Merger-Enabled Innovation and Upward Pricing Pressure: A Local Theory-of-Benefit Test

Kirsten Edwards-Warren* Francisco Espinosa† Salvatore Piccolo‡

August, 2026

Abstract

Merger review often quantifies unilateral pricing incentives while treating claimed innovation efficiencies as exogenous. We develop a local theory-of-benefit test for horizontal mergers said to enable process innovation. The test separates benefit generation from benefit transmission. A project is merger-enabled only if the transaction raises its expected operating return above implementation cost while realistic less restrictive arrangements do not. The resulting product-level efficiency vector then enters benefit-adjusted upward pricing pressure for each merging product. Common ownership does not create a direct investment incentive when an investment only lowers the investor's own cost and that benefit was already internalized. Moreover, efficiencies on products receiving diverted demand can strengthen the incentive to raise another product's price, so aggregate savings are not sufficient. We extend the framework to uncertainty, multiproduct ownership, demand-enhancing innovation, and local price–investment feedback. The method provides an implementable first screen based on pre-merger quantities, margins, diversion, pass-through, and contemporaneous project evidence.

Corresponding author: Salvatore Piccolo, University of Bergamo, via dei Caniana 2, 24127 Bergamo, Italy. Email: salvatore.piccolo@unibg.it

Keywords: horizontal mergers; efficiencies; innovation; upward pricing pressure; merger specificity; theory of benefit.

JEL codes: L13, L41, O31, D43.

1 Introduction

The 2026 Draft EU Merger Guidelines give efficiency claims a more explicit place in merger review. Where merging parties rely on efficiencies, they are expected to articulate a "theory of benefit" explaining how the transaction generates gains and how those gains benefit customers. The draft places particular emphasis on dynamic efficiencies, including changes in the parties' ability or incentive to invest and innovate. Yet, this groundbreaking opening creates a practical challenge. If theories of benefit are to enter the competitive assessment, authorities need a disciplined way to distinguish innovation genuinely enabled by the merger from projects that would occur under the counterfactual, and benefits likely to reach

*Econic Partners, LLC
†Econic Partners, LLC
‡University of Bergamo and Econic Partners, LLC

customers from gains that remain private or reinforce the merged firm's pricing incentives (European Commission, 2026).

A robust first-screening test is therefore essential. It must be demanding enough to reject speculative claims, but parsimonious enough to be used before a full merger simulation or a detailed dynamic model. It should rely on evidence available early in the review, make the causal assumptions visible, and identify the claims that warrant deeper investigation. Without such a structure, credible innovation benefits may be discounted because they are difficult to quantify, while weak claims may receive undue weight because the link between the transaction, the project, and customer benefit is left implicit.

Yet, on the harm side, merger economics already provides practical local tools. Compensating-efficiency tests and upward pricing pressure translate pre-merger margins, diversion, and demand information into a first-order prediction of post-merger pricing incentives (Werden, 1996; Farrell and Shapiro, 2010; Jaffe and Weyl, 2013). Innovation claims are usually handled less rigorously. The parties identify a project, estimate the resulting cost reduction, and ask whether it offsets the merger's price effect. The efficiency is thus treated as an input even though the first factual question is whether the transaction will cause the project to be undertaken.

This paper develops a local theory-of-benefit test that answers two questions in sequence. The first stage concerns *benefit generation*. It evaluates a specified project at the observed pre-merger sales base, compares its expected operating return with implementation expenditure, and then performs the same calculation for realistic alternatives, including standalone investment, licensing, supply agreements, joint ventures, and narrower asset combinations. A strong local theory of benefit requires the project to cross the investment threshold under the merger while the best realistic less restrictive alternative does not. The second stage concerns *benefit transmission*. It inserts the predicted product-level efficiency vector into the post-merger price gradient and produces a benefit-adjusted upward-pricing-pressure measure for each product in the merging portfolio.

The two stages translate the familiar requirements of verifiability, merger-specificity, and customer benefit into separate economic objects (Competition and Markets Authority, 2021; U.S. Department of Justice and Federal Trade Commission, 2023; European Commission, 2026). Project evidence and implementation costs discipline whether the claimed gain is credible. Comparison with alternative arrangements operationalizes merger specificity. Product-level pricing pressure then asks whether the resulting gain is transmitted strongly enough to offset the unilateral effect of common ownership. The test is not a substitute for a full investigation. It is an early screen for whether a theory of benefit is economically coherent and where richer evidence or simulation is required.

The framework is general enough and allows an arbitrary number of differentiated-product firms and any subset of them to merge. This matters because both harm and benefit are portfolio objects. A price increase on one merging product can divert sales to several products that become jointly owned, while a common project can generate different cost reductions across those products. The analysis therefore remains product-specific and aggregates only when the economic structure permits it.

Two results sharpen the assessment. First, common ownership does not by itself create a direct investment incentive. If an investment lowers only the marginal cost of the investor's own product, and the firm already appropriated that benefit before the merger, the local investment condition is unchanged. A merger-induced incentive requires a newly internalized cross-product or cross-firm benefit, a common capability that can be deployed over a larger sales base, a change in appropriability, or a genuine feasibility

change. Second, the location of efficiencies matters as much as their total amount. An own-product cost reduction lowers the incentive to raise that product's price, whereas a cost reduction on another product in the merged portfolio can make diverted sales more profitable. A project can therefore be genuine and merger-specific, yet still be too small, too delayed, or too unevenly distributed to offset unilateral pricing pressure.

The methodology is purposefully local in the sense that its baseline inputs are pre-merger quantities, margins, diversion ratios, expected unit savings, project expenditure, success evidence, and realistic alternatives. When richer evidence is available, the screen can incorporate uncertainty, pre-existing multiproduct ownership, quality or demand effects, pass-through, and price–investment feedback. The contribution is to endogenize the efficiency input within conventional first-order merger analysis and provide an implementable screen for the generation and transmission of dynamic efficiencies.

We also show that the main insights of the framework are robust to natural extensions of the baseline set-up. First, we allow project success to be uncertain, which makes it possible to distinguish the expected contribution of the claimed efficiencies from the risk that they fail to materialize in the states in which prices are ultimately set. Second, we account for pre-existing multiproduct ownership. This clarifies that only ownership links newly created by the merger generate additional direct pricing pressure, while efficiencies on all products included in the post-merger portfolio may still affect how those benefits are transmitted across products. We then allow merger-induced price changes to alter the quantity base over which innovation benefits are earned. This provides a simple correction for the possibility that higher prices reduce output and thereby weaken the merged firm's incentive to undertake the project. We also extend the framework to innovations that improve product quality or shift demand, rather than merely reducing marginal costs. In this case, the analysis must account not only for cost savings but also for demand expansion, cannibalization across products, changes in diversion patterns, and the direct benefits that quality improvements confer on consumers. Finally, when the relevant derivatives are available, we solve for prices and innovation jointly. This makes it possible to assess the feedback running in both directions: price changes affect the returns to innovation, while innovation changes firms' pricing incentives. Together, these extensions identify which conclusions are robust and which depend on additional evidence about project risk, ownership structure, demand, quality, and price–investment feedbacks.

The remainder of the paper is organized as follows. Section 2 places the approach in the relevant literature. Section 3 develops the project-generation and merger-specificity tests. Section 4 derives benefit-adjusted upward pricing pressure. Section 5 brings the two stages together in a compact merger-to-monopoly illustration. Section 6 develops the main extensions, and Section 7 concludes. Appendices A and B contain the proofs of all formal results stated in the paper. Online Resource 1 retains the dynamic normalization, comparative statics, expanded practical applications, proofs of supplemental results, and additional implementation derivations.

2 Related literature

2.1 UPP, compensating efficiencies, and first-order merger analysis

The paper builds directly on the literature that replaces a full equilibrium merger simulation with a local comparison of the merger's first-order incentives. Werden (1996) derives a robust compensating

marginal-cost-reduction test for differentiated-product mergers. Farrell and Shapiro (2010) formulate upward pricing pressure as the difference between the value of diverted sales and an efficiency credit. The attractiveness of the measure is informational economy: it can often be constructed from margins, diversion, and a transparent efficiency assumption.

In the familiar two-product presentation, direct pressure on one product depends on diversion to the other product and the other product's margin. The same logic extends naturally to an arbitrary merging set. For any product controlled by the merged firm, direct pressure is the sum of diversion-weighted margins on every other newly internalized product, net of the own-product efficiency. The simplified implementation is not the entire post-merger profit gradient. As Farrell and Shapiro (2010) note, a reduction in the cost of a product that receives diverted demand raises the value of those diverted sales and therefore strengthens the incentive to raise the originating product's price. The full local object must thus include diversion-weighted efficiencies on all other products in the merged portfolio. This cross-efficiency term becomes especially important when a common project generates uneven savings across several products.

Jaffe and Weyl (2013) develop a general first-order approach that maps local profit incentives into predicted price changes through a pass-through matrix. Willig (2011) extends upward-pricing-pressure reasoning to product quality and other competitive variables. The simulation evidence in Miller et al. (2017) shows that upward pricing pressure can be informative about merger price effects, while also emphasizing the importance of pass-through and functional-form choices. Our contribution is upstream of these tools: rather than accepting the efficiency vector as given, we derive a local prediction of whether the efficiency project is undertaken and how large it is.

The treatment of efficiencies has a much older lineage. Williamson (1968) formalizes the tradeoff between market-power losses and production efficiencies, while Farrell and Shapiro (1990) studies the equilibrium welfare effects of cost-saving horizontal mergers. Choné and Linnemer (2008) analyze merger assessment when efficiency gains are uncertain. We retain uncertainty but focus on a different margin: the probability and intensity of an investment are themselves chosen, and the merger may change whether the project exists at all.

2.2 Mergers, innovation, and investment incentives

A large literature studies the effect of a reduction in the number of competitors on innovation. The classic tension runs from the escape-rivalry logic associated with Arrow (1962) to the appropriability and scale logic associated with Schumpeter (1942). Modern merger models show that the sign depends on the form of innovation, strategic interaction, business stealing, spillovers, and synergies.

Federico et al. (2017) and Federico et al. (2018) show in a tractable product-innovation model that a merger can reduce innovation incentives through the internalization of business stealing, although efficiencies can reverse the result. Motta and Tarantino (2021) analyze firms that compete in both prices and investments and highlight the output channel: a merger-induced reduction in output can weaken the return to cost-reducing investment. Bourreau et al. (2026) provide a unified treatment of incremental cost-reducing and demand-enhancing innovation and distinguish price diversion from innovation diversion. Their decomposition clarifies when cannibalization and margin-expansion forces increase or reduce post-merger investment. Reviews by Jullien and Lefouili (2018), Régibeau and Rockett (2019), and Lefouili and Madio (2026), together with the unified modeling work of Moraga-González and

Motchenkova (2026), emphasize that no general monotonic relationship exists between concentration and innovation. Empirical evidence is similarly heterogeneous; for example, Haucap et al. (2019) find reductions in R&D and patenting following pharmaceutical mergers in their sample.

Our question is narrower from a positive perspective, but complementary from a normative angle. We do not seek a general sign for the effect of merger on an innovation technology that is already available to both firms. Instead, we study a specified project whose productivity, implementation cost, or appropriability changes with the organizational arrangement. Examples include integrating two production systems, combining complementary datasets, redesigning a common platform, reallocating capacity across formerly separate networks, or deploying a technology across both product bases. The relevant counterfactual is therefore not merely "two firms rather than one," but a menu of feasible arrangements under which the project may or may not cross its investment threshold.

Two contributions are particularly close. Salinger (2019) proposes net innovation pressure, an analogue to UPP in which the effect of a merger on innovation depends on the relationship between product diversion and R&D spillovers. That framework is designed to assess how horizontal internalization changes an existing R&D incentive. We instead construct an arrangement-specific project hurdle and feed the predicted project output into price UPP. Pietola et al. (2026) explicitly study merger-specific innovation synergies and show, in a procurement model, that a merger can increase investment and buyer surplus if the innovation gain exceeds the merger premium. Our objective is a more reduced-form implementation for differentiated-product pricing markets with arbitrary merging sets. To the best of our knowledge, the existing literature does not combine an arrangement-specific project-generation hurdle with product-level UPP transmission in a single general local screen for arbitrary merging sets.

2.3 Efficiencies and the evidentiary theory of benefit

Merger policy generally requires efficiencies to be verifiable, merger-specific, and likely to benefit customers. Competition and Markets Authority (2021) and U.S. Department of Justice and Federal Trade Commission (2023) make that structure explicit in their merger guidelines, and OECD (2025) provides a recent policy synthesis. Innovation-specific implementation issues are discussed by, among others, Kokkoris and Valletti (2020). European Commission (2026) uses the language of a theory of benefit and explicitly links dynamic efficiencies to evidence about the merged firm's ability and incentive to invest or innovate relative to the counterfactual. It also identifies standalone expansion, joint ventures, licensing, and other contractual arrangements as relevant alternatives in the merger-specificity inquiry.

The proposed test translates these criteria into separate economic objects. *Ability* is represented by the arrangement-specific efficiency technology. *Incentive* is represented by the local project score. *Merger specificity* is a comparison of scores and, where necessary, project net present values across alternatives. *Consumer benefit* is assessed through product-level benefit-adjusted UPP and, when pass-through is available, a first-order price or consumer-surplus prediction. The framework therefore disciplines the sequence of the argument: a claimed benefit must first be generated, then transmitted.

3 Model

This section develops the first stage of the screening exercise. The goal is not to predict the full post-merger equilibrium. It is to ask a more basic causal question: does the transaction change the private

return to a concrete innovation project enough to make that project occur, and would a realistic less restrictive arrangement generate the same result? Keeping this question separate from pricing is useful because the required evidence is different. Project plans, engineering estimates, implementation budgets, internal hurdle rates, and contractual alternatives discipline benefit generation; margins, diversion, and pass-through discipline benefit transmission. The analysis is local in the same sense as conventional UPP. Project returns are first evaluated at the observed pre-merger price vector and sales base. This normalization makes the screen transparent and avoids embedding an unverified efficiency claim in a full equilibrium model. It does not assume that quantities will literally remain fixed after the merger. Sections 4.3 and 4.5 show how to add price–quantity feedback when the necessary derivatives are available.

Environment. There are $N \geq 2$ firms indexed by $i \in \mathcal{N} \equiv \{1, \ldots, N\}$. Each firm initially sells one differentiated product and chooses its price. A subset $\mathcal{M} \subseteq \mathcal{N}$, with $r \equiv |\mathcal{M}| \geq 2$, proposes to merge. Firms outside $\mathcal{M}$ remain independently owned. The bilateral merger is the special case $r = 2$; allowing $r \geq 2$ keeps the analysis product-specific when a transaction combines larger portfolios.

Let $p = (p_1, \ldots, p_N)'$ and let demand for product i be

$$q_i = q_i(p),$$

where each demand function is twice continuously differentiable. Write $q_{j,i} \equiv \partial q_j / \partial p_i$. Own-price effects satisfy $q_{i,i} < 0$. Cross-price effects may be positive or negative, so the general formulas use signed diversion; results that invoke nonnegative matrices assume that the relevant products are substitutes.

The observed pre-merger equilibrium is denoted by

$$p^0 = (p_1^0, \ldots, p_N^0)', \qquad q_i^0 = q_i(p^0), \qquad m_i^0 = p_i^0 - c_i^0,$$

where c_i^0 is firm i's constant marginal cost. Because each firm is single-product before the merger, its first-order condition is

$$q_i^0 + m_i^0 q_{i,i}^0 = 0, \qquad i \in \mathcal{N}.$$

For $i \neq j$, signed diversion from product i to product j is

$$D_{ij} = \frac{q_{j,i}^0}{-q_{i,i}^0}, \qquad D_{ii} = 0.$$

Let $\mathbf{D}_{\mathcal{M}}$ denote the $r \times r$ matrix containing the diversion ratios between products in the merging set. Its rows need not sum to one: substitution to products outside $\mathcal{M}$ and to the outside good is not internalized by the merger.

