What's in a Margin? Marginal Costs and Merger Efficiencies in Contracting Settings

Christopher Lau and Andrew E Sfekas

Feb 20, 2026 12 min read

Summary

- Marginal cost is an important input for modern merger analysis, and can differ based on how competitive prices are set.
- When contracts are negotiated, costs typically considered "fixed" may be appropriately classified as margins.
- In such settings, reductions in overhead costs may constitute marginal cost reductions that could be credited to a merger.

Introduction

Marginal cost—the cost of increasing production to serve an additional customer or produce an additional unit of a good—is an important input for modern merger analysis, as it enters both the analysis of competitive effects and the evaluation of potential efficiencies. Marginal costs are relevant because they affect pricing decisions. However, it is not often obvious which costs should be considered "marginal" and which should not. Costs that vary with production in the short run are usually marginal costs, while costs that remain fixed in the short run are usually not. However, analysis of what costs constitute marginal costs is case-specific and can depend on the structure of competition.

One setting where the dividing line as to which costs are marginal can be unclear occurs when buyers and sellers negotiate contracts for substantial unit quantities, i.e., where purchases are both large and lumpy. In these settings, such as healthcare provision and many business-to-business services, production could change substantially depending on whether a contract is won or lost, or a negotiation succeeds or fails. In such cases, production from the sale to an additional customer is different from the production of an additional unit because customers may purchase many units of the product. Such a sale could thus entail adding or removing salaried employees or setting up a new production line—costs that the firm's decision-makers could take into account when bargaining or bidding for a contract. As a result, whether costs should be considered as marginal or not may depend upon the margin being analyzed. A cost that may be categorized as fixed in some settings (e.g., because it does not change with producing an additional unit) may be considered marginal in other settings (e.g., because the transaction may involve large changes in unit production due to large contracts, so the margin is not a single unit but a contract or customer). In other words, marginal costs in large-quantity contracting settings may also need to be evaluated with respect to an additional *contract*, not just an additional *unit produced*—and strategic decision-making within the firm likely reflects this key difference.

Economists often consider the implications of such settings when considering how to model the dynamics of competition. However, even with the appropriate choice of model, it is still necessary to make an appropriate choice of marginal cost. For example, economists often use models such as the second-score auction, which assumes that buyers select among multiple bidders to supply the quantity needed to fulfill a contract, or make adjustments to approaches like upward pricing pressure to reflect that buyers and sellers may bargain over price. [1] However, these approaches are agnostic as to how to appropriately define marginal costs and are generally implemented using margins that reflect small scale, or incremental, changes in production, rather than the larger-scale

changes that a new contract may entail. Different costs may vary at different increments of change: some may vary at the unit production level while others may vary only at the contract level. For example, a firm could enter into a contractual relationship (e.g., via a procurement auction or a negotiated agreement) to be a supplier to a large-scale customer. In that case, the change in total production cost from the increment of a single unit may only tell part of the story. Parts of the organization, such as sales and administrative staff, factory capacity, etc. could be scaled if the company is faced with large changes in production—and depending on circumstances, an individual "sale" may involve such a large change.

Marginal costs also matter for merger efficiencies as economics predicts that changes in marginal costs will be passed through to price. An implication is that merger efficiencies could include changes in certain costs that we might traditionally consider to be fixed because they do not vary with unit production, but do vary with large-quantity production. Changes in such costs have historically not been considered as a cognizable efficiency because they are not viewed as reductions to marginal costs. For example, extra administrators needed for a large contract could be a source of efficiencies when treated as variable costs if the combined entity can reduce the number of administrators after a merger (since there are fewer contracts to oversee). However, under the logic outlined above, such cost savings may benefit customers in the same way as other marginal cost savings that are likely to be credited. As a result, such redundancies in these settings may be a candidate for cognizable efficiencies.

Competitive conditions help determine marginal costs

Regulatory guidance on estimating marginal or incremental costs recognizes that the change in output may vary across different settings in response to a small change in price. The 2023 Merger Guidelines suggest that the agencies can use accounting data, but that ordinary course margins may not align with "the concept of incremental cost that is relevant in economic analysis of a merger." [2] Relevant incremental costs might incorporate more or less than ordinary course margins, depending on the relevant change in output. [3]

The notion of "marginal" differs in settings where a small increase in price can potentially cause a significant reduction in sales, such as bargaining and procurement settings, compared to those where producers sell to individual consumers. In a traditional model of pricing, such as the Nash Bertrand models frequently used for consumer products, the "margin" is typically in reference to the production associated with a single sale: If a firm raised its prices by a small amount, a customer who exactly values the product at the old price may choose to no longer purchase the product, thus the decision to raise price will

lose that customer. That is, if the seller raises prices in a bargaining setting, the buyer may decide to no longer partner with the seller, which would reduce their sales (and thus production) to zero. Or, if a seller raises its bid in a procurement setting, the seller may lose the auction and will be unable to provide the good or service to the buyer at all. In these settings, the marginal cost is not defined by just the incremental customer, it is instead defined by the incremental *contract*. If the cost of increasing output from 100 to 101 units is the same as the average cost associated with increasing output from 100 to 1,000 units, then this would not matter—large changes to production could be treated similarly to small changes. However, this is not likely to be the case in a wide range of circumstances.

