A Sufficient Statistics Approach to Merger Evaluation *

Vincenzo Alfano[†] Christoph Carnehl[‡] Konrad O. Stahl[§]

January 26, 2026

Abstract

We develop a sufficient-statistics method to evaluate the consumer surplus effects of mergers at scale. Building on the theory of multiproduct-firm oligopoly, we demonstrate that a merger increases consumer surplus if and only if the merging parties' combined revenue rises relative to its pre-merger level. The result accommodates efficiencies, quality upgrading, and changes in product scope, and extends to mergers in an upstream input market. We implement this test by assembling a novel dataset linking European Commission merger decisions to firm accounts and estimating counterfactual revenue paths using synthetic difference-in-differences. Analyzing 152 approved mergers, we find that while the average merger is consumer-surplus neutral, approximately 20% of cases significantly lowered consumer surplus. We also compute compensating cost synergies. Even mergers in our data that do not undergo an in-depth investigation would require roughly 6–11% cost reductions to avoid consumer harm.

JEL Classification Numbers: D43, G34, L41
Keywords: Merger Policy, Consumer Surplus, Ex-Post Evaluation

---

*We thank Andrea D'Amico, Marisol Diaz, Silvia Gatteschi, Nadine Hahn, Lion Holste, Katerina Simkova, and Philip vom Baur for excellent research assistance. We are grateful to Massimo Motta and Tommaso Valletti for helpful discussions. Konrad Stahl gratefully acknowledges support by the Deutsche Forschungsgemeinschaft through CRC TR224 (project B04).

†Bocconi University. Email: vincenzo.alfano@phd.unibocconi.it

‡Bocconi University and IGIER. Email: christoph.carnehl@unibocconi.it

§University of Mannheim, CEPR, CESifo, ZEW. Email: konrad.stahl@uni-mannheim.de

1 Introduction

The recent debate on the rise of market power has renewed interest in merger policy, both because acquisitions are a key channel through which industries consolidate and because merger review is the main ex-ante instrument competition authorities have to prevent market power from being created in the first place. But the market power debate is ultimately a welfare question: whether consolidation has reduced the benefits of competition for consumers. Merger review is explicitly framed in those terms, protecting consumers through lower prices, higher quality, greater variety, and innovation (European Commission, 2004; U.S. Department of Justice and Federal Trade Commission, 2023). Yet, broad academic and policy retrospectives still tend to evaluate mergers through concentration, prices, or markups, objects that are often convenient to measure but do not map cleanly into consumer surplus. A merger that generates efficiencies, improves quality, or changes product scope can raise measured markups and prices while increasing consumer surplus. As a result, the existing evidence provides only limited systematic guidance on how effective merger policy has been at protecting consumers.

The scarcity of direct evidence on consumer surplus reflects a trade-off between rigor and scalability. In principle, structural merger analysis can recover welfare changes by estimating demand and marginal costs. But the data requirements (such as detailed product-market information on prices, quantities or shares, characteristics, and credible instruments) make this approach difficult to apply consistently across a broad set of mergers spanning many industries and markets. The scaling problem is especially acute because mergers typically involve large multiproduct firms and affect multiple markets simultaneously. As a consequence, the literature often ends up speaking either in the language of structurally grounded case studies in a handful of industries, or in the language of large-sample retrospectives built on concentration, prices, or markups that lack a transparent welfare mapping.

In this paper, we propose a way to bridge structural discipline and large-scale measurement. Building on recent theoretical advances in multiproduct-firm oligopoly and aggregative games (Nocke and Schutz, 2018), we derive a sufficient-statistic sign test for a merger's consumer-surplus effect that relies on a single observable: the change in the merging parties' combined revenues. We show that, under a CES demand assumption, a merger increases consumer surplus if and only if the merged entity's revenue rises relative to the pre-merger sum of the merging firms' revenues. Intuitively, with a fixed market scale, consumer-surplus gains correspond to a shift of expenditure toward inside goods; because the merger leaves outsiders' fundamentals unchanged, such a shift must show up as a higher revenue share for the merged entity. In an extension of the baseline model, we show that the same logic applies to mergers in an upstream input-good market.

We implement our test by estimating merger-induced revenue changes relative to a counterfactual no-merger path using a synthetic difference-in-differences design (Arkhangelsky et al., 2021). We assemble a novel dataset that links European Commission (EC) merger decisions to firm-level accounts. From EC decision reports we code merging parties, competitors, affected product and geographic markets, market shares, and procedural features; we then match firms to ORBIS balance-sheet data via Bureau van Dijk identifiers. Because the ORBIS coverage is incomplete, we cannot follow every reviewed merger. Still, we recover the revenue panels required for our analysis for 202 mergers between 2002 and 2021, corresponding to 42.5% of the cases we coded from decision reports, spanning 58% of two-digit NACE industries, with coverage that is stable over time. Moreover, using the observables available in the decision reports, we show that included and excluded cases look broadly comparable, supporting the external validity of the matched sample.

The theoretical framework also yields a transparent ex-ante welfare diagnostic, the compensating cost synergies required for a merger to be consumer-surplus neutral. Using only the market shares reported in the EC decisions and a standard range of substitution elasticities, we compute the cost improvements needed to offset the market-power effect of a given merger. Even Phase 1 mergers require economically meaningful synergies to be consumer-surplus neutral (roughly 6–11% cost reductions on the merger level), while Phase 2 mergers require substantially larger efficiencies (about 14–24%).

We then bring the sufficient-statistic approach to the data. The average merger is consumer-surplus neutral, but heterogeneity is substantial. Around 20% of approved mergers exhibit significantly negative consumer-surplus effects. The incidence of harm is higher among Phase 2 cases (six out of 15 in our baseline sample). This pattern is partly consistent with selection. Phase 2 cases are, by design, the deals raising greater prima facie concerns. At the same time, it is striking because these are exactly the mergers receiving the most intensive scrutiny.

Beyond consumer surplus, the ORBIS linkage lets us study outcomes that matter in policy debates but are rarely available at scale. We examine employment effects and find that average impacts are close to zero, though a non-trivial minority of cases, especially among Phase 2 mergers, exhibit significant employment reductions. Finally, the combination of detailed ex-ante market structure measures from decision reports with ex-post welfare classifications creates a natural platform for prediction: we can evaluate whether the screening metrics embedded in merger guidelines (e.g., HHI levels and ΔHHI thresholds, market shares, and procedural indicators) actually forecast consumer harm, and how such tools could be improved to better target intensive review.

Beyond consumer surplus, the ORBIS linkage lets us study other outcomes that matter in policy debates but are rarely available at scale. We examine employment effects and find that most mergers have no statistically detectable impact, while a non-trivial minority show employment reductions, especially among Phase 2 cases. Finally, combining rich ex-ante market structure measures from decision reports with ex-post welfare classifications creates a natural platform for evaluating whether the screening metrics embedded in merger guidelines (e.g., HHI and ΔHHI thresholds, market shares, and procedural indicators) predict consumer harm, and for designing improved targeting of intensive review.

Related Literature We connect to three strands of literature: (i) Recent theoretical advances in the analysis of multiproduct-firm mergers, (ii) the literature on merger retrospectives, and (iii) empirical research on European merger control.

