Analyzing the Elimination of Double Marginalization

A manufacturer makes clocks and sells them to a retailer, which resells them to customers. The manufacturer sells only to this retailer, and the retailer buys clocks only from this manufacturer. The retailer sells every clock it buys, so the manufacturer's sales equal the retailer's sales.[1]

The manufacturer charges the retailer $20 per clock. Making each clock costs the manufacturer $8. The retailer sells each clock for $30. In addition to the $20 it pays per clock, the retailer spends $4 per clock to sell it. The retailer currently sells 300 clocks per month. The two firms deal only through the $20 per-clock price: there are no fees, quantity discounts, or agreements about the retail price.

The retailer estimates that, over the relevant price range, each $1 decrease in its retail price will increase its monthly sales by 50 clocks.

Assume that for each $1 reduction in the retailer's cost per clock (the price it pays for a clock plus its $4 selling cost), the retailer reduces its retail price by $0.50. This is a 50% pass-through rate.

The manufacturer acquires the retailer. The costs of making and selling each clock do not change. After the merger, the two businesses operate to maximize their combined profit.

Use these facts to calculate the prices, margins, quantities, and profits before and after the merger. Explain how eliminating double marginalization (EDM) affects customers and the merged firm. The completed worksheet and explanations follow.[2]

The idea in brief

Each firm adds its own markup. The manufacturer marks up its $8 cost to $20, and the retailer marks up its $24 cost to $30. Two markups in a row is double marginalization. When the retailer sets its price, it treats the manufacturer's $12 markup as part of the cost of selling one more clock. At $30, each clock sold earns the two firms $18 together ($30 − $8 − $4), but the retailer keeps only $6 of that, and the retailer sets the price. So it sells fewer clocks than would maximize the two firms' combined profit. After the merger, the $12 markup is no longer paid to an outside firm. At the original $30 price, the merged firm counts the full $18 that each sale earns, so it cuts the price and sells more. Customers pay less, and the merged firm earns more.

---

[1] This is a "no inventory" restriction.

[2] Technical assumptions: Each cost is the same for every clock, so the cost per clock is also the cost of one more clock (marginal cost). Before the merger, the manufacturer sets its per-clock price anticipating how the retailer will respond, and each firm maximizes its own profit. There are no fixed costs, inventory changes, losses, or capacity limits. Each customer buys at most one clock per month. Demand, quality, and the firms' other business activities are unchanged. Treat the stated pass-through rate and sales response as accurate over the relevant range, which here means retail prices from $20 to $31.

Completed worksheet

Prices, costs, and margins are per clock; quantities, profits, and customer gains are per month.

Calculation order. Line 4 depends on line 14, so work the after-merger column in this order: lines 1-3, then lines 6-14, then lines 4-5 and 15-19.

Line	Item	Before merger	After merger	Change	How calculated
Manufacturer					
1	Price charged to retailer (transfer price after merger)	$20	$8	−$12	Before: given. After: equals line 2.
2	Cost of making a clock	$8	$8	None	Given
3	Manufacturer's margin per clock	$12	$0	−$12	Line 1 − line 2
4	Clocks sold to the retailer	300	600	+300	Equals line 14
5	Manufacturer's profit	$3,600	$0	−$3,600	Line 3 × line 4
Retailer					
6	Price paid for a clock (transfer price after merger)	$20	$8	−$12	Equals line 1
7	Other selling cost per clock	$4	$4	None	Given
8	Retailer's cost per clock	$24	$12	−$12	Line 6 + line 7
9	Reduction in retailer's cost per clock	—	$12	—	Line 8 before − line 8 after
10	Pass-through rate	—	50%	—	Given
11	Reduction in retail price	—	$6	—	Line 9 × line 10
12	Retail price	$30	$24	−$6	Before: given. After: $30 − line 11.
13	Retailer's margin per clock	$6	$12	+$6	Line 12 − line 8
14	Clocks sold to customers	300	600	+300	Before: given. After: 300 + (50 × line 11).
15	Retailer's profit	$1,800	$7,200	+$5,400	Line 13 × line 14
Combined					
16	Combined profit	$5,400	$7,200	+$1,800	Line 5 + line 15
Customers					
17	Savings to existing customers	—	—	+$1,800	300 × line 11
18	Gain to new customers (value above price)	—	—	+$900	300 × line 11 ÷ 2
19	Total gain to customers (increase in consumer surplus)	—	—	+$2,700	Line 17 + line 18

Reductions in lines 9 and 11 are shown as positive numbers. Lines 17-19 measure the increase in consumer surplus: the value customers receive above what they pay. Because they measure only the change caused by the merger, they appear in the Change column.

The chain of effects

Figure 1 traces the chain of effects behind the worksheet. It uses two general labels that also fit other vertical mergers: the supplier, the firm that sells the input (here, the manufacturer), and the partner, the downstream firm the supplier acquires (here, the retailer). If the merged firm prices on its real cost, the supplier's markup is no longer a cost to it (A), so the partner's cost of selling one more unit falls by that markup (B). The partner passes on half of the cost reduction as a lower price (C), and sales rise as new customers buy (D). Customers gain, both on purchases they were already making and on new ones (E). The merged firm gains too, because its margin on the added sales outweighs the lower price on the sales it was already making (F).