Merger-enabled process innovation. The transaction is said to make available, improve, or accelerate an R&D or implementation project with choice variable $x \geq 0$. In applications, x is most naturally read as investment expenditure or as an engineering intensity that can be mapped into expenditure. The project produces expected marginal-cost reductions on product $i \in \mathcal{M}$,

$$e_i^M(x) \geq 0, \qquad e_i^M(0) = 0.$$

Set $e_i^M(x) = 0$ for products outside the merging set. The superscript M denotes implementation under the

merger. If x is expenditure, the resource cost $K_M(x)$ can simply be set equal to x. More generally, whenever K_M is increasing and convex, the model can be written in expenditure units by defining $x = K_M(z)$ and $\widehat{e}_i^M(x) = e_i^M(K_M^{-1}(x))$. A more difficult project then appears as a flatter expenditure-to-saving function or as a larger expenditure requirement for a given saving.

Notice that the efficiency functions may already incorporate engineering risk. For example,

$$e_i^M(x) = s_M(x)\Delta_i^M,$$

where $s_M(x)$ is the probability of successful implementation and Δ_i^M is the cost reduction conditional on success. The scalar formulation describes a common project whose intensity affects several products. Product-specific workstreams can be handled by replacing x with a vector; the economic logic is unchanged, but the scalar version keeps the screen and its evidentiary inputs more transparent.

At fixed market prices, the merged firm's operating profit net of project cost is

$$\Pi^M(p,x) = \sum_{i\in\mathcal{M}} \left[p_i - c_i^0 + e_i^M(x)\right] q_i(p) - K_M(x).$$

Evaluating the incremental project return at the observed price vector gives the local objective

$$V_M^L(x) = \sum_{i\in\mathcal{M}} q_i^0 e_i^M(x) - K_M(x),$$

where terms independent of x have been omitted without loss of generality. Each product contributes its observed quantity multiplied by the expected saving generated on that product. This is the project-side counterpart to using observed margins and diversion in a local pricing test.

3.1 The local project problem

The first basic question is whether the merger moves the project away from the no-investment boundary. For a smooth project—i.e., a project for which the investment choice is a continuous variable—the relevant first screen is the slope of the local objective at $x = 0$. It compares the marginal operating return generated across the merged portfolio with the marginal resource cost of beginning or expanding the project. Under concavity, this single number identifies the extensive decision at the boundary and, when positive, points toward a positive local investment. Fixed launch costs and genuinely indivisible projects require a total-NPV test and are treated separately below.

Definition 1 (Local innovation incentive)**.** *The merger's local innovation incentive is*

$$\mathrm{LII}_M \equiv \left.\frac{\mathrm{d}V_M^L(x)}{\mathrm{d}x}\right|_{x=0} = \sum_{i\in\mathcal{M}} q_i^0 e_i^{M\prime}(0) - K_M'(0).$$

The contribution $q_i^0 e_i^{M\prime}(0)$ has a direct interpretation. It is the incremental operating-profit gain from a marginal increase in project intensity on product i, evaluated on the observed sales base. The term $K_M'(0)$ is marginal project expenditure; it equals one when x is measured directly in expenditure units. The score adds the product-level returns that the merged firm internalizes, but it does not discard their incidence across products. That incidence will matter when the same project is placed in the pricing test.

Assumption 1 (Local concavity)**.** *For each feasible organizational arrangement a, $e_i^a(x)$ is increasing and concave and $K_a(x)$ is increasing and convex on the relevant local domain. Moreover, the arrangement-specific local objective $V_a^L(x)$ is strictly concave and attains a maximum on a domain containing a right-neighborhood of zero.*

The following then holds.

Proposition 1 (Local project generation)**.** *Under Assumption 1, the merger generates a strictly positive local project level if and only if* $\mathrm{LII}_M > 0$*. If* $\mathrm{LII}_M \leq 0$*, $x = 0$ maximizes the local project objective. If* $\mathrm{LII}_M > 0$*, every maximizer of the local project objective is strictly positive.*

The proposition is a boundary result, not a complete prediction of the post-merger R&D program. It says that, given the observed commercial base, the merger-side project objective initially rises only when the marginal portfolio benefit exceeds marginal implementation cost. It does not yet allow prices and quantities to reoptimize, and it does not by itself establish merger specificity. Those are separate steps rather than hidden assumptions inside the score.

As an illustrative example, suppose the merger combines two products with the same observed quantity q^0, the project creates the same marginal expected saving η on each product, and $K'_M(0) = k_M$. Then

$$\mathrm{LII}_M = 2q^0\eta - k_M.$$

The smooth project is locally generated exactly when $2q\eta > k_M$. This is the simplest operational reading of the generation test: the merger-side sales base multiplied by marginal project productivity must exceed the marginal implementation hurdle.

3.2 When does the merger change investment incentives?

A positive merger-side project score should not be read as saying that common ownership mechanically creates an investment incentive. The correct comparison is with the return already internalized before the transaction. In the simplest own-cost model, common ownership creates no direct local increment at all. Suppose that, before the merger, firm i chooses expenditure x_i that lowers only the marginal cost of product i:

$$\pi_i(p, x_i) = \big[p_i - c_i + e_{ii}(x_i)\big]q_i(p) - K_i(x_i).$$

At fixed prices, the investment first-order condition is

$$e'_{ii}(x_i)q_i(p) - K'_i(x_i) = 0.$$

After the merger, the merged firm maximizes the sum of the profits of the products it owns. If x_i affects neither the cost, quality, nor demand of any other merging product, then

$$\frac{\partial \Pi^M}{\partial x_i} = \frac{\partial \pi_i}{\partial x_i}.$$

Thus an investment level satisfying the pre-merger own-product condition also satisfies the fixed-price local post-merger condition. Prices may subsequently change and alter the output base, but there is no *direct* newly internalized investment benefit at the pre-merger point.

Hence, a direct merger-induced effect requires a benefit that was not fully internalized before the transaction. Let $e_{ki}(x_i)$ denote the cost reduction on product k created by expenditure x_i. The post-merger investment derivative is

$$\frac{\partial \Pi^M}{\partial x_i} = \sum_{k \in \mathcal{M}} e'_{ki}(x_i) q_k^0 - K'_i(x_i).$$

Using the pre-merger own-product condition, the direct residual is

$$\text{UIP}_i \equiv \sum_{k \in \mathcal{M} \setminus \{i\}} e'_{ki}(x_i) q_k^0.$$

This is the investment-side analogue of UPP. UPP captures the newly internalized value of diverted sales. UIP captures the newly internalized value of project benefits on other merging products. If a common technology previously deployed only on product i can, after the merger, be deployed on products i and j, the residual at the pre-merger investment level is $e'_i(x_i) q_j^0$. The gain comes from the additional sales base, not from relabeling an own-product saving as merger-enabled.

3.3 Merger specificity and alternative arrangements

A project is merger-enabled only relative to a counterfactual. The screen therefore repeats the project calculation for the best realistic less restrictive arrangements rather than comparing the merger with an undefined "no project" benchmark. It is useful to distinguish a feasibility claim from an incentive claim. Under a feasibility claim, the project is technically, legally, or organizationally unavailable absent the merger. The counterfactual then restricts the feasible choice set. Under an incentive claim, the project is feasible but would be unprofitable, smaller, slower, riskier, or less productive. The analyst must then identify the cost, productivity, timing, risk, or appropriation mechanism that changes the project return.

Let $\mathcal{A}$ be the set of feasible realistic alternatives, such as standalone investment by one or more parties, a joint venture, licensing, a supply agreement, or a partial asset combination. Arrangement a is characterized by three objects:

(i) an expected technical cost-reduction function $e_i^a(x)$ for each affected product;

(ii) an implementation expenditure schedule $K_a(x)$, or simply $K_a(x) = x$ when x is measured as expenditure; and

(iii) an appropriation weight $\omega_i^a \in [0, 1]$, equal to the fraction of the operating gain on product i internalized by the project decision maker under arrangement a.

For the merger, $\omega_i^M = 1$ for every $i \in \mathcal{M}$. Under standalone investment, the sponsor may internalize gains only on its own product. A joint venture or licensing arrangement can have intermediate weights and a different technology, cost schedule, timing profile, or probability of success.

The arrangement-specific local objective is

$$V_a^L(x) = \sum_{i \in \mathcal{M}} \omega_i^a q_i^0 e_i^a(x) - K_a(x),$$

and its local innovation incentive is

$$\text{LII}_a = \sum_{i \in \mathcal{M}} \omega_i^a q_i^0 e_i^{a\prime}(0) - K'_a(0).$$

The same formula is used for every arrangement; only the technology, cost, and appropriation inputs change. This makes the counterfactual contestable: a party cannot obtain merger specificity merely by assuming a worse standalone technology without identifying why it is worse.

Definition 2 (Strong local merger specificity)**.** *The project is strongly merger-specific in the local sense if*

$$\mathrm{LII}_M > 0 \quad \textit{and} \quad \max_{a \in \mathcal{A}} \mathrm{LII}_a \leq 0. \tag{1}$$

The following then holds.

Proposition 2 (Local theory of benefit)**.** *Under Assumption 1, condition eq.* (1) *is equivalent to the following statement: the merger produces a positive local investment, while every modeled less restrictive alternative selects zero local investment.*

The condition separates channels that are often bundled under the word "synergy." The merger may improve technical productivity, reduce implementation cost, accelerate deployment, lower risk, or improve appropriability. The decomposition matters for evidence. Engineering records speak to technical productivity, integration budgets to implementation cost, project calendars to timing, and contracts, bargaining records, intellectual-property rights, and leakage risks to appropriability. If scale alone drives the difference, the technology should not be made artificially more productive under the merger; the larger quantity exposure should do the work.

3.4 A linear-quadratic specification

A screening tool needs a transparent benchmark that can be implemented with a small number of inputs. The following linear-quadratic specification offers that benchmark. The restrictions apply to the project technology and implementation cost, not to product demand. Demand can remain fully nonlinear; only observed quantities enter the project calculation at this stage.

For every arrangement a, suppose that

$$e_i^a(x) = \eta_i^a x, \qquad K_a(x) = k_a x + \frac{\kappa_a}{2} x^2, \qquad \kappa_a > 0.$$

The coefficient η_i^a is the expected marginal saving on product i from one unit of project intensity. It may combine engineering productivity and the probability of successful implementation. The parameter k_a is the marginal implementation hurdle at the origin; it is a variable cost of starting to scale the project, not a fixed launch cost. The curvature κ_a captures increasing marginal difficulty as the project becomes more ambitious.

Define the arrangement's effective project base

$$A_a = \sum_{i \in \mathcal{M}} \omega_i^a q_i^0 \eta_i^a.$$

This scalar is the quantity- and appropriation-weighted marginal productivity of the project. A product contributes more when it has a larger observed sales base, when the project is more effective on that product, or when the decision maker captures a larger share of the resulting operating gain. The local objective becomes

$$V_a^L(x) = (A_a - k_a)x - \frac{\kappa_a}{2} x^2,$$

so that

$$\mathrm{LII}_a = A_a - k_a, \qquad x_a^L = \left[\frac{A_a - k_a}{\kappa_a}\right]_+.$$

Hence, there are two relevant margins. The project is activated only when $A_a > k_a$. Conditional on activation, its intensity rises with the gap between marginal operating benefit and marginal implementation cost and falls with the rate at which implementation becomes more difficult. For an active project, larger quantities, stronger technical productivity, and broader appropriation all increase x_a^L at rate proportional to $1/\kappa_a$; a larger k_a or κ_a reduces it.

The maximized variable project surplus, before any fixed launch cost, is

$$W_a^L \equiv \max_{x \geq 0} V_a^L(x) = \frac{\left[A_a - k_a\right]_+^2}{2\kappa_a}.$$

This expression will be useful below because a fixed cost must be compared with total project surplus, not with the derivative at the origin. In the smooth baseline, strong local merger specificity reduces to the transparent set of inequalities

$$A_M > k_M, \qquad A_a \leq k_a \quad \text{for every } a \in \mathcal{A}.$$

The conditions show exactly where the merger-specificity claim must come from: a larger effective project base, a lower marginal implementation hurdle, or both.

The workhorse also clarifies the separation between generation and transmission. Once the project is selected, its product-level efficiency profile is

$$e_i^M(x_M^L) = \eta_i^M \left[\frac{A_M - k_M}{\kappa_M}\right]_+.$$

The scalar term determines how much innovation is generated; the vector $(\eta_i^M)_{i \in \mathcal{M}}$ determines where the savings appear. Two projects can therefore have the same total private return and the same investment intensity but different pricing implications because their savings are distributed differently across products.

The specification is useful for three reasons. First, it is the natural second-order local approximation to a smooth project objective: marginal benefits can be held constant over a small neighborhood and the first relevant curvature placed in implementation cost. Second, it yields closed forms that can be audited in a spreadsheet and linked directly to quantities, technical estimates, appropriation assumptions, and project-cost records. Third, it keeps the causal structure visible. The merger may affect the project base A_a, the marginal hurdle k_a, the curvature κ_a, or the product-incidence vector η^a, and each channel has a different evidentiary meaning.

3.5 Discrete projects and fixed launch costs

The baseline model presented above can fail for two independent reasons. A project may be nonconvex because a fixed cost must be paid before any benefit is realized, or because the project is intrinsically indivisible. Moreover, its benefits may persist for several periods, possibly after an implementation delay. Fixed costs change the extensive launch decision; persistence changes the discounted operating base on which the project is valued. Conflating the two can lead to the wrong hurdle and, in applications, to inconsistent comparisons between project cost, price effects, and consumer benefit.

3.5.1 Fixed launch costs and project indivisibilities

Suppose arrangement a requires a nonrecoverable launch cost $F_a \geq 0$ whenever $x > 0$. Conditional on launching, the project can still be adjusted smoothly. The relevant local net present value is then

$$\mathrm{NPV}_a^L = \sup_{x>0} \left\{ \sum_{i \in \mathcal{M}} \omega_i^a q_i^0 e_i^a(x) - K_a(x) \right\} - F_a. \tag{2}$$

The project is undertaken when $\mathrm{NPV}_a^L > 0$; at equality the firm is indifferent and we adopt the conservative convention of no launch. The derivative LII_a remains informative about whether the variable part of the project initially creates value, but it is no longer sufficient for launch: a project can have $\mathrm{LII}_a > 0$ and still fail to recover F_a. Conditional on launch, a pure fixed cost does not change the optimal positive intensity; it changes only whether that intensity is selected.

In the linear-quadratic workhorse,

$$\mathrm{NPV}_a^L = \frac{\left[A_a - k_a\right]_+^2}{2\kappa_a} - F_a.$$

Hence $A_a > k_a$ is the smooth marginal hurdle, while $[A_a - k_a]_+^2/(2\kappa_a) \geq F_a$ is the launch hurdle. The distinction has a direct evidentiary implication. Marginal engineering returns and implementation curvature determine the variable project surplus; the one-time integration budget, platform migration cost, certification expense, or plant reconfiguration cost belongs in F_a and must be supported separately.

For a genuinely indivisible project there may be no economically meaningful path from $x = 0$ to a small positive intensity. If the available choice is simply whether to implement a specified scale $\bar{x}_a$, the correct test is

$$\mathrm{NPV}_a^L(\bar{x}_a) = \sum_{i \in \mathcal{M}} \omega_i^a q_i^0 e_i^a(\bar{x}_a) - K_a(\bar{x}_a) - F_a.$$

A local derivative should not be used to decide such a lumpy project. The analyst should compare the total expected operating gain with all incremental and fixed costs at the specified scale.

With fixed costs or indivisibilities, strong merger specificity is therefore an NPV comparison:

$$\mathrm{NPV}_M^L > 0, \qquad \max_{a \in \mathcal{A}} \mathrm{NPV}_a^L \leq 0.$$

Relative to the smooth baseline, the merger may enable the project not only by raising marginal productivity or lowering marginal implementation cost, but also by spreading or avoiding a fixed cost. The claimed fixed-cost advantage must still be merger-specific; a cost that would also be incurred under a licensing agreement or joint venture cannot simply be assigned to the standalone counterfactual.