Our empirical strategy builds on the aggregative-games approach to multiproduct-firm oligopoly developed by Nocke and Schutz (2018), which delivers a tractable characterization of equilibrium in differentiated-product industries and a type-aggregation property that collapses rich product-portfolio heterogeneity into a one-dimensional firm index. Building on these insights, we derive a new ex-post welfare diagnostic: under CES demand, the sign of the merger's consumer-surplus effect is identified by the sign of the change in the merging parties' combined revenues. This result provides a structural bridge between standard accounting outcomes and welfare, enabling welfare-relevant evaluation without estimating demand and marginal cost primitives.

Our approach also complements recent work clarifying the welfare content of the concentration screens used in practice. Nocke and Whinston (2022) and Nocke and Schutz (2025) provide microfoundations for why merger-induced changes in concentration measures (rather than levels) can be informative about welfare consequences in these models. We use these insights in two ways: first, to compute compensating marginal-cost reductions that would render each merger consumer-surplus neutral as a transparent ex-ante benchmark; and second, to evaluate how well such screening metrics forecast the ex-post outcomes produced by our revenue-based test.

Empirically, our paper contributes to the merger retrospective literature that measures the effects of mergers on observable outcomes (Asker and Nocke, 2021, provide an excellent survey). This literature can be divided into two groups: (i) detailed industry studies, such as Luo (2014), Das (2019), and Peters (2006) on airlines and Ashenfelter et al. (2015) and Miller and Weinberg (2017) on beer. In some cases, a larger set of mergers, still in a single industry have been studied (for example, Kim and Singal, 1993; Focarelli and Panetta, 2003; Bhattacharya et al., 2023). The majority of papers focusing on larger merger retrospectives studies price, quantity or markup effects (Bhattacharya et al., 2023; Stiebale and Szücs, 2022, is a recent example). The papers focused on smaller studies sometime provide structural estimates (for example, Miller and Weinberg, 2017). We complement this literature with a theory-based sufficient-statistics approach that is applicable at larger scale.

Finally, we contribute to the literature on European merger control. Affeldt (2019) provides an excellent overview of the data that can be obtained from the EC. Our data set takes the same source as initial basis but includes more recent years and links merging entities to financial data. Stiebale and Szücs (2022) measure the effect of mergers on markups following De Loecker and Warzynski (2012) and building on the data in Affeldt (2019) using information on merging firms' competitors.

2 Oligopoly, Mergers and Consumer Surplus

In this section, we first set up the multiproduct firm oligopoly game and characterize the equilibrium following Nocke and Schutz (2018). Second, we formalize mergers in this environment and define the compensating synergies; specifically, the cost synergies necessary to ensure that a merger is consumer-surplus neutral. Third, we prove our main theoretical result: that the revenue changes of the merging firms constitute a sufficient statistic for the consumer surplus effects of a merger. This result provides the structural foundation for our empirical analysis.

2.1 Oligopoly Model and Equilibrium

Consider an industry with a finite set of differentiated products $\mathcal{N}$ and a set of firms $\mathcal{F}$ with $|\mathcal{F}| \geq 3$. The set of firms forms a partition of the set of products, such that each product $j \in \mathcal{N}$ is supplied by exactly one firm. We identify each firm $f \in \mathcal{F}$ with the subset of products it produces, denoted $\mathcal{N}_f \subset \mathcal{N}$. We assume constant returns to scale, where $c_j > 0$ denotes the constant marginal cost of product j.

Preferences and Demand. We derive the market demand from a representative consumer with quasilinear preferences and an indirect utility function $V(\mathbf{p}) = y + \frac{1}{\eta} \ln H(\mathbf{p})$, where y is the numeraire and $H(\mathbf{p})$ is the industry aggregator. We assume demand exhibits Constant Elasticity of Substitution (CES), implying an aggregator of the form

$$H(\mathbf{p}) \equiv H^0 + \sum_{j \in \mathcal{N}} b_j p_j^{1-\sigma}, \tag{1}$$

where $H^0 \geq 0$ represents the value of the outside option, b_j denotes the vertical quality of product j, and $\sigma > 1$ is the elasticity of substitution.

Applying Roy's Identity, the demand for product j at price vector $\mathbf{p}$ is

$$D_j(\mathbf{p}) = \frac{b_j p_j^{-\sigma}}{H(\mathbf{p})} M, \tag{2}$$

where $M \equiv (\sigma - 1)/\eta$ is an exogenous market scale parameter. In this framework, consumer surplus is strictly increasing in the aggregator $H(\mathbf{p})$, and is given by the indirect utility of the representative consumer

$$CS(\mathbf{p}) = y + \frac{1}{\eta} \ln H(\mathbf{p}). \tag{3}$$

Firms' profits and choices. Firms compete by simultaneously setting prices for their respective product portfolios to maximize profit. Using the demand specification in (2), the profit function of firm f is given by:

$$\Pi_f(\mathbf{p}) = \sum_{j \in \mathcal{N}_f} (p_j - c_j) D_j(\mathbf{p}). \tag{4}$$

In the following, we denote firm f's revenues by $R_f(\mathbf{p}) = \sum_{j \in \mathcal{N}_f} p_j D_j(\mathbf{p})$. As established in Nocke and Schutz (2018), this pricing game is *aggregative*: the profit of firm f depends on the pricing decisions of its rivals solely through the value of the aggregator $H(\mathbf{p})$.

This structural property allows for *type aggregation*. Specifically, the optimal pricing strategy implies that all relevant information regarding firm f's product portfolio (including the number of products, their vertical qualities b_j, and their marginal costs c_j) can be summarized by a uni-dimensional sufficient statistic, the firm's type T_f:

$$T_f \equiv \sum_{j \in \mathcal{N}_f} b_j c_j^{1-\sigma}. \tag{5}$$

Economically, the type T_f represents the firm's aggregate contribution to consumer surplus under perfect competition. A higher type reflects a portfolio with higher quality products or more efficient production technologies.

Equilibrium. Nocke and Schutz (2018) show that, in the unique Nash equilibrium, firm f's market share s_f (defined as its share of the structural market scale M, i.e., $s_f \equiv R_f/M$) is determined exclusively by its own type T_f and the equilibrium aggregator H^*. That is,

$$s_f(H^*) = S\left(\frac{T_f}{H^*}\right), \tag{6}$$

where $S(\cdot)$ is the *market share fitting-in function*. Nocke and Schutz (2018) prove that $S(z)$ is strictly increasing in z. This yields the following monotonicity property, which is central to our empirical identification strategy.

Lemma 1. *Holding a firm's type T_f constant, its equilibrium revenue market share $s_f(H) = S(T_f/H)$ is strictly decreasing in the level of the aggregator H.*

Moreover, Nocke and Schutz (2018) show that, for any vector of firm types $\mathbf{T} = (T_f)_{f\in\mathcal{F}}$, a unique equilibrium exists, and that the equilibrium aggregator H^* is uniquely pinned down by the market-clearing condition that the sum of the firms' market shares and the share of the outside option ($s_0 \equiv H^0/H^*$) must equal unity

$$\sum_{f\in\mathcal{F}} S\left(\frac{T_f}{H^*}\right) + \frac{H^0}{H^*} = 1. \tag{7}$$

By equation (3), consumer surplus is proportional to the level of the aggregator, and hence, type vectors $\mathbf{T}$ that lead to a greater equilibrium aggregator H^* feature higher consumer surplus.

