is $V_H(1000) - p$ and to type L consumers is $V_L(1000) – p$
  - suppose a superior technology comes along with value $V_H^{new}(Q) > V_H(Q)$ and $V_L^{new}(Q) > V_L(Q)$ for all Q. Hence, for the same number of users (and price), all consumers prefer the new technology
  - if $V_H^{new}(0) < V_H(1000)$ and $V_L^{new}(0) < V_L(1000)$, then (at the same price) no consumer would want to switch technologies if she thought the other consumers would remain with the old technology
  - Unless the new technology is priced sufficiently lower than the current technology (which may make it unprofitable, especially since the existing technology can also lower its price) or consumers manage to somehow coordinate a shift, this new and superior technology will fail

An example of *sustained dominance:* IBM vs Microsoft

- consider when IBM introduced OS/2 in 1987 as an OS for the personal computer. At the time, the dominant OS was sold by Microsoft
- IBM spent about $2 billion developing OS/2, which was generally considered to be superior to Microsoft's OS.
- In spite of its technological appeal and that IBM was the largest computer manufacturer in the world, OS/2 was a failure
- IBM never convinced personal computer manufacturers and consumers that it would become popular
- Consumers did not buy it because they did not expect others to do so, and if not enough people bought it, then there would not be much software written for it.

Network effects and optimal price decisions: Building the installed base of customers

- Consider two firms competing in a new market with network effects

Hypotheses:

- The value that a consumer attaches to the product of firm 1 (2) depends on the installed base of firm 1 (2), where the installed base is the set of consumers who bought in the past and are currently using the product
- New consumers are flowing into the market each period, and their purchasing decisions depend on firms' prices, the characteristics of firms' products, and firms' installed bases. A consumer is more likely to buy a firm's product when its price is lower and network effects are stronger
- Let $D_1 (p_1, p_2, B_1, B_2)$ be the demand function of firm 1, and $D_2 (p_1, p_2, B_1, B_2)$ be that of firm 2
  - p_i is the price of firm *i*
  - B_i is the (current) installed base of firm *i*
- The demand function $D_1 (p_1, p_2, B_1, B_2)$ is decreasing in its own price p_1, increasing in its rival price p_2, increasing in its own installed base B_1 (so network effects are stronger), and decreasing in its rival's installed base B_2. Analogous properties hold for firm 2's demand function
- With no network effects, firm 1 would choose its price to maximize its profit $(p_1 - c_1)q_1$, where c_1 is its unit cost, and q_1 is how many units it sells (and depends only on p_1 and p_2)

Network effects and optimal price decisions: dynamic considerations

- The presence of network effects inserts a dynamic consideration into a firm's pricing decision:
  - The price it sets today influences how many consumers buy today, which affects the installed based it will have tomorrow and thereby affects its demand tomorrow
  - Selling more in the current period will then raise demand and profit in the next period and, in fact, every period thereafter, as all those periods will have a higher installed base
- We summarize the effect of the installed base on a firm's future profits with the expression $W_1(B_1, B_2)$, which is the sum of discounted future profits when the initial installed bases are B_1 and B_2. $W_1(B_1, B_2)$ is increasing in B_1 and decreasing in B_2
- The firm's pricing problem can then be posed as choosing a price to maximize the sum of current profit and the discounted sum of future profits:
  - Choose p_1 to maximize $(p_1 - c_1)q_1 + W_1(B_1 + q_1, B_2 + q_2)$, where $q_1 = D_1 (p_1, p_2, B_1, B_2)$, $q_2 = D_2 (p_1, p_2, B_1, B_2)$.
- A firm will take into account how its price affects current profit $(p_1 - c_1)q_1$ but also its future profit stream $W_1 (B_1 + q_1, B_2 + q_2)$ in that price influences the future installed base through its effect on the current amount sold q_1

With no network effects, a firm would choose price p_1'
to maximize its current profit

