Merger-enabled process innovation and UPP

A first-order theory-of-benefit screen for merger review

Purpose of this paper

This is a non-technical companion to Edwards-Warren, Espinosa and Piccolo (2026), which derives a first-order screen for evaluating a process innovation theory of benefit. It translates the methodology into a practical framework for competition lawyers, agency teams and economists. The assessment is sequential and contains three distinct tests. **Step 1** asks whether the proposed merged firm has an incentive to increase investment on a certain innovation project. **Step 2** asks whether that incentive is merger-specific by testing whether a different ownership allocation or a realistic less restrictive arrangement would provide similar incentives. **Step 3** asks whether the resulting product-level benefits of the project are transmitted into prices and are sufficient to offset the unilateral pricing concern measured by UPP. The principal inputs are pre-merger quantities, margins and diversion ratios, together with evidence on project costs, expected benefits, timing and risk.

Executive summary

Horizontal merger review has a disciplined framework for assessing unilateral pricing incentives. Margins and diversion ratios can be combined in an upward pricing pressure (UPP) calculation, providing a transparent first-order indication of whether the merger makes a price increase attractive.

Innovation and efficiency claims are often less disciplined. They may be expressed as broad claims about scale, data, complementary assets or integration, without an equally precise account of why the transaction changes the incentive to undertake a defined project, why that change is attributable to the proposed ownership structure, and how the resulting benefits reach the products and customers affected by the loss of competition.

This note proposes a bounded three-step screen for assessing a process innovation theory of benefit. It is not a merger simulation and it is not a complete welfare analysis. Its purpose is to organise the evidence around three separate causal questions:

1. **Incentive to innovate.** Does the proposed merged firm find it privately profitable to start or expand the project?
2. **Merger specificity.** Would the investment incentives also arise under a different realistic ownership allocation or a less restrictive arrangement?
3. **Transmission to prices.** At the merger-induced investment level, do the product-level benefits offset the unilateral pricing incentives identified by UPP?

The three tests are sequential but not interchangeable. Increasing investment on the project may be profitable under the merger yet not merger-specific. It may be profitable and merger-specific yet still fail to offset pricing pressure on one or more affected products. Conversely, a strong theory of benefit requires a coherent chain running from investment incentive, through counterfactual ownership, to product-level price transmission.

An important practical advantage of the local approach is that it can be implemented largely using observed pre-merger evidence. Rather than requiring the analyst to construct the merged firm's full post-merger investment problem, the screen evaluates marginal investment incentives at an observable benchmark, using familiar merger inputs such as pre-merger quantities, margins and diversion ratios together with evidence on incremental project costs and benefits.

Core legal takeaway

A credible innovation theory of benefit should do more than claim that the merged firm will be larger, better resourced or more capable. It should identify the project and show, first, that the proposed merged owner has an incentive to undertake it; second, that the best realistic alternative does not generate the incentives; and third, that the resulting benefits are timely, product-specific and sufficient to offset the pricing pressure identified on the harm side.

1. Introduction: why a theory-of-benefit screen is useful

Authorities appear increasingly open to a more symmetric assessment of competitive harms and efficiencies. In Europe, the draft Merger Guidelines[1] invite parties to articulate a theory of benefit. This requires more than an assertion that the merger will be efficient. It requires an explanation of the mechanism through which the benefit will be generated, evidence that the merged firm will have the incentive to realise it, and a credible account of

[1] European Commission, Draft Guidelines on the Assessment of Mergers under Council Regulation (EC) No 139/2004, published 30 April 2026 (hereinafter 'Draft Merger Guidelines').

how the benefit will reach affected customers. Similar evidentiary questions arise in other jurisdictions whenever parties rely on dynamic efficiencies to offset concerns about unilateral effects.

The practical problem is one of evidentiary symmetry. UPP asks a precise first-order question on the harm side: after the merger, will a firm find it profitable to raise the price of one product because some lost sales are recaptured by another product it now owns? Innovation claims should, where the evidence permits, be subjected to an analogous discipline. The agency or court needs to know whether a defined project will actually be undertaken, whether the investment is caused by the proposed ownership structure rather than by the counterfactual, and whether its benefits affect the same products, customers and time periods as the alleged harm.