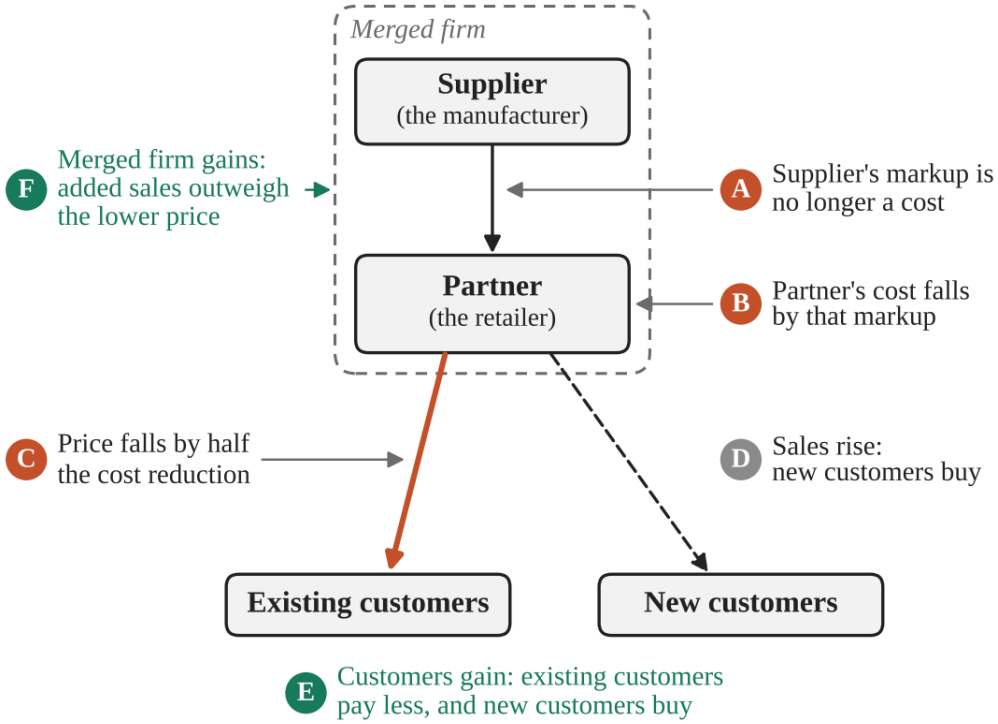


Figure 1. How eliminating double marginalization works. Coral marks the cost and price changes, gray the change in sales, and teal the gains.

Explanation of the worksheet

Letters in parentheses refer to the callouts in Figure 1.

1. Price charged to retailer (transfer price after merger)

Before the merger, the manufacturer charges the retailer $20 per clock. After the merger, clocks move from one division of the same company to another, and the company records an internal transfer price for them. That transfer price matters because the retail division treats it as its cost when it sets the retail price. If the retail division were still charged $20, it would still charge customers $30, and the merger would change nothing. To maximize the merged firm's profit, the transfer price should equal what a clock actually costs the firm to make: $8. (We use this rule without proving it.) The transfer price itself only moves money from one division to the other, so it does not change combined profit. The result would be the same if headquarters set the retail price directly, using the combined $12 cost to make and sell a clock.

2. Cost of making a clock

Given. Each clock costs $8 to make, before and after the merger. Because every clock costs the same $8, this is also the cost of making one more clock, which economists call marginal cost. The merger does not make clocks cheaper to produce: EDM removes a markup, not a production cost.

3. Manufacturer's margin per clock (A)

Line 1 − line 2. Before the merger, the manufacturer earns $20 − $8 = $12 on each clock. After the merger, the transfer price equals the cost of making a clock, so the manufacturing division records $8 − $8 = $0. That does not mean manufacturing earns nothing for the firm. The firm's profit on each clock now appears in the retail division's margin (line 13).

4. Clocks sold to the retailer

Equals line 14, because every clock sold to a customer is bought from the manufacturer. Before the merger: 300 clocks. After the merger: 600 clocks, as calculated in line 14.

5. Manufacturer's profit

Line 3 × line 4. Before the merger: $12 × 300 = $3,600. After the merger: $0 × 600 = $0. With the transfer price set at cost, all of the merged firm's profit is recorded in the retail division.

6. Price paid for a clock (transfer price after merger)

Equals line 1, seen from the retailer's side: $20 paid to an outside supplier before the merger, and an $8 internal transfer price after it.

7. Other selling cost per clock

Given. The retailer spends $4 to sell each clock, in addition to what it pays for the clock. The merger does not change this cost.

8. Retailer's cost per clock

Line 6 + line 7. Before the merger: $20 + $4 = $24. After the merger: $8 + $4 = $12. This is the cost the retailer uses to set its price. After the merger, it equals what it actually costs the firm to make and sell one more clock.

9. Reduction in retailer's cost per clock (B)

Line 8 before the merger minus line 8 after: $24 − $12 = $12. This equals the manufacturer's old $12 markup (line 3). The real cost of making and selling a clock was $12 before the merger and is $12 after it. What disappears is the $12 markup the retailer used to pay an outside firm.

10. Pass-through rate

Given: for each $1 drop in its cost per clock, the retailer cuts its price by $0.50. The rate refers to dollar changes, not percentage changes. The 50% figure is consistent with the sales response in the problem: when each $1 price cut adds the same 50 clocks, a retailer maximizing its own profit passes through exactly half of any change in its cost. (The section "Why the merged firm cuts its price" below illustrates this, and the Appendix explains why.) With other demand patterns, pass-through can be higher or lower, so 50% is not a general rule for EDM.