Consideration of where to draw the line between marginal and fixed costs affects both competitive effects analyses and analyses of efficiencies. From an economic perspective, efficiencies that reduce marginal costs are likely to be passed on to the consumer because they directly impact price setting behavior—when marginal costs fall, firms can increase profits by lowering their price such that they make the same margin on each sale, but are making more unit sales due to increased demand. Reductions to fixed costs, on the other hand, are unlikely to have an immediate impact on price setting—if a firm tried to lower its price, it would earn a lower margin on each sale since the cost that was reduced does not vary with the number of sales. As a result, because marginal cost reductions have a theoretical justification for benefits that are passed through to consumers, economists tend to emphasize these efficiencies when evaluating procompetitive effects of a merger. [4]

We provide an example below that suggests ways of considering what should be classified as marginal vs. fixed and consider how this classification could affect analysis of efficiencies. The example focuses on a "contracting" pricing setting, where a seller and a group of buyers agree to terms for the seller to provide the product to customers who are members of the buyer group, and compares marginal cost in that setting to how we might think about marginal cost in the same setting from a "retail" point of view—i.e., considering each customer as the relevant unit. Settings such as this occur in a range of industries, such as health insurance, employee vetting, and wholesale supply. In retail terms, marginal cost applies to production of a single unit, while in contracting, we argue that the definition of "marginal" should reflect all sales associated with the agreement, such that costs associated with an additional agreement are marginal, even if they may be considered fixed in a retail setting.

Example. **Hospital administration costs and merger efficiencies.** Consider an example of a hypothetical hospital merger. Hospitals negotiate with insurers over network inclusion; if a hospital is in an insurer's network, the hospital receives an agreed upon reimbursement for providing services to the insurer's members. Conversely, if the hospital is not included in the insurer's network, the insurer's members are likely to face

considerably higher costs for using it and will tend to choose other, in-network hospitals. Thus, the price negotiated between the hospital and insurer will hinge on the extent to which the failure of the negotiation results in a large swing in patient volume. Likewise, the cost that influences the negotiated price will reflect the costs associated with that large swing in patient volume.

The example in Table 1 illustrates this type of scenario. To begin, we focus on calculating marginal costs for Hospital 1, as shown in column (1) of the table. For ease of exposition, assume that each procedure done by the hospital is rendered to a single patient, such that demand volume can be measured in patients (i.e., patients do not receive multiple procedures). Hospital 1 has a total patient volume of 150, where 100 patients are insured by Insurer A and 50 patients are insured by Insurer B (rows [a]-[c]). For simplicity, we assume that the only cost to serving these patients is the cost of hiring administrators, and that one administrator can handle at most 65 patients. Each administrator earns a yearly salary of $100,000 that does not depend on the number of patients they manage. Thus, if the hospital contracts with both insurers and serves all 150 patients, it will need 3 administrators (see row [d]) for a total cost of $300,000 (see row [f]).

Table 1. Example of Marginal Cost Efficiencies in Hospitals

	(1) Hospital 1	(2) Hospital 2	(3) Combined
Procedures rendered (patients)			
[a] Total	150	100	250
[b] Insured by A	100	65	165
[c] Insured by B	50	35	85
Number of administrators			
[d] In-network with A and B	3	2	4
[e] Out-of-network with A, in-network with B	1	1	2
Total cost			
[f] In-network with A and B	$300,000	$200,000	$400,000

[g] Out-of-network with A, in-network with B	$100,000	$100,000	$200,000
Marginal Cost with A (when B is in-network)			
[h] Total: [f]—[g]	$200,000	$100,000	$200,000
[i] Per-patient: [h]/[b]	$2,000	$1,538	$1,212

If Hospital 1 comes to a disagreement with Insurer A such that it is no longer in A's network, the hospital contracts only with Insurer B and sees 50 patients (the patients insured by Insurer B, row [c]) assuming no recapture from customers who switch from Insurer A to Insurer B to retain access to the hospital. As a result, it would only need one administrator (see row [e]) so that the total cost is $100,000 (see row [g]). This means that from the perspective of Insurer A's contract, $200,000 from the 2 administrators who would scale up or down with Insurer A's contract would be the marginal cost of serving that contract (see row [h]). Since reimbursement rates between hospitals and insurers are often set on a per-patient basis, it is also useful to represent costs on a per-patient basis—marginal costs associated with Insurer A are $2,000 ($200,000/100) per patient (see row [i]). [5] This is significantly higher than the "traditional" marginal cost of $0 that one would get if considering the change in costs from scaling up or down by a single patient—i.e., the retail perspective.