2.2 Mergers and Compensating Synergies

We now analyze a horizontal merger between a subset of firms $\mathcal{M} \subset \mathcal{F}$. Let $\mathcal{O} = \mathcal{F} \setminus \mathcal{M}$ denote the set of non-merging "outsider" firms. Applying the type aggregation property, we can model mergers in a reduced form as affecting the market equilibrium by replacing the independent pre-merger types $\left(T_f\right)_{f\in\mathcal{M}}$ with a single post-merger type $T_{\mathcal{M}}$. This way of modeling mergers allows for many distinct effects a merger might have on the merging parties. We can capture cost efficiencies (lower c_j for $j \in \mathcal{N}_{\mathcal{M}}$) or quality improvements (higher b_j for $j \in \mathcal{N}_{\mathcal{M}}$), but also can have heterogeneous impacts across products in the merged firm's product portfolio. Overall, the post-merger type $T_{\mathcal{M}}$ aggregates the effect on the entire portfolio into a scalar statistic.

A merger without any synergies would correspond to a post-merger type equal to the sum of the pre-merger types, $T^0_{\mathcal{M}} = \sum_{f \in \mathcal{M}} T_f$. Such a merger only reduces competition and necessarily harms consumers (see Proposition 2 in Nocke and Schutz, 2025). Hence, to be consumer-surplus neutral, a merger requires sufficiently large synergies; in particular, there exists a unique threshold type $\overline{T}_{\mathcal{M}} > T^0_{\mathcal{M}}$ such that a merger is consumer-surplus neutral if and only if $T_{\mathcal{M}} = \overline{T}_{\mathcal{M}}$ (see Proposition 2 in Nocke and Whinston, 2022). This threshold type is determined from the observation that (i) consumer surplus neutrality requires the same level of the equilibrium aggregator H^* pre- and post-merger, and (ii) that the market clearing condition (7) has to hold, implying

$$S\left(\frac{\overline{T}_{\mathcal{M}}}{H^*}\right) = \sum_{f \in \mathcal{M}} S\left(\frac{T_f}{H^*}\right). \tag{8}$$

Due to the concavity of $S(\cdot)$ (see Lemma 5 in Nocke and Schutz, 2025), the neutral type must strictly exceed the sum of pre-merger types: $\overline{T}_{\mathcal{M}} > \sum_{f \in \mathcal{M}} T_f$.

While the threshold type is an abstract object without immediate empirical analogue, we can map it into a straightforwardly computable and interpretably "compensating marginal cost reductions" as analyzed by Nocke and Whinston (2022). Suppose that the synergies take the form of a uniform percentage change in marginal costs κ, so that for all $j \in \mathcal{N}_{\mathcal{M}}$, the post-merger marginal cost are $c^{\mathcal{M}}_j = (1+\kappa)c_j$. Then, the required cost synergy $\overline{\kappa}$ ensuring that the merger is consumer-surplus neutral can be expressed in terms of the merging firms' pre-merger market shares (see Corollary 3 in Nocke and Whinston, 2022)

$$\overline{\kappa} = 1 - \left(\frac{s_{\mathcal{M}} \left(\sigma + \frac{s_{\mathcal{M}}}{1 - s_{\mathcal{M}}}\right)^{\sigma - 1}}{\sum_{f \in \mathcal{M}} s_f \left(\sigma + \frac{s_f}{1 - s_f}\right)^{\sigma - 1}} \right)^{\frac{1}{1-\sigma}} \tag{9}$$

where $s_{\mathcal{M}} \equiv \sum_{f \in \mathcal{M}} s_f$ is the naively computed market share of the merged entity.

2.3 Revenues as Sufficient Statistic

Compensating synergies provide an interpretable ex-ante benchmark for the required merger effects on merging firms' post-merger types to ensure that a merger is CS-neutral. Yet, they do not provide clear guidance on evaluating a merger's consumer surplus effects ex post. The type threshold $\overline{T}_{\mathcal{M}}$ is more permissive to realistic heterogeneous effects, but is not directly observable. Having a simple sufficient statistic for consumer surplus that can be straightforwardly computed is therefore of significant value to assess realized mergers ex post. We now show that the change in the sum of the merging firms' revenues from pre-merger to the post-merger entity provides such a simple sufficient statistic for the consumer surplus effect of that merger.

Let $R_{\mathcal{M}} = \sum_{f\in\mathcal{M}} R_f$ be the combined pre-merger revenue of the merging parties, and let $\overline{R}_{\mathcal{M}}$ be the post-merger revenue of the merged entity.

Proposition 1. *The sign of the change in the merging firms' revenues identifies the sign of the change in consumer surplus:*

1. *The merger strictly increases consumer surplus if and only if* $\overline{R}_{\mathcal{M}} > R_{\mathcal{M}}$.
2. *The merger is consumer surplus neutral if and only if* $\overline{R}_{\mathcal{M}} = R_{\mathcal{M}}$.
3. *The merger strictly decreases consumer surplus if and only if* $\overline{R}_{\mathcal{M}} < R_{\mathcal{M}}$.

Proof. Recall that $CS(\mathbf{p})$ is strictly increasing in the equilibrium aggregator H^*. Thus, it suffices to show that the sign of the change in the equilibrium aggregator coincides with the sign of the change in revenues, $\text{sgn}(\overline{H} - H^*) = \text{sgn}(\overline{R}_{\mathcal{M}} - R_{\mathcal{M}})$, where $\overline{H}$ denotes the post-merger equilibrium aggregator.

Recall that, in any equilibrium, the sum of the market shares of the insiders ($s_{\mathcal{M}}$), the outsiders (s_g for $g \in \mathcal{O}$), and the outside option (s_0) must equal unity

$$s_M(H) + \underbrace{\sum_{g\in\mathcal{O}} s_g(H) + \frac{H^0}{H}}_{\equiv\Omega(H)} = 1. \tag{10}$$

Consider the function $\Omega(H)$, which represents the aggregate market share of all non-merging entities. First, the share of the outside option, H^0/H, is strictly decreasing in H. Second, by Lemma 1, the share of any outsider firm g, $s_g(H) = S(T_g/H)$, is strictly decreasing in H, as their types T_g are invariant to the merger. Consequently, $\Omega(H)$ is strictly decreasing in H. To satisfy the identity $s_M(H) = 1 - \Omega(H)$, the equilibrium share of the merging firms must be strictly increasing in the aggregator H.

