Merger-enabled innovation: making the EU theory of benefit operational

Kirsten Edwards-Warren, Francisco Espinosa and Salvatore Piccolo

Econic Partners LLC; University of Bergamo

The views expressed are those of the authors. Any errors are their own.

1. Introduction

Horizontal merger review has become increasingly precise about measuring competitive harm and increasingly demanding about the assessment of claimed dynamic benefits. The mismatch is not merely academic. The European Commission's draft Merger Guidelines, published on 30 April 2026, are the broadest overhaul of the EU framework in more than two decades. They expressly require merging parties that rely on efficiencies to articulate a 'theory of benefit' and devote substantial attention to non-price dimensions of competition such as investment and innovation.[1]

This is a groundbreaking advance. The draft recognises that mergers can combine complementary capabilities, spread an indivisible project over a larger customer base, relieve hold-up or financial constraints, and make investments possible that would not otherwise occur. Building on recent academic literature, it also insists that such claims be concrete, merger-specific, timely, verifiable and capable of benefiting the consumers exposed to the merger's harm. Yet, one practical question remains underdeveloped: how should an authority move from those general requirements to an operational first assessment of a particular innovation claim?

The objective of this article is to help answer this question. We start from a simple asymmetry. On the harm side, margins, diversion ratios and pricing-pressure tools often provide a disciplined route from pre-merger evidence to an assessment of local unilateral incentives.[2] On the benefit side, the parties may invoke more scale, more data, complementary engineering teams or faster deployment without showing which project will be undertaken, why the merger is needed, where the resulting benefit will arise or how it will reach consumers. The problem is not that those mechanisms are implausible. It is that several distinct propositions are being compressed into one, and often poorly articulated, narrative.

The central results of this article help minimise both types of enforcement errors in assessing efficiency claims. A merger may be genuinely necessary to generate an innovation and the resulting benefit may nevertheless be too small, too late or located on the wrong products to offset the loss of competition. Conversely, an innovation claim need not be rejected merely because a complete dynamic merger simulation is unavailable. A bounded local analysis can identify whether the merger changes the investment decision and whether the expected product-level outcome points against, or reinforces, the immediate pricing incentive.

This article provides a non-technical account of the methodology developed in Edwards-Warren, Espinosa and Piccolo (2026)[3] and relates it to the draft Merger Guidelines. We ask what a credible innovation theory of benefit should establish at the first-screening stage, which analytical steps are needed to meet the required standard of proof, and what the Commission could clarify as it finalises the Merger Guidelines. In that respect, the article follows a broader recent effort to connect new economic results to specific operational gaps in the draft framework.[4]

2. What the draft Merger Guidelines say about innovation benefits

The draft Merger Guidelines already contain most of the right ingredients. Paragraph 25 defines a theory of benefit as an explanation of how specific merger efficiencies arise and maintain or enhance effective competition to the benefit of consumers. The framework then distinguishes direct efficiencies, such as cost savings or quality improvements, from dynamic efficiencies that operate through investment and innovation. Both categories must be verified, attributed to the merger and shown to benefit consumers on a lasting basis.[5]

For direct efficiencies, the draft asks whether cost savings affect the costs that matter for pricing and whether they are likely to be passed on. It also recognises compensating marginal-cost analysis as a way of comparing cost efficiencies with price harm. For quality, the Commission proposes to appraise consumer valuation, where possible through willingness-to-pay evidence, and to deduct any part of the gain that the merged firm claws back through higher prices.[6]

For dynamic efficiencies, the draft is even more explicit. The parties should identify the investment or innovation as concretely as possible; explain how the merger changes both their ability and their incentive to pursue it; account for timing and uncertainty; and support the claim with ordinary-course documents, business plans, historical performance and other contemporaneous evidence. Merger specificity is comparative: standalone investment, licensing, joint ventures and other realistic arrangements must be considered rather than dismissed in the abstract.[7]

The remaining gap is not one of legal principle but of analytical nature. Profitability under the merger and merger specificity are related, but they are not the same question. Nor does the draft explain how an authority should use local evidence when exact quantification of a new technology or a long-run quality improvement is not yet possible. Without a clearer sequence, the theory of benefit risks becoming either too permissive - because a broad innovation narrative is accepted without a causal chain - or too demanding - because only a complete dynamic model is regarded as informative.