11. Reduction in retail price (C)

Line 9 × line 10: \$12 × 50% = \$6. Customers get part of the cost reduction as a lower price; the merged firm keeps the rest as a larger margin per clock.

12. Retail price

Before the merger: \$30 (given). After the merger: \$30 − \$6 = \$24. The new price comes from the pass-through rate; the starting prices and costs alone do not determine it.

13. Retailer's margin per clock

Line 12 − line 8. Before the merger: \$30 − \$24 = \$6. After the merger: \$24 − \$12 = \$12. The margin rises because the retailer's cost falls by \$12 while its price falls by only \$6.

14. Clocks sold to customers (D)

Before the merger: 300 clocks (given). The retailer sells 50 more clocks for each \$1 price cut, and line 11 gives a \$6 cut, so sales rise by 6 × 50 = 300. After the merger: 300 + 300 = 600 clocks per month. This step needs both the price cut and the sales response; the pass-through rate alone does not tell you how much sales grow.

15. Retailer's profit

Line 13 × line 14. Before the merger: \$6 × 300 = \$1,800. After the merger: \$12 × 600 = \$7,200. This example has no fixed costs to subtract.

16. Combined profit (F)

Line 5 + line 15. Before the merger: \$3,600 + \$1,800 = \$5,400. After the merger: \$0 + \$7,200 = \$7,200. Combined profit rises by \$1,800 per month.

17. Savings to existing customers (E)

300 × line 11. The 300 customers who bought at \$30 now pay \$24, saving \$6 each: 300 × \$6 = \$1,800 per month. This is cash saved on purchases customers were already making.

18. Gain to new customers (value above price)

The 300 new customers value a clock at between \$24 and \$30: enough to buy at the new price, but not at the old one. Because sales rise steadily (50 clocks for each \$1 cut), their values are spread evenly over that range. On average, each new customer gets a clock worth \$3 more than the \$24 price: 300 × \$3 = \$900 per month. This is not cash saved. It is the "consumer surplus" these customers receive above the \$24 they pay, on purchases they did not make before the merger. Dividing by 2, as the worksheet does, is the same as using that \$3 average.

Courts—and, to my knowledge, the agencies—do not attempt to account for this part of the gain. In merger review, the benefit to customers is measured by the cash savings on purchases they already make (line 17), so the gain to customers who buy only because of the lower postmerger price is ignored. A compelling showing of this additional consumer surplus might persuade the agencies, or at least the agencies' economists, to credit it, but I am not aware of any precedent for doing so.

19. Total gain to customers (increase in consumer surplus)

Line 17 + line 18: $1,800 + $900 = $2,700 per month. This is the increase in consumer surplus. Customers gain more than the merged firm does (+$1,800 in line 16), though in merger review only their cash savings in line 17 would ordinarily count (see explanation 18). Like the other results, these figures depend on the pass-through rate and sales response given in the problem.

Why the merged firm cuts its price

The merged firm does not cut its price out of generosity. It costs the firm $12 to make and sell a clock. If it kept the $30 price, it would sell 300 clocks at an $18 margin and earn $5,400, exactly what the two firms earned separately. The table below uses the problem's sales response (50 more clocks for each $1 price cut) to compute the merged firm's profit at other prices.

Retail price	Clocks sold per month	Margin per clock (price − $12)	Monthly profit
$30	300	$18	$5,400
$28	400	$16	$6,400
$26	500	$14	$7,000
$24	**600**	**$12**	**$7,200**
$22	700	$10	$7,000
$20	800	$8	$6,400

Among the prices shown, profit is highest at $24, the same price the worksheet reached using the 50% pass-through rate: a $12 cost reduction led to a $6 price cut. Before the merger, the retailer had no reason to make this cut, because it ignored the $12 the manufacturer earned on each extra clock.

The table also shows why the merged firm stops at $24. A price cut adds sales, but it also lowers the price of every clock the firm was already selling. Cutting the price from $26 to $24 adds 100 clocks at a $12 margin (+$1,200) and gives up $2 on the 500 clocks already selling at $26 (−$1,000), so profit rises by $200. Cutting further, from $24 to $22, adds 100 clocks at a $10 margin (+$1,000) but gives up $2 on the 600 clocks already selling at $24 (−$1,200), so profit falls by $200.

Where the gains come from

The merger does not save any resources: making and selling a clock still costs $12. On the 300 clocks that were already being sold, the lower price shifts money from the firm to customers. The firm's margin on each of those clocks falls from $18 to $12, and customers save the same $1,800 a month (line 17).

All of the net gain comes from the 300 additional clocks the merged firm sells postmerger. The firm earns $3,600 a month on them (300 × $12), and the new customers gain $900 (line 18). The firm's profit rises by $1,800 (line 16) because the $3,600 it earns on the new clocks more than offsets the $1,800 it gives up on the old ones. Together, the firm and its customers gain $4,500 a month ($1,800 + $2,700), all from the increase in output.

In short, EDM does not reduce the resources needed to make and sell each clock. Its economic benefit comes from changing pricing incentives and increasing output. The lower price on existing purchases transfers money to customers; the additional purchases create a net economic gain because customers value those clocks more than they cost to produce and sell.[3]

Four assumptions that matter

Four facts in the problem do important work. Each matters when a court or agency evaluates an EDM claim.