Since, for any firm $f \in \mathcal{F}$, revenues are given by $R_f = M \cdot s$, and the structural parameter M is constant, we have

$$\overline{H} > H^* \iff s_{\mathcal{M}}(\overline{H}) > s_{\mathcal{M}}(H^*) \iff \overline{R}_{\mathcal{M}} > M \sum_{f \in \mathcal{M}} s_f(H^*), \tag{11}$$

where $M \sum_{f \in \mathcal{M}} s_f(H^*)$ is the sum of pre-merger revenues of the merging parties. □

Intuitively, the result obtains from three consecutive observations. First, under CES demand, the market scale parameter M is exogenously fixed and measures consumers' total expenditures on the goods supplied by the firms in the market $\mathcal{F}$ as well as on the outside option H^0. Second, any increase in consumer surplus (which corresponds to a higher level of the aggregator H) must come from additional expenses on the inside goods $\mathcal{N}$, as the value of the outside option remains fixed. Third, the merger affects only the merging firms structurally; that is, the types of non-merging firms T_f for $f \in \mathcal{O}$ remain constant. Thus, any increase in consumer surplus implies an increase in the revenue share, and hence, in the revenues of the merging parties.

Proposition 1 provides us with a simple empirical test for the consumer surplus effect of a merger. If the merging firms' revenues increase post-merger, the merger was procompetitive. We leverage this novel theoretical result to evaluate realized mergers in the European Union.

2.4 Vertical Supply Chains

While our baseline model focuses on firms selling directly to consumers, some of the mergers in our sample involve upstream firms producing intermediate inputs for downstream manufacturers. We briefly illustrate how our theoretical insights extend to such cases and show that the revenue sufficient statistic can, under appropriate restrictions on the downstream market, be applied to such vertical markets as well.

Consider an upstream sector supplying differentiated inputs to a downstream market, which in turn sells to final consumers with CES demand. We assume downstream firms produce using a common CES bundle of these upstream inputs. In this setting, an upstream merger affects downstream firms symmetrically via a change in the input price index. As we show in Appendix E, this acts as a uniform multiplicative cost shock to the downstream sector. We prove that such uniform shocks are "neutral": they scale downstream prices but leave equilibrium market shares, and crucially, total downstream expenditure on inputs invariant.

This neutrality implies that the upstream pricing game is isomorphic to our baseline model and that the upstream aggregator is passed on to final good consumers monotonically. Hence, we obtain the analogous result to the case in Proposition 1. We formally show that both under downstream (i) monopolistic competition and (ii) CES oligopoly, an increase in the merging upstream firms' revenue identifies an improvement in the upstream aggregator. This improvement passes through to final consumers, increasing consumer surplus.

3 Data & Descriptive Evidence

In this section, we describe the construction of the merger-level and firm-level data that underlie our analysis and present key descriptive statistics on the mergers in our data. In particular, we compute the model-based compensating cost synergies for the mergers in our sample.

3.1 Data Sources and Sample Construction

The generation of the complex data set itself is one of the main contributions of the paper: we link the European Commission's merger decisions to firm-level balance-sheet data and construct a panel that follows the merging entities before and after the merger. We use three data sources to obtain our final data set; the merger decision reports provided by the European Commission, the Zephyr M&A database, and ORBIS.

Commission reports. Our starting point is the universe of horizontal mergers reviewed by the European Commission (EC) between 2002 and 2021. In constructing our data, we apply the following four filters: (i) We exclude cases handled under the EC's simplified procedure, which are cleared without an in-depth competitive assessment. (ii) We drop mergers decided before 2002, because financial coverage in ORBIS is sparse in earlier years. (iii) For Phase-1 mergers, we restrict attention to cases in the twenty most frequently affected two-digit NACE industries, and to product markets whose geographic scope is at least national (national, EEA-wide, or worldwide). (iv) We exclude mergers that were prohibited or withdrawn, as our objective is to estimate the consumer-surplus effects of mergers that were approved by the EC. The resulting set of cases covers the large majority of non–simplified EC merger decisions over this period.

For each merger case, we use the published decision to extract: the identities of the merging parties and the main competitors; the dates of the merger announcement and the final decision; whether the case was cleared in Phase 1 or Phase 2, and whether remedies were imposed;[1] the list of affected product and geographic markets as specified by the EC; for each market, the reported market-share intervals of the merging firms and of the main competitors; an indicator of whether the EC raises explicit competition concerns in that market ("problematic" market).

Using the reported market-share intervals, we construct pre- and post-merger Hirschman-Herfindahl indices (HHIs) and the naively computed merger-induced changes in these indices ΔHHI (obtained by adding up the merging firms' market shares) for each product × geographic market. Because the EC reports market shares as intervals and only for the main competitors, computing HHIs requires assumptions about how to allocate mass within and outside the reported ranges. We adopt a simple and transparent rule that we discuss in Appendix B. Therein, we also discuss possible alternatives and their consequences.

[1] In Phase 1 the EC has 25 working days to decide whether to clear a merger (with or without remedies which may be offered by the merging parties), reject it, or open a Phase 2 investigation on the basis of competition concerns. If a Phase 2 investigation is opened, the EC has up to 90 working days to carry out an in-depth assessment and decide whether to clear the merger, possibly with remedies, or to prohibit it.

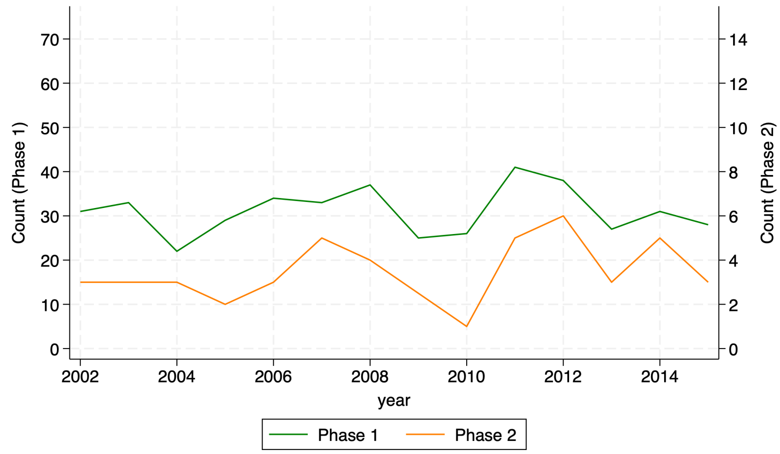

Figure 1: Count of Cases by Year of Application

Zephyr and ORBIS. [2] To link merger decisions to firm-level accounts in ORBIS, we link the EC cases to Zephyr using the names of the parties, deal characteristics, and dates reported in the EC decisions. This yields Zephyr deal identifiers and, crucially, Bureau van Dijk (BvD) firm identifiers for acquirers and targets.

Using BvD firm identifiers, we extract annual accounting and financial data from ORBIS for all merging firms and for a large pool of potential control firms. For each merger case m, we assemble a firm–level panel from (whenever available) five years before until five years after the merger decision year T_m, i.e., $t \in [T_m - 5, T_m + 5]$.

Final sample and coverage. To estimate dynamic treatment effects, we require at least two years of outcome data before and after the merger. Specifically, we retain only cases for which we observe complete information on operating revenues from (at least) $T_m - 2$ through $T_m + 2$ for the merging firms in ORBIS.[3] Applying this requirement yields a final sample of 204 out of 487 collected merger cases. Figure 1 shows the distribution of cases over our sample period.