Edwards-Warren, Espinosa and Piccolo (2026)[2] address this asymmetry by deriving a measure of the investment incentive created by the merger and then combining the merger-induced product benefits with UPP. The resulting benefit-adjusted UPP, or BAUPP, uses inputs that are familiar to merger practitioners: pre-merger quantities, margins and diversion ratios, together with estimates of unit cost savings or quality values, project costs, probability of success, timing and implementation risk.

The screen is most useful for a concrete project whose productivity, cost, risk or appropriability changes with ownership. Examples include integrating production systems, combining complementary datasets, redesigning a common platform, reallocating capacity across a single network, or deploying a technology across both merging product bases. These are precisely the settings in which a broad innovation narrative can be converted into a sequence of testable propositions.

2. The three-step assessment logic

Each step answers a different question. Step 1 concerns the private investment decision under the proposed merger. Step 2 is a counterfactual arrangement test. Step 3 concerns the incidence and transmission of the merger-induced efficiency. Although Steps 1 and 2 may use closely related calculations, they answer different causal questions: Step 1 asks whether the investment occurs under the merger; Step 2 asks whether the merger causes that investment relative to the relevant counterfactual.

Methodological step	What the analyst asks	Typical evidence needed	Failure mode
Step 1: Incentive to innovate	Would the proposed merged entity choose a positive project intensity?	Pre-merger quantities, expected unit benefits, success probabilities, project costs, internal hurdle rates and contemporaneous plans.	The project is described, but the merged entity's investment incentive is not established.
Step 2: Merger specificity	Would any realistic alternative ownership allocation, standalone strategy or less restrictive arrangement also generate the investment incentive?	Alternative ownership maps, licensing and joint-venture feasibility, data and IP rights, governance constraints, implementation costs and counterfactual business plans.	The project or the same project intensity may occur in the counterfactual.
Step 3: Transmission to prices	At the merger-induced investment level, do the product-level efficiencies offset the unilateral pricing incentive identified by UPP?	Margins, diversion ratios, benefit allocation by product and customer, timing, quality valuation.	A real and merger-specific project does not offset pricing pressure on the affected products or customers.

How to read the screen

Step 1 is an investment test: does the proposed owner have an incentives to undertake the project? **Step 2 is an attribution test:** is the investment incentive caused by the proposed ownership allocation rather than by a realistic counterfactual? **Step 3 is a price-transmission test:** do the merger-induced efficiencies offset the same competitive pressure identified on the harm side? Passing one step does not imply that the next step is satisfied. The fact that Steps 1 and 2 may look numerically similar is deliberate: Step 2 reruns the investment logic under counterfactual ownership in order to separate generation (incentives) from attribution (merger-specificity).

3. Step 1: does the proposed merged firm have an incentive to innovate?

Consider a proposed merger between two firms supplying differentiated products. The transaction makes available a defined process innovation project. The project may reduce marginal costs, improve reliability, increase quality or permit a technology to be deployed across both product bases. Let project intensity be denoted

[2] Kirsten Edwards-Warren, Francisco Espinosa and Salvatore Piccolo, 'Merger-Enabled Innovation and Upward Pricing Pressure: A Local Theory-of-Benefit Test' (August 2026, manuscript).

by *x*. A higher value of *x* may represent additional engineering effort, greater investment expenditure, a broader rollout or a more ambitious technical specification.

In the baseline cost-reduction formulation, the expected value of the project to the merged firm is the operating benefit generated across the products it owns, minus the implementation cost:

Equation 1: project value under the proposed merger

$$V_M(x) = \sum_i q_i^0 \, e_i(x) - K_M(x).$$

Here, q_i^0 is the pre-merger quantity of product i; $e_i(x)$ is the expected per-unit benefit on that product at project intensity x; and $K_M(x)$ is the expected project cost under the proposed merger. Benefits and costs should be expressed over a common horizon and on a comparable probability-adjusted and discounted basis.