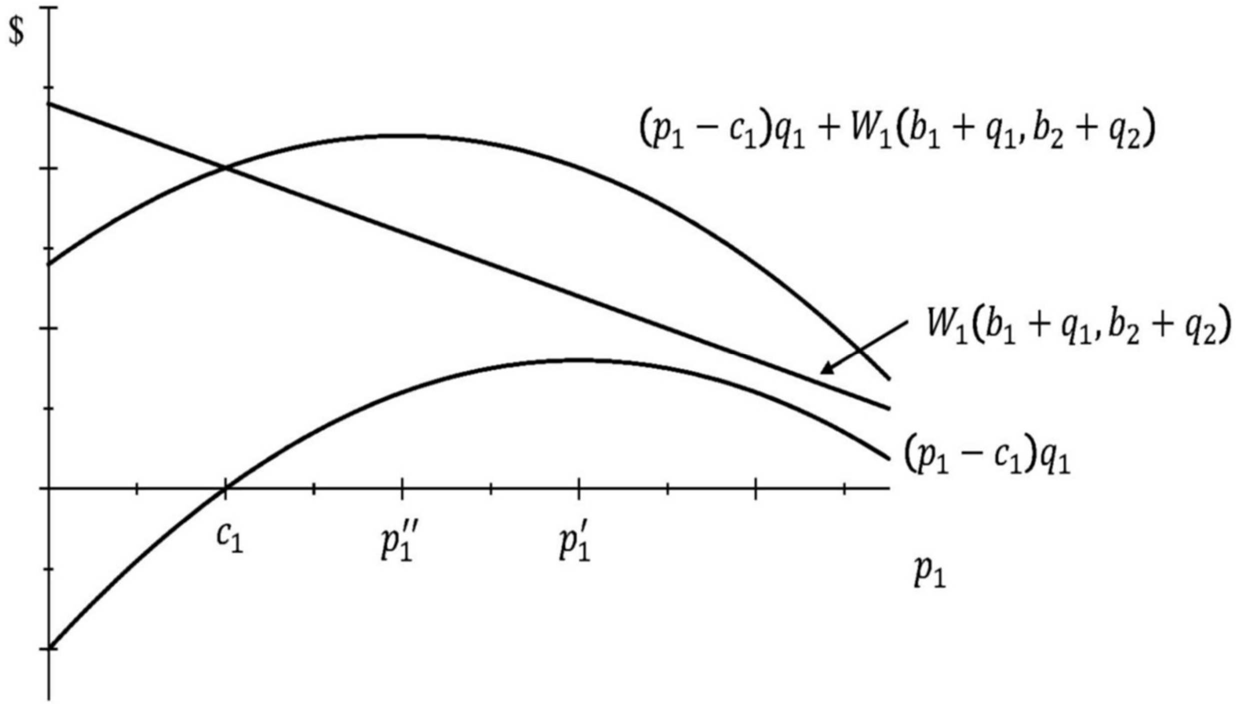

With network effects, the current price impacts the future profit stream W_1

W_1 is decreasing in the current price because the lower is price, the more consumers will buy and, therefore, the higher is the installed base in the next period. This implies stronger future demand and higher profits

A lower price not only raises current demand but also future demand. Optimal dynamic pricing has the firm price below that which maximizes current profit. That price is p_1''

Pricing below p_1' forgoes some current profit, but it is more than compensated by higher future profits by building the installed base

The consequences of the dynamic pricing incentive: *Penetration pricing* and market power

- The following figure depicts the optimal price for firm 1 (which is given by the height of the plotted surface) depending on both firms' current installed bases
- Firm 1 (and firm 2) prices very low when firms have comparable installed bases (the "trench" around the diagonal).
- Referred to as *penetration pricing*, it is intended to build a firm's installed base and thus is a form of investment that, in expectation, pays off in terms of future profit.
- Penetration pricing is intended to spark the positive feedback associated with network effects. A firm that gets a slight edge in its installed base has a significantly higher chance of dominating the market and earning high profits
- Though price could even be below marginal cost, it is not predatory pricing, because the intent is to build a firm's installed base rather than drive its rival from the market
- When installed bases are far enough away from the diagonal, prices are much higher because the intense phase of competition is over. The firm with a smaller installed base is resigned to having a smaller share of the market because supplanting the current market leader would take a sustained policy of pricing low, which is too costly. Thus, firms settle down to more standard competition though with the firm with a higher installed base having higher demand, a higher price, and higher profits