Both firms mark up their prices. In this example, double marginalization arises because both firms independently charge a markup over marginal cost. Suppose instead that identical retailers competed on price and competition eliminated their markup, so that the retail price equaled the wholesale price plus the $4 selling cost. The manufacturer's best wholesale price would be $20, the retail price would be $20 + $4 = $24, and 600 clocks would sell, the same result the merged firm reaches. Similarly, if clocks were sold to the retailer at their $8 cost, the retailer alone would choose $24. With only one markup, the price is the same $24 whichever level takes it, and a merger has no second markup to remove.

The firms deal only through a per-clock price. Before the merger, the retailer pays $20 for each clock, with no fees, quantity discounts, or agreements about the retail price. If the firms could reach the same result by contract, without merging, the benefit of EDM may not be specific to the merger. The next section explains how a contract can do so.

The manufacturer sells only to this retailer. If the manufacturer also supplied rival retailers, some of the retailer's added sales might come from those retailers. The merged firm would then lose the manufacturing profit on each clock those retailers no longer sell, and it would count that lost profit as a cost of cutting its retail price. That opportunity cost could reduce its incentive to lower its retail price. What matters is not whether the manufacturer has other customers, but whether the retailer's added sales would come at the expense of retailers that sell the manufacturer's clocks. Sales won from retailers selling other brands do not matter here: the manufacturer earns nothing on those clocks, and those customers are already counted in the sales response of 50 clocks per $1. Winning sales from rivals by charging lower prices is not, by itself, a competitive harm. Supplying rivals does raise a separate question this example does not address: whether the merged firm would raise its price to them.

The merged firm prices on its real cost. The problem assumes that after the merger the two businesses maximize their combined profit. The retail division therefore bases its pricing decision on the combined $12 cost of making and selling a clock, including an $8 internal transfer price. Common ownership does not by itself ensure that divisional pricing decisions immediately maximize combined profit.[4] Charging the retail division only $8 for a clock leaves the manufacturing division with no reported profit. If the manufacturing division's managers are

---

[3] This is an economic point. Whether EDM counts as a cognizable efficiency in merger review is a separate legal question. Kwoka and Slade, for example, describe EDM as only a "pecuniary" economy because it saves no resources, while acknowledging that the resulting increase in output improves economic efficiency. *See* John Kwoka & Margaret Slade, *Second Thoughts on Double Marginalization*, ANTITRUST, Spring 2020, at 51, 53.

[4] For evidence from one industry, *see* Gregory S. Crawford et al., *The Welfare Effects of Vertical Integration in Multichannel Television Markets*, 86 ECONOMETRICA 891 (2018), which estimates that the divisions of integrated firms in multichannel television take into account about $0.79 of each dollar of profit earned by their affiliated divisions. *See* Kwoka & Slade, *supra* note 3, at 55.

paid or evaluated on their own division's profit, they have reason to keep charging something close to the old $20, and the retail division would then keep pricing as it does today. This is a principal–agent problem: the firm's owners want to maximize combined profit, but the managers who set the internal price may be rewarded for something else. Unless the merged firm aligns those incentives, for example by having headquarters set the internal price or by tying managers' pay to combined profit, the price reduction and customer benefits would be smaller, or slower to arrive, than this example shows.

Eliminating double marginalization by contract

A merger is not the only way to eliminate double marginalization. The manufacturer and the retailer could instead agree on a two-part tariff: a per-clock price equal to the manufacturer's $8 cost of making a clock, plus a fixed monthly fee that does not depend on how many clocks the retailer buys. Suppose the fee is $4,500 a month.

Under that contract, the retailer's cost of selling one more clock is $8 + $4 = $12, the same as the merged firm's. Because the fee is the same however many clocks the retailer sells, it does not affect the retailer's pricing decision. The retailer therefore does what the merged firm would do: it cuts its price from $30 to $24 and sells 600 clocks. Before paying the fee, the retailer earns $12 × 600 = $7,200 a month, the same as the merged firm. The manufacturer, by contrast, earns nothing on the clocks themselves, because it sells them at its $8 cost. All of its profit comes from the fee. In effect, the $7,200 is the two firms' joint profit, earned on the retailer's sales to customers, and the fee divides it between them: the manufacturer receives $4,500 and the retailer keeps $2,700. Each firm earns $900 a month more than under the $20 per-clock price, and customers get the same $24 price and $2,700 gain that the merger would produce. Any fee between $3,600 and $5,400 a month would leave both firms better off than they are today. Where the fee falls in that range depends on how the firms bargain. It changes how they split the $7,200, but not the retail price or what customers gain. That is why a two-part tariff is the standard contractual fix for double marginalization.

	Separate firms, $20 per clock	Separate firms, two-part tariff	Merged firm
Price per clock charged to the retailer	$20	$8	$8 (internal)
Fixed monthly fee	—	$4,500	—
Retailer's cost per clock	$24	$12	$12
Retail price	$30	$24	$24
Clocks sold per month	300	600	600
Manufacturer's monthly profit	$3,600	$4,500	$0
Retailer's monthly profit	$1,800	$2,700	$7,200
Combined monthly profit	$5,400	$7,200	$7,200
Gain to customers (consumer surplus)	—	$2,700	$2,700

In the two-part-tariff column, all of the manufacturer's profit is the fee. In the merged firm's column, all of the profit appears in the retail division because the internal transfer price equals

the manufacturer's cost, as in lines 5 and 15 of the worksheet. In both columns, the whole $7,200 is earned on the retailer's sales to customers. Under the tariff, the fee decides how two separate firms share it. In the merged firm, the transfer price decides only which division reports it.