Relative to the set of the EC merger cases we coded from the decision reports, this represents 42.5% of mergers. These cases span 58% of the two–digit NACE

[2]These databases, provided by Bureau van Dijk, contain arguably the best available data on firms located and active in Europe.

[3]Note that in our main specification, we apply a more restrictive criterion of a complete time series from $T_m - 2$ through $T_m + 4$ to ensure sufficient time for merger effects to materialize.

industries covered in the merger case sample, and the coverage rate is quite stable over time. In Appendix A we document the time series of coverage and compare the observable characteristics of included and excluded cases.[4]

Our ORBIS data contain both unconsolidated and consolidated accounts: For a subset of cases, we observe unconsolidated financial statements for the legal entities involved in the merger; in others, only consolidated accounts are available, combining the acquirer with its subsidiaries (and, post-merger, the acquired target). In our baseline analysis, we include both types of cases and later show that results are robust to restricting attention to cases with unconsolidated accounts or to analyzing consolidated and unconsolidated cases separately.

3.2 Descriptive Statistics

In the descriptive statistics of our data, we emphasize expected properties of mergers and their investigation. Mergers involve large multiproduct firms active in many markets. Phase 2 mergers tend to feature larger firms and larger changes in concentration. There is vast heterogeneity across mergers in their ex-ante observables, however.

In Table 1, we summarize information on the merger cases in our final data set.[5] We report the statistics separately for Phase 1 and Phase 2 mergers. Three important observations emerge from the descriptive statistics.

First, by the product markets classified by the EC in their decision reports, mergers have a substantial relevant multiproduct component. The average Phase 1 merger affects 5 product markets, while for Phase 2 mergers this average goes up to 8 product markets. Even when restricting attention to only the (more relevant) problematic markets, we find that the average Phase 2 merger features almost 4 problematic markets. Hence, taking the multiproduct nature of competition seems crucial in assessing mergers.

Second, merging parties in our data set have substantial market shares. In

[4]Note that our sample includes only few cases after the 2015 notification date. This is due to many missing observations within the ORBIS database for recent years. See Figure 8 in Appendix A for the share of complete cases after 2015 that we obtain.

[5]The post-merger HHI and ΔHHI are computed using the naive aggregation of merging firms' market shares. As the same product market can be served in several geographic markets, the total number of markets corresponds to at most the number of product × geographic markets.

	Phase 1				Phase 2			
	Affected		Problematic		Affected		Problematic	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD
N Markets								
Geographic Markets	3.22	3.83	0.18	1.60	6.52	8.12	3.11	5.24
Product Markets	5.09	6.23	0.16	0.78	8.48	10.12	3.74	7.50
Combined	8.37	11.10	0.01	0.18	23.83	38.80	8.41	20.11
Market Shares								
Acquirers	0.17	0.17	0.58	0.33	0.24	0.19	0.29	0.19
Targets	0.11	0.15	0.00	0.00	0.26	0.25	0.32	0.26
All	0.28	0.18	0.58	0.33	0.50	0.26	0.61	0.23
HHI								
Pre-Merger HHI	0.15	0.15	0.40	0.30	0.28	0.17	0.32	0.18
Post-Merger HHI	0.17	0.16	0.40	0.30	0.38	0.22	0.46	0.21
Delta HHI	0.02	0.03	0.00	0.00	0.09	0.09	0.14	0.09
Revenues (Billion Euros)								
Operating Revenues	14.18	22.98	21.41	24.16	22.96	36.61	19.92	37.18

Table 1: Summary Statistics of Merger Cases. We report descriptive statistics for the markets associated with the mergers in our sample, separately for Phase 1 and Phase 2 cases, distinguishing between all affected markets and those classified as problematic. For each subsample, we report the mean and standard deviation of the number of geographic and product markets (and their sum), the market shares of acquirers and targets (and their combined share), as well as the pre-merger, post-merger, and change in the Herfindahl-Hirschman Index. The bottom row shows the operating revenues (in billion euros) of the merging parties.

markets affected by Phase 1 cases, the merging parties together account on average for about 28% of market revenues; in markets affected by Phase 2 cases, their combined share rises to about 50%, and to more than 60% in problematic Phase 2 markets. As one would expect, Phase 2 cases typically involve larger merging parties than Phase 1 cases.

Third, the HHI panel shows that markets flagged as problematic are, on average, more concentrated than the broader set of affected markets, and that Phase 2 cases tend to feature both higher pre-merger concentration levels and larger merger-induced increases. In affected markets, the mean ΔHHI is about 0.02 for Phase 1 mergers and 0.09 for Phase 2 mergers; in problematic Phase 2 markets it is around 0.14, indicating that the largest increases in concentration occur in the subset of

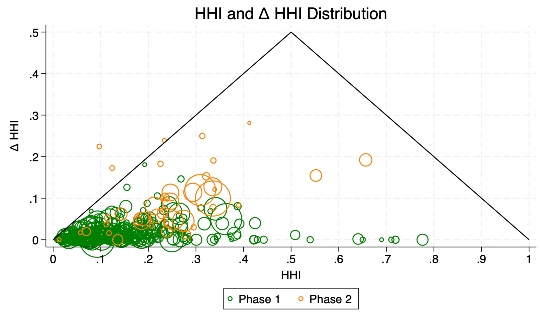

(a) Affected markets.

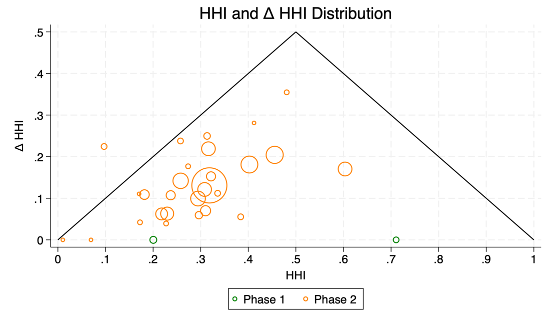

(b) Problematic markets.

Figure 2: HHI–ΔHHI combinations for affected and problematic mergers. Each bubble represents a merger case in our data set. The size of the bubble depends on the number of affected markets in panel (a) and on the number of problematic markets within the merger in panel (b).

Phase 2 markets where the Commission ultimately identifies serious concerns.

These somewhat surprising observations are detailed in Figure 2a. We document both the pre-merger HHI and the ΔHHI for each merger. These pre-merger statistics play an important role in practical merger policy (see, for example, U.S. Department of Justice and Federal Trade Commission, 2023; European Commission, 2004). The ΔHHI has been identified in theoretical work to be a good indicator for the welfare consequences of mergers (see Nocke and Schutz, 2025). Unsurprisingly, Phase 2 mergers are more frequently found to be located higher along the vertical axis, relative to Phase 1 mergers; that is, they tend to have higher ΔHHI, implying larger naively computed changes in the post-merger HHI. Similarly, in Figure 2b, we see that markets for which the EC raises competitive concerns tend to have a higher ΔHHI rather than a high pre-merger concentration.

3.3 Compensating Synergies

Our theoretical framework provides us with a simple, but powerful ex-ante assessment tool: the synergies required to ensure that a merger is at least consumer-surplus neutral. In particular, we can compute the compensating synergies in equation (9) directly using the merging firms' market shares from the decision reports.