For a project with adjustable investment intensity starting from no investment,[3] a useful first-order diagnostic is the local innovation incentive:

Equation 2: local innovation incentive

$$\mathrm{LII}_M = V_M'(0) = \sum_i q_i^0 \, e_i'(0) - K_M'(0).$$

A positive value means that, at the no-investment boundary, the marginal operating gain exceeds the marginal project cost. Under the usual concavity condition (which would need to be proven for each specific case), the merged firm then chooses a *positive* project intensity.

In practice, the local derivative will rarely be observed directly. The analyst can approximate it using a sufficiently small observed or planned increment in project intensity around the relevant benchmark (for example, an incremental engineering work package) and compare the associated incremental product benefits with the incremental project cost. The increment is an empirical finite-difference approximation to the local derivative.

The local score is deliberately modest. It uses contemporaneous pre-merger evidence to answer a threshold question: does the proposed owner have a private incentive to start or expand the project starting *at* the benchmark? However, it does not by itself predict the full post-merger equilibrium once prices, quantities and rival responses adjust. Where the project is discrete and requires a fixed launch cost, the corresponding test is whether expected net present value is positive. The local derivative and the fixed-cost NPV answer related but distinct questions: the former concerns the intensive margin, while the latter determines whether the project is launched at all.

A positive score at the benchmark shows that the merger changes investment incentives in the direction of greater investment. Where the Parties put forward a larger investment programme consistent with that direction of change, the local test provides evidence supporting that claimed investment response without requiring the analyst to solve for the merged firm's optimal investment level.

What practitioners should establish at Step 1

The project should be defined in operational terms. The parties should identify the products that receive the efficiency, estimate the efficiency as a function of project intensity, place implementation costs on the same time and risk-adjusted basis, and explain how the estimates relate to contemporaneous budgets, engineering plans and internal decision rules. Step 1 helps turn a generic assertion that scale improves innovation into a measure of investment incentives.

4. Step 2: is the investment incentive merger-specific?

A project that is profitable under the proposed merger is not automatically merger-specific. The relevant counterfactual is not simply "no project". It may involve a different allocation of ownership, standalone investment by one of the parties, licensing, a supply or data-access agreement, a network-sharing arrangement, a joint venture or another less restrictive structure. Step 2 asks whether any realistic alternative would generate the same investment incentive.

Let $\mathcal{A}$ denote the set of ownership allocations and arrangements that are sufficiently realistic to form part of the counterfactual. The proposed merger, denoted by M, is one element of this set. For each arrangement $a \in \mathcal{A}$, project value can be written as:

[3] The framework also easily accommodates the case of expanding investment above some benchmark value different than zero.

Equation 3: project value under an ownership allocation or alternative arrangement

$$V_a(x) = \sum_i \omega_i^a \, q_i^0 e_i^a(x) - K_a(x).$$

The weight $\omega_i^a \in [0,1]$ measures the share of the product-level benefit that the project sponsor can appropriate under arrangement a; $e_i^a(x)$ allows technical productivity to vary across ownership structures; and $K_a(x)$ captures arrangement-specific cost, timing and implementation risk. The proposed merger is included by setting $a = M$ and, for the benefits fully internalised by the merged firm, normalising $\omega_i^M = 1$.

Alternative ownership can weaken the investment incentive through three related mechanisms. Fragmented control may reduce the technical productivity of the project because assets, data or engineering teams cannot be integrated fully. It may raise implementation costs or delay the rollout because coordination and monitoring are more difficult. It may also limit appropriability because some benefits accrue to a separately owned product and cannot be captured through an incomplete contract. These mechanisms should be evidenced rather than assumed.