3.5.2 Benefits that persist over time

Even with no fixed cost and a perfectly smooth project choice, the project should be valued over the periods in which its savings are expected to apply. Let $t = 0, \ldots, T$, denote by $\delta \in (0, 1]$ the firm's discount factor, and let q_{it}^0 and $e_{it}^a(x)$ define the counterfactual quantity and expected unit saving in period

t. The dynamic local objective is

$$V_a^{D,L}(x) = \sum_{t=0}^{T} \delta^t \sum_{i \in \mathcal{M}} \omega_i^a q_{it}^0 e_{it}^a(x) - K_a(x),$$

and the dynamic local innovation incentive is

$$\mathrm{LII}_a^D = \sum_{t=0}^{T} \delta^t \sum_{i \in \mathcal{M}} \omega_i^a q_{it}^0 e_{it}^{a\prime}(0) - K_a'(0).$$

Thus persistence changes the marginal project return itself. It affects both whether a smooth project is launched and, when it is launched, how much is chosen.

If the unit saving is constant after an arrangement-specific implementation date τ_a, define discounted quantity exposure by

$$Q_{ia}^0 \equiv \sum_{t=\tau_a}^{T} \delta^t q_{it}^0.$$

The one-period formulas then continue to apply after replacing q_i^0 with Q_{ia}^0. When quantities are constant over time, the replacement amounts to multiplying the project base by the annuity factor $\sum_{t=\tau_a}^{T} \delta^t$. A longer duration, earlier implementation, or slower depreciation increases the effective project base; delay, discounting, technical obsolescence, or a declining quantity path reduces it.

The dynamic comparison can also change merger specificity. The merger and a less restrictive arrangement may have the same steady-state saving but different implementation dates, ramp-up paths, or useful lives. In that case, merger specificity comes from the discounted exposure profile rather than from a different one-period technology. These timing assumptions should be reported explicitly rather than hidden in a single annualized efficiency number.

Fixed costs and dynamics can of course be combined:

$$\mathrm{NPV}_a^{D,L} = \sup_{x>0} \left\{ \sum_{t=0}^{T} \delta^t \sum_{i \in \mathcal{M}} \omega_i^a q_{it}^0 e_{it}^a(x) - K_a(x) \right\} - F_a.$$

But the two ingredients should still be diagnosed separately. The discounted operating stream determines the value created conditional on implementation; F_a determines whether that value is sufficient to justify the discrete launch. For the transmission stage, the same discipline applies period by period: the time-t efficiency vector enters the time-t pricing-pressure calculation, and consumer effects should be discounted over the same relevant horizon. Online Resource 1, Section A, gives the full dynamic normalization.

3.6 Benefit-adjusted upward pricing pressure

The project analysis produces two outputs: a predicted project level and a product-level vector of expected benefits. This section asks what those benefits do to the merged firm's incentive to change prices. We return to the observed pre-merger price vector, evaluate the post-merger pricing first-order conditions at that point, and compare the incentive created by common ownership with the incentive created by the project.

Three forces play a role in this exercise. First, after the merger a price increase on product i is less

costly because some lost sales are recaptured by other products in the merged portfolio. This is the conventional unilateral-effect term. Second, a cost reduction on product i makes sales of that product more profitable and therefore pushes its price incentive downward. Third, a cost reduction on a product that receives diverted demand makes recaptured sales more valuable and can push the incentive to raise product i's price upward. The third effect is why an aggregate efficiency credit is generally insufficient.

The section proceeds as follows. We first derive the normalized local profit gradient, which we call benefit-adjusted UPP. We then identify the product-level efficiencies required to neutralize direct upward pressure. Next, we substitute the project level predicted in Section 3, so that efficiencies are an outcome rather than an assumption. Finally, we use pass-through to translate the incentive measure into local price and consumer-surplus effects. The distinction between a pressure measure and a price prediction is central throughout.

3.6.1 The local price gradient

Fix a project level x and consider an infinitesimal increase in the price of product $i \in \mathcal{M}$, holding all other prices at their observed pre-merger levels. Differentiating merged profit gives

$$\left.\frac{\partial \Pi^M(p,x)}{\partial p_i}\right|_{p=p^0} = q_i^0 + \left[m_i^0 + e_i^M(x)\right] q_{i,i}^0 + \sum_{j \in \mathcal{M}\setminus\{i\}} \left[m_j^0 + e_j^M(x)\right] q_{j,i}^0.$$

The three pieces have a simple economic meaning. The first is the extra revenue earned on the units still sold by product i. The second is the margin lost because the higher price reduces demand for product i; the project raises that margin by $e_i^M(x)$. The final line is the margin earned or lost on the other products in the merged portfolio as customers respond to the price of product i. Before the merger, those cross-product margins were not part of firm i's pricing decision.

The observed pre-merger first-order condition, $q_i^0 + m_i^0 q_{i,i}^0 = 0$, cancels the revenue and own-margin terms that were already balanced before common ownership. What remains is exactly the change in incentives created by the merger and the project:

$$\left.\frac{\partial \Pi^M(p,x)}{\partial p_i}\right|_{p=p^0} = e_i^M(x) q_{i,i}^0 + \sum_{j \in \mathcal{M}\setminus\{i\}} \left[m_j^0 + e_j^M(x)\right] q_{j,i}^0.$$

Because $-q_{i,i}^0 > 0$, dividing by the absolute own-price slope preserves the sign of the profit gradient and expresses it in the same monetary units as a margin or a marginal-cost reduction. Using signed diversion yields

$$\left.\frac{\partial \Pi^M(p,x)}{\partial p_i}\right|_{p=p^0} = (-q_{i,i}^0) \left\{ \sum_{j \in \mathcal{M}\setminus\{i\}} D_{ij}\left[m_j^0 + e_j^M(x)\right] - e_i^M(x) \right\}.$$

With this characterization in place, we can now introduce a central concept of the analysis.

Definition 3 (Benefit-adjusted UPP)**.** *For every $i \in \mathcal{M}$, benefit-adjusted upward pricing pressure is*

$$\text{BAUPP}_i(x) = \sum_{j \in \mathcal{M}\setminus\{i\}} D_{ij}\left[m_j^0 + e_j^M(x)\right] - e_i^M(x). \tag{3}$$

A positive value means that, at the observed price vector and holding other prices fixed, the merged firm's profit rises when product i's price is increased. A negative value means that the immediate local incentive points downward. The measure is therefore a normalized first-order condition, not a forecast of the equilibrium price change. Demand curvature and the reactions of all products enter only in the pass-through step below.

Equation (3) nests familiar cases. With no project, it reduces to standard UPP: the diversion-weighted margins on products whose ownership is newly internalized. If the project lowers only product i's cost, the benefit enters as the conventional own-product efficiency credit. If it also lowers the cost of product j, that saving raises the margin earned on sales diverted from i to j and therefore partly offsets the credit. In a two-product merger (say a merger involving only products i and j), the above formula simplifies to

$$\text{BAUPP}_i(x) = D_{ij}m_j^0 - e_i^M(x) + D_{ij}e_j^M(x),$$

which makes the three forces immediately transparent.

Let $\boldsymbol{m}_{\mathcal{M}}$ and $\boldsymbol{e}_{\mathcal{M}}(x)$ collect margins and efficiencies in the merging set. Stacking the product-level conditions gives

$$\mathbf{BAUPP}(x) = \mathbf{D}_{\mathcal{M}}\boldsymbol{m}_{\mathcal{M}} - \left(\mathbf{I} - \mathbf{D}_{\mathcal{M}}\right)\boldsymbol{e}_{\mathcal{M}}(x).$$

The first vector is gross pricing pressure from common ownership. The second is the net transmission content of the efficiency profile: each own-product saving is credited once, while diversion-weighted savings on destination products are subtracted from that credit. The following then holds.

Proposition 3 (Exact local decomposition)**.** *At the pre-merger price vector, the post-merger profit gradient for product $i \in \mathcal{M}$ has the sign of* $\text{BAUPP}_i(x)$*, where*

$$\text{BAUPP}_i(x) = \underbrace{\sum_{j \neq i} D_{ij}m_j^0}_{\text{newly internalized diversion}} - \underbrace{e_i^M(x)}_{\text{own-product benefit}} + \underbrace{\sum_{j \neq i} D_{ij}e_j^M(x)}_{\text{benefits on diverted products}},$$

with all sums restricted to $j \in \mathcal{M}$.

The proposition formalizes the product-incidence effect. A productive investment can reduce total cost and still strengthen the incentive to raise a particular price when its largest savings accrue to products receiving diverted demand. For substitute products, the full measure therefore exceeds a simplified formula that credits only the own-product saving by

$$\sum_{j \in \mathcal{M}\setminus\{i\}} D_{ij}e_j^M(x) \geq 0.$$

This does not make destination-product savings socially undesirable. It says that their effect on the source product's pricing incentive is different from the effect of an own-product marginal-cost reduction.

Notably, if both products have margin m, bilateral diversion D, and the project creates the same

expected unit saving e on each product, then

$$\text{BAUPP} = D(m + e) - e = Dm - (1 - D)e.$$

Only the leakage share $1 - D$ of a common saving disciplines pricing pressure: the fraction D is recaptured through the higher efficiency-enhanced margin on the partner product.

3.6.2 Critical efficiencies

Benefit-adjusted UPP can be used in two related ways. The first is directional: does an additional unit of the project reduce or increase pressure on a particular product? The second is a level question: how large must the complete efficiency vector be to neutralize the merger's direct pricing incentive on every product? The first question depends on the marginal incidence of the project. The second produces a product-level compensating-efficiency threshold.

Corollary 1 (Product-incidence condition)**.** *The project reduces pricing pressure on product i at project level x if and only if*

$$e_i^{M\prime}(x) > \sum_{j \in \mathcal{M} \setminus \{i\}} D_{ij} e_j^{M\prime}(x).$$

It increases pricing pressure if the inequality is reversed.

The left-hand side is the marginal own-product saving. The right-hand side is the marginal increase in the value of sales diverted to the other merging products. The condition is therefore not a comparison between "efficiency" and "harm" in the aggregate. It asks whether the next unit of the project is located more heavily on the product whose price is being assessed than on its diversion destinations. A fixed project profile may reduce pressure on some products and increase it on others.

Corollary 2 (Compensating expected efficiency)**.** *There is no upward pricing pressure on any product in the merging set if and only if*

$$\left(\mathbf{I} - \mathbf{D}_{\mathcal{M}}\right) \boldsymbol{e}_{\mathcal{M}}(x) \geq \mathbf{D}_{\mathcal{M}} \boldsymbol{m}_{\mathcal{M}} \tag{4}$$

componentwise. If $\mathbf{D}_{\mathcal{M}}$ is nonnegative and $\rho(\mathbf{D}_{\mathcal{M}}) < 1$, the exact-balance vector is

$$\boldsymbol{e}^{\text{crit}} = \left(\mathbf{I} - \mathbf{D}_{\mathcal{M}}\right)^{-1} \mathbf{D}_{\mathcal{M}} \boldsymbol{m}_{\mathcal{M}}. \tag{5}$$

The right-hand side of (4) is the gross pressure that must be offset. The left-hand side is not simply the efficiency vector. It is the own-product saving net of the diversion-weighted savings that make recaptured sales more profitable. The componentwise inequality therefore requires a separate pass for each product; a favorable portfolio average can conceal a product with positive pressure.

The inverse in (5) has a network interpretation. The first round of efficiency offsets the direct value of newly internalized diversion. But that efficiency also raises the margins on destination products, which creates another round of recapture that must itself be offset. The inverse sums these direct and indirect within-portfolio feedbacks. The spectral-radius condition ensures that some demand leaks outside the internal diversion network quickly enough for the sequence to converge.

In the symmetric case, if each product has margin m, efficiency e, and total diversion $\delta < 1$ to the

other merging products, exact balance requires

$$e = \frac{\delta m}{1-\delta}.$$

The familiar direct term is δm. The denominator $1-\delta$ reflects the fact that the same efficiencies credited on one product raise the value of diverted sales on the others. As internal recapture approaches one, a pure marginal-cost efficiency must become increasingly large to neutralize pressure across the entire portfolio.

3.6.3 Endogenous-efficiency UPP

The preceding formulas treat the efficiency vector as given. Yet, the central purpose of the paper is to remove that assumption. We now feed the project level predicted by the generation and merger-specificity analysis into the pricing test. This produces a single local diagnostic while preserving its two economically distinct components: why the project occurs and how the resulting benefits are transmitted.

Definition 4 (Endogenous-efficiency UPP)**.** *Let x_M^L be the project level predicted by the local investment model. Endogenous-efficiency UPP is*

$$\mathrm{EEUPP}_i = \mathrm{BAUPP}_i(x_M^L), \qquad i \in \mathcal{M}.$$

In the linear-quadratic workhorse,

$$\mathrm{EEUPP}_i = \underbrace{\sum_{j\neq i} D_{ij} m_j^0}_{\text{gross merger pressure}} + \underbrace{\left(\sum_{j\neq i} D_{ij}\eta_j^M - \eta_i^M\right)}_{\text{product-incidence coefficient}} \underbrace{\left[\frac{A_M - k_M}{\kappa_M}\right]_+}_{\text{generated project scale}},$$

where the sums are over the other products in $\mathcal{M}$. The formula is useful precisely because it separates scale from incidence. The project base and implementation parameters determine whether innovation is generated and how much. The incidence coefficient determines whether each unit of that innovation mitigates or amplifies pressure on product i.

If $A_M \leq k_M$, the project is not generated and EEUPP collapses to ordinary UPP. If $A_M > k_M$ but $\eta_i^M > \sum_{j\neq i} D_{ij}\eta_j^M$, a larger project lowers pressure on product i. If the inequality is reversed, a larger project raises that product's local pricing incentive because the destination-product margin effect dominates the own-product cost credit. We can thus state the following.

Proposition 4 (Generation does not imply transmission)**.** *For any merging set with at least two products, there are parameter values for which the project is strongly merger-specific but* $\mathrm{EEUPP}_i > 0$ *for at least one product. There are also parameter values for which a larger merger-enabled project increases* EEUPP_i *on one or more products.*

A positive project score is therefore evidence about private project generation, not a conclusion about customer benefit. The project may be genuine, incremental, and merger-specific yet remain too small, too delayed, or too concentrated on diversion destinations to reach the compensating threshold in (4).

Combining the linear-quadratic project choice with the symmetric price condition gives the single diagnostic

$$\text{EEUPP} = Dm - (1 - D)\eta \left[\frac{2q\eta - k_M}{\kappa_M} \right]_+ .$$

The bracketed term is generated project scale; the coefficient $(1 - D)\eta$ is its net transmission content. A positive generation score can therefore coexist with positive endogenous pricing pressure.

3.6.4 Pass-through and local price effects

BAUPP is a normalized profit-gradient shock. It answers whether the merged firm would like to move a price up or down at the pre-merger point, holding other prices fixed. Consumer surplus depends instead on the equilibrium price response of the entire market. Translating pressure into prices therefore requires one additional object: the local pass-through system. This step is important for policy because two projects with the same BAUPP vector can have different consumer effects when price curvature, strategic interaction, or rival responses differ.

Embed the direct pressure vector in the full N-product market by defining

$$\widetilde{b}_i(x) = \begin{cases} \text{BAUPP}_i(x), & i \in \mathcal{M}, \\ 0, & i \notin \mathcal{M}. \end{cases}$$

The zeros for non-merging products mean only that the merger does not create a direct ownership shock in their first-order conditions. Their equilibrium prices can still change because they respond strategically to the prices of the merging products.

To see the mapping, let $F_p(p)$ collect the market's price first-order conditions and let $A = \partial F_p / \partial p'$ be their Jacobian at the pre-merger equilibrium. Let

$$S = \text{diag}(-q^0_{1,1}, \ldots, -q^0_{N,N}).$$

The merger and project shift the price conditions by approximately $S\widetilde{b}(x)$. A first-order expansion around p^0 gives

$$A\Delta p + S\widetilde{b}(x) \approx 0.$$

If A is nonsingular, the local price prediction is

$$\Delta p \approx \mathcal{P}\widetilde{b}(x), \qquad \mathcal{P} \equiv -A^{-1}S. \tag{6}$$

The matrix $\mathcal{P}$ is local pass-through. Its kth column records how every market price responds to one monetary unit of normalized pressure on product k. Diagonal elements capture own-price pass-through; off-diagonal elements capture strategic and cross-product responses. The formula also makes clear why BAUPP and pass-through should be reported separately. BAUPP is determined by ownership, diversion, margins, and the project. Pass-through adds demand curvature and the structure of the pricing game.