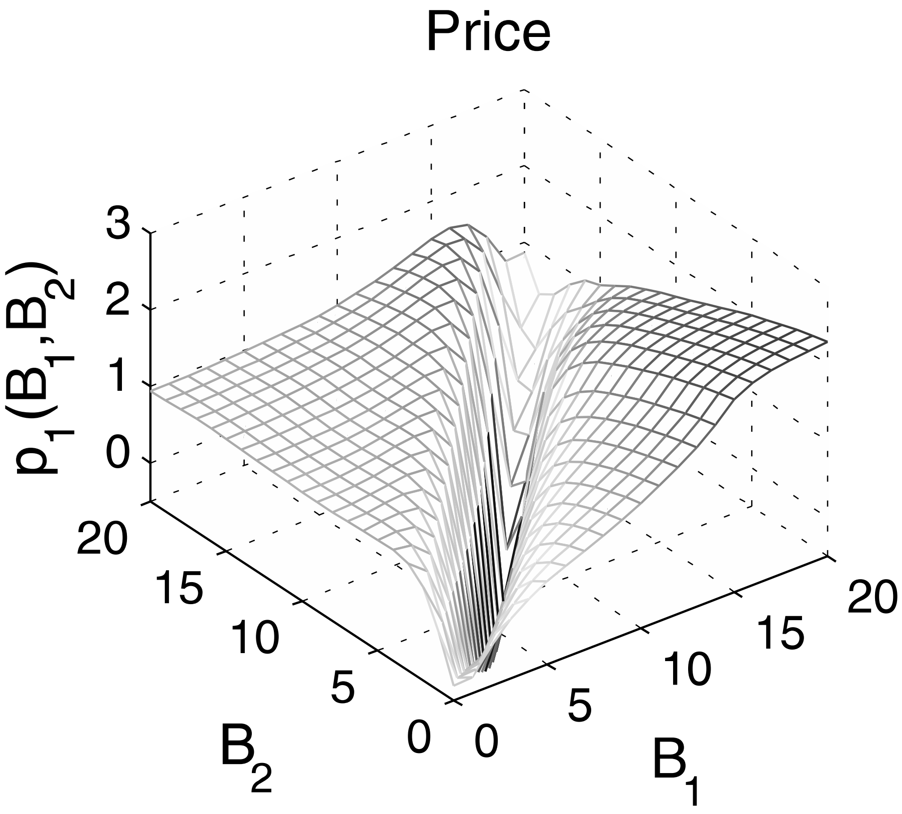

Price
3
2
1
0
$p_1(B_1,B_2)$
20
15
10
5
0
B_2
0
5
10
15
20
B_1

The consequences of the dynamic pricing incentive: *The tendency for market dominance*

- The following figure plots the average movement in firms' installed bases, based on their current installed bases and given equilibrium prices
- When installed bases are equal (so they are on the diagonal), the average tendency is for them to remain equal; either rising when their initial bases are low or falling when they are high. (They can fall over time because some consumers who bought in the past stop using the product.) With identical installed bases and identical prices in equilibrium, no firm has an advantage in expectation
- However, this only shows what happens on average, and the actual realization can be different. If installed bases start on the diagonal, they can move away from it
- Once firms' installed bases are different, the tendency is for them to move in the direction of expanding the advantage of the firm with the larger installed base. For example, if installed bases are to the right of the diagonal, so firm 1 has the larger base, the arrows point to the right which means the average tendency is for installed bases to move to the right, which means firm 1's advantage grows
- In this case firm 2 has to invest relevant resources and sell below its cost for a long period to change the average direction toward which market evolves