A similar calculation applies to Hospital 2, shown in column (2) of the table. When Hospital 2 is in-network with both Insurer A and Insurer B, it sees 100 patients (see row [a]) and requires 2 administrators (see row [d]). If Hospital 2 goes out-of-network with Insurer A, it would see only 35 patients (see row [c]) and would only need one administrator (see row [e]). Thus, $100,000 are categorized as marginal (see row [h]), yielding a per-patient marginal cost in the negotiation between Hospital 2 and Insurer A of $1,538 ($100,000/65) per patient (see row [i]). [6]

One could make the example more realistic by including costs that vary at the patient level, which would contribute to marginal costs regardless of whether one were considering a change of one patient or 100 patients. At the other end of the spectrum, there could also be costs such as physical plant costs that cannot be scaled up or down with the number of patients and would contribute only to fixed costs. The addition of that level of complexity would not change the basic point here that some costs do not change with small movements in volume but do change with the movements in volume that result from changes in contracts.

Treating the costs of administrators as marginal has implications for merger effects. If a merger between Hospital 1 and Hospital 2 is consummated, as shown in column (3), the combined entity will see 250 patients (see row [a]), requiring four administrators (see row [d]). The administrative cost when contracting with all insurers post-merger is therefore $400,000 (4 administrators times $100,000 each; see row [f]). If the combined entity were to go out-of-network with Insurer A, they would see 85 patients (see row [c]), requiring only two administrators (see row [e]) at a cost of $200,000 (see row [g]). As a result, the marginal cost to serve Insurer A's contract is $200,000 (see row [h]). On a per-patient basis, marginal costs in negotiations with Insurer A are $1,212 ($200,000/165) per patient (see row [i]), lower than the marginal cost of both Hospital 1 and Hospital 2 when negotiating with Insurer A. Relative to the discharge-weighted average marginal cost of the parties before the merger of $1,815.20 ($2,000 * 0.6 + $1,538 * 0.4), this is a reduction in marginal cost of 33.2% that could potentially be credited as a cognizable efficiency even though the "traditional" marginal cost that varies by each patient did not change from the merger. [7]

Determining a reasonable value for marginal cost

When estimating marginal cost in an environment where large-scale changes to production could occur, practitioners will want to bear in mind that retrospective accounting data and the financial reports that come from them may not adequately reflect the types of changes that are expected to occur—but firms may have performed the types of analyses that would be of interest for identifying an appropriate marginal cost. Finding the appropriate financial analyses will be important for determining how the firms think about costs in their decision-making, and thus how costs should be treated in economic modeling of the proposed merger. For example, the firms may have forecasts that include scenarios where they gain or lose large customers, which may reflect the adjustments to administrative personnel, manufacturing footprint, or other considerations that would be fixed for smaller changes to production. One could find additional evidence in, e.g., natural experiments that identify the incremental cost associated with a firm winning (or losing) a contract, or the total firm costs before and after. An economist could examine the costs incurred by the firm before and after the contract was active (or terminated) to gain a better understanding of which costs are "marginal."

On the other hand, readily available financial reports may lack the information necessary to account for large-scale changes. For example, gross profit margins are often reported in public filings but may be of limited utility. These margins are calculated as the difference between total revenue and the cost of goods sold ("COGS"), which are the direct costs for producing the product, which may include raw materials, wages of workers,

manufacturing overhead (e.g., factory rent, utilities, and depreciation), and packaging and shipping costs. These could be adjusted to exclude overhead as fixed and to include operating costs that are also scalable with production. However, it could be challenging to use them to predict a production change of sufficient size as to affect additional categories of costs.

Alternatively, some academic studies in the industrial organization literature have assumed that the marginal cost relevant to decision-making is not possible to extract from available financial data, and have instead used models to infer marginal costs from consumer demand, prices, and assumptions about the nature of competition. This approach has been developed in several seminal articles in industrial organization. [8] In particular, one can infer the marginal cost of a product or service by estimating the elasticity of demand, then using that elasticity, as well as a set of assumptions about competitive conditions, to derive the marginal cost that would yield the observed price. While many academic studies have taken this approach (in part because of an inability to obtain detailed cost information from private firms), it is often impractical for merger evaluation due to data availability, the computational burden of estimating reliable demand elasticities, or the use of additional assumptions that may not always be appropriate for the setting being evaluated.