For a process innovation that leaves utility and demand fixed, the first-order change in consumer surplus is

$$\Delta \text{CS} \approx -(q^0)'\Delta p \approx -(q^0)'\mathcal{P}\widetilde{b}(x). \tag{7}$$

The first expression is the usual envelope approximation: a small price increase of one unit on product i lowers consumer surplus by approximately q_i^0 units. The second expression combines this quantity weighting with the market-wide price response. All products belong in the calculation, including outsiders whose direct shock is zero but whose prices respond through cross-pass-through.

Several implications follow. First, nonpositive BAUPP on every merging product is a strong direct-incentive result, but it is not mechanically the same object as a consumer-surplus prediction. When $\mathcal{P}$ is componentwise nonnegative, a nonpositive embedded pressure vector implies weakly lower prices throughout the market and therefore a nonnegative local consumer-surplus effect. Without that sign structure, strategic responses can make some prices move in the opposite direction. Conversely, a positive BAUPP component does not by itself establish an aggregate consumer-surplus loss if other prices fall sufficiently. Product-level pressure and aggregate consumer incidence answer related but distinct questions.

Second, the allocation of savings matters twice. It determines BAUPP through the own-product and destination-product terms, and it determines consumer incidence through heterogeneous pass-through and baseline quantities. A dollar of marginal-cost saving on a high-volume product with strong own-pass-through may generate a larger consumer benefit than the same dollar on a low-volume product. At the same time, placing that saving on a product that receives substantial diversion can raise pressure elsewhere. A policy assessment should therefore retain the product-level efficiency vector until after both the pressure and pass-through calculations have been performed.

Third, the pass-through step disciplines what can be concluded from a first screen. When an estimated demand system is available, $\mathcal{P}$ can be computed from the Jacobian of the pricing conditions. When only empirical cost pass-through or a range of plausible elasticities is available, the analyst can report sensitivity ranges or bounds for (6) and (7). Assuming zero pass-through discards the claimed benefit; assuming complete pass-through can overstate it. Uncertainty about pass-through should therefore lead to a qualified or inconclusive consumer-benefit conclusion, not to silently replacing the missing object with a convenient constant. Online Resource 1, Section D.5, provides simple consumer-surplus bounds when a full matrix is unavailable.

Fourth, timing must be aligned across the generation and transmission tests. If the project produces the period-t efficiency vector $e_t(x)$, the corresponding local prediction is

$$\Delta p_t \approx \mathcal{P}_t \widetilde{b}_t(x).$$

A discounted consumer-surplus calculation is then

$$\Delta \mathrm{CS}^D \approx -\sum_{t=0}^{T} \beta^t (q_t^0)' \mathcal{P}_t \widetilde{b}_t(x),$$

where β is the welfare discount factor. Delayed efficiencies should not be credited as though they reduced prices immediately, and an upfront project NPV should not be compared with a one-period consumer-surplus number. The same horizon, implementation lag, and expected efficiency path should be used on both sides of the theory of benefit.

Finally, (7) captures only the price channel. That is complete for a pure process innovation with fixed demand. A quality, reliability, variety, or service improvement can benefit consumers directly at

unchanged prices and may also change the demand system itself. Section 4.4, therefore, adds the direct utility term and, when quality is money-metric, carries out the pass-through exercise in hedonic rather than nominal prices.

3.7 A compact merger-to-monopoly illustration

This section combines the generation, merger-specificity, and transmission tests in the smallest operational example. Consider two symmetric single-product firms that are the only suppliers of the two inside products and merge into a two-product monopolist. Consumers may still switch to an outside good, so bilateral diversion satisfies $D < 1$. Let

$$q_1^0 = q_2^0 = q, \qquad m_1^0 = m_2^0 = m, \qquad D_{12} = D_{21} = D,$$

and suppose the merger enables the common process project

$$e_1^M(x) = e_2^M(x) = \eta x, \qquad K_M(x) = kx + \frac{\kappa}{2}x^2.$$

A standalone firm can deploy the project only on its own product and faces $K_S(x) = k_S x + (\kappa/2)x^2$. The local choices are therefore

$$x_M^L = \left[\frac{2q\eta - k}{\kappa}\right]_+, \qquad x_S^L = \left[\frac{q\eta - k_S}{\kappa}\right]_+.$$

The project is locally merger-enabled when $2q\eta > k$ and $q\eta \le k_S$. This is the generation side of the theory of benefit.

The transmission side is equally compact. At any project level,

$$\mathrm{BAUPP}_1(x) = \mathrm{BAUPP}_2(x) = Dm - (1 - D)\eta x.$$

Hence the generated project eliminates direct upward pricing pressure only if

$$x_M^L \ge x^{\mathrm{UPP}} \equiv \frac{Dm}{(1 - D)\eta}, \qquad \text{or equivalently} \qquad \eta x_M^L \ge e^{\mathrm{UPP}} \equiv \frac{Dm}{1 - D}.$$

The denominator shows why leakage matters. If $D = 1$, a common process saving raises the partner-product recapture margin by exactly as much as it lowers the own-product marginal cost, so no finite common saving can offset a positive pre-merger margin in this local screen.

Table 1 reports two compact calibrations. Online Resource 1, Section B.1, develops the associated comparative statics and parameter frontiers. Both use $q = 10$, $m = 20$, $D = 0.5$, $k = 20$, and $\kappa = 10$; in both cases the standalone choice is $x_S^L = 0$.

Table 1: Compact merger-to-monopoly calibrations

Calibration	η	k_S	x_M^L	e^L	x^{UPP}	BAUPP	Screening result
Generation only	2	25	2	4	10	8	Fails transmission
Both hurdles	4	45	6	24	5	−2	Passes both

In the first calibration the transaction changes the project decision, but the generated unit saving is far

below the compensating efficiency of 20. The theory of benefit therefore fails at transmission rather than at project causation. In the second calibration the merger again uniquely generates the project and the resulting unit saving of 24 exceeds the same compensating efficiency. The negative BAUPP indicates a local incentive to reduce both prices.

4 Extensions

The analytical framework developed above provides a robust initial screen for assessing mergers that are claimed to promote innovation. However, the model rests on specific assumptions, and its application may require incorporating additional elements or considering alternative hypotheses regarding the underlying innovation process. The purpose of this section is to introduce the most plausible extensions one at a time, in order to identify which conclusions change and which additional evidence drives those changes.

The first extension allows project success to be uncertain and distinguishes expected pressure from state-contingent failure risk. The second admits pre-existing multiproduct ownership, separating margins already internalized before the merger from ownership links created by the transaction. The third adds a one-step quantity correction when merger-induced price changes alter the sales base over which savings are earned. The remaining extensions allow the project to shift demand or quality and, when the required derivatives are available, solve prices and investment jointly. Each subsection begins from the baseline object, states the new terms, and explains how the interpretation of the screen changes.

4.1 Uncertain success and expected pricing pressure

The baseline analysis treats the project-induced cost reductions as known conditional on the chosen project intensity. In many merger investigations, however, this is too strong. A claimed efficiency may depend on a technical milestone, integration timing, supplier adoption, regulatory approval, or the ability to scale a prototype across products. It is therefore useful to separate three objects: the *ex ante incentive* to launch the project, the *expected* price-pressure effect, and the *state-contingent* risk that the claimed efficiencies do not materialize when prices are set.

Let $\widetilde{e}_i(x, \theta)$ denote the realized marginal-cost reduction on product i in state θ, and define its ex ante mean by

$$\bar{e}_i(x) \equiv \mathbb{E}[\widetilde{e}_i(x, \theta)].$$

The project is assessed at the same local demand system and pre-merger margins as in the baseline model; only the efficiency vector is random. This restriction isolates success risk from changes in demand, pass-through, capacity, or the continuation pricing game.

The first implication concerns project generation. If the merged firm chooses project intensity before the success state is realized, the local benefit of raising x is the expected value of the product-level cost savings, evaluated at pre-merger quantities. Thus the stochastic analogue of the local innovation incentive is

$$\mathrm{LII}^{U}_{M} = \sum_{i \in \mathcal{M}} q_i^0 \left. \frac{\partial \bar{e}_i(x)}{\partial x} \right|_{x=0} - K'_M(0),$$

where the superscript U stands for uncertainty. This expression says that success risk discounts the private project base through expected marginal cost savings. If a project succeeds with probability π and, conditional on success, creates the deterministic local efficiency profile $e_i(x)$, then $\bar{e}_i(x) = \pi e_i(x)$

and the project base is multiplied by π. If project intensity also changes the probability of success, the relevant marginal term is the derivative of $\pi(x)e_i(x)$ at the origin. In particular, for a binary project with conditional saving $\Delta_i(x)$,

$$\left.\frac{\partial \bar{e}_i(x)}{\partial x}\right|_{x=0} = \pi(0)\Delta_i'(0) + \pi'(0)\Delta_i(0).$$

When the project has no saving at zero intensity, the second term vanishes; when the first unit of investment mainly raises the probability of completing a lumpy project, the second term captures that extensive-margin effect.

Proposition 5 (Expected project generation under success risk)**.** *Suppose the merged firm's expected local project payoff is concave in x and equals*

$$\mathbb{E}\left[\sum_{i\in\mathcal{M}} q_i^0 \widetilde{e}_i(x,\theta)\right] - K_M(x).$$

Then the merger generates a positive local project if and only if $\mathrm{LII}_M^U > 0$. If $\mathrm{LII}_M^U \le 0$, the local optimum is $x = 0$.

The second implication concerns transmission. Conditional on a state, realized benefit-adjusted UPP is obtained by replacing the deterministic efficiency vector with the realized one. Because the local BAUPP index is affine in that vector, expected pricing pressure depends only on the mean efficiency vector.

Proposition 6 (Expected benefit-adjusted UPP)**.** *If local demand derivatives and pre-merger margins are state-independent, then for every $i \in \mathcal{M}$,*

$$\mathbb{E}\big[\mathrm{BAUPP}_i(x,\theta)\big] = \sum_{j\in\mathcal{M}\setminus\{i\}} D_{ij}\big[m_j^0 + \bar{e}_j(x)\big] - \bar{e}_i(x).$$

Thus BAUPP evaluated at expected efficiencies equals expected state-contingent BAUPP.

A useful special case is binary success. Let the project deliver $\Delta_i(x)$ on product i with probability $\pi(x)$ and zero otherwise. Define the no-efficiency upward pressure and the net incidence of successful cost savings by

$$u_i^0 = \sum_{j\in\mathcal{M}\setminus\{i\}} D_{ij}m_j^0, \qquad G_i(x) = \Delta_i(x) - \sum_{j\in\mathcal{M}\setminus\{i\}} D_{ij}\Delta_j(x).$$

Expected BAUPP can then be written compactly as

$$\mathbb{E}\big[\mathrm{BAUPP}_i(x,\theta)\big] = u_i^0 - \pi(x)G_i(x).$$

This form is useful because it separates the gross unilateral pressure created by the merger, u_i^0, from the risk-adjusted net efficiency credit, $\pi(x)G_i(x)$. If $G_i(x) > 0$, product i passes the expected local screen only when

$$\pi(x) \ge \frac{u_i^0}{G_i(x)}.$$

If $G_i(x) \le 0$, the claimed efficiency pattern cannot cure expected upward pressure on product i: successful savings on other products increase the value of diverted sales by at least as much as the own-product saving reduces product i's marginal cost. The whole merging set passes the expected test only if the success probability exceeds the largest product-level critical probability among products for which $G_i(x) > 0$.

This critical-probability representation is a convenient device. It shows whether failure is due to a low probability of success or to a poor incidence pattern of savings across products. A high conditional saving on one product may still be insufficient for the portfolio if the products receiving diverted sales also become much more profitable. Conversely, a moderate but well-distributed efficiency vector can pass with a lower success probability.

Expected BAUPP is the appropriate ex ante screen when the authority evaluates the project before uncertainty is resolved and when local primitives are not expected to change across states. It is not a complete risk assessment. A merger may pass the expected test while failing in a large set of states. For that reason, the analyst should also report state-contingent diagnostics, such as

$$\Pr\left(\text{BAUPP}_i(x,\theta) \leq 0 \text{ for all } i \in \mathcal{M}\right),$$

together with the success-state and failure-state BAUPP vectors. In the binary case, this probability is especially transparent: it equals one if both the success and failure states pass, equals $\pi(x)$ if only the success state passes, and equals zero if even the success state fails. These numbers answer a different question from the expected screen: not "what is the average pressure?", but "how often do the claimed efficiencies actually cure unilateral price pressure product by product?"

Finally, the linearity result should not be overinterpreted. Once success changes demand, capacity, marginal costs outside the local efficiency vector, pass-through, or the continuation pricing game, expected prices need not equal prices evaluated at expected efficiencies. In those cases, the recommended implementation is to compute the local screen separately by state, attach probabilities to the states, and treat expected BAUPP as a first-order summary rather than as a substitute for a state-contingent price analysis.

4.2 Multiproduct firms

The single-product baseline keeps the sufficient statistics transparent, but pre-merger portfolios matter whenever a merging firm already owns more than one differentiated product. In that setting, diversion to a sister product is already internalized before the transaction. It therefore does not create new direct upward pricing pressure. The same diversion nevertheless matters for benefit transmission because a cost reduction on the sister product raises the value of sales diverted to it after the merger. This subsection states the general ownership-matrix formulation. A compact symmetric four-product specialization closes the subsection; the expanded calibration is in Online Resource 1, Section B.2.

4.2.1 General ownership matrices

The matrix notation in this subsection is only a bookkeeping device. The economic calculation remains product by product. It has three steps. First, record which product margins were already jointly internalized before the transaction. Second, record which products are jointly owned afterward; the difference identifies the ownership links that create direct merger pressure. Third, use post-merger ownership to determine which efficiency-enhanced margins are earned when demand diverts. Keeping the second and third steps separate prevents pre-existing portfolio diversion from being incorrectly counted as a new unilateral effect.

Let there be J products. In state $s \in \{0, 1\}$, let F_s denote the number of independent owners and

define the product-owner incidence matrix $\mathbf{G}^s \in \{0,1\}^{J \times F_s}$ by

$$G^s_{if} = \begin{cases} 1, & \text{if firm } f \text{ controls product } i \text{ in state } s, \\ 0, & \text{otherwise.} \end{cases}$$

When every product has one controlling owner, the ownership matrix is

$$\mathbf{\Omega}^s = \mathbf{G}^s(\mathbf{G}^s)', \qquad \Omega^s_{ik} = \begin{cases} 1, & \text{if products } i \text{ and } k \text{ are jointly owned in state } s, \\ 0, & \text{otherwise.} \end{cases}$$

Thus $\mathbf{\Omega}^s$ is a $J \times J$ symmetric matrix with ones on the diagonal. After products are ordered by owner, it consists of blocks of ones, one block for each multiproduct firm. The matrix $\mathbf{\Omega}^0$ records which margins were already internalized in the pre-merger first-order conditions. The matrix $\mathbf{\Omega}^1$ records the portfolio whose margins are internalized after the transaction. Their difference $\Delta\mathbf{\Omega} = \mathbf{\Omega}^1 - \mathbf{\Omega}^0$ identifies ownership links created by the merger. For a pure consolidation without divestitures its off-diagonal entries are zero or one; a broader restructuring can also generate negative entries.

The role of the two matrices is different. Direct merger pressure depends on $\Delta\mathbf{\Omega}$ because only a newly internalized margin changes the pricing first-order condition relative to the pre-merger benchmark. Efficiency recapture depends on $\mathbf{\Omega}^1$ because every other product jointly owned after the merger contributes its efficiency-enhanced margin when demand diverts to it, whether that product was acquired in the transaction or was already in the same portfolio.