A strong local merger-specificity result requires the proposed merger to generate investment while every realistic alternative fails:[4]

Equation 4: strong local merger-specificity test

$$\mathrm{LII}_M > 0 \qquad \text{and} \qquad \mathrm{LII}_a \leq 0 \text{ for every } a \in \mathcal{A} \text{ with } a \neq M.$$

where $\mathrm{LII}_a = V_a{}'(0)$. Strong local merger-specificity requires a positive score under the merger and a non-positive score under every relevant counterfactual at the same benchmark. The further inference that a non-positive counterfactual local score rules out investment at higher intensities relies on the usual concavity condition. If that condition cannot be supported, the result could be supplemented with evidence at additional investment levels.
A weaker, but potentially relevant, claim arises when an alternative also generates investment but the proposed merger materially improves project intensity, timing, probability of success or the products over which the benefit can be deployed. In that case, only the incremental advantage over the best counterfactual is merger-specific.

What practitioners should establish at Step 2

The analysis should identify the best realistic alternative ownership allocation, not merely compare the merger with inaction. For each realistic counterfactual, it should explain who controls the necessary assets, which benefits that owner can capture, whether the project can be contracted over, and how cost, timing and risk differ. The merger-specific credit is the incremental project value relative to the best supported alternative.

5. Step 3: are the merger-induced benefits transmitted into prices?

Passing Steps 1 and 2 provides first-order evidence that the merger creates an investment incentive that is absent in the relevant counterfactual. Step 3 then takes the investment claim put forward by the Parties and asks whether the associated product-level efficiencies are sufficient to offset the post-merger pricing incentive.[5]

Standard UPP captures the lost-competition effect. When the merged firm raises the price of product i, some lost sales divert to product j, which is now under common ownership. If the project reduces marginal costs, the saving on product i lowers the incentive to raise its price. By contrast, a saving on product j makes diverted sales to that product more profitable and can strengthen upward pressure on product i.

For two single-product firms, the benefit-adjusted UPP on product i, evaluated at the merger-induced project intensity x_M, is:

Equation 5: Benefit-adjusted UPP (BAUPP)

[4] In the case of investment expansion above some benchmark level x^*, it must be that x^* is the optimal investment level under an alternative setting $a^* \in \mathcal{A}$, with $a^* \neq M$. Then, the question becomes whether the merger provides an incentive to increase investment above x^*. The conditions become $V_M'(x^*) > 0$ and $V_a'(x^*) \leq 0$ for every $a \in \mathcal{A}$ with $a \neq M$ and, in particular, $V_{a^*}'(x^*) = 0$. This case is less straightforward to implement than the no-investment benchmark when a^* cannot be constructed with observed pre-merger data. However, this case is important when the test fails to prove strong local merger-specificity at no investment (i.e., at $x = 0$). In this case, there is an alternative setting under which investment is also privately profitable for the merging parties. A further additional step would then consist in finding a new investment level, x'', at which $V_M'(x'') > 0$ and $V_a'(x'') \leq 0$ for every $a \in \mathcal{A}$ with $a \neq M$. In this way, the merging parties would prove that investment level $x_M' > x''$ passes Steps 1 and 2.

[5] The framework does not require the analyst to establish that x_M is the merged firm's *optimal* investment level. Steps 1 and 2 instead provide a first-order screen of whether the merger changes investment incentives in the direction of the claimed investment and whether that change is merger-specific. Step 3 evaluates the supported product-level benefits associated with the investment response put forward by the Parties.

$$\text{BAUPP}_i(x_M) = D_{ij}\left[m_j^0 + e_j(x_M)\right] - e_i(x_M).$$

Here, D_{ij} is diversion from product i to product j; m_j^0 is the pre-merger margin on product j; $e_i(x_M)$ is the own-product benefit; and $e_j(x_M)$ is the benefit on product j, the product receiving diverted sales. A negative value indicates downward first-order pricing pressure; a positive value indicates residual upward pressure.

Component	Contribution to BAUPP on product i	Interpretation
Lost competition	$+D_{ij}m_j^0$	Raising the price of product i becomes more attractive because some lost sales are recaptured on product j, which yields unit margins m_j^0 pre-merger.
Own-product benefit	$-e_i(x_M)$	A lower cost or higher customer value on product i reduces the incentive to raise its price.
Diverted-product benefit	$+D_{ij}e_j(x_M)$	If product j becomes more profitable due to lower costs, diverted sales to product j become more valuable, which increases upward pricing pressure on i.