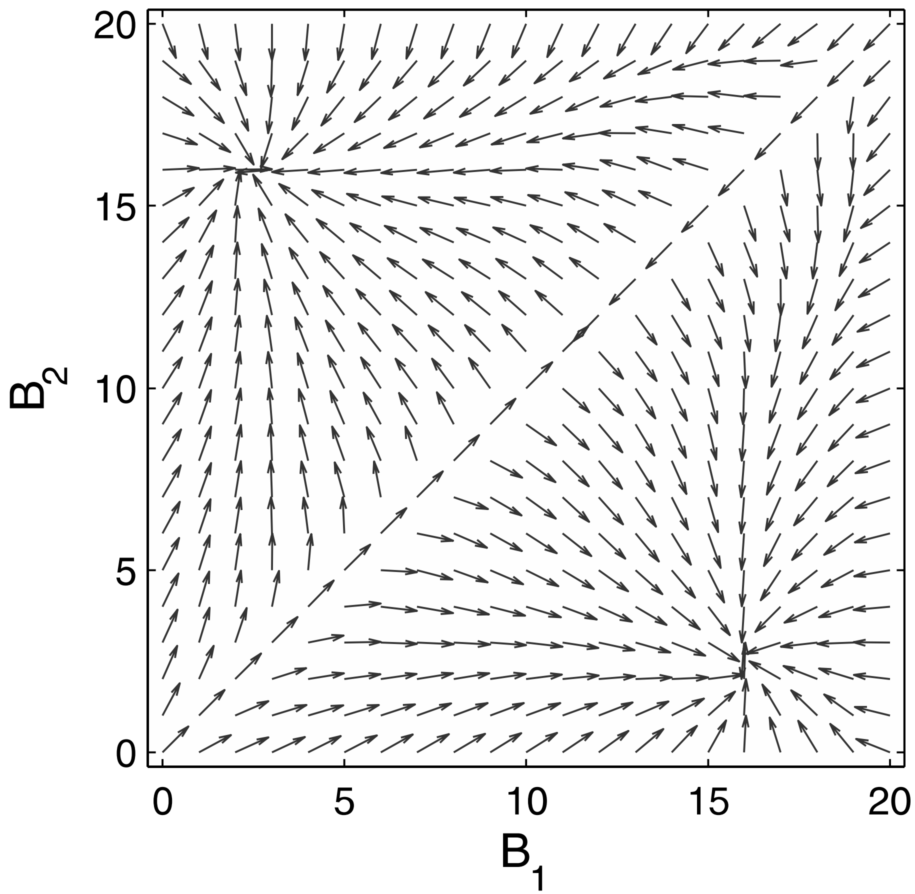

20
15
10
5
0
0
5
10
15
20
B_1
B_2

The consequences of the dynamic pricing incentive: some conclusions

- Market dominance is a natural competitive outcome when there are strong network effects, and with it comes two positive features
  - First, competition to become the dominant firm (“competition for the market”) is intense, which benefits consumers by ensuring low prices
  - Second, consumers benefit from having a dominant firm, because that results in a higher valued product due to stronger network effects (for example, more software being written for an OS)
- Of course, the dominant firm will eventually take advantage of its position by charging a high price, but consumers are still likely to fare well

Network effects, market dominance and exclusionary prices (1)

- Exclusionary contracts can keep a new firm's installed base from reaching the critical mass required for the product to be attractive to consumers. This can drive out a more efficient firm or deter one from entering
- Suppose an upstream monopolist 1 provides input A to a downstream firm, which combines it with other inputs and sell it to final consumers. Upstream firm 2 has developed product B as a substitute for A. These upstream products (e.g. computer OS) are subject to network effects
  - *Example*: firm 1 is Microsoft and firm 2 has a competing operating system (for example, IBM and OS/2). The downstream firms are original equipment manufacturers (OEMs) of personal computers, such as Dell and Hewlett-Packard, who install an OS and then sell it to consumers

Network effects, market dominance and exclusionary prices (2)

- Assume there is just one downstream firm.
- Suppose that the network effects for the product of firm 1 have largely been realized, so that additional customers will not add any more value. The value of input A to the downstream firm is assumed to be 150, and firm 1 charges a price of 100. Hence, a downstream firm receives a surplus of 50 on its M customers.
- Now, a new input B appears, which is supplied by firm 2.
- The downstream firm's M customers differ in how they value B. A fraction $(1-\vartheta)$ attach zero value to B, because there are no network effects (for example, no software has been developed), while a fraction ϑ attach a value of 200 (for example, they prize the OS's stability and develop their own software)
- Assume marginal cost is 0 for both inputs A and B, and firm 2 must earn revenue of at least F to profitably supply its input.
- The surplus maximizing solution is to have the $(1-\vartheta)M$ customers buy A, which they value at 150 compared to 0 from B, and the ϑM customers buy B, which they value at 200 compared to 150 for A
- If firm 1 continues to price A at 100, that outcome will occur if, for example, firm 2 prices B at 100. The resulting revenue for firm 2 is $\vartheta M 100$, which we assume exceeds F
- Firm 2 and consumers are better off, but firm 1 is harmed in two ways
  - First, its profit is lower by $\vartheta M100$ because of weaker demand
  - Second, firm 2 may eventually attract the other $(1-\vartheta)$ M consumers as it builds up network effects for its product.