3. The economic literature and the missing bridge

Three strands of economic literature have already touched on the problem, but they answer different questions. The first is the classic literature on merger efficiencies and unilateral effects. Williamson (1968) framed the trade-off between market power and production efficiencies. Werden (1996), Farrell and Shapiro (2010), and Jaffe and Weyl (2013) later developed compensating-efficiency, upward-pricing-pressure and first-order methods that can be implemented

[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'). The public consultation closed on 26 June 2026; the Commission expects to finalise the review in Q4 2026.

[2] Throughout this article, "local" refers to a first-order analysis of a variable, such as price or investment intensity, around a specified benchmark value, typically the observed pre-merger position or, where appropriate, the relevant absent-merger counterfactual. A local analysis assesses the direction and immediate magnitude of the incentives associated with a small change in that variable departing from its benchmark value, while holding other relevant factors fixed at their benchmark values as well. It therefore does not, by itself, incorporate the full adjustment of quantities, capacities and other dimensions of competition, or the responses of rivals and other market participants. Importantly, "local" describes the analytical method, not the scale of the project or its expected effects: the project being assessed may be implemented at full scale, and its effects need not themselves be small. In academic works "local" is sometimes referred to as "first-order" as the analysis is based on first-order conditions.

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

[4] See, for example, Volker Nocke, Martin Peitz and Nicolas Schutz, 'Merger control amid market evolutions and shocks: what the EU Merger Guidelines should say' (2026) Journal of European Competition Law & Practice, published 22 July 2026, for a parallel effort to connect new theoretical results to operational gaps in the draft Merger Guidelines.

[5] Draft Merger Guidelines, paras 25, 305-320 and 324-338. Paras 327-330 address verifiability, the distinction between ability and incentive, timing, magnitude and ordinary-course evidence; paras 331-333 address standalone and contractual alternatives; paras 334-338 address benefits to consumers.

[6] Draft Merger Guidelines, paras 317–323. Para 317 addresses the pass-on of variable or marginal-cost savings and the use of compensating marginal-cost reduction; paras 319–323 address the valuation and pass-on of quality efficiencies, including willingness-to-pay evidence and the deduction of benefits clawed back through higher prices.

[7] Draft Merger Guidelines, paras 327–334. Paras 327–330 address verifiability, ability and incentive, timing, magnitude and contemporaneous evidence; paras 331–333 address standalone investment and less anticompetitive arrangements, including licensing and joint ventures; para 334 addresses uncertainty and expected consumer benefits.

with relatively limited information. Their attractive is "informational" economy: they identify the immediate change in incentives without requiring a complete model of the post-merger equilibrium.[8]

Those tools normally take the efficiency credit as given. They can ask how a stated reduction in marginal cost changes pricing pressure, but not whether the merger will induce the project that generates that reduction. This omission is innocuous when the efficiency is technologically certain and directly caused by the change in ownership. It is central when the claimed benefit depends on the success of an uncertain R&D project, the existence of a common platform or the implementation of a reorganisation that might also be achieved through another arrangement.

The second strand studies how mergers affect investment and innovation incentives. The modern literature shows why no general presumption is available/justifiable. A merger can weaken innovation by internalising business stealing or reducing output, but it can also improve appropriability, spread fixed costs, combine capabilities or relieve contracting frictions. The sign depends on the form of innovation, spillovers, cannibalisation, the strategic environment and the presence of genuine synergies.[9]

The third strand is the policy and evidentiary literature. It asks whether claimed efficiencies are verifiable, merger-specific and likely to reach consumers. This is where the draft Merger Guidelines make their largest contribution. Yet evidentiary criteria alone do not produce an economic test. An authority still needs to know which project to evaluate, what counterfactual to use, how to treat risk and timing, how to allocate the benefit across products, and how to compare it with the mechanism of harm.[10]

The local theory-of-benefit approach developed in Edwards-Warren, Espinosa and Piccolo (2026) is intended to provide that formal bridge. It does not seek a universal relationship between concentration and innovation. It asks narrower, case-specific questions: does this transaction move a defined project over the investment threshold relative to realistic alternatives, and does the expected output of that project counter the product-level pricing incentives through which the merger harms consumers?