If the firms already have such a contract, a merger has nothing left to eliminate. The retailer already prices on the $12 cost, and the merged firm would charge the same $24. The parties would then have no EDM efficiency to claim. More generally, EDM is specific to the merger only to the extent that contracting could not achieve it. The 2023 Merger Guidelines ask whether contracts short of a merger have eliminated, or could eliminate, double marginalization.[5] If the parties do not already use a two-part tariff, the question is whether they could, and it is up to the merging parties to show why a contract would not work.

The contractual fix works only if three conditions hold.

The per-clock price equals the manufacturer's cost. If the contract instead charged $14 a clock plus a fee, the retailer's cost would be $18, and it would cut its price only to $27, passing on half of the $6 reduction from $24. Part of the double marginalization would remain, and a merger could eliminate that part: the remaining $6 markup, which would lower the price by another $3. Only that remaining part would count as a merger-specific efficiency.

The fee is fixed over the sales that matter. A fee that rises with the number of clocks sold works like a per-clock charge and brings the markup back. Contracts often apply their terms only up to some volume. The fix holds as long as the retailer's current sales, and the sales it would reach after cutting its price, stay within that limit.

The contract will last. If the contract is about to expire, can be repriced, or would end on a change in control, the firms might return to a per-clock markup. A merger that prevents that future markup could then claim some EDM benefit. If the contract has years to run, survives the deal, and is expected to be renewed on similar terms, it cannot.

Finally, a two-part tariff with one retailer does nothing for the manufacturer's other customers. If the manufacturer also sells to rival retailers at a per-clock price above its $8 cost, double marginalization remains on their sales. Acquiring one retailer does not integrate the others with the manufacturer, so it cannot remove that markup. What the acquisition may do instead is give the merged firm an incentive to raise the price it charges those rivals: the raising-rivals'-costs theory taken up in the Class 23 note.

Checking the starting prices

The starting prices fit the sales response: each firm's price maximizes its own profit. The comparisons below illustrate this: for each firm, profit is lower at $1 above or below its actual price.

Retailer. With a $24 cost per clock, the retailer earns $1,750 at $29 (350 clocks × $5), $1,800 at $30 (300 × $6), and $1,750 at $31 (250 × $7).

Manufacturer. The manufacturer knows that a $1 change in its price changes the retail price by $0.50 and the retailer's sales by 25 clocks. It earns $3,575 at $19 (325 clocks × $11), $3,600 at $20 (300 × $12), and $3,575 at $21 (275 × $13).

[5] U.S. Dep't of Justice & Fed. Trade Comm'n, MERGER GUIDELINES § 2.5 n.31 (2023).

For students with an economics background. At $30, the retailer's demand elasticity is 50 × 30/300 = 5, and its margin satisfies the Lerner condition: (30 − 24)/30 = 1/5. At $20, the elasticity of the manufacturer's demand from the retailer is 25 × 20/300 = 5/3, and (20 − 8)/20 = 3/5. Because the sales response is linear, elasticity changes with price. At $24, it is 50 × 24/600 = 2, and the merged firm's margin satisfies (24 − 12)/24 = 1/2.

A caution about the numbers

Several numbers repeat in this example, and you may wonder whether the repetitions mean something.

By definition. The $12 reduction in the retailer's cost (line 9) equals the manufacturer's old $12 margin (line 3), because the reduction is the markup the merger removes.

Because demand is a straight line. Output doubles. The new $24 retail price equals the retailer's old $24 cost. The $6 price cut equals the retailer's old $6 margin. The retailer's new $12 margin equals the manufacturer's old $12 margin. And $1,800 is the retailer's old profit, the merged firm's profit gain, and the existing customers' savings. These patterns follow from the straight-line sales response combined with each firm's choosing its profit-maximizing price. The note on linear demand below explains when they hold.

By coincidence. The firm's real cost per clock after the merger (line 8) is also $12. That is simply a result of the numbers chosen.

A note on linear demand

This course assumes linear demand throughout: each $1 change in price changes sales by the same number of units. In this example, linear demand, combined with constant costs and the manufacturer–retailer relationship described above, produces several useful patterns. These patterns recur when the same assumptions apply.

Pass-through is 50%. With constant marginal cost and an unchanged linear residual demand curve facing the firm, the profit-maximizing price changes by half the change in marginal cost (see the Appendix). That is why the problem's 50% pass-through rate is consistent with its sales response of 50 clocks per $1. If rivals responded by changing their own prices, the demand facing the firm could shift, and pass-through could differ.

Profit is symmetric around the best price. Profit falls by the same amount whether the price is set a given amount above or below the best price. So when profit is the same at two prices, the best price lies exactly halfway between them. Equal profits at $1 above and below each starting price therefore show that $30 and $20 are exactly the best starting prices, and equal profits at $22 and $26 show that $24 is exactly the merged firm's best price.

Eliminating a double markup follows a fixed pattern. Under the assumptions of this example—including the manufacturer setting its wholesale price first and anticipating the retailer's response—the patterns described in the previous section follow. Two additional profit relationships also hold: before the merger, the manufacturer earns twice what the retailer earns ($3,600 versus $1,800), and the two firms together earn three-quarters of what the merged firm earns ($5,400 versus $7,200).