Network effects, market dominance and exclusionary prices (3)

- How can firm 1 exclude firm 2 from the market? Instead of requiring a downstream firm to pay 100 for each unit of input A it buys, it requires to pay 100 for each unit of output that it sells
- Let P_B denote the price charged by firm 2 for input B.
- The downstream firm has three options
  1. it can agree to this contract from firm 1 and buy all its input from firm 1, which yields a profit of $50M$
  2. it can agree to firm 1's contract and buy $(1-\vartheta)M$ units from firm 1 and ϑM units from firm 2. In that case, its profit is $(1-\vartheta)M(150 - 100) + \vartheta M(200 - P_B - 100)$. It buys $(1-\vartheta)M$ units of A at a price of 100, which delivers value of 150, so the resulting per unit profit is 50. It buys ϑM units of B, which delivers per unit value of 200 and which requires the downstream firm to pay P_B to firm 2 and 100 to firm 1
  3. It can decline the contract from firm 1 and buy ϑM units from firm 2, which delivers profit of $\vartheta M(200 - P_B)$

Network effects, market dominance and exclusionary prices (4)

a) The first option is more profitable than the second option when:

$$50M > (1 - \vartheta)M(150 - 100) + \vartheta M(200 - P_B - 100)$$

This condition can be simplified to $P_B > 50$. Hence, if firm 2's price for B exceeds 50, then the downstream firm prefers buying all its inputs from firm 1 than buying inputs from both firms 1 and 2

b) The first option is more profitable than the third when:

$$50M > \vartheta M(200 - P_B)$$

Even if $P_B = 0$, so firm 2 gives its product away, it is more profitable to have firm 1 as the exclusive supplier when $\vartheta < 1/4$; that is, the fraction of consumers who value B (when there are no network effects) is less than 25%

Network effects, market dominance and exclusionary prices (5)

Summarizing:

- The downstream firm will agree to the contract from firm 1 and buy all its inputs from firm 1 when $P_B > 50$
- Firm 2 can have a positive demand only if $P_B \leq 50$, but…
  - If firm 2 were to charge a price of 50, then its profit is $\vartheta M50$. Recalling that it must earn at least F to be profitable, if $\vartheta M50 < F$, then the highest price that firm 2 can charge to have a positive demand is not sufficient for it to be profitable, and it will not supply
  - If $\vartheta < 1/4$, then the downstream firm will agree to the contract from firm 1 and buy all its input from firm 1

Concluding:

- This kind of contract can produce exclusionary effects towards competitors with better products. Firm 1 is allowed to continue to be the exclusive supplier at its original price of 100. Furthermore, it is able to prevent firm 2 from getting a foothold in the market, which, after building its installed base and generating network effects, might have allowed firm 2 to effectively compete for the remainder of firm 1’s customers

The Microsoft case: Maintenance of monopoly in the operating systems market (1)

- An OS runs software applications by using application programming interfaces (APIs), which allow the application to interact with the OS. A platform is a collection of such APIs
- When writing software, a developer needs to write it so that it works with an OS's APIs. If the software is also to work on a different OS, then the code must be rewritten for its APIs. Such a process is known as "porting" and can be costly
- Given the large number of Windows users, software developers generally write their programs for the Windows platform; such is the advantage of having the largest installed base, as we know from our analysis of network effects. Given fewer Mac users, not all software written for Windows' APIs would get ported to the Mac OS.