4. The economics: from project generation to consumer benefit

The methodology we propose for operationalising an innovation theory of benefit involves three sequential steps. Step 1 concerns the incentive to innovate: would the proposed merged firm find it privately profitable to undertake, expand or accelerate the defined project? Step 2 concerns merger specificity: would the merging parties also have an incentive to undertake, expand or accelerate the same project under the relevant absent-merger counterfactual, including under a different realistic ownership allocation or a less restrictive arrangement? Step 3 concerns competitive transmission and consumer benefit: what incremental effects does the merger-specific project produce, where and when do they arise, and do they offset the competitive incentive created by the merger? For process innovation, benefit-adjusted upward pricing pressure provides a first-order screen at Step 3. Pass-through and other market responses may require further analysis before reaching a complete conclusion on consumer outcomes.

4.1 Step 1: Incentive to innovate

The first question is whether the project will be undertaken. For a cost-reducing project, the private gain comes from the expected saving on the sales to which the new process will apply, adjusted for technical success, delay, rollout and useful life. For a demand- or quality-enhancing project, the relevant private gain is the profit on net additional or preserved sales, together with any cost effects. In both cases, the gain must be compared with development and implementation expenditure.

The form of the hurdle depends on the project. A project with adjustable intensity can be screened by asking whether a small expansion of effort is profitable at the observed pre-merger position. An indivisible project with a fixed launch cost requires an expected net-present-value comparison. The local incentives test we propose is not meant to replace the firm's business case. It translates that business case into a transparent statement about the margin on which the merger is claimed to change behaviour.

Ability and incentive should also remain separate. The merger may make an engineering solution technically feasible because the parties control complementary assets, thus providing the parties with the ability to implement it. Alternatively, the project may already be feasible, while fragmented ownership, spillovers or hold-up prevent the investor from capturing enough of the return. A persuasive ability claim identifies the constraint that ownership removes rather than treating scale as a self-executing explanation. The identification of these constraints will then help defining the right formal definition of the problem, and the correct form of the screening test(s) for incentives.

An important implication is that common ownership does not, by itself, create a direct incentive to increase investment. If a project only reduces the cost of the investing firm's own product, and the firm already captures that benefit before the transaction, the local investment incentive is unchanged. A direct merger-induced incentive therefore requires something additional, such as a benefit on the other party's products, a common capability that can be deployed across a broader sales base, improved appropriability or a genuine change in feasibility.

4.2 Step 2: Merger specificity

Proof of incentives is not yet proof of merger specificity. The relevant counterfactual is not necessarily "no project". The same project assessment must be repeated for the best realistic alternatives: standalone investment, a licence, a research joint venture, a supply or data-sharing agreement, a minority investment or a narrower asset acquisition.

These alternatives may differ in technical productivity, cost, delay, governance and the fraction of the return the decision-maker can appropriate. The relevant question is whether the same or a materially equivalent project outcome would arise under the best realistic alternative—not whether an alternative is technically conceivable in the abstract.

A strong merger-specificity result arises where the project is undertaken under the merger but not under any realistic alternative. Where an alternative would also generate some investment, only the incremental improvement attributable to the merger (such as earlier implementation, greater intensity, a higher probability of success or deployment across a broader product base) should be treated as merger-specific.

4.3 Step 3: Transmission and consumer benefit

Passing Steps 1 and 2 provides evidence that the merger generates an investment incentive that would not arise, or would arise only to a materially lesser extent, in the relevant counterfactual. The next task is to translate that incremental difference into the expected effects relevant to the competitive assessment. These effects should be mapped by product and period: the expected unit-cost reduction,

[8] See Oliver E. Williamson, 'Economies as an Antitrust Defense: The Welfare Tradeoffs' (1968) 58 American Economic Review 18; Gregory J. Werden, 'A Robust Test for Consumer Welfare Enhancing Mergers among Sellers of Differentiated Products' (1996) 44 Journal of Industrial Economics 409; Joseph Farrell and Carl Shapiro, 'Antitrust Evaluation of Horizontal Mergers: An Economic Alternative to Market Definition' (2010) 10 B.E. Journal of Theoretical Economics, Article 9; Sonia Jaffe and E. Glen Weyl, 'The First-Order Approach to Merger Analysis' (2013) 5 American Economic Journal: Microeconomics 188; Nathan H. Miller, Marc Remer, Conor Ryan and Gloria Sheu, 'Upward Pricing Pressure as a Predictor of Merger Price Effects' (2017) 52 International Journal of Industrial Organization 216.