New customers gain half the price cut, on average. Because each $1 price cut adds the same number of sales, the new customers' values are spread evenly between the new and old prices. That is why line 18 divides by 2.

Outside the course. Actual demand need not be linear, and costs, contracts, and competitive conditions may differ from those assumed here. If so, the patterns above need not hold, and pass-through can be more or less than 50%; the table at the end of the Appendix shows in which direction. Pass-through and the size of any price reduction are therefore empirical questions in an actual merger.

EDM: Step by Step

The steps below restate the analysis as a method that applies to any set of facts. They use the labels in Figure 1: the supplier sells the input (the manufacturer in the example), and the partner is the downstream firm the supplier acquires (the retailer). Letters refer to the callouts in Figure 1; line numbers and section titles point to the worked numbers.

Set up the problem

1. Identify the roles and the markup at issue. Find the supplier, the partner, and the price the supplier charges the partner for the input. Every unit the partner sells uses one unit from the supplier, so the partner's price sets both firms' sales. (Lines 1–4.)

2. Find the supplier's margin. It is the supplier's price to the partner minus the supplier's cost of making one more unit. That margin is the most the merger can remove from the partner's cost (A). (Lines 1–3.)

3. Check the starting point. The method assumes that, before the merger, each firm already charged its most profitable price, given the other's. If not, part of any predicted change would have happened without the merger. ("Checking the starting prices.")

Find the incentive

4. Ask whether there are two markups. Double marginalization requires that both firms mark up their prices. If competition already holds the partner's markup to zero, or the supplier already sells at cost, there is only one markup, and the merger has nothing to remove. ("Four assumptions that matter.")

5. Ask whether a contract already removes the markup, or could. A per-unit price equal to the supplier's cost, plus a fixed fee, removes double marginalization without a merger. Only the part of the markup that a contract could not remove is specific to the merger. ("Eliminating double marginalization by contract.")

6. Find the partner's new cost. After the merger, the partner's cost of selling one more unit falls by the supplier's margin (B). Before the merger, the partner treated that margin as a cost and ignored what the supplier earned on each added sale. The cost falls only if the merged firm actually prices on its real cost, for example by setting the internal price at cost. (Lines 6–9; "Four assumptions that matter.")

7. Check where the added sales come from. If some of the partner's added sales would come from other retailers the supplier serves, the supplier loses its margin on those sales, and that lost margin offsets part of the cost reduction. ("Four assumptions that matter.")

Trace the effects

8. Predict the price. The partner cuts its price by half the cost reduction (C). (Lines 10–12; Appendix.)

9. Find the added sales. Apply the sales response to the price cut. Because only half the cost reduction is passed on, sales rise by half as much as they would if all of it were (D). (Line 14.)

10. Measure the customers' gain. Existing customers save the price cut on each purchase, and new customers gain, on average, half the price cut each (E). Merger review ordinarily counts only the savings to existing customers. (Lines 17–19.)

11. Measure the merged firm's gain. The firm earns its new margin on the added sales but gives up the price cut on the sales it was already making. The first outweighs the second, which is why the firm cuts its price (F). (Line 16; "Where the gains come from.")

Test the conclusion

12. Test the inputs and assumptions against the evidence. The margins, the pass-through rate, the sales response, whether the merged firm will price on its real cost, and whether a contract could do the same are all questions of evidence. If the supplier also sells to the partner's rivals, the benefit must be weighed against the merged firm's incentive to raise their costs. ("Four assumptions that matter"; "Eliminating double marginalization by contract.")

These results are only as reliable as the estimates and assumptions behind them, and they do not by themselves make a merger lawful.

Appendix: Cost Pass-Through with Linear Demand

The note's central step is that the merged firm cuts the retail price by half of the manufacturer's $12 markup, from $30 to $24. This appendix explains where the "half" comes from and when it holds. It first shows the idea in the example's numbers and then gives two proofs, one using only algebra and one using calculus. Neither proof is needed to follow the note.

Pass-through

The *pass-through rate* is the change in a firm's price divided by the change in its marginal cost, the cost of making and selling one more unit. In the example, marginal cost falls by $12 and the price falls by $6, so pass-through is $6 ÷ $12 = 50%. Pass-through determines how much of a cost saving reaches customers as a lower price. The same rule turns pricing pressure, such as a GUPPI, into a predicted price increase, and it will reappear when we take up raising rivals' costs.

The result

Suppose that:

- the firm sells one product at a single price per unit;
- over the range of prices at issue, the demand curve the firm faces is a straight line, and it does not shift when the firm's cost changes (economists call this the firm's *residual* demand curve: the demand the firm faces with its rivals' prices held fixed);
- the firm's marginal cost is the same at every quantity; and
- both before and after the cost change, the firm charges the price that maximizes its profit, that price lies within the range where demand is a straight line, the firm makes positive sales, and no contract, regulation, or capacity limit prevents the price change.

Then a change in marginal cost, up or down, changes the profit-maximizing price in the same direction by half as much. Pass-through is 50%. Sales change by half as much as they would if the entire cost change were passed on.

The idea in the example's numbers

In the example, the retailer sells 300 clocks a month at $30 and would sell 50 more for each $1 price cut, or 50 fewer for each $1 increase. If that straight line is extended upward, sales reach zero at $36, because 300 ÷ 50 = 6. The sales response covers only retail prices from $20 to $31 (see note 2), and nothing here depends on what sales would be outside that range. The $36 is simply a convenient way to describe the line.