Let $\mathbf{D}$ be the $J \times J$ diversion matrix with zero diagonal and entries defined above. Let $\boldsymbol{m}^0 = (m^0_1, \ldots, m^0_J)'$ and let $\boldsymbol{e}(x) = (e_1(x), \ldots, e_J(x))'$, assigning zero to unaffected products.

Proposition 7 (Multiproduct benefit-adjusted UPP)**.** *For a product i controlled by the merged entity after the transaction, the post-merger price gradient at pre-merger prices has the sign of*

$$\text{BAUPP}_i(x) = \sum_{k \neq i} (\Omega^1_{ik} - \Omega^0_{ik}) D_{ik} m^0_k - e_i(x) + \sum_{k \neq i} \Omega^1_{ik} D_{ik} e_k(x).$$

The first term is the value of diversion that becomes newly internalized because ownership changes. The second is the own-product efficiency. The final term contains efficiencies on every other product in the post-merger portfolio that receives diverted demand. This formula can therefore be implemented with two ownership matrices, a diversion matrix, pre-merger margins, and a product-level efficiency vector.

For a compact matrix representation, let $\circ$ denote the Hadamard (element-by-element) product and define

$$\mathbf{H} \equiv (\mathbf{\Omega}^1 - \mathbf{\Omega}^0) \circ \mathbf{D}, \qquad \mathbf{R} \equiv \mathbf{\Omega}^1 \circ \mathbf{D}.$$

The vector of normalized price gradients is

$$\boldsymbol{b}(x) = \mathbf{H}\boldsymbol{m}^0 - (\mathbf{I} - \mathbf{R})\boldsymbol{e}(x).$$

The exact portfolio-wide transmission condition is therefore

$$(\mathbf{I} - \mathbf{R})\boldsymbol{e}(x) \geq \mathbf{H}\boldsymbol{m}^0$$

componentwise.

Corollary 3 (A fixed product-level efficiency profile)**.** *Suppose the project generates* $\boldsymbol{e}(x) = x\boldsymbol{\eta}$*, where* $\boldsymbol{\eta} \geq \mathbf{0}$*, and define*

$$\boldsymbol{h} \equiv \mathbf{H}\boldsymbol{m}^0, \qquad \boldsymbol{g} \equiv (\mathbf{I} - \mathbf{R})\boldsymbol{\eta}.$$

If $g_i > 0$ *for every product under review, then all local price gradients are nonpositive if and only if*

$$x \geq x^{\text{UPP}} \equiv \max_i \frac{h_i}{g_i}.$$

If $h_i > 0$ *and* $g_i \leq 0$ *for some product i, no project intensity along that fixed profile can eliminate upward pricing pressure on product i.*

The vector $\boldsymbol{g}$ is the net transmission content of the project profile. Its ith component equals the own-product efficiency coefficient minus the diversion-weighted efficiency coefficients on all other products jointly owned after the merger. This makes clear why aggregate savings are not sufficient: two projects with the same total cost reduction can have opposite implications if their product incidence differs.

When $\rho(\mathbf{R}) < 1$, the matrix $\mathbf{I} - \mathbf{R}$ is invertible and

$$\boldsymbol{e}^{\text{bal}} = (\mathbf{I} - \mathbf{R})^{-1}\mathbf{H}\boldsymbol{m}^0 = \sum_{n=0}^{\infty} \mathbf{R}^n \mathbf{H}\boldsymbol{m}^0$$

is the unique efficiency vector that sets every component of $\boldsymbol{b}$ exactly equal to zero. The Neumann series gives an economic interpretation: the first term offsets newly internalized diversion, while each successive term offsets another round of efficiency recapture through the post-merger diversion network. Importantly, $\boldsymbol{e}^{\text{bal}}$ is the exact-balance vector. When $\mathbf{R}$ is nonnegative, every passing vector must lie componentwise above it. Componentwise dominance is not, however, sufficient for passage: increasing efficiency on one destination product can raise pricing pressure on another source product. Thus the passing set described above need not be an upper set.

Proposition 8 (Why a portfolio monopoly needs diversion leakage)**.** *Suppose the merged firm owns all J inside products, so that* $\boldsymbol{\Omega}^1 = \mathbf{1}\mathbf{1}'$ *and* $\mathbf{R} = \mathbf{D}$*. Assume* $\mathbf{D}$ *is a nonnegative, irreducible, row-stochastic diversion matrix and* $\mathbf{H}\boldsymbol{m}^0 \geq \mathbf{0}$ *is nonzero. Then no finite product-level efficiency vector can make every component of* $\boldsymbol{b}$ *nonpositive.*

The proposition describes the limiting case in which every unit lost by one inside product is captured by another inside product. A portfolio-wide pass then cannot be obtained from marginal-cost efficiencies alone: the weighted sum of the direct merger terms remains positive. A passing merger-to-monopoly example therefore requires diversion leakage to an outside good or market contraction, zero newly internalized diversion, or a richer benefit channel such as quality or demand expansion. A simple sufficient condition for leakage is that every row sum of $\mathbf{D}$ is strictly below one, which implies $\rho(\mathbf{D}) < 1$.

To gain further insights, consider the case of two firms and suppose each pre-merger firm owns two products. Let d_W be diversion from a product to its pre-merger sister product, let d_C be diversion to each of the two products owned by the other firm, and let $d_0 = 1 - d_W - 2d_C$ be outside leakage. With a

common margin m and common unit saving e, all four products have

$$\text{BAUPP} = 2d_C m - d_0 e.$$

Thus, when $d_0 > 0$, transmission requires $e \geq 2d_C m/d_0$. Within-firm diversion does not create direct merger pressure, but it weakens transmission because savings on the sister product make diverted sales more profitable.

4.3 A one-step quantity-feedback correction

The baseline project score values a unit of cost reduction on pre-merger quantities. This is the natural local starting point, but it can be optimistic when the no-efficiency merger would raise prices and contract the very output base over which savings are earned. The issue is conceptually different from pass-through to consumers. Here pass-through matters because it changes the merged firm's private return to innovation. The output channel is emphasized in investment models such as Motta and Tarantino (2021).

A useful robustness check can be implemented without solving the full price–investment equilibrium. It proceeds in three steps. First, calculate the no-efficiency local price response. Second, use the demand Jacobian to update the quantity base. Third, recompute the project score at those updated quantities. This preserves the transparency of the sequential method while showing whether the project-generation conclusion depends materially on holding output fixed.

Let J_q denote the full demand Jacobian at p^0, and let $\widetilde{b}^0$ be the embedded no-efficiency pressure vector. The first two steps are

$$\Delta p^0 \approx \mathcal{P}\widetilde{b}^0, \qquad \widehat{q} = q^0 + J_q \Delta p^0.$$

The first equation maps the ownership shock into prices. The second is a first-order demand update: each product's quantity changes in response to all market prices, not only its own. The feedback-adjusted local innovation score is then

$$\widehat{\text{LII}}_M = \sum_{i \in \mathcal{M}} \widehat{q}_i e_i^{M\prime}(0) - K_M'(0).$$

Equivalently, the correction relative to the baseline score is

$$\widehat{\text{LII}}_M - \text{LII}_M \approx \left(\sum_{i \in \mathcal{M}} e_i^{M\prime}(0) \nabla_p q_i^0 \right)' \Delta p^0.$$

This expression identifies the mechanism directly. The vector in parentheses measures how the quantity-weighted marginal return to the project changes with prices. Its inner product with the no-efficiency price response is the first-order loss or gain in project incentives caused by output adjustment.

For ordinary goods, a broad price increase usually lowers the effective project base, so the corrected score is typically smaller than the pre-merger-quantity score. The sign need not be negative product by product: a price change can redirect demand toward a product on which the project is especially productive. The relevant object is therefore the change in the whole quantity-weighted project base, not simply the merged firm's aggregate output.

The correction is most useful as a diagnostic. If LII_M and $\widehat{\text{LII}}_M$ have the same sign and similar magnitude, the baseline screen is robust to the first round of output feedback. If the sign changes or the magnitude falls sharply, the merger-specific project claim depends on endogenous quantities and

should be evaluated with the joint local system in Section 4.5 or a fuller simulation. The calculation can be iterated, but repeated iteration progressively turns the screen into a merger model and should be accompanied by the corresponding additional demand and curvature evidence.

4.4 Demand-enhancing innovation and quality

This subsection extends both halves of the screening test when the project changes demand, quality, or willingness to pay. We first explain why a demand shift cannot generally be treated as an ordinary cost reduction, then derive the generation, pricing, and consumer-incidence formulas. A money-metric quality case provides a more parsimonious implementation when the evidence supports it.

4.4.1 Why the extension requires a separate treatment

The baseline project lowers marginal cost while leaving demand fixed. This makes the two stages of the exercise particularly simple. At the innovation stage, each unit of expected cost reduction is valued on the observed pre-merger sales base. At the pricing stage, the demand derivatives and diversion ratios remain fixed, so the predicted efficiency vector can be inserted directly into benefit-adjusted upward pricing pressure.

Yet, many merger-enabled projects have a different structure. A common data platform may improve matching. A redesigned product may raise reliability. A network integration may improve coverage. A combined product-development team may increase variety or accelerate feature introduction. These projects can change demand even when marginal cost is unaffected. They can also have both effects at once: quality may rise while unit cost falls, or quality may rise while unit cost increases. Demand innovation creates three logically distinct questions.

First, does the project create enough private operating value to be launched? A demand gain is valuable to the firm only through the margin earned on the additional sales. If the project shifts customers from one product in the merged portfolio to another, the gain on the expanding product must be netted against the margin lost on the cannibalized product.

Second, how does the project change pricing incentives? An outward demand shift can induce a firm to charge a higher nominal price even when consumers are better off. The project may also change own-price sensitivity, cross-price sensitivity, and the diversion matrix. A fixed-diversion formula is therefore not generally sufficient.

Third, how much of the benefit reaches consumers? With process innovation, the relevant first-order channel is the change in prices. With quality innovation, consumers may receive a direct benefit at unchanged prices. The consumer calculation must combine the direct quality gain with the induced nominal-price response. A nominal price increase, by itself, is not a sufficient statistic for consumer harm.

The analysis below first gives a general formulation in which the project can change the entire demand system. It then develops a more parsimonious special case in which quality is measured as a product-specific willingness-to-pay increment. The latter produces a quality-adjusted UPP test in the spirit of Willig (2011).

4.4.2 Environment and notation

The extension changes one baseline object: demand now depends on project intensity as well as prices. Everything else is introduced so that the new terms can be compared directly with the process-only formulas. The project may lower marginal cost, shift demand, or do both; allowing the cost effect to be signed also covers quality improvements that are more expensive to produce.

There are $N \geq 2$ differentiated products indexed by $i \in \mathcal{N} \equiv \{1, \ldots, N\}$. A set $\mathcal{M} \subseteq \mathcal{N}$, with $|\mathcal{M}| \geq 2$, is brought under common ownership by the merger. Before the merger, each product is owned by a single-product firm. The observed pre-merger equilibrium is

$$p^0 = (p_1^0, \ldots, p_N^0)', \qquad q_i^0 = q_i(p^0, 0), \qquad m_i^0 = p_i^0 - c_i^0,$$

where $x = 0$ denotes the no-project state and c_i^0 is pre-project marginal cost. Demand is now written as

$$q_i = q_i(p, x).$$

The project intensity $x \geq 0$ may change demand directly. It may also change marginal cost. Let

$$e_i(x) = c_i^0 - c_i(x), \qquad e_i(0) = 0.$$

The cost effect is signed. A positive $e_i(x)$ is a marginal-cost reduction; a negative $e_i(x)$ is a marginal-cost increase associated, for example, with a higher-quality input. The non-volume-dependent resource cost of the project is $K_M(x)$, with $K_M(0) = 0$.

Write

$$q_{j,i}^0 \equiv \left.\frac{\partial q_j(p,x)}{\partial p_i}\right|_{(p^0,0)}, \qquad q_{i,i}^0 < 0,$$

and define signed pre-merger diversion by

$$D_{ij} \equiv \frac{q_{j,i}^0}{-q_{i,i}^0}, \qquad D_{ii} = 0.$$

Because each firm is single-product before the merger, the pre-merger first-order condition is

$$q_i^0 + m_i^0 q_{i,i}^0 = 0, \qquad i \in \mathcal{N}.$$

The merged firm's profit is

$$\Pi^M(p, x) = \sum_{i \in \mathcal{M}} \left[p_i - c_i^0 + e_i(x)\right] q_i(p, x) - K_M(x).$$

This expression makes the distinction between the two project channels transparent. The project changes the margin on existing units through $e_i(x)$ and changes the number and allocation of units through $q_i(p, x)$.

Assumption 2 (Local regularity)**.** *Demand and project functions are twice continuously differentiable in a neighborhood of* $(p^0, 0)$*. Own-price derivatives remain negative in that neighborhood. The local project objective defined below is strictly concave and attains a maximum on the relevant interval of*

project intensities, which contains a right-neighborhood of zero.

4.4.3 Project generation with demand effects

We first repeat the generation exercise while holding prices at their observed levels. The only new issue is that a marginal project now changes operating profit through two channels: it changes the margin earned on existing units and it changes the number and allocation of units sold. The formula below keeps those channels separate before subtracting project cost.

As in the baseline methodology, the local screen evaluates the project at the observed pre-merger price vector. The merged firm's incremental local project objective is

$$V_M^{Q,L}(x) = \sum_{i \in \mathcal{M}} \left\{ \left[m_i^0 + e_i(x) \right] q_i(p^0, x) - m_i^0 q_i^0 \right\} - K_M(x).$$

The subtraction of pre-project operating profit is only a normalization. It ensures $V_M^{Q,L}(0) = 0$ and does not affect the project choice.

Definition 5 (Local innovation incentive with quality or demand effects)**.** *The merger's local quality-and-demand innovation incentive is*

$$\mathrm{LII}_M^Q \equiv \left. \frac{\mathrm{d}V_M^{Q,L}(x)}{\mathrm{d}x} \right|_{x=0} = \sum_{i \in \mathcal{M}} \left[q_i^0 e_i'(0) + m_i^0 \frac{\partial q_i(p^0, 0)}{\partial x} \right] - K_M'(0).$$

The displayed quality local innovation incentive has three components. The first is the *process-efficiency channel*, $q_i^0 e_i'(0)$. It values a marginal unit-cost change on the observed sales base. This term is positive for a cost reduction and negative if the quality improvement raises marginal cost. The second is the *margin-weighted demand channel*, $m_i^0 q_{i,x}^0$, where $q_{i,x}^0 \equiv \partial q_i(p^0, 0)/\partial x$. It is positive when the project expands demand for product i at unchanged prices and negative when that product is cannibalized. The merged firm sums this term across its portfolio. Hence, a demand shift from one merging product to another is not a pure gain: the additional margin on the destination product is netted against the margin lost on the origin product. The third is the marginal resource cost $K_M'(0)$. This term should contain the marginal cost of developing and implementing the common project. Any unit cost that varies with sales belongs in $e_i(x)$ instead. Keeping the two objects separate avoids counting a quality-related cost increase twice.

Proposition 9 (Local generation with quality or demand effects)**.** *Under Assumption 2, the merger chooses a strictly positive local project level if and only if* $\mathrm{LII}_M^Q > 0$*. If* $\mathrm{LII}_M^Q \leq 0$*,* $x = 0$ *maximizes the local project objective.*

The proposition has the same threshold structure as the process-only model, but the content of the threshold is different. A project can be technically valuable and nevertheless fail the private hurdle because it cannibalizes a high-margin sister product. Conversely, a project with a modest engineering cost effect can cross the hurdle when it expands demand on products with large margins.

4.4.4 Merger specificity with demand innovation

Merger specificity is assessed in exactly the same sequence as in the process-only model. For each realistic alternative, the analyst asks how the arrangement changes technical performance, implementation cost, demand, and the share of the resulting operating gain captured by the project sponsor. The additional demand term does not relax the counterfactual requirement; it makes the appropriation mechanism more important.