[9] See Giulio Federico, Gregor Langus and Tommaso Valletti, 'A Simple Model of Mergers and Innovation' (2017) 157 Economics Letters 136, and 'Horizontal Mergers and Product Innovation' (2018) 59 International Journal of Industrial Organization 1; Massimo Motta and Emanuele Tarantino, 'The Effect of Horizontal Mergers, When Firms Compete in Prices and Investments' (2021) 78 International Journal of Industrial Organization 102774; Marc Bourreau, Bruno Jullien and Yassine Lefouili, 'Horizontal Mergers and Incremental Innovation' (2026) 57 RAND Journal of Economics 122; Michael A. Salinger, 'Net Innovation Pressure in Merger Analysis' (2019) SSRN Working Paper 3051249; Matias Pietola, Emanuele Tarantino and Hans Zenger, 'Mergers, Innovation Efficiencies, and the Investment Channel' (2026) CEPR Discussion Paper 21529.

[10] See US Department of Justice and Federal Trade Commission, Merger Guidelines (2023); OECD, Efficiencies in Merger Control, OECD Roundtables on Competition Policy Papers No 321 (2025); Ioannis Kokkoris and Tommaso Valletti, 'Innovation Considerations in Horizontal Merger Control' (2020) 16 Journal of Competition Law & Economics 220.

quality improvement or demand effect; the probability of success; the implementation date; the rollout path; and the duration of the gain. A single, portfolio-wide synergy percentage may be informative only if the evidence supports a genuinely uniform effect across products.

This product- and time-specific mapping is necessary because the subsequent pricing-pressure analysis does not compare aggregate project value with aggregate competitive harm. It asks how the expected effect on each product changes the pricing incentive for that product and for other products receiving diverted demand.

The final question is whether the expected product-level efficiency counters the pricing incentive created by the change in ownership. For process innovation (that is, innovation that can be eventually mapped into a reduction in costs), the natural first screen embeds the expected cost saving in the same local pricing-pressure logic used on the harm side. We derive a benefit-adjusted UPP formula (BAUPP) for this purpose. For quality innovation, the relevant comparison is between the merger's changed pricing incentive and demand's monetary valuation of the quality improvement, net of any increase in the cost of delivering it.

Evidence of benefit transmission is not a forecast of the final price path. Pass-through, rival responses, capacity, repositioning, entry and coordinated effects may all matter. The local inquiry asks a prior and narrower question: at the point where common ownership changes the merged firm's pricing incentives, does the claimed innovation operate in the direction and on the scale needed to offset that change?

5. Why product-level allocation of efficiencies changes the result

The most important implication concerns the location of the benefit within the merged entity's product portfolio. Consider two substitutable products. Before the transaction, the seller of the first product does not take account of the profit earned on the second when a customer switches between them. After the merger, a price increase on the first product becomes less costly because some lost sales are recaptured by the second. This is the familiar source of unilateral upward pricing pressure.

A cost reduction on the product whose price is being assessed works against that pressure. It makes an additional sale of that product more profitable and therefore strengthens the incentive to keep its price low. A cost reduction on the second product, however, can work differently. It makes the sales diverted to that product more valuable and can therefore strengthen the incentive to raise the first product's price. The full local profit gradient contains both effects.[11]

This is not a claim that an efficiency on the second product is socially harmful. The project may create substantial value. The point is that its contribution to a particular unilateral pricing incentive depends on where it arises. The own-product benefit and the diverted-product benefit enter the local decision through different channels.

Three conclusions follow. First, project generation does not imply benefit transmission. A project can be genuinely merger-specific and still fail to neutralise the loss of competition. Second, total portfolio savings are not a sufficient statistic. The same aggregate saving can produce very different customer outcomes depending on its allocation across the merged entity's products. Third, the same project can reduce pricing pressure on one product while increasing pressure on another. A complete analysis should therefore match each synergy with every relevant product according to the portfolio's diversion map, rather than netting, for each product, one headline synergy number against one headline measure of harm.