To compute the cost synergies, we need to choose a value for the elasticity of substitution σ. Following Nocke and Whinston (2022), we offer a range from $\sigma = 4$ to $\sigma = 7$, which is in line with estimated substitution elasticities in the trade literature.

In Table 2 we report the average compensating cost synergies, split up across Phase 1 and Phase 2 mergers.[6] The top panel shows the merger-level average compensating synergies; that is, we first average the compensating synergies for all markets affected by a given merger and then average across all mergers. The bottom panel shows the market-level average compensating synergies; that is, we average all affected markets in our data set. The main picture is similar. Even Phase 1 mergers require substantial cost synergies to be consumer-surplus neutral, between 6 and 11% on the merger level. On the market level, the average compensating synergies are higher (between 9 and 20 %). This is due to a clustering of markets with relatively large compensating synergies in some mergers. The compensating statistics that we estimate show that even Phase 1 mergers, that only undergo a relatively brief assessment, require non-negligible synergies to overcome their market-power effect.

Phase 2 mergers require, unsurprisingly, larger compensating synergies, between 14 and 24% on the merger level. This statistic shows that the EC indeed identifies mergers that appear detrimental from an ex-ante perspective for consumer surplus and that require careful consideration before being approved.

Overall, the large compensating synergies in Table 2 suggest that an ex-post evaluation of merger decisions is warranted. Synergies from a merger are difficult to verify ex ante, but are crucial in ensuring that merger activities do not come at the expense of consumers. An ex-post evaluation helps evaluation whether the current decision processes ensure that only (or at least predominantly) mergers with sufficient procompetitive components are approved.

Remark. We compute the compensating synergies on a market-by-market basis even within a given merger. The theoretical setting suggests a single statistic based on the aggregated revenue shares of the merging firms. While our market-by-market computation is due to data constraints, it also provides a meaningful

[6]We report the more abstract and harder to interpret *type synergies* in Appendix C.

	Phase 1		Phase 2	
	Mean	SD	Mean	SD
Merger Level				
$\sigma = 4$	−0.11	0.31	−0.24	0.24
$\sigma = 5$	−0.08	0.22	−0.19	0.19
$\sigma = 6$	−0.07	0.18	−0.16	0.17
$\sigma = 7$	−0.06	0.15	−0.14	0.15
Market Level				
$\sigma = 4$	−0.20	1.15	−0.22	0.46
$\sigma = 5$	−0.14	0.75	−0.18	0.38
$\sigma = 6$	−0.10	0.55	−0.16	0.33
$\sigma = 7$	−0.09	0.43	−0.14	0.29

Table 2: Compensating Cost Synergies. The compensating cost synergies are ex-ante specified as the percentage by which pre-merger marginal costs must decrease, so that the merger is consumer surplus-neutral. The merger-level panel displays the mean of within-merger average compensating synergies across all mergers. The market-level panel displays the mean compensating synergies across affected markets without pre-averaging on the merger-level.

benchmark.[7] Our statistics provide the compensating synergy for consumers not to be harmed in any individual market and do not allow for cross-subsidization of markets, as the merger-level synergy would. Such compensations across markets may be considered depending on the legislation (see Werden, 2017, for a discussion of balancing anticompetitive effects across markets).

4 Assessment of Merger Cases

In this section, we investigate the consumer surplus effects of the merger cases in our data set. To do so, we employ our novel empirical test based on *realized* revenue changes of the merging firms in Proposition 1.

[7]As we only have access to revenue market shares on a product-geography level, we cannot compute a merger-level market share without additional information on the relative sizes of the markets.

4.1 Empirical Strategy: Synthetic Difference-in-Differences

Our empirical objective is to estimate the CS effect of approved mergers. Employing our Proposition 1, this amounts to estimating the treatment effect τ^m of a merger case m on the merging parties' revenues relative to a counterfactual no-merger revenue path. We obtain the treatment effects using a synthetic difference-in-differences (SDID) approach following Arkhangelsky et al. (2021).

Let i index firms, m merger cases and t calendar years. For each merger case m, we denote by $\mathcal{M}_m$ the set of firms involved in the merger and by t_m the merger notification year. We work in event time, so that for each merger m and year t, we compute $k = t - t_m$ such that $k = 0$ denotes the merger year. For each firm-year combination i, t, we use the year-specific operating revenues from ORBIS as the variable of interest, $rev_{i,t}$.[8] We obtain post-merger revenues by considering unconsolidated accounts on ORBIS for the merging parties whenever available. If these are not available, we consider the consolidated revenues of the acquirers and drop the targets accordingly to avoid double-counting.[9] In Appendix D.3, we also estimate the treatment effects not as average treatment effect for each treated unit, but instead consider the sum of merging firms' revenues as variable of interest. We find that both approaches lead to comparable outcomes.

In our baseline specification, we focus on the time window from two years before ($k = -2$) until four years after ($k = +4$) the merger.[10] We denote the treatment indicator $D^m_{i,t} = \mathbf{1}_{i \in \mathcal{M}_m} \cdot \mathbf{1}_{k \geq 0}$.

For each merger case m, we construct the set of control firms starting from a dataset of firms that are not directly involved in any mergers within our sample, and then progressively narrow this set along product-market and size dimensions. First, we exclude all firms that are explicitly mentioned as competitors in the corresponding European Commission (EC) merger report. Second, among the remaining firms with a complete time series for the period considered, we retain

[8] We obtain $rev_{i,t}$ from ORBIS' operating revenues after deflation. In the empirical estimation, we use the logged deflated revenues.

[9] In Appendix D.2, we show that our results are not driven by only consolidated or unconsolidated accounts.

[10] In Appendix D.1, we provide the estimated treatment effects also for alternative time windows and show that they are robust. Our choice of the $k = -2$ to $k = 4$ window strikes a reasonable balance between sample size and a sufficiently long post-merger period to let merger effects materialize.

only those operating in the same 2-digit NACE industry as that of the merging parties. Third, we apply three filters in sequence: one based on a finer industry definition (3-digit NACE code), and two based on firm size (financial assets and employment).

The size filters are defined as follows. For each firm, we compute the average size (in terms of financial assets or employment) over the pre-merger years. We then pool the candidate control firms with the merging firms and sort all firms into quintiles of the pre-merger average. Candidate control firms are flagged if they belong to the same quintile as at least one merging firm. If at least 11 firms are flagged, we retain only these firms as potential controls and proceed to the next filter. If fewer than 11 firms are flagged, we do not apply this filter and keep the broader set of candidate controls. In our baseline specification, we apply all three filters. We impose this minimum donor-pool size to avoid the counterfactual revenue evolution to be driven by few control firms with potentially extreme weights.

Our donor-pool construction (same two- and three-digit NACE codes and size filters) is designed to net out common industry and size-related shocks in firm-wide operating revenues. However, because ORBIS revenues may aggregate activities beyond the affected markets, we cannot rule out that there are shocks that affect the merging firms entirely.[11]

With the SDID method we combine the advantages of synthetic control and difference-in-differences approaches.[12] As in synthetic control approaches, SDID uses reweighting and matching based on pre-trends to construct a counterfactual revenue evolution. As in difference-in-differences approaches, SDID controls for time- and industry-specific fixed effects. For each merger case m, the SDID estimator

[11]We carry out additional robustness checks with different donor-pool constructions in Appendix D.4.