The relevant counterfactual remains a menu of realistic organizational arrangements. Let $a \in \mathcal{A}$ index standalone investment, licensing, a joint venture, a supply agreement, a narrower asset acquisition, or another feasible alternative. Arrangement a can have its own demand technology $q_i^a(p, x)$, cost effect $e_i^a(x)$, project cost $K_a(x)$, and appropriation weight $\omega_i^a \in [0, 1]$. The weight is the fraction of the incremental operating gain on product i internalized by the project decision maker under arrangement a.

The arrangement-specific local objective is

$$V_a^{Q,L}(x) = \sum_{i \in \mathcal{M}} \omega_i^a \left\{ \left[m_i^0 + e_i^a(x) \right] q_i^a(p^0, x) - m_i^0 q_i^0 \right\} - K_a(x),$$

and its local score is

$$\mathrm{LII}_a^Q = \sum_{i \in \mathcal{M}} \omega_i^a \left[q_i^0 e_i^{a\prime}(0) + m_i^0 q_{i,x}^{a,0} \right] - K_a'(0).$$

Proposition 10 (Strong local merger specificity)**.** *Suppose every arrangement-specific local objective is strictly concave. The quality or demand project is strongly merger-specific in the local sense if and only if*

$$\mathrm{LII}_M^Q > 0 \qquad \textit{and} \qquad \max_{a \in \mathcal{A}} \mathrm{LII}_a^Q \leq 0.$$

Under this condition, the merger selects a positive local project while every modeled less restrictive alternative selects zero.

The demand term makes appropriation especially important. A standalone firm may create consumer value that spills over to a partner's product or a platform. A licensing contract may allow the innovator to capture only part of that value. The merger can therefore enable a demand project through technical complementarity, lower implementation cost, broader appropriation, or any combination of the three.

4.4.5 The exact nominal-price gradient

The objective is to preserve the logic of BAUPP while allowing the project to change the demand system used to construct it. The derivation starts from the raw nominal-price gradient, identifies the terms that are no longer held fixed, and only then normalizes by the project-dependent own-price slope. Fix a project level x and hold nominal prices at p^0. For $i \in \mathcal{M}$, define

$$G_i^Q(x) \equiv \left. \frac{\partial \Pi^M(p, x)}{\partial p_i} \right|_{p=p^0}.$$

Differentiating the merged-profit expression gives

$$G_i^Q(x) = q_i(p^0, x) + \left[m_i^0 + e_i(x)\right] q_{i,i}(p^0, x) + \sum_{j \in \mathcal{M} \setminus \{i\}} \left[m_j^0 + e_j(x)\right] q_{j,i}(p^0, x).$$

There are now two reasons why the process-only formula is insufficient. First, the project can change the level of demand at the evaluation prices. Second, it can change the derivatives that determine own-price sensitivity and diversion.

For a local project level at which $q_{i,i}(p^0, x) < 0$, define the project-dependent signed diversion ratio

$$D_{ij}^Q(x) \equiv \frac{q_{j,i}(p^0, x)}{-q_{i,i}(p^0, x)}, \qquad j \neq i,$$

and define the own-demand repricing term

$$R_i^Q(x) \equiv \frac{q_i(p^0, x) + m_i^0 q_{i,i}(p^0, x)}{-q_{i,i}(p^0, x)}.$$

By the pre-merger first-order condition, $R_i^Q(0) = 0$.

Definition 6 (Quality-and-demand adjusted UPP)**.** *For $i \in \mathcal{M}$, nominal quality-and-demand adjusted upward pricing pressure is*

$$\mathrm{QBAUPP}_i(x) \equiv R_i^Q(x) + \sum_{j \in \mathcal{M} \setminus \{i\}} D_{ij}^Q(x) \left[m_j^0 + e_j(x)\right] - e_i(x).$$

The following then holds.

Proposition 11 (Exact decomposition of nominal pricing pressure)**.** *For every $i \in \mathcal{M}$, the nominal post-merger price gradient at (p^0, x) has the sign of* $\mathrm{QBAUPP}_i(x)$ *and satisfies*

$$G_i^Q(x) = \left[-q_{i,i}(p^0, x)\right] \mathrm{QBAUPP}_i(x).$$

Equivalently,

$$\mathrm{QBAUPP}_i(x) = \underbrace{\sum_{j \in \mathcal{M} \setminus \{i\}} D_{ij} m_j^0}_{\text{baseline merger pressure}} + \underbrace{R_i^Q(x)}_{\text{own-demand repricing}} + \underbrace{\sum_{j \in \mathcal{M} \setminus \{i\}} \left[D_{ij}^Q(x) - D_{ij}\right] m_j^0}_{\text{diversion reconfiguration}} - \underbrace{e_i(x)}_{\text{own cost effect}} + \underbrace{\sum_{j \in \mathcal{M} \setminus \{i\}} D_{ij}^Q(x) e_j(x)}_{\text{cost effects on diverted products}}.$$

The first term is the familiar value of newly internalized diversion in the absence of a project. The second term is present even for a standalone firm: it measures the incentive to reprice product i because

the project changes that product's demand level or own-price slope. The third term records the fact that quality may redirect customers across products and thereby change the value of diversion. The fourth is the own-product marginal-cost credit. The final term is the cross-product feedback emphasized in the process-only model: a cost reduction on a destination product makes diverted sales more profitable.

The displayed decomposition explains why a single fixed diversion formula is not generally available. A demand-enhancing project can change not only the numerator of the UPP calculation but the demand system that generates the diversion ratios themselves.

4.4.6 A first-order incidence test

The exact formula is appropriate when the analyst can estimate demand at several project states. Often the available evidence supports only the marginal demand shift and the marginal change in price sensitivities at the no-project point. A first-order expansion then provides a parsimonious incidence test and shows exactly which derivatives are needed. Define

$$\eta_i \equiv e_i'(0), \qquad g_i \equiv q_{i,x}^0, \qquad h_{ji} \equiv \left.\frac{\partial^2 q_j(p,x)}{\partial p_i \partial x}\right|_{(p^0,0)}.$$

The object g_i is the direct demand-level effect. The object h_{ji} measures how the project changes the response of demand for product j to the price of product i; it therefore captures changes in own-price sensitivity and in diversion. Normalize the nominal gradient by the pre-project own-price slope:

$$B_i^Q(x) \equiv \frac{G_i^Q(x)}{-q_{i,i}^0}.$$

Define the demand-system contribution

$$\Gamma_i^Q \equiv \frac{g_i + \sum_{j\in\mathcal{M}} m_j^0 h_{ji}}{-q_{i,i}^0}.$$

The following holds.

Proposition 12 (Marginal project effect on nominal pricing pressure)**.** *For every $i \in \mathcal{M}$,*

$$B_i^Q(x) = \sum_{j\in\mathcal{M}\setminus\{i\}} D_{ij} m_j^0 + x\left[\Gamma_i^Q - \eta_i + \sum_{j\in\mathcal{M}\setminus\{i\}} D_{ij}\eta_j\right] + o(x).$$

A marginal increase in project intensity reduces nominal pricing pressure on product i if and only if

$$\eta_i > \sum_{j\in\mathcal{M}\setminus\{i\}} D_{ij}\eta_j + \Gamma_i^Q.$$

The condition separates the process and demand channels. The left-hand side is the own-product marginal-cost benefit. The first term on the right is the diversion-weighted cost benefit on other merging products. The final term is the change in the demand system.

The linear-demand application in Online Resource 1, Section B.3 clarifies the sign of the demand-system contribution.

Remark 1 (Generation and transmission use different demand objects)**.** *The generation score uses $m_i^0 g_i$: it asks whether the project creates private operating value at unchanged prices. The nominal transmission test uses g_i and the slope effects h_{ji}: it asks how the project's demand change alters the incentive to reprice each product. A positive contribution to project generation need not reduce nominal pricing pressure.*

4.4.7 Money-metric quality and hedonic prices

The general demand-shift model is necessary when quality changes substitution patterns in a way that cannot be summarized by a willingness-to-pay shift. It can be simplified when quality has a credible money-metric interpretation. The hedonic representation expresses a quality gain and a cost reduction in the same units from the consumer's perspective. It does not require nominal prices to fall; it asks whether the quality-adjusted price rises or falls. Let $v_i(x)$ be the willingness-to-pay increment for product i, with $v_i(0) = 0$, and suppose demand depends on the quality-adjusted, or hedonic, price

$$h_i = p_i - v_i(x).$$

Specifically,

$$q_i(p, x) = \bar{q}_i\big(h\big) = \bar{q}_i\big(p - v(x)\big).$$

Set $v_i(x) = 0$ for products not affected by the project. This specification means that, from the consumer's perspective, a one-dollar increase in willingness to pay offsets a one-dollar increase in nominal price. It is the natural case for implementing quality-adjusted UPP when quality can be translated into willingness-to-pay units.

Let

$$\nu_i \equiv v_i'(0).$$

At fixed nominal prices,

$$q_{i,x}^0 = -\sum_{k \in \mathcal{M}} q_{i,k}^0 \nu_k.$$

This identity permits a sharp decomposition of the private quality incentive.

Proposition 13 (Generation of money-metric quality)**.** *Under the hedonic-demand specification, the local innovation incentive is*

$$\begin{aligned} \mathrm{LII}_M^H = &\sum_{i \in \mathcal{M}} q_i^0 \eta_i \\ &+ \sum_{k \in \mathcal{M}} \big(-q_{k,k}^0\big) \left[m_k^0 - \sum_{i \in \mathcal{M} \setminus \{k\}} D_{ki} m_i^0 \right] \nu_k - K_M'(0). \end{aligned}$$

The bracket in the hedonic local innovation incentive is the net margin value of improving product k. The term $(-q_{k,k}^0) m_k^0 \nu_k = q_k^0 \nu_k$ is the value of the additional demand attracted to product k. The subtraction

$$(-q_{k,k}^0) \sum_{i \in \mathcal{M} \setminus \{k\}} D_{ki} m_i^0 \nu_k$$

is the margin lost when the quality improvement draws customers away from other products in the merged

portfolio. Quality is privately valuable to the extent that it expands the merged firm's total margin base rather than merely reshuffling sales among its products.

Define the net quality-appropriation margin

$$\chi_k^Q \equiv m_k^0 - \sum_{i \in \mathcal{M}\setminus\{k\}} D_{ki} m_i^0.$$

For substitute products, high diversion to high-margin sister products reduces χ_k^Q. It can even make the local private demand return negative. This does not mean that consumers dislike the quality improvement. It means that, at unchanged prices, the improvement cannibalizes more internal margin than it creates.

4.4.8 Quality-adjusted upward pricing pressure

Having expressed quality in money units, we can repeat the transmission exercise in hedonic prices. A dollar of willingness-to-pay improvement and a dollar of marginal-cost reduction both relax the quality-adjusted pricing constraint by one dollar on the product where the benefit occurs. Product incidence and diversion remain essential. Rewrite merged profit as a function of hedonic prices:

$$\Pi^M(h, x) = \sum_{i \in \mathcal{M}} \left[h_i - c_i^0 + v_i(x) + e_i(x)\right] \bar{q}_i(h) - K_M(x).$$

Define the total money-metric product benefit

$$b_i(x) \equiv e_i(x) + v_i(x).$$

A cost reduction and a willingness-to-pay increase enter the hedonic margin in exactly the same way. A one-dollar cost reduction permits a one-dollar lower nominal price at a given margin. A one-dollar willingness-to-pay gain permits a one-dollar higher nominal price without increasing the hedonic price faced by consumers.

Evaluating the price gradient at $h = p^0$ holds consumers' hedonic prices at their pre-merger levels. Equivalently, nominal prices are evaluated at $p = p^0 + v(x)$, so the nominal price adjustment exactly matches the quality increment before the merger incentive is added.

Definition 7 (Hedonic benefit-adjusted UPP)**.** *For $i \in \mathcal{M}$, define*

$$\text{HBAUPP}_i(x) \equiv \sum_{j \in \mathcal{M}\setminus\{i\}} D_{ij}\left[m_j^0 + b_j(x)\right] - b_i(x).$$

Then, we have.

Proposition 14 (Hedonic equivalence)**.** *Under the hedonic-demand specification, the post-merger profit gradient with respect to the hedonic price h_i, evaluated at the pre-merger hedonic price vector $h^0 = p^0$, has the sign of* $\text{HBAUPP}_i(x)$. *In particular,*

$$\left.\frac{\partial \Pi^M(h, x)}{\partial h_i}\right|_{h=p^0} = (-q_{i,i}^0)\,\text{HBAUPP}_i(x).$$

In vector form,

$$\mathbf{HBAUPP}(x) = \mathbf{D}_{\mathcal{M}}\left[\mathbf{m}^0_{\mathcal{M}} + \mathbf{b}_{\mathcal{M}}(x)\right] - \mathbf{b}_{\mathcal{M}}(x).$$

Hence, there is no upward pressure on any hedonic price if and only if

$$\left(\mathbf{I} - \mathbf{D}_{\mathcal{M}}\right)\mathbf{b}_{\mathcal{M}}(x) \geq \mathbf{D}_{\mathcal{M}}\mathbf{m}^0_{\mathcal{M}}$$

componentwise. If $\mathbf{D}_{\mathcal{M}}$ *is nonnegative and its spectral radius is less than one, the minimal componentwise compensating benefit vector is*

$$\mathbf{b}^{\text{crit}} = \left(\mathbf{I} - \mathbf{D}_{\mathcal{M}}\right)^{-1}\mathbf{D}_{\mathcal{M}}\mathbf{m}^0_{\mathcal{M}}.$$

The result is the quality counterpart to benefit-adjusted UPP. Once quality is expressed as a willingness-to-pay increment, the analyst can add it to the marginal-cost effect product by product. The sum $b_i = e_i + v_i$, rather than either component in isolation, is the relevant own-product benefit for hedonic pricing pressure.

The cross-product term remains essential. A quality improvement on product j raises the value of customers diverted from product i to product j. It therefore increases the incentive to raise the hedonic price of product i. The project reduces hedonic pressure on product i only when its own money-metric benefit dominates the diversion-weighted benefits on the other products in the merged portfolio.

Corollary 4 (Hedonic product-incidence condition)**.** *A marginal increase in project intensity reduces hedonic pricing pressure on product* i *if and only if*

$$e_i'(x) + v_i'(x) > \sum_{j \in \mathcal{M}\setminus\{i\}} D_{ij}\left[e_j'(x) + v_j'(x)\right].$$

This special case is empirically attractive because it does not require the analyst to estimate how every diversion ratio changes with project intensity. It does require credible evidence that the quality change can be represented as a money-metric shift and credible estimates of the product-level values $v_i(x)$. If quality changes the ranking of products, creates a new product, or changes substitution in a way that is not equivalent to a price shift, the general formula in Proposition 11 should be used instead.

4.4.9 A linear-quadratic workhorse

This benchmark combines the generation and transmission formulas in a form that can be implemented with product-level process coefficients, willingness-to-pay coefficients, and a quadratic project cost. As in Section 3.4, the restrictions are local approximations to the project, not restrictions on the entire demand system. Their value is to separate the scalar amount of innovation generated from the product-level mix of process and quality benefits.

For implementation, suppose

$$e_i(x) = \eta_i x, \qquad v_i(x) = \nu_i x, \qquad K_M(x) = k_M x + \frac{\kappa_M}{2}x^2, \qquad \kappa_M > 0.$$

Define the effective quality-and-process project base

$$A_M^Q \equiv \sum_{i \in \mathcal{M}} q_i^0 \eta_i + \sum_{k \in \mathcal{M}} \left(-q_{k,k}^0\right) \left[m_k^0 - \sum_{i \in \mathcal{M} \setminus \{k\}} D_{ki} m_i^0 \right] \nu_k .$$

The predicted local project intensity is

$$x_M^{Q,L} = \left[\frac{A_M^Q - k_M}{\kappa_M} \right]_+ .$$

The endogenous hedonic pricing-pressure measure is then

$$\text{HEEUPP}_i = \sum_{j \in \mathcal{M} \setminus \{i\}} D_{ij} m_j^0 + \left[\sum_{j \in \mathcal{M} \setminus \{i\}} D_{ij} (\eta_j + \nu_j) - (\eta_i + \nu_i) \right] x_M^{Q,L} .$$

The two displayed formulas preserve the two-hurdle logic of the baseline model. The effective project base A_M^Q determines whether the combined process and quality project clears the implementation hurdle. Its process part values unit-cost savings on observed quantities. Its quality part values additional demand using the net margin retained after cannibalization within the merged portfolio. Conditional on generation, $x_M^{Q,L}$ scales every product's total money-metric benefit $\eta_i + \nu_i$.