Product-incidence principle

An aggregate synergy figure cannot be mechanically subtracted from an aggregate UPP figure. Product allocation is part of the transmission test. The innovation component dampens pricing pressure on product i when $e_i(x_M) > D_{ij}e_j(x_M)$. It fully offsets UPP only when $e_i(x_M) \geq D_{ij}\left[m_j^0 + e_j(x_M)\right]$. The same project can therefore reduce pressure on one product while leaving positive pressure on another.

In the symmetric two-product benchmark, with common margin m, diversion D and unit saving e, BAUPP is non-positive when $e \geq Dm/(1-D)$. This threshold makes clear that the required efficiency rises with diversion and with the margin on the diverted-to product.

Quality-improving innovation can be incorporated by translating reliability, functionality, speed or other quality improvements into a customer-value or willingness-to-pay equivalent. Where pass-through evidence is available, the BAUPP vector can then be translated into approximate price effects. Without such evidence, BAUPP should be reported as a first-order incentive measure rather than as a complete price prediction.

6. A running numerical example

The following example applies all three steps using one internally consistent set of assumptions. Suppose four single-product firms sell differentiated products. Firms 1 and 2, selling Product 1 and Product 2, respectively, propose to merge. The innovation project requires access to complementary engineering systems controlled by the two firms. Project intensity x is measured by the number of equivalent full-time engineering teams assigned to the project. Quantities are in millions of units, and all benefits and project costs are expressed in comparable € million-equivalent units over the assessment horizon.

Product	Pre-merger quantity	Margin (€ per unit)	Expected unit savings	Diversion to the other merging product
Product 1	10	30	$e_1(x) = 3x$	$D_{12} = 0.25$
Product 2	8	25	$e_2(x) = 2x$	$D_{21} = 0.20$

Assume that the project cost under the proposed merger is $K_M(x) = 16x + 5x^2$.

Step 1: incentive to innovate

The gross operating benefit under the merger is $10(3x) + 8(2x) = 46x$. The first-order project value is therefore

$$V_M(x) = 46x - (16x + 5x^2) = 30x - 5x^2.$$

The local innovation incentive is $\text{LII}_M = V_M'(0) = 30$, so the proposed merged passes Step 1. Within the workhorse approximation, $V_M'(x) = 30 - 10x$, which gives an optimal intensity of $x_M = 3$. At that intensity, the expected unit savings are €9 on Product 1 and €6 on Product 2. The gross operating benefit is €138 million-equivalent, project cost is €93 million-equivalent, and project value is €45 million-equivalent.

Step 2: merger specificity

Now consider the best alternative ownership allocation in the relevant counterfactual set. Because all alternatives leave the complementary engineering systems under separate control, suppose the best-supported engineering

estimate is that the alternative can achieve only 30 per cent of the technical benefit generated under common ownership, while facing the same project-cost schedule. This is a case-specific evidentiary assumption, not a general benchmark.

The value of the project under the best alternative is then

$$V_A(x) = 0.30(46x) - (16x + 5x^2) = -2.2x - 5x^2.$$

Its local innovation incentive is $\mathrm{LII}_A = -2.2$. Because $V_A(x)$ is decreasing for every $x \geq 0$, the best alternative chooses $x_A = 0$. The different ownership allocation does not generate the investment, so the example passes Step 2 in the strong local sense.