This also clarifies the role of simplified pricing-pressure tools. A simplified calculation that credits only the cost saving on the repriced product may be useful for transparency. But when a common project changes costs across several products, omitting the effect on destination products can make the merger appear more favourable than the full local incentive warrants. The choice between the simplified and full calculation should be explicit, not accidental.

6. Quality-enhancing innovation

Many merger-enabled projects do not primarily reduce marginal cost. They improve reliability, coverage, matching, safety, interoperability, speed, variety, functionality or customer support. The same causal sequence applies, but private project value and consumer benefit must be kept particularly distinct.

For the firm, a quality improvement is valuable because it creates or preserves profitable sales, changes costs or reduces churn. The private return is not the full increase in consumers' willingness to pay. It is the part the firm can appropriate through margins on net additional sales, after accounting for cannibalisation within the portfolio and any increase in the cost of serving customers. Using the full willingness-to-pay gain as private project revenue would overstate the incentive to invest.

A comparison of nominal prices alone to assess consumer value can be misleading. A better product may command a higher price while still offering more value. The relevant question is whether the additional value received exceeds the associated price increase, so that the quality-adjusted price falls or rises by less than it otherwise would. The draft Merger Guidelines recognise this by requiring consumer valuation to be appraised, often through willingness-to-pay evidence, and by deducting any part of the gain clawed back through higher prices.[12]

The allocation of the effects remains central. A quality improvement on the product whose price is being assessed directly benefits its customers. An improvement on another product in the merged portfolio can also make diverted sales more attractive to the merged firm. The analysis should therefore allocate quality benefits product by product rather than treating an aggregate or portfolio-wide quality (e.g., "willingness-to-pay") increase as an offset to aggregate harm.

The local shortcut is most reliable when quality can be represented as a credible product-specific monetary gain and the project does not fundamentally alter substitution. If the innovation creates a new product, changes product rankings or materially changes demand sensitivities, the pre-merger diversion evidence must be updated and a richer demand model may be required. A first screen should reveal its limitations and identify when that additional modelling is necessary.

7. Why local methods are often necessary for first screening

Merger review is prospective, evidence-intensive and conducted under tight procedural deadlines. Innovation claims often arise before the project has been completed, when technical success, rollout and customer valuation remain uncertain. Requiring a full dynamic equilibrium model at the outset can be disproportionate and, in some cases, impossible. Accepting a qualitative narrative without a disciplined comparison creates the opposite error.

Local methods fill the space between those extremes. They use the observed pre-merger position and a limited set of inputs that are already familiar in merger analysis: quantities, margins, diversion, project-level savings or quality effects, implementation costs, success probabilities and timing. They can be implemented with ranges and sensitivity analysis when the evidence does not support a single point estimate.

Their fundamental value is diagnostic. A first screen can show that the alleged project is not defined; that a realistic alternative appears capable of producing it; that the expected benefit is too uncertain or delayed to matter over the relevant horizon; or that the benefit is concentrated on products that do not face the identified harm. It

[11] Farrell and Shapiro (2010) note that the full local profit gradient includes the effect of an efficiency on the product receiving diverted demand, although simplified UPP implementations often omit it for transparency. Edwards-Warren, Espinosa and Piccolo (2026) develop this term as the product-location component of benefit-adjusted UPP.

[12] Draft Merger Guidelines, paras 319-323 and 336-338. See also Robert D. Willig, 'Unilateral Competitive Effects of Mergers: Upward Pricing Pressure, Product Quality, and Other Extensions' (2011) 39 Review of Industrial Organization 19.

can also identify a credible claim that deserves a more complete pass-through analysis, demand estimation or merger simulation.

The local nature of the test is therefore both its strength and its limitation. It is a strength because it makes the causal assumptions visible and focuses investigative resources on the inputs that can change the conclusion. It is a limitation because the test does not incorporate all post-merger feedback. If higher prices shrink the sales base on which savings are earned, the investment incentive may be weaker than the initial screen suggests. If the project changes demand, capacity or rival strategy, the local calculation may need to be updated. That said, a local analysis may more directly reflect the investment incentives underlying the claims put forward by the merging parties.