[12]We adopt the SDID approach as it more flexibly constructs the counterfactual revenue evolution. In particular, we expect the synthetic control approach to not do as well, as we are considering the largest firms in the European Union that engage in merger activities under the scrutiny of the EC. In Appendix D.5, we compare our estimation results to those of standard difference-in-differences and synthetic control approaches. We find that the DID results are comparable to the SDID results, but that the synthetic control estimates systematically overestimate the revenue effect of mergers.

is specified as

$$(\hat{\tau}^m, \hat{\mu}^m, \hat{\alpha}^m, \hat{\beta}^m) = \arg \min_{\tau,\mu,\alpha,\beta} \left\{ \sum_{i=1}^{N} \sum_{k=-2}^{4} \left(rev_{i,k} - \mu - \alpha_i - \beta_k - D_{i,k}^m \tau \right)^2 \hat{\omega}_i \hat{\lambda}_k \right\} \tag{12}$$

where $\hat{\omega}_i$ is an estimated control-unit weight to match trends in the pre-merger revenues of the merging firms, $\hat{\lambda}_k$ is an estimated time-weight to balance pre-merger time periods, and N denotes the size of the control firm pool plus merging firms. Within the SDID implementation, we first compute the weights $\hat{\omega}_i$ for the control firms and the time-weights $\hat{\lambda}_k$ for the pre-merger periods, weighing periods relatively more that predict post-merger outcomes better.

To obtain standard errors for each estimated merger-specific treatment effect $\hat{\tau}^m$, we follow the placebo procedure outlined in Arkhangelsky et al. (2021), that builds on Abadie et al. (2010) and Abadie et al. (2015). This procedure remains valid if the number of treated units is small. It approximates the estimator's variance by iteratively assigning placebo treatments to control units and computing the effect distribution when the true effect is known to be zero.

Whenever we find a significantly negative (positive) treatment effect $\hat{\tau}^m$, we conclude that the merger decreased (increased) consumer surplus.

4.2 Consumer Surplus Effects of Mergers

In Table 3 we report our estimated treated effects across the 152 merger cases for which we have complete data. The results indicate that the average treatment effect is slightly negative, and this with substantial heterogeneity. More specifically, we find that there are only few cases for which we obtain a significantly positive treatment effect.[13] In contrast, we obtain for around 20% of mergers a significantly negative treatment effect, suggesting that a substantial share of mergers in our sample reduced consumer surplus. The share of anticompetitive merger cases goes up further if we restrict the sample to Phase 2 mergers only; here, 40 % or six out of the 15 mergers in our sample appear to have harmed consumers. Note that we interpret only the *sign* of the treatment effects. This is due to our reduced-form

[13] Note that at our 5% significance level, finding around 3% of positively significant treatment effects is to be expected even when the true treatment effect is zero.

Sample	TE	SE	% Impact	% Pos Sign.	% Neg. Sign.	N
Full Sample	-0.166	0.297	-0.712	3.29	20.39	152
Phase 1	-0.143	0.304	-0.606	3.65	18.25	137
Phase 2	-0.375	0.233	-1.681	0.00	40.00	15

Table 3: Summary table of estimated treatment effects. TE reports the average treatment effect $\tau = \frac{1}{M}\sum_{m=1}^{M}\tau^m$ and SE the corresponding standard errors computed using the placebo approach. % Impact translates the treatment effect into percentage changes in pre-merger operating revenues. % Pos Sign. and % Neg Sign. indicate the share of cases in which we find positively and negatively significant treatment effects at the five percent level, respectively.

approach, by which we infer the consumer surplus effect from the sign of the revenue change of the merging parties. To quantify the *magnitude* of the losses in consumer surplus, a more structural approach is necessary.

Event study decomposition. While the estimator $\hat{\tau}^m$ captures the average treatment effect over the post-merger period, we also estimate dynamic effects to trace the evolution of revenues over time. Following the decomposition in Ciccia (2024), we define the dynamic treatment effect $\hat{\tau}_k^m$ for each event-year $k \geq 0$. This estimator compares the observed revenues of the merging firms to their synthetic counterfactual at time k, adjusted for the time-weighted pre-merger fit.

$$\hat{\tau}_k^m = \left(\frac{1}{|\mathcal{M}_m|}\sum_{i\in\mathcal{M}_m} rev_{i,t_m+k} - \sum_{i\notin\mathcal{M}_m}\hat{\omega}_i \cdot rev_{i,t_m+k}\right) - \sum_{t<t_m}\hat{\lambda}_t\left(\frac{1}{|\mathcal{M}_m|}\sum_{i\in\mathcal{M}_m} rev_{i,t} - \sum_{i\notin\mathcal{M}_m}\hat{\omega}_i \cdot rev_{i,t}\right)$$

The first term represents the cross-sectional difference between the revenues of the merging parties and the weighted control units in the post-merger year k. The second term acts as a time-invariant adjustment; it subtracts the weighted average gap between the treatment and control groups during the pre-merger period. By using the estimated time-weights $\hat{\lambda}_k$, this adjustment ensures that pre-existing level differences are netted out, specifically those in pre-merger years most relevant for predicting the outcome.

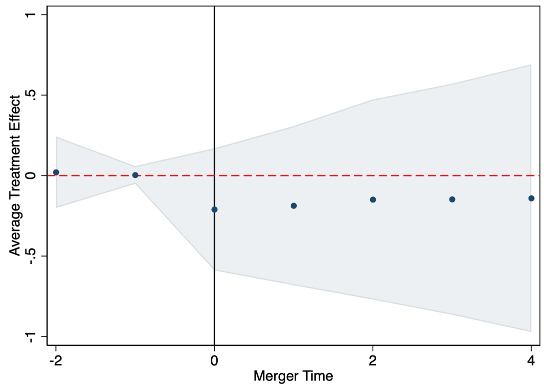

(a) Full sample

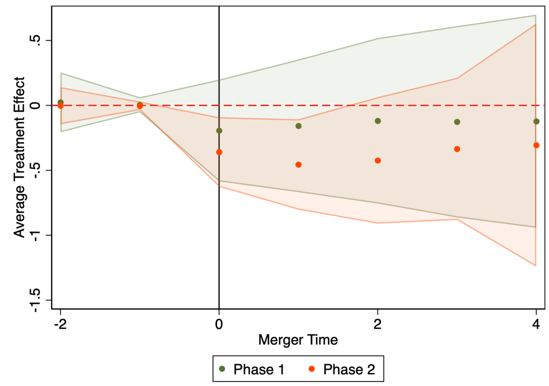

(b) Sample Split by Procedure

Figure 3: Event-Study Estimation of Merger Effects

Figure 3a shows the event-study decomposition of the SDID treatment effects for the full merger sample. It illustrates that the merger effects on the revenues realize quickly and are stable over the post-merger periods. However, the standard errors are large relative to the estimated treatment effects. In Figure 3b, we show the treatment effects for Phase 1 and Phase 2 mergers separately. In line with the aggregated results from Table 3, we see a stronger negative effect throughout for Phase 2 mergers.