Microsoft: Maintenance of monopoly in the operating systems market (2)

- A threat to Windows' dominant position would be a technology that allows the same programs to run on all OSs. In that case, if a superior OS came along, it would not be at a disadvantage to Microsoft, because the existing software applications could run on it
- This was the threat that Netscape Navigator and Java posed. Written by Sun Microsystems, Java was a cross-platform language that allowed a program to run on many different OSs. Navigator relied on Java and could run software applications independently of Windows (it ran on seventeen different OSs, including Windows, Mac, and various versions of the UNIX operating system)
- Referred to as "middleware", Navigator with Java was a potential challenge to Windows' position as the dominant platform.

The Microsoft case: the assessment of market power

- To establish a monopolization claim, the plaintiffs must argue that
  - 1) the accused firm has monopoly power in a relevant market, and
  - 2) it has sought to maintain that monopoly through anticompetitive behavior (for example, harming a rival's product as opposed to making one's product better)
- Windows was found to have monopoly power for two reasons:
  - It had the highest market share in the market of OSs (from 95 to 80% depending on the inclusion of Mac OS and the OSs for mobile devices)
  - Network effects and a large installed base of users provided a relevant barrier to entry. A new OS would lack such an installed base, which meant that not much software might be written for it, and if there is not much software, consumers will not be inclined to buy it, even if the OS is superior

The Microsoft case: the anticompetitive conducts

- As for the anticompetitive conduct, it was observed that the two primary avenues for distributing browsers are
  - OEMs, which install programs on the computers they sell
  - ISPs (such as America Online), which offer browsers when someone signs up for Internet access
- In both cases, Microsoft was found to have used contracts that severely disadvantaged competing browser
  - Microsoft prohibited OEMs from altering the Windows desktop and the initial boot sequence. So, for example, an OEM could not take IE off the desktop and replace it with Navigator, even if the customer so desired it (in 1996 Compaq wanted to load the more popular Netscape browser on its machines and remove the icon for IE from Windows 95. Microsoft informed Compaq that if it removed IE, Compaq would lose its license for Windows. Compaq complied with Microsoft's wishes)
  - With regard to ISPs, Microsoft agreed to provide easy access to an ISP's services on the Windows desktop in exchange for them exclusively promoting IE and keeping shipments of Navigator under 25% (AOL and CompuServe had contracts with Microsoft that restricted them in their ability to promote non-IE browsers. The market share of IE for those two ISPs rose from 20% to 87% in less than two years)
- The government also argued that the integration of Windows and IE was anticompetitive because it involved excluding IE from the "Add/Remove Programs" utility, commingling code so that deletion of IE would harm the OS, and that Windows could override the user's default browser when it was not IE.

The Microsoft case: the conducts toward Java

- The conducts adopted by Microsoft toward Java were also very aggressive.
- Java was a potentially serious threat to the Windows platform. Though Microsoft announced its intent to promote Java, it actually crippled it as a cross-platform threat by designing a version of Java that ran on Windows but was incompatible with the one written by Sun Microsystems, and then using its Windows monopoly to induce parties (including Intel) to adopt its version

The Microsoft case: the remedies

- Three remedies were discussed. The first was behavioral:
  1. To restrict conduct prohibiting some or all of the anticompetitive practices, such as limiting what Microsoft could put in a contract with an OEM
- The other two remedies were structural and draconian. They aimed at reducing monopoly power by breaking up the company
  2. One plan was to create three identical Microsoft companies (dubbed “Baby Bills”), each with Windows and applications. The objective was to inject competition into the OS market
  3. The second structural remedy was to create two companies, one with Windows and one with applications (such as Word and Excel)
- The latter remedy was proposed by the DOJ and ordered by the district court. However, the circuit court remanded the remedy, because the district court had not adequately considered the facts and justified its decision

The Microsoft case: the settlement

- At that point, the government stopped pursuing a structural remedy, and a settlement restricting the conduct of Microsoft was reached in November 2001.
  - The remedy prohibited Microsoft from retaliating against an OEM for using or promoting software that competes with Microsoft products
  - Required uniform licensing agreements to OEMs (so as to make it more difficult for Microsoft to hide rewards and punishments in discriminatory agreements)
  - It also sought to make it more difficult for Microsoft to hamper the development of middleware and more generally to promote interoperability by mandating that it make available the APIs and related documentation that are used by middleware to interact with a Windows OS product.