The hedonic EEUPP expression then asks whether that generated profile is located on the right products. A larger project lowers pressure on product i only when its own process-plus-quality coefficient exceeds the diversion-weighted coefficients on the destination products. Thus a high private return to quality can coexist with weak consumer transmission when the project mainly improves products that receive diverted demand. The first equation is a generation test; the second is a product-incidence and transmission test. The compact symmetric specialization below combines the two hurdles.

4.4.10 Consumer transmission

The final step moves from the firm's pricing incentive to the consumer object. With demand innovation there are two first-order channels: a direct utility gain at unchanged prices and the indirect effect of the induced nominal-price response. Reporting only nominal prices omits the first; reporting only willingness to pay omits the second. Suppose demand is generated by a differentiable consumer-surplus function $S(p, x)$ satisfying

$$\frac{\partial S(p,x)}{\partial p_i} = -q_i(p,x).$$

Let Δp be the local equilibrium nominal-price change induced by the merger and project.

Proposition 15 (First-order consumer transmission)**.** *Around $(p^0, 0)$, the first-order change in consumer surplus is*

$$\Delta \text{CS} \approx S_x^0 x - (q^0)' \Delta p, \qquad S_x^0 \equiv \left. \frac{\partial S(p,x)}{\partial x} \right|_{(p^0,0)} .$$

Under the money-metric quality specification $S(p,x) = \bar{S}(p - v(x))$,

$$S_x^0 = \sum_{i \in \mathcal{N}} q_i^0 v_i,$$

and therefore

$$\Delta \mathrm{CS} \approx -(q^0)' \Delta h, \qquad \Delta h \equiv \Delta p - vx.$$

The displayed general consumer-surplus expression separates direct and indirect consumer transmission. The term $S_x^0 x$ is the direct value of improved quality, reliability, variety, or service at unchanged prices. The term $-(q^0)'\Delta p$ is the ordinary first-order price effect.

In the hedonic special case, the two channels combine into the change in quality-adjusted prices. A nominal price increase equal to the willingness-to-pay gain leaves the hedonic price unchanged and has no first-order consumer-surplus effect. A smaller nominal price increase lowers the hedonic price and benefits consumers. A larger increase raises the hedonic price and harms consumers.

Let $\mathcal{P}^H$ be a local pass-through matrix that maps embedded hedonic pricing pressure into equilibrium hedonic-price changes. Define

$$\widetilde{b}_i^H(x) = \begin{cases} \mathrm{HBAUPP}_i(x), & i \in \mathcal{M}, \\ 0, & i \notin \mathcal{M}. \end{cases}$$

Then a first-order implementation is

$$\Delta h \approx \mathcal{P}^H \widetilde{b}^H(x), \qquad \Delta \mathrm{CS} \approx -(q^0)' \mathcal{P}^H \widetilde{b}^H(x).$$

This is the quality-adjusted analogue of the baseline pass-through step. The direct quality benefit is not omitted; it is embedded in the hedonic price change.

Symmetric two-product benchmark. Let both products have quantity q, margin m, and bilateral diversion $D < 1$. Suppose the project creates the common process saving ηx and the common money-metric quality gain vx, with $K_M(x) = k_M x + (\kappa_M/2)x^2$. Then

$$x_M^{Q,L} = \left[\frac{2q[\eta + (1-D)v] - k_M}{\kappa_M} \right]_+,$$

and common hedonic pricing pressure is

$$\mathrm{HBAUPP}(x) = Dm - (1-D)(\eta + v)x.$$

Generation therefore requires $2q[\eta + (1-D)v] > k_M$, while hedonic transmission requires $x_M^{Q,L} \geq Dm/[(1-D)(\eta+v)]$. Diversion reduces both the private return to a common quality improvement and the amount of that improvement that disciplines quality-adjusted prices.

4.5 A joint local system for prices and innovation

The logic followed in the baseline model is sequential. It first evaluates the project at the observed pre-merger price vector, obtains a predicted project level, and then maps the resulting benefit-adjusted

pricing pressure into price changes. This "triangular" structure keeps the evidentiary burden low and preserves the distinction between benefit generation and benefit transmission. It is an approximation, however. A merger-induced price change can alter the quantities on which cost savings are earned; innovation can change pricing incentives; and non-merging rivals can respond to both. When the relevant derivatives are available, these first-order feedbacks can be incorporated in a single local system without moving immediately to a full nonlinear merger simulation.

The construction can be read as one Newton step around the observed pre-merger, no-project point. The right-hand side contains the two shocks already used by the sequential analysis: ordinary merger pricing pressure and the local project score. The Jacobian on the left describes how price conditions and the project condition respond to prices and innovation. The solution therefore shows both the direct effects and the first round of feedback between them.

Let

$$z = \begin{pmatrix} p \\ x \end{pmatrix}, \qquad z_0 = \begin{pmatrix} p^0 \\ 0 \end{pmatrix},$$

and write the post-merger first-order-condition system as

$$F(z) = \begin{pmatrix} F_p(z) \\ F_x(z) \end{pmatrix}.$$

For $i \in \mathcal{M}$, the ith component of F_p is the derivative of the merged firm's profit with respect to p_i. For $i \notin \mathcal{M}$, it is the derivative of the independent owner of product i with respect to its own price. The final component is the merged firm's project first-order condition,

$$F_x(z) = \frac{\partial \Pi^M(p, x)}{\partial x}.$$

The value of the system at the pre-merger, no-project point has a direct economic interpretation. Let

$$S \equiv \operatorname{diag}(-q^0_{1,1}, \ldots, -q^0_{N,N})$$

and let $\widetilde{b}^0 \equiv \widetilde{b}(0)$ denote the embedded no-efficiency BAUPP vector defined above. Under the baseline single-product ownership structure,

$$F_p(z_0) = S\widetilde{b}^0, \qquad F_x(z_0) = \ell_0,$$

where $\ell_0 = \mathrm{LII}_M$ in the process-only model and $\ell_0 = \mathrm{LII}^Q_M$ when the project also changes demand. Thus the joint system starts from exactly the two objects used by the sequential method: the no-project pricing-pressure vector and the local innovation score.

Partition the Jacobian at z_0 as

$$J_F(z_0) = \begin{pmatrix} A & B \\ C & d \end{pmatrix}, \qquad A \equiv \frac{\partial F_p}{\partial p'}, \quad B \equiv \frac{\partial F_p}{\partial x}, \quad C \equiv \frac{\partial F_x}{\partial p'}, \quad d \equiv \frac{\partial F_x}{\partial x},$$

with all derivatives evaluated at z_0. The four blocks have distinct roles. The matrix A contains own-price curvature and strategic interaction among all products. The vector B records how a marginal increase in project intensity changes each price first-order condition. The row vector C records how price changes

alter the marginal return to the project. Finally, d is the local curvature of the project objective and is negative under strict local concavity.

The first-order joint approximation is therefore

$$\begin{pmatrix} A & B \\ C & d \end{pmatrix} \begin{pmatrix} \Delta p \\ \Delta x \end{pmatrix} = - \begin{pmatrix} S\widetilde{b}^0 \\ \ell_0 \end{pmatrix}.$$

If the Jacobian is nonsingular, the unconstrained Newton step is

$$\begin{pmatrix} \Delta p^J \\ \Delta x^J \end{pmatrix} = -J_F(z_0)^{-1} \begin{pmatrix} S\widetilde{b}^0 \\ \ell_0 \end{pmatrix}. \tag{8}$$

Because project intensity satisfies $x \geq 0$, the joint local system is an interior candidate. If it yields $\Delta x^J \leq 0$, the boundary candidate sets $x = 0$ and resolves the price block with the project fixed; the analyst then verifies the local Kuhn–Tucker condition $F_x(p^0 + \Delta p, 0) \leq 0$. The same active-set logic applies if project intensity has an upper bound.

The block form makes the relation to pass-through especially transparent. If A is invertible, the price equation gives

$$\Delta p = -A^{-1}\big(S\widetilde{b}^0 + B\Delta x\big).$$

Hence $\mathcal{P} = -A^{-1}S$ is the local pass-through matrix when project intensity is held fixed, and $S^{-1}B$ is the marginal change in normalized pricing pressure generated by one unit of the project. Substitution into the project equation gives

$$\Delta x^J = -\big[d - CA^{-1}B\big]^{-1}\big[\ell_0 - CA^{-1}S\widetilde{b}^0\big].$$

The second bracket is the innovation score after allowing the no-project price response to change the project's operating base. The first bracket is an effective curvature term that also incorporates the feedback loop from innovation to prices and back to innovation.

The sequential method is recovered from the same system by setting $C = 0$: prices are allowed to respond to innovation through B, but the innovation decision is evaluated at p^0. In a locally quadratic project problem, $\Delta x = -\ell_0/d$ is then the project choice, and the price block of the joint system becomes the usual pass-through step conditional on that choice. The one-step quantity correction in Section 4.3 retains the term $-CA^{-1}S\widetilde{b}^0$ but not the full feedback term $CA^{-1}B$.

An expanded product-incidence application is collected in Online Resource 1, Section B.5.

5 Conclusion

This paper develops a local framework for a recurring merger claim: the transaction enables an innovation that would not otherwise be generated and the resulting benefit offsets unilateral price effects. The framework separates that claim into two tests. The generation test asks whether the merger changes project feasibility or raises the expected return above the investment hurdle relative to realistic less restrictive arrangements. The transmission test places the predicted product-level efficiency vector in the post-merger price gradient.

The separation produces two practical cautions. A horizontal merger does not mechanically create a direct incentive to invest when the relevant own-product benefit was already internalized before the

transaction. A credible merger-specific story requires a newly internalized cross-product benefit, a common capability, an appropriation change, or a genuine feasibility constraint. Moreover, aggregate savings are not enough to assess price effects. Efficiencies on products receiving diverted demand raise the value of diversion and can strengthen the incentive to raise another product's price. The amount, timing, and product incidence of the benefit all matter.

The method is designed as a first screen. It uses the same pre-merger evidence that supports conventional UPP–quantities, margins, diversion, and pass-through–together with contemporaneous evidence on project technology, variable and fixed cost, risk, timing, and alternative arrangements. Smooth projects can be screened at the margin; projects with fixed launch costs or indivisibilities require an NPV test; benefits that last for several periods must be valued on discounted quantity exposure. These are distinct departures from the one-period baseline and should be reported separately.

On the consumer side, BAUPP is only the direct incentive object. The local pass-through system maps that object into market-wide prices and consumer-surplus effects, including the responses of non-merging products. Section 4 shows how uncertainty, multiproduct ownership, quality, and price–investment feedback can be incorporated when the record supports the additional structure. The broader implication is methodological: dynamic efficiencies should enter merger analysis as the output of an incentive and counterfactual test, and then be subjected to the same product-level and consumer-incidence discipline used to assess harm.

Acknowledgements

We are grateful to Michele Bisceglia, Raffaele Fiocco, Shiva Shekhar, Emanuele Tarantino, and Federico Vaccari for useful comments. The views expressed in this paper are solely the responsibility of the authors and do not necessarily reflect the views of the organizations with which they are affiliated.

Statements and declarations

Data availability. The paper is theoretical and does not use an external dataset.

Supplementary information

Online Resource 1. The dynamic normalization, comparative statics for the compact merger-to-monopoly example, expanded multiproduct, quality, and feedback applications, proofs of supplemental results, and additional implementation derivations. Proofs of all formal results stated in the paper are contained in Appendices A and B below.

A Proofs of the baseline results

A.1 Proof of Proposition 1

The local project objective is

$$V_M^L(x) = \sum_{i \in \mathcal{M}} q_i^0 e_i^M(x) - K_M(x).$$

By Assumption 1, V_M^L is concave and attains a maximum. Its right derivative at zero is LII_M. If $\mathrm{LII}_M \leq 0$, concavity implies

$$V_M^L(x) \leq V_M^L(0) + x\,\mathrm{LII}_M \leq V_M^L(0)$$

for every feasible $x \geq 0$, so $x = 0$ is optimal. If $\mathrm{LII}_M > 0$, then $V_M^L(x) > V_M^L(0)$ for all sufficiently small positive x. Hence no maximizer can be zero. □

A.2 Proof of Proposition 2

Apply Proposition 1 separately to the merger and to every arrangement $a \in \mathcal{A}$, using the arrangement-specific objective defined above. The merger chooses a positive local project if and only if $\mathrm{LII}_M > 0$. Each alternative chooses zero if and only if $\mathrm{LII}_a \leq 0$. Combining these conditions gives the stated strong-specificity condition. □

A.3 Proof of Proposition 3

For $i \in \mathcal{M}$, differentiating the merged-profit expression with respect to p_i and evaluating at p^0 gives

$$q_i^0 + \left[m_i^0 + e_i^M(x)\right] q_{i,i}^0 + \sum_{j \in \mathcal{M}\setminus\{i\}} \left[m_j^0 + e_j^M(x)\right] q_{j,i}^0.$$

The pre-merger first-order condition is $q_i^0 + m_i^0 q_{i,i}^0 = 0$. Subtracting it yields

$$e_i^M(x) q_{i,i}^0 + \sum_{j \in \mathcal{M}\setminus\{i\}} \left[m_j^0 + e_j^M(x)\right] q_{j,i}^0.$$

Dividing by $-q_{i,i}^0 > 0$ gives the stated BAUPP expression; rearranging gives the stated decomposition. □

A.4 Proof of Corollary 1

Differentiating the BAUPP expression while holding the local diversion ratios fixed gives

$$\frac{\mathrm{d}\,\mathrm{BAUPP}_i(x)}{\mathrm{d}x} = \sum_{j \in \mathcal{M}\setminus\{i\}} D_{ij} e_j^{M\prime}(x) - e_i^{M\prime}(x).$$

The derivative is negative exactly under the stated incidence condition, and positive when the inequality is reversed. □

A.5 Proof of Corollary 2

For a single product, $\mathrm{BAUPP}_i(x) \leq 0$ is equivalent to the stated componentwise inequality. Stacking these inequalities and using the vector BAUPP expression gives the stated passing condition. If $\mathbf{D}_{\mathcal{M}}$ is

nonnegative and $\rho(\mathbf{D}_{\mathcal{M}}) < 1$, then

$$(\mathbf{I} - \mathbf{D}_{\mathcal{M}})^{-1} = \sum_{n=0}^{\infty} \mathbf{D}_{\mathcal{M}}^{n}$$

exists and is componentwise nonnegative. Premultiplying any feasible inequality in eq. (4) by this inverse yields

$$\boldsymbol{e}_{\mathcal{M}}(x) \geq (\mathbf{I} - \mathbf{D}_{\mathcal{M}})^{-1} \mathbf{D}_{\mathcal{M}} \boldsymbol{m}_{\mathcal{M}}.$$

The vector on the right satisfies the passing condition with equality. It is therefore the least componentwise vector in the passing set. The last step does not imply that every vector componentwise above it passes, because $\mathbf{I} - \mathbf{D}_{\mathcal{M}}$ has nonpositive off-diagonal entries. □

A.6 Proof of Proposition 4

In the linear-quadratic model, strong merger-specific generation requires $A_M - k_M > 0$ and $A_a - k_a \leq 0$ for the modeled alternatives. These restrictions do not constrain the incidence coefficient

$$\sum_{j \in \mathcal{M} \setminus \{i\}} D_{ij} \eta_j^M - \eta_i^M.$$

Choose $A_M - k_M > 0$ but sufficiently small that the generated efficiency does not offset the no-efficiency pressure on at least one product. Then $\text{EEUPP}_i > 0$. Alternatively, choose the incidence coefficient to be positive for some product i. The displayed EEUPP expression is then increasing in x_M^L, so a larger generated project raises pricing pressure on that product. □