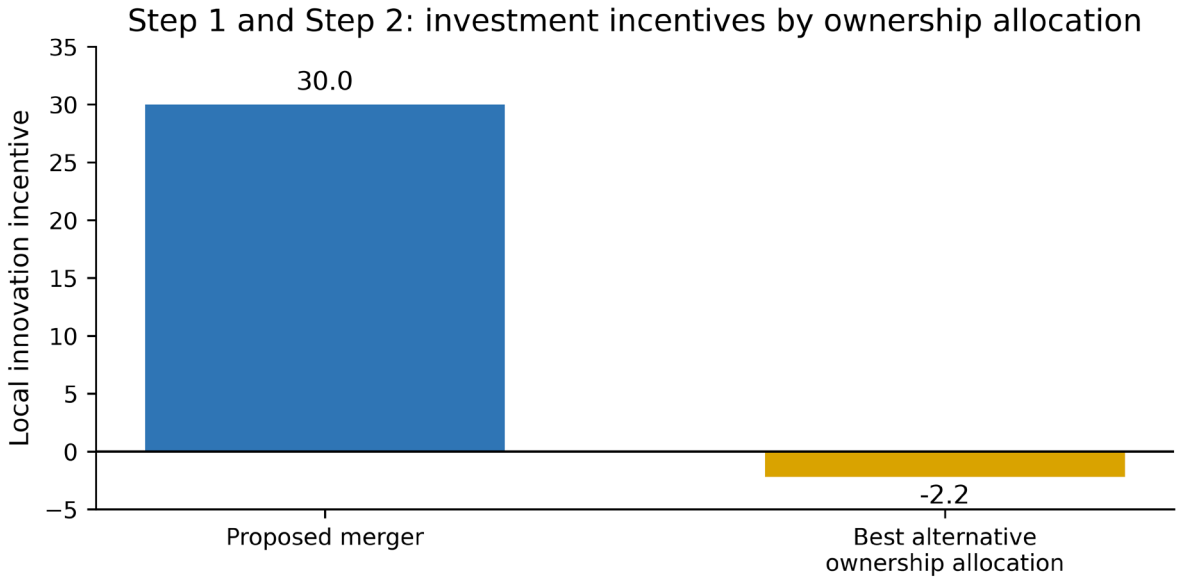


Figure 1. The merger induces investment; the best alternative ownership allocation does not.

Step 3: transmission to prices

At the merger-induced intensity $x_M = 3$,[6] the relevant unit savings are €9 on Product 1 and €6 on Product 2. Applying Equation 5 gives:

Product repriced	Lost competition	Own-product benefit	Diverted-product benefit	BAUPP
Product 1	$0.25 \times 25 = 6.25$	-9.00	$0.25 \times 6 = 1.50$	-1.25
Product 2	$0.20 \times 30 = 6.00$	-6.00	$0.20 \times 9 = 1.80$	$+1.80$

[6] As noted in footnote 5, x_M need not be, or be shown to be, optimal.