Profit at any price is the margin per clock times the number of clocks sold. Before the merger, the retailer's cost is $24 a clock. Along the extended line, profit is zero at two prices: at $24, where the margin is zero, and at $36, where sales are zero. Between those prices, profit rises and then falls in a symmetric hill, so the peak lies exactly halfway between them, at $30. After the merger, the firm prices on a cost of $12. Profit is now zero at $12 and at $36, and the peak lies halfway between, at $24. Figure 2 shows the two hills.

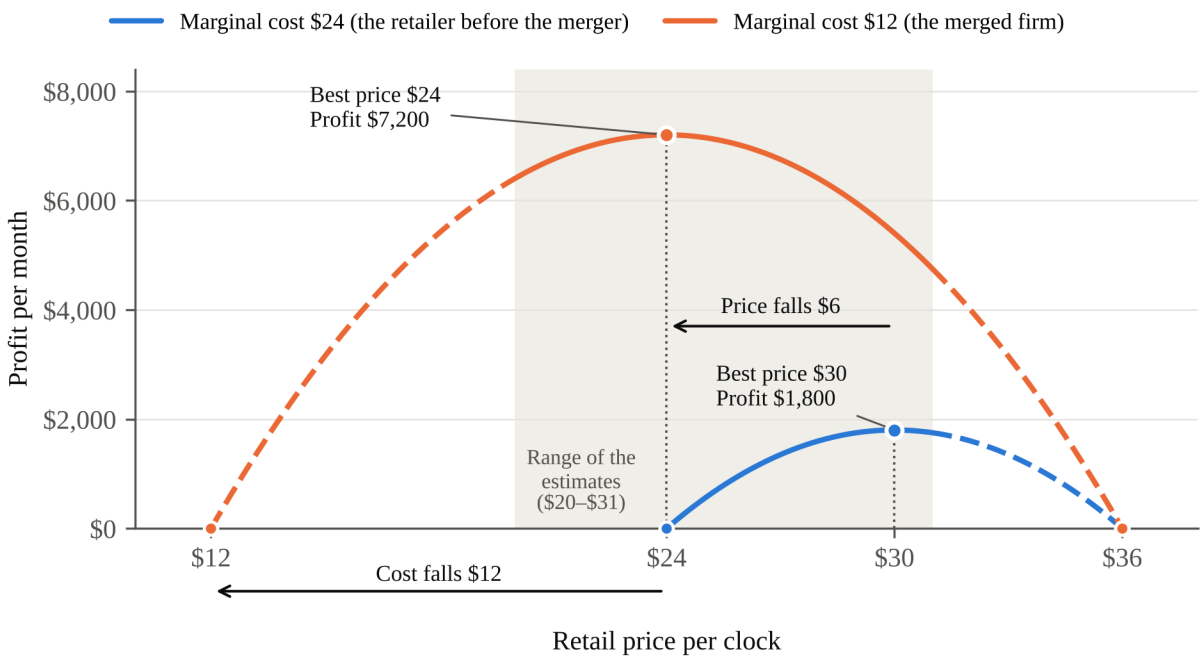


Figure 2. Profit per month at each retail price when marginal cost is $24 (the retailer before the merger) and $12 (the merged firm). Each curve is zero at the cost and at $36 and peaks halfway between. The curves are solid over the $20 to $31 range covered by the example's sales response and dashed where they extend the demand line beyond it.

The cost reduction moves the left end of the hill by $12. The right end stays at $36, because demand has not changed. The point halfway between the two ends therefore moves by half as much: $6. That is the whole result. The proofs below show that with straight-line demand the hill is always symmetric, so the best price is always halfway between cost and the price at which the line reaches zero sales.

Sales follow from the price. The $6 cut raises sales by 50 × 6 = 300, from 300 to 600 clocks a month. Passing on the entire $12 would have added 600 sales; with 50% pass-through, sales rise by half that. "Why the merged firm cuts its price" shows in dollars why the firm stops at half.

Proof 1: Algebra

This proof generalizes the reasoning in the example and uses only algebra. Let p be the price, q the quantity sold, and c the marginal cost. Write the demand line as

$$p = a - bq, \quad \text{where } b > 0.$$

Here a is the price at which the line reaches zero sales, and b is how much the price must fall to sell one more unit. Equivalently, the firm sells 1/b more units for each $1 price cut. In the example, 1/b = 50 clocks, so b = $0.02, and a = $36. Assume that c is less than a, so that the firm makes positive sales. Solving the demand line for quantity gives $q = (a - p)/b$.

Step 1. Write profit in terms of price.

Profit, written π, is the margin per unit times the number of units sold:

$$\pi = (p - c)q = (p - c)(a - p)/b. \qquad (1)$$

Profit is zero at p = c, where the margin is zero, and at p = a, where sales are zero. It is positive between those prices.

Step 2. Find the price that maximizes profit.

Let m = (a + c)/2 be the price halfway between c and a, and let h = (a − c)/2 be the distance from m to either end. Then p − c = h + (p − m) and a − p = h − (p − m). Multiplying these two expressions gives (p − c)(a − p) = h² − (p − m)², so equation (1) becomes

$$\pi = [h^2 - (p - m)^2]/b. \qquad (2)$$

The term (p − m)² is never negative, and it is zero only when p = m. Profit is therefore highest at

$$p^* = (a + c)/2, \qquad (3)$$

the price halfway between the cost and the price at which sales reach zero. Equation (2) also shows why the hill is symmetric: moving a given distance above or below m reduces profit by the same amount. In the example before the merger, h = \$6 and m = \$30, so profit is 50 × [36 − (p − 30)²]: \$1,800 at \$30 and \$1,750 at either \$29 or \$31.