B Proofs of the extensions

B.1 Uncertain project success

B.1.1 Proof of Proposition 5

The merged firm's expected local payoff is concave by assumption. Its right derivative at the origin is

$$\sum_{i \in \mathcal{M}} q_i^0 \left. \frac{\partial \bar{e}_i(x)}{\partial x} \right|_{x=0} - K_M'(0) = \text{LII}_M^U,$$

where expectation and differentiation can be interchanged under the maintained local regularity assumptions. If this derivative is positive, concavity implies that the objective initially increases and every local maximizer is strictly positive. If it is weakly negative, concavity implies that the objective is maximized at the boundary $x = 0$. □

B.1.2 Proof of Proposition 6

State-contingent benefit-adjusted UPP is

$$\text{BAUPP}_i(x, \theta) = \sum_{j \in \mathcal{M} \setminus \{i\}} D_{ij} \left[m_j^0 + \widetilde{e}_j(x, \theta) \right] - \widetilde{e}_i(x, \theta).$$

The diversion ratios and margins are state-independent by assumption. Taking expectations and applying linearity yields the expected-BAUPP expression. □

B.2 Multiproduct ownership

B.2.1 Proof of Proposition 7

The pre-merger first-order condition for product i is

$$q_i^0 + \sum_k \Omega_{ik}^0 m_k^0 q_{k,i}^0 = 0.$$

After the merger and project implementation, the local derivative of the relevant owner's profit is

$$q_i^0 + \sum_k \Omega_{ik}^1 \left[m_k^0 + e_k(x)\right] q_{k,i}^0.$$

Subtracting the pre-merger first-order condition gives

$$\sum_k (\Omega_{ik}^1 - \Omega_{ik}^0) m_k^0 q_{k,i}^0 + \sum_k \Omega_{ik}^1 e_k(x) q_{k,i}^0.$$

The ownership difference is zero for $k = i$, while $\Omega_{ii}^1 = 1$. Dividing by $-q_{i,i}^0 > 0$ makes the own-efficiency term $-e_i(x)$ and every other term signed diversion times the relevant margin or efficiency. This is the stated multiproduct BAUPP expression. □

B.2.2 Proof of Corollary 3

With $\boldsymbol{e}(x) = x\boldsymbol{\eta}$, the vector BAUPP expression becomes

$$\boldsymbol{b}(x) = \boldsymbol{h} - x\boldsymbol{g}.$$

If every $g_i > 0$, the inequality $b_i(x) \leq 0$ is equivalent to $x \geq h_i/g_i$. All components are nonpositive exactly when $x \geq \max_i h_i/g_i$. If $h_i > 0$ and $g_i \leq 0$ for some product, then $b_i(x) = h_i - xg_i \geq h_i > 0$ for all $x \geq 0$, so no intensity along the fixed profile can pass on that product. □

B.2.3 Proof of Proposition 8

Because $\mathbf{D}$ is nonnegative, irreducible, and row-stochastic, the Perron–Frobenius theorem provides a strictly positive left stationary vector $\boldsymbol{\pi}$ satisfying

$$\boldsymbol{\pi}'\mathbf{D} = \boldsymbol{\pi}'.$$

Let $\boldsymbol{h} = \mathbf{H}\boldsymbol{m}^0$. If a finite efficiency vector made every component of $\boldsymbol{b}$ nonpositive, then

$$(\mathbf{I} - \mathbf{D})\boldsymbol{e} \geq \boldsymbol{h}.$$

Premultiplication by $\boldsymbol{\pi}'$ would give

$$0 = \boldsymbol{\pi}'(\mathbf{I} - \mathbf{D})\boldsymbol{e} \geq \boldsymbol{\pi}'\boldsymbol{h} > 0,$$

a contradiction. The final strict inequality follows because $\boldsymbol{\pi} > \mathbf{0}$ and $\boldsymbol{h} \geq \mathbf{0}$ is nonzero. □

B.3 Demand-enhancing innovation and quality

B.3.1 Proof of Proposition 9

Differentiate the local objective at $x = 0$. Since $e_i(0) = 0$ and $q_i(p^0, 0) = q_i^0$, the right derivative is

$$\sum_{i \in \mathcal{M}} \left[q_i^0 e_i'(0) + m_i^0 q_{i,x}^0 \right] - K_M'(0) = \mathrm{LII}_M^{\mathcal{Q}} .$$

Under Assumption 2, the objective is strictly concave and attains a maximum. The same supporting-line argument used in the proof of Proposition 1 implies that $x = 0$ is optimal when the right derivative is nonpositive, whereas a positive right derivative makes every maximizer strictly positive. □

B.3.2 Proof of Proposition 10

Apply Proposition 9 to the merger objective and to each alternative arrangement using its arrangement-specific local objective. The merger selects a positive project if and only if $\mathrm{LII}_M^{\mathcal{Q}} > 0$, and arrangement a selects zero if and only if $\mathrm{LII}_a^{\mathcal{Q}} \leq 0$. Taking the maximum over the modeled alternatives gives the stated specificity condition. □

B.3.3 Proof of Proposition 11

Starting from the raw gradient expression, separate the own-product pre-project margin:

$$\begin{aligned} G_i^{\mathcal{Q}}(x) = q_i(p^0, x) + m_i^0 q_{i,i}(p^0, x) + e_i(x) q_{i,i}(p^0, x) \\ + \sum_{j \in \mathcal{M} \setminus \{i\}} \left[m_j^0 + e_j(x) \right] q_{j,i}(p^0, x). \end{aligned}$$

Factoring out $-q_{i,i}(p^0, x) > 0$ and using the project-diversion and repricing definitions gives the factored gradient expression. Adding and subtracting $\sum_{j \neq i} D_{ij} m_j^0$ inside the quality-adjusted BAUPP expression yields the stated decomposition. □

B.3.4 Proof of Proposition 12

Around $x = 0$,

$$\begin{aligned} q_i(p^0, x) &= q_i^0 + g_i x + o(x), \\ q_{j,i}(p^0, x) &= q_{j,i}^0 + h_{ji} x + o(x), \\ e_j(x) &= \eta_j x + o(x). \end{aligned}$$

Substitution into the raw gradient expression, followed by the pre-merger first-order condition, gives

$$\begin{aligned} G_i^{\mathcal{Q}}(x) = \sum_{j \in \mathcal{M} \setminus \{i\}} m_j^0 q_{j,i}^0 \\ + x \left[g_i + \sum_{j \in \mathcal{M}} m_j^0 h_{ji} + \eta_i q_{i,i}^0 + \sum_{j \in \mathcal{M} \setminus \{i\}} \eta_j q_{j,i}^0 \right] + o(x). \end{aligned}$$

Divide by $-q_{i,i}^0$ and use $D_{ij} = q_{j,i}^0 / (-q_{i,i}^0)$ together with the definition of $\Gamma_i^{\mathcal{Q}}$. This gives the local-pressure expansion. Its first-order coefficient is negative exactly under the stated incidence condition. □

B.3.5 Proof of Proposition 13

Under the hedonic-demand specification,

$$q^0_{i,x} = -\sum_{k\in\mathcal{M}} q^0_{i,k}\nu_k.$$

The demand component of the quality local innovation incentive can therefore be rearranged as

$$\begin{aligned}\sum_{i\in\mathcal{M}} m^0_i q^0_{i,x} &= -\sum_{k\in\mathcal{M}}\left(\sum_{i\in\mathcal{M}} m^0_i q^0_{i,k}\right)\nu_k \\ &= \sum_{k\in\mathcal{M}}(-q^0_{k,k})\left[m^0_k - \sum_{i\in\mathcal{M}\setminus\{k\}} D_{ki} m^0_i\right]\nu_k.\end{aligned}$$

The second equality uses $q^0_{i,k} = (-q^0_{k,k})D_{ki}$ for $i \neq k$. Adding the process component $\sum_i q^0_i \eta_i$ and subtracting $K'_M(0)$ gives the stated expression. □

B.3.6 Proof of Proposition 14

Because $p = h + v(x)$, merged profit can be written as

$$\Pi^M(h,x) = \sum_{i\in\mathcal{M}}\left[h_i - c^0_i + b_i(x)\right]\bar{q}_i(h) - K_M(x).$$

At $h = p^0$, differentiation with respect to h_i yields

$$q^0_i + \left[m^0_i + b_i(x)\right]q^0_{i,i} + \sum_{j\in\mathcal{M}\setminus\{i\}}\left[m^0_j + b_j(x)\right]q^0_{j,i}.$$

Subtract the pre-merger first-order condition and divide by $-q^0_{i,i} > 0$ to obtain the hedonic BAUPP formula. Stacking the product conditions gives the vector hedonic condition, and rearrangement gives the hedonic critical-benefit condition. If $\mathbf{D}_\mathcal{M}$ is nonnegative and has spectral radius below one, its resolvent is componentwise nonnegative. Thus every passing benefit vector is componentwise at least the hedonic critical vector, while that vector itself satisfies the condition with equality. □

B.3.7 Proof of Corollary 4

Differentiate the hedonic BAUPP expression and use $b'_i(x) = e'_i(x) + v'_i(x)$. Hedonic pressure falls exactly when

$$b'_i(x) > \sum_{j\in\mathcal{M}\setminus\{i\}} D_{ij} b'_j(x),$$

which is the hedonic incidence condition. □

B.3.8 Proof of Proposition 15

A first-order Taylor expansion of $S(p,x)$ around $(p^0, 0)$ gives

$$\Delta\,\mathrm{CS} \approx \nabla_p S(p^0,0)'\Delta p + S^0_x x.$$

By Roy's identity, $\nabla_p S(p^0, 0) = -q^0$, which yields the stated general consumer-surplus expression. Under $S(p, x) = \bar{S}(p - v(x))$, the chain rule gives

$$S_x^0 = \sum_{i \in \mathcal{N}} \frac{\partial \bar{S}}{\partial h_i}(-\nu_i) = \sum_{i \in \mathcal{N}} q_i^0 \nu_i.$$

Substitution into the general expansion yields

$$\Delta \mathrm{CS} \approx -(q^0)'\big(\Delta p - \nu x\big) = -(q^0)'\Delta h,$$

which is the stated hedonic consumer-surplus expression. □

References

Arrow, Kenneth J. (1962). "Economic Welfare and the Allocation of Resources for Invention". In: *The Rate and Direction of Inventive Activity: Economic and Social Factors*. Princeton, NJ: Princeton University Press, pp. 609–626. DOI: 10.1515/9781400879762-024.

Bourreau, Marc, Bruno Jullien, and Yassine Lefouili (2026). "Horizontal Mergers and Incremental Innovation". In: *The RAND Journal of Economics* 57.1, pp. 122–139. DOI: 10.1111/1756-2171.70045.

Choné, Philippe and Laurent Linnemer (2008). "Assessing Horizontal Mergers under Uncertain Efficiency Gains". In: *International Journal of Industrial Organization* 26.4, pp. 913–929. DOI: 10.1016/j.ijindorg.2007.08.002.

Competition and Markets Authority (Mar. 18, 2021). *Merger Assessment Guidelines*. Guidance CMA129. London: Competition and Markets Authority. URL: https://www.gov.uk/government/publications/merger-assessment-guidelines.

European Commission (Apr. 30, 2026). *Draft Communication from the Commission: Guidelines on the Assessment of Mergers under Council Regulation (EC) No 139/2004 on the Control of Concentrations between Undertakings*. Brussels: European Commission. URL: https://competition-policy.ec.europa.eu/document/download/46dde10f-85c1-4590-a3f4-2b71f85685ef_en.

Farrell, Joseph and Carl Shapiro (1990). "Horizontal Mergers: An Equilibrium Analysis". In: *American Economic Review* 80.1, pp. 107–126.

— (2010). "Antitrust Evaluation of Horizontal Mergers: An Economic Alternative to Market Definition". In: *The B.E. Journal of Theoretical Economics* 10.1, 9, pp. 1–41. DOI: 10.2202/1935-1704.1563.

Federico, Giulio, Gregor Langus, and Tommaso Valletti (2017). "A Simple Model of Mergers and Innovation". In: *Economics Letters* 157, pp. 136–140. DOI: 10.1016/j.econlet.2017.06.014.

— (2018). "Horizontal Mergers and Product Innovation". In: *International Journal of Industrial Organization* 59, pp. 1–23. DOI: 10.1016/j.ijindorg.2018.03.001.

Haucap, Justus, Alexander Rasch, and Joel Stiebale (2019). "How Mergers Affect Innovation: Theory and Evidence". In: *International Journal of Industrial Organization* 63, pp. 283–325. DOI: 10.1016/j.ijindorg.2018.10.003.

Jaffe, Sonia and E. Glen Weyl (2013). "The First-Order Approach to Merger Analysis". In: *American Economic Journal: Microeconomics* 5.4, pp. 188–218. DOI: 10.1257/mic.5.4.188.

Jullien, Bruno and Yassine Lefouili (2018). "Horizontal Mergers and Innovation". In: *Journal of Competition Law & Economics* 14.3, pp. 364–392. DOI: 10.1093/joclec/nhy010.

Kokkoris, Ioannis and Tommaso Valletti (2020). "Innovation Considerations in Horizontal Merger Control". In: *Journal of Competition Law & Economics* 16.2, pp. 220–261. DOI: 10.1093/joclec/nhaa008.

Lefouili, Yassine and Leonardo Madio (2026). "Mergers and Investments: Where Do We Stand?" In: *International Journal of Industrial Organization* 105, p. 103269. DOI: 10.1016/j.ijindorg.2026.103269.

Miller, Nathan H., Marc Remer, Conor Ryan, and Gloria Sheu (2017). "Upward Pricing Pressure as a Predictor of Merger Price Effects". In: *International Journal of Industrial Organization* 52, pp. 216–247. DOI: 10.1016/j.ijindorg.2017.01.010.

Moraga-González, José-Luis and Evgenia Motchenkova (2026). *Mergers and R&D Investment: A Unified Approach*. CEPR Discussion Paper 20986. Paris and London: CEPR Press. URL: https://cepr.org/publications/dp20986.

Motta, Massimo and Emanuele Tarantino (2021). "The Effect of Horizontal Mergers, When Firms Compete in Prices and Investments". In: *International Journal of Industrial Organization* 78, p. 102774. DOI: 10.1016/j.ijindorg.2021.102774.

OECD (May 5, 2025). *Efficiencies in Merger Control*. OECD Roundtables on Competition Policy Papers 321. Paris: OECD Publishing. DOI: 10.1787/f4ce548f-en.

Pietola, Matias, Emanuele Tarantino, and Hans Zenger (May 22, 2026). *Mergers, Innovation Efficiencies, and the Investment Channel*. CEPR Discussion Paper 21529. Paris and London: CEPR Press. URL: https://cepr.org/publications/dp21529.

Régibeau, Pierre and Katharine E. Rockett (2019). "Mergers and Innovation". In: *The Antitrust Bulletin* 64.1, pp. 31–53. DOI: 10.1177/0003603X18822576.

Salinger, Michael A. (2019). "Net Innovation Pressure in Merger Analysis". Available at SSRN 3051249. DOI: 10.2139/ssrn.3051249.

Schumpeter, Joseph A. (1942). *Capitalism, Socialism and Democracy*. New York: Harper & Brothers.

U.S. Department of Justice and Federal Trade Commission (Dec. 18, 2023). *2023 Merger Guidelines*. Washington, DC: U.S. Department of Justice and Federal Trade Commission. URL: https://www.justice.gov/atr/merger-guidelines.

Werden, Gregory J. (1996). "A Robust Test for Consumer Welfare Enhancing Mergers among Sellers of Differentiated Products". In: *The Journal of Industrial Economics* 44.4, pp. 409–413.

Williamson, Oliver E. (1968). "Economies as an Antitrust Defense: The Welfare Tradeoffs". In: *American Economic Review* 58.1, pp. 18–36.

Willig, Robert D. (2011). "Unilateral Competitive Effects of Mergers: Upward Pricing Pressure, Product Quality, and Other Extensions". In: *Review of Industrial Organization* 39.1, pp. 19–38. DOI: 10.1007/s11151-011-9307-7.