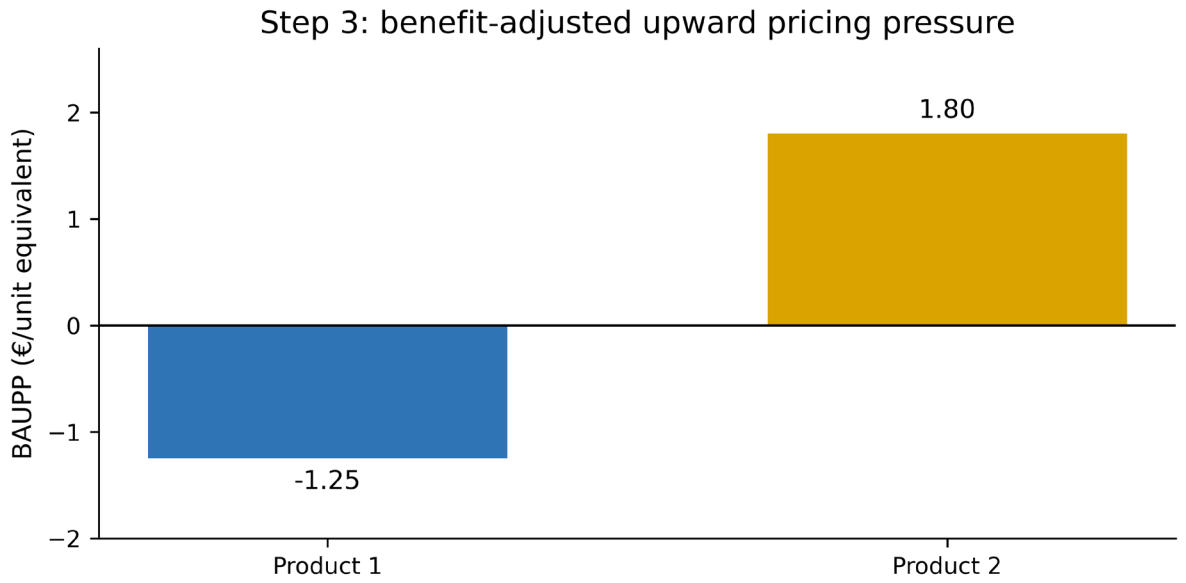


Figure 2. BAUPP is negative for Product 1 and positive for Product 2.

The proposed merger creates the investment and the investment is merger-specific, but the project does not fully offset pricing pressure on both products. Product 1 receives the larger own-product saving and has negative BAUPP. Product 2 receives the smaller own saving, while diversion to the now more profitable Product 1 remains valuable, leaving positive BAUPP. A pass-through analysis could translate these two incentive measures into approximate price effects, but it would not change the logical conclusion of the three-step screen.

7. Extensions for legal and agency practice

The baseline screen is intentionally simple, but the three-step logic extends naturally to common features of innovation claims.

Issue	How the analysis changes	Primary step affected
Discrete projects and fixed launch costs	Replace the local derivative with an expected NPV test for launch; where project scale is also adjustable, report both launch and intensity decisions.	Step 1
Uncertainty of success	Probability-weight benefits and costs consistently, and report outcome states where failure changes demand, capacity or strategic incentives.	Steps 1 and 3
Quality-improving innovation	Translate quality, reliability, functionality or speed into customer-value equivalents before comparing them with pricing pressure.	Step 3
Delayed benefits	Discount benefits and costs over a common horizon, report ramp-up and delay risks, and keep interim pricing pressure visible.	Steps 1 and 3
Multiproduct firms	Compare realistic portfolio ownership allocations at Step 2 and use product- or segment-level diversion and benefit vectors at Step 3.	Steps 2 and 3
Quantity feedback and pass-through	Where price changes materially alter the sales base, update project value and translate BAUPP into approximate prices using supported pass-through evidence.	Step 3
Licensing, joint ventures and contractual alternatives	Model how the alternative changes technical productivity, implementation cost, timing, risk and appropriability rather than treating “less restrictive” as a label.	Step 2

Introducing timing is important. A project may be privately profitable because it creates future benefits while customers experience upward pricing pressure immediately after closing. A submission should therefore identify the expected launch date, ramp-up profile and probability of delay, and should ask whether the customers exposed to interim harm are the same customers expected to receive the later benefit.

For quality innovation, nominal price alone is an incomplete outcome measure. A higher nominal price can coexist with customer benefit if the quality-adjusted value rises by more than the price. The quality benefit nevertheless

requires the same incidence analysis: which customers value the improvement, on which product, and at what point in time?

8. Practical implementation workflow

A case team can implement the methodology as three linked workstreams corresponding exactly to the three analytical steps.

Workstream	Core tasks	Required output
1. Incentive to innovate	Define the project, estimate product-level benefit functions and project costs, incorporate timing and success probabilities, and calculate the merger LII or project NPV. Approximate derivatives with small discrete changes in project intensity around the benchmark.	Project definition, local investment score (or launch NPV).
2. Merger specificity	Identify realistic alternative arrangements; estimate their productivity, cost, timing, risk and appropriability; rerun the investment test under the relevant counterfactuals and compare the results with the merger.	Counterfactual set, best-alternative local score or NPV, conclusion on whether the merger-induced investment incentive is locally merger-specific.
3. Transmission to prices	Allocate the merger-specific benefit by product, customer and time; calculate BAUPP.	Product-level benefit vector, BAUPP by product, price-effect ranges where available, and break-even thresholds.

Suggested reporting dashboard
Report the merger investment score, the best alternative ownership score, the predicted project intensity or launch NPV, the merger-specific benefit vector, BAUPP by affected product, timing and customer incidence, and critical values for the most uncertain inputs. Where supported input ranges straddle a relevant zero or break-even threshold, label that result inconclusive and report the critical value rather than assigning a pass/fail conclusion.