Step 3. Change the cost.

Let p_0 be the best price at cost c, and let p_1 be the best price after cost changes by Δc, which is negative for a reduction. Demand does not change, so a stays the same, and equation (3) gives

$$p_0 = (a + c)/2 \quad \text{and} \quad p_1 = (a + c + \Delta c)/2.$$

Subtracting, the a and c terms cancel:

$$p_1 - p_0 = \Delta c/2.$$

Because q = (a − p)/b, sales change by −Δc/(2b), half the −Δc/b change that passing on the full cost change would produce. In the example, Δc = −\$12, so the price falls by \$6 and sales rise by 50 × 6 = 300 clocks.

Proof 2: Calculus

This proof reaches the same result by the standard economic route, the rule that marginal revenue equals marginal cost. It is worth knowing because it does not need a formula for the best price. It is therefore the method that still works when demand is not a straight line, and the one economists use to show how pass-through changes when demand curves.

Marginal revenue. Marginal revenue is the change in total revenue from selling one more unit. It is less than the price, because selling one more unit requires lowering the price on all units. Total revenue is

$$pq = (a - bq)q = aq - bq^2.$$

Differentiating with respect to q gives marginal revenue: MR = a − 2bq. Marginal revenue falls by 2b for each additional unit, twice as fast as the price on the demand line, which falls by b. Here the firm is described as choosing a quantity rather than a price. Because each price on the demand line corresponds to exactly one quantity, choosing the best quantity and choosing the best price are the same problem.

Marginal revenue equals marginal cost. Profit is π(q) = aq − bq² − cq. At the top of the profit hill, a small change in quantity neither raises nor lowers profit, so the derivative of profit is zero:

$$d\pi/dq = a - 2bq - c = 0, \quad \text{or} \quad a - 2bq = c. \qquad (4)$$

Equation (4) says that marginal revenue equals marginal cost. If marginal revenue were higher, selling more would add to profit; if it were lower, selling less would. Because $d^2\pi/dq^2 = -2b < 0$, profit curves downward, so equation (4) identifies the maximum.

Step 1. Find how the best quantity responds to cost.

Write $q^*(c)$ for the profit-maximizing quantity at cost c. The star marks the best choice, and (c) shows that it depends on cost. Equation (4) becomes

$$a - 2b \cdot q^*(c) = c.$$

Differentiate both sides with respect to c. The derivative of a is zero because demand does not change, and the derivative of the right-hand side is 1. The quantity term changes because the best quantity changes with cost:

$$-2b \cdot dq^*/dc = 1, \quad \text{so} \quad dq^*/dc = -1/(2b).$$

A $1 increase in marginal cost reduces the best quantity by 1/(2b) units, half the 1/b units that a $1 price increase would cost. In the example, that is 25 clocks.

Step 2. Translate the quantity response into a price response.

The best price is the price on the demand line at the best quantity: $p^* = a - b \cdot q^*$. Differentiate with respect to c and substitute the result of Step 1:

$$dp^*/dc = -b \cdot dq^*/dc = -b \times [-1/(2b)] = 1/2.$$

Step 3. Apply the result to the cost change.

The derivative equals 1/2 at every cost, not just at the starting point. Anything whose rate of change never varies changes along a straight line, so a cost change of Δc changes the best price by exactly $\Delta c/2$.

The calculation has a simple geometric reading. When marginal cost falls by Δc, the point where marginal revenue meets marginal cost moves out by $\Delta c/(2b)$ units, because marginal revenue falls 2b per unit. Moving that many units along the demand line, which falls only b per unit, lowers the price by $b \times \Delta c/(2b) = \Delta c/2$. Because marginal revenue is twice as steep as demand, price moves half as far as cost.

What 50% pass-through does and does not mean

Dollars, not percentages. In the example, the $12 cost reduction is 50% of the retailer's $24 cost, but the $6 price cut is only 20% of the $30 price.

Any per-unit cost, but not fixed costs. The rule applies to anything that works like a cost of each additional sale, including profit the firm gives up elsewhere when it makes the sale, as long as that amount is the same for every additional sale. It does not apply to fixed costs. A saving that does not vary with the number of units sold leaves marginal cost, and therefore the best price, unchanged.

The assumptions. The result depends on each assumption listed at the start. The table shows what happens when one does not hold.

Assumption	If it does not hold
The demand curve is a straight line	Pass-through can be above or below 50%. If the demand curve bows inward toward the origin (it is convex), pass-through is above 50% and can exceed 100%. If it bows outward (it is concave), pass-through is below 50%.
Marginal cost is the same at every quantity	If marginal cost rises as output grows, pass-through is below 50%, because the added sales cost more to supply. If marginal cost falls as output grows, pass-through is above 50%.
Rivals' prices stay fixed	Rivals may change their prices in response, which shifts the firm's residual demand curve. Applying the 50% rule then gives only a first-round effect. A cost change that affects all competitors at once is usually passed on more than one that affects a single firm.
The firm is free to change its price	Contracts, regulation, or capacity limits can delay or prevent the price change.