Regardless of whether one relies on economic models or income statements, one would rarely end up with margins that match those reported in income statements. In fact, the 2023 Merger Guidelines suggest that "[t]he Agencies may use accounting data to measure incremental costs, but they do not necessarily rely on accounting margins recorded by firms in the ordinary course of business because such margins often do not align with the concept of incremental cost that is relevant in economic analysis of a merger." [9] As a result, careful consideration of industry dynamics and how to think about the relevant "margin" in which prices are being set should play an important role in any antitrust analysis.

Conclusion

When industry-specific factors indicate that there may be large production shifts, it may be appropriate to include costs that are variable for a large shift, but that might be fixed for a small shift, in production—such as the cost of a contract administrator who would be unaffected by unit changes in production but may be affected by the gain or loss of a large contract. Practitioners should evaluate the impact that industry and firm structure have on what costs should be considered marginal—potentially leading to significant implications for predicting merger price effects and efficiencies.

Endnotes

1. Upward pricing pressure (UPP) has been included as part of the FTC and DOJ's jointly issued merger guidelines since 2010. UPP can be applied to bargaining settings; see Christopher Garmon, The Accuracy of Hospital Merger Screening Methods, 48 RAND J. Econ. 1068 (2017). Extensions to UPP that also use margins as an input have been applied to raising rivals' costs in vertical mergers and to bilateral bargaining; see Serge Moresi & Steven C. Salop, vGUPPI: Scoring Unilateral Pricing Incentives in Vertical Mergers, 79 Antitrust L.J., 185 (2013) and Keith Brand & Christopher V. Lau, Upward Pricing Pressure in Mergers Where Prices Are Determined Through Joint Bargaining, 73 J. Ind. Econ, 523–539 (2025), henceforth Brand & Lau (2025)

2. 2023 Merger Guidelines, Endnote 72.

3. The 2010 Horizontal Merger Guidelines make this notion more explicit. See 2010 Horizontal Merger Guidelines, § 4.1.3 ("The profit margin on incremental units is the difference between price and incremental cost on those units. The Agencies often estimate incremental costs, for example using merging parties' documents or data the merging parties use to make business decisions. Incremental cost is measured over the change in output that would be caused by the price increase under consideration.").

4. For example, there is a line of academic research known as "compensating marginal cost reduction" that evaluates mergers by calculating the amount marginal costs must reduce to offset upward pricing pressure from a merger; see Gregory J. Werden, A Robust Test for Consumer Welfare Enhancing Mergers Among Sellers of Differentiated Products, 44 J. Ind. Econ., 409–413 (1996); and Joseph Farrell and Carl Shapiro, Antitrust Evaluation of Horizontal Mergers: An Economic Alternative to Market Definition, 10 BE J. of Theoretical Econ, 1-39 (2010). Also see Brand and Lau (2025).

5. If the Hospital 1 were to go out-of-network with Insurer B, it would still need two administrators—and therefore only $100,000 would be marginal ($300,000 when A and B are in-network minus $200,000 when B is out-of-network). This again yields a marginal cost of $2,000 per patient. The fact that the Hospital 1's marginal cost is the same for Insurer A and B is a coincidence—under a different setup, it could have been the case that marginal cost was different from the perspective of Insurer A than from Insurer B.

6. If the Hospital 2 were to go out-of-network with Insurer B, it would also need only one administrator. This means Hospital 2's marginal cost when negotiating with Insurer B, $2,857 ($100k/35) per patient.

7. The combined entity would still require 3 administrators if it were to go out-of-network with Insurer B. This means the total cost when going out-of-network with Insurer B is $300,000, and the total marginal cost is $100,000. Because the combined entity sees 85 patients from Insurer B, the post-merger per-patient marginal cost is $1,176.47 ($100,000/85). As a result, the merger reduces marginal cost by 54% when negotiating with Insurer B relative to the discharge-weighted average of pre-merger marginal costs with Insurer B.

8. See Steven Berry, James Levinsohn, and Ariel Pakes, Automobile prices in market equilibrium, 63 Econometrica, 841–890 (1995). Also see Aviv Nevo, Measuring market power in the ready-to-eat cereal industry, 69 Econometrica, 307–342 (2001)

9. 2023 Merger Guidelines, footnote 72.