The draft Merger Guidelines already provide a natural home for this approach. They require quantification where reasonably possible, accept a reasoned assessment of magnitude when exact quantification is unavailable, and expressly recognise pricing-pressure and compensating-efficiency tools in the balancing exercise. What is missing is a clear statement that local techniques can serve as an initial theory-of-benefit screen, not merely as the final stage of a fully quantified case.[13]

That clarification would improve symmetry between harm and benefit. A local theory of harm is not treated as invalid merely because it is not a complete price forecast. The same should be true of a carefully bounded local theory of benefit. The important condition is that the screen be used for what it can establish – the incentive to innovate, merger specificity and immediate transmission - and not be presented as a complete welfare conclusion.

8. Where the draft Merger Guidelines could be strengthened

The draft is directionally right, but its operational sequence could be made more explicit. Three additions would materially improve the assessment of merger-enabled innovation without turning the Merger Guidelines into a technical manual.

A. Require a project-level counterfactual

A theory of benefit should identify a defined project and the mechanism through which the merger changes it. Claims based on scale, data or complementary capabilities should explain whether the transaction changes technical feasibility, project cost, probability of success, timing or appropriability. The best realistic alternative should be analysed against the same friction. Contracting should not be assumed costless, but it should not be declared implausible without compelling explanation either.

The finding should be stated narrowly: the project is expected to proceed under the merger, while the best realistic alternative would not produce the same outcome or would do so materially later, at smaller scale or with a lower probability of success. This would give practical content to the draft's distinction between ability and incentive and prevent a generic claim of greater capability from substituting for merger specificity.

B. Match the location of benefit and harm

In the case of process innovation, a local price concern should be compared with the expected unit-cost effect on the affected product, while accounting for benefits on destination products that alter the value of diversion.

For quality, the Merger Guidelines should emphasise that two separate calculations are required. Private project generation depends on appropriable profit, while consumer benefit depends on the value received net of any price clawback and additional delivery cost. Keeping these objects separate avoids both overstating the investment case and understating the consumer benefit.

C. Recognise the role of a staged local screen

The Merger Guidelines should state that a first-order, local assessment can be an appropriate first screen when the objective is to determine whether a claim warrants deeper investigation. The screen should report three findings separately: whether the project is expected to be undertaken; whether it depends on the merger relative to realistic alternatives; and whether its expected product-level output offsets the immediate unilateral incentive in the channel analysed.

A favourable result would not create a safe harbour. Pass-through, market adjustment, implementation risk and other theories of harm may still require analysis. An unfavourable result would not imply that the project has no social value. It would mean that the evidence does not establish a merger-specific benefit that compensates the identified harm. This staged approach would make the evidentiary burden more transparent while preserving the Commission's case-by-case discretion.

9. Conclusion

The draft Merger Guidelines represent a major advance in the treatment of non-price competition, such as the one occurring through investment and innovation. Their theory-of-benefit framework correctly asks whether efficiencies are caused by the merger, supported by evidence and capable of reaching consumers. What the framework still lacks is an operational bridge from those principles to the first assessment of a concrete innovation project.

The economic analysis provided in Edwards-Warren, Espinosa and Piccolo (2026) provides a first attempt to supply that bridge. A credible claim should be assessed in three steps. First, the project must be privately worthwhile under the proposed merger. Second, the improved project outcome must be merger-specific relative to the relevant absent-merger counterfactual and realistic alternative arrangements. Third, the project's incremental effects must be mapped across products and periods and tested against the same customer-facing incentive through which the merger creates harm.

The product-location result is the central discipline. A benefit on the repriced product can offset pressure; a benefit on a product receiving diverted demand can increase the value of that diversion. The strongest theory of benefit is therefore not the one with the largest aggregate synergy estimate. It is the one in which the project counterfactual, the allocation of benefits and the consumer-facing mechanism tell the same story.

Clearly, local methods cannot answer every question, but that is not the standard they should be judged against. When the objective is first screening, they are often necessary precisely because they reveal what has been established, what remains speculative and where a fuller model would add value. Making that role explicit would help the final Merger Guidelines treat investment and innovation benefits with the same discipline - and the same practical realism - that merger control already brings to competitive harm.

[13] Draft Merger Guidelines, paras 316, 329 and 343-344. Footnote 414 of the draft expressly refers to merger simulation and pricing-pressure tools including GUPPI, compensating marginal-cost reduction and vertical GUPPI.