9. Evidence map and red flags

Step	Key inputs and corroborating evidence	Common red flag
1. Incentive to innovate	Product-level quantities; engineering or customer-value estimates; budgets, staffing and vendor quotes; stage-gate probabilities; contemporaneous plans and internal hurdle rates.	The benefit is described only as scale, data or improved capability, with no defined project technology, cost or investment decision.
2. Merger specificity	Alternative ownership maps; standalone plans; licensing or JV evidence; data and IP rights; prior negotiations; governance, monitoring, hold-up and leakage constraints; comparative timing and cost.	The analysis compares the merger only with inaction or dismisses alternatives without modelling the best realistic counterfactual.
3. Transmission to prices	Margins and diversion; product- and customer-level benefit allocation; rollout dates; quality valuation; output and customer-switching evidence.	A single aggregate synergy is subtracted from UPP, benefits arrive after material interim harm, or the claimed benefit accrues to different products or customers from those affected by the concern.

10. Decision table for a theory of benefit

The three-step structure produces a diagnostic matrix. It is not a legal decision rule; it prevents one evidentiary claim from silently substituting for another.

Finding	Step 1: incentive	Step 2: specificity	Step 3: transmission	Implication
No investment generation	Fails	Not reached	Not reached	No innovation credit should be counted within the screen.

Finding	Step 1: incentive	Step 2: specificity	Step 3: transmission	Implication
Investment generated but not merger-specific	Passes	Fails because an alternative also invests[7]	Not determinative	The investment incentive may be real, but it should not be entirely credited to the merger; only any incremental difference may qualify.
Investment generated and merger-specific, but weak transmission	Passes	Passes	Positive or mixed BAUPP	The investment incentive is genuine and attributable to the merger, but the benefit does not fully offset pricing pressure on all affected products or customers.
Strong theory of benefit	Passes	Passes	Non-positive BAUPP on the relevant products, with supported timing and incidence	The evidence forms a coherent chain from ownership to investment and from investment to customer-facing price incentives.

11. Scope and limitations

The framework is a screen. It is designed to structure evidence, identify weak links in an innovation theory of benefit and provide a first-order comparison between merger-induced benefits and unilateral pricing incentives. It does not replace a full merger simulation, a complete welfare analysis or a detailed assessment of coordinated effects, entry, repositioning, capacity constraints, bargaining or dynamic rival responses. In particular, a local score is a first-order incentive measure at a specified benchmark. Like UPP on the harm side, it is not intended to solve for the firm's post-merger optimum. Its role is to identify whether the transaction changes incentives in the direction relevant to the theory of benefit.

The value of the methodology is its separation of the evidentiary tasks. Parties must show why the proposed transaction would incentivise investment, why a different ownership allocation does not generate the same investment incentive, and how the incremental project output affects the customers and products implicated by the theory of harm. Agencies can then distinguish merger-enabled innovation from benefits that are speculative, achievable in the counterfactual, delayed beyond the relevant horizon or poorly aligned with the competitive concern.

12. Conclusion

A theory of benefit should be more than a general claim that the merged firm will be larger, more capable or better placed to innovate. It should identify a project, estimate its expected benefits and costs, show that the proposed ownership structure changes the investment decision relative to realistic alternatives, and test whether the merger-specific benefits offset the same pricing incentives that create concern.

The proposed three-step screen provides a practical way to do this. Step 1 disciplines the claim that the merger generates innovation. Step 2 turns merger specificity into an explicit comparison across ownership allocations and less restrictive arrangements. Step 3 links the resulting product-level benefits to UPP and, where evidence permits, to prices. The practical appeal of the test is that the investment side can be disciplined using the same type of observed, pre-merger evidence that makes first-order pricing screens useful on the harm side. For lawyers, agency teams and economists, the central contribution is evidentiary: the framework converts an innovation claim into a coherent sequence of propositions that can be supported, challenged and stress-tested in merger review.

[7] See footnote 4.