Next, we investigate the heterogeneity of the treatment effects further. To do so, we follow Bhattacharya et al. (2023) and split the sample into quartiles based on the estimated treatment effects.[14] Figure 4 illustrates the heterogeneity clearly. The lowest quartile depicted in blue shows an immediate and stable and significantly negative treatment effect. For the medium quartiles the effect is negligible throughout. Only for the highest quartile, there appears to be a positive (though insignificant) effect that grows over time.

The heterogeneity in the results highlights that there is no clear evidence of the approval policy in the EU being systematically too lax. Taking our results on compensating synergies (Section 3.3) and our estimated CS effects together, it appears as if the vast majority of mergers generates compensating synergies sufficient to overcome the negative CS effects of weaker competition. Yet, there is

[14]Note that this split of the sample is based on the estimated treatment effects and therefore purely descriptive.

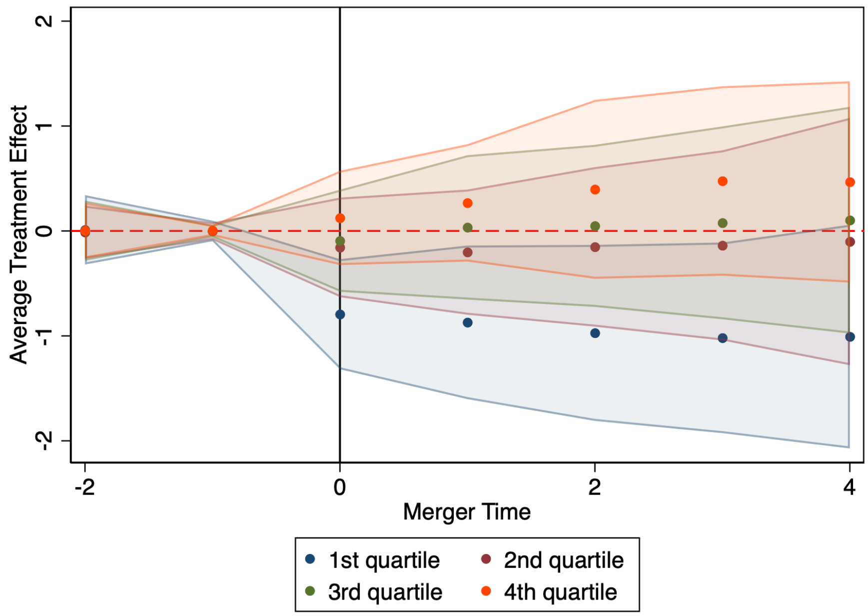

Figure 4: Event-Study Estimation of Merger Effects by Quartiles

no non-negligible set of mergers, particularly Phase 2 mergers, which, despite the Commission's investigations, were harmful to consumers.

Remark. A central challenge in interpreting the treatment effects for Phase 2 mergers is the prevalence of structural remedies. The EC frequently conditions merger approval on the divestiture of assets. In such cases, a negative treatment effect on the merging firms' revenues may be a mechanical accounting artifact of the divestiture. To address this, we plan to use our detailed data extracting the EC decision reports. We will re-estimate our baseline model by (i) examining the subsample of mergers cleared without structural remedies, and (ii) including the full set of competitors whenever available, as these are frequently the purchasers of the divested assets.

5 Predicting Consumer Surplus Effects

In this section, we use our estimated CS effects to gain additional insights on the fundamentals that correlate with negative outcomes. In particular, we assign each merger to one of three groups: (i) anticompetitive, i.e., mergers with negative and significant CS effect, (ii) neutral, i.e., mergers with insignificant CS effect, and (iii) procompetitive, i.e., mergers with positive and significant CS effect.

We then regress these discrete outcomes on ex-ante observables of the mergers to see which, if any of them predict the outcome on consumer surplus. Natural candidates for the explanatory variables are the screening measures applied by competition authorities, such as pre- and post-merger market shares and HHIs as well as naively-computed changes in the HHI, as suggested in the academic literature (see, for example, Nocke and Whinston, 2022). Both the European Commission's and the U.S. authorities' (FTC and DOJ) horizontal merger guidelines specify thresholds above which they may have competition concerns.

With our broad ex-post assessment of merger cases and the detailed data on pre-merger characteristics of markets and merging firms, we can test which of these thresholds are suited to actually predict CS effects. However, we also have information beyond the simple thresholds and data on the merger procedure that we can include in order to assess the effectiveness of various tools in preventing harmful outcomes for consumers.

5.1 The Effectiveness of Concentration Thresholds

In the following, we will run probit regressions at the case level of a binary variable indicating a significant negative consumer surplus effect on various pre-merger observables. Due to the substantial heterogeneity in the number of affected markets, we include the number of affected markets as a control in all specifications and use different percentile levels of the explanatory variables in the regressions.

Table 4 shows the results of these regressions on HHI and market share statistics available before the merger. The table shows no significant results, indicating that, conditional on approval by the EC, these frequently-used statistics do not provide additional explanatory power for negative consumer surplus effects.

Table 5 shows the results for the pre-merger classifications of the EU and US authorities. The results indicate that the screening measures employed by the

Negative CS Effect	(1) HHI Median	(2) HHI 90th	(3) MS Median	(4) MS 90th
Median Post-Merger HHI	-0.0000230	-0.0000458		
	(0.0000493)	(0.0000446)		
Median Post-Merger ΔHHI	0.000233	0.000115		
	(0.000320)	(0.000133)		
Joint Market Share			0.180	-0.137
			(0.618)	(0.500)
No. Affected Markets	-0.00228	-0.00219	-0.000842	-0.000199
	(0.00510)	(0.00497)	(0.00528)	(0.00531)
Constant	-0.821***	-0.666**	-0.939***	-0.828***
	(0.257)	(0.286)	(0.235)	(0.254)
N	176	176	156	156

Standard errors in parentheses
* p<0.10, ** p<0.05, *** p<0.01

Table 4: Pre-merger concentration measures and post-merger consumer outcomes.

authorities pick some relevant anti-competitive merger properties up. While the EU's safe harbor indicator (a joint post-merger market share of less than 25%) is indeed associated with a lower likelihood of negative consumer surplus effects, markets that would be classified as highly concentrated according to the 1982 DOJ merger guidelines ($\Delta HHI > 100$ and post-merger $HHI > 1800$) are associated with a higher likelihood of negative consumer surplus effects. The revised 2010 classification is more lenient ($\Delta HHI > 200$ and post-merger $HHI > 2500$) and, in fact, does not pick up any association anymore.

Overall, these results indicate that screening mergers based on simple indicators and without an in-depth investigation is non-trivial. In particular, as many mergers affect a wide class of product and geographic markets. However, there is some evidence that a combination of the EU safe harbor requirement and a combination of post-merger HHI and ΔHHI statistics can give indications on a merger's relative safety or potential to harm consumers.

Negative CS Effect	(1) EU	(2) US 1982	(3) US 2010
EU Safe	-0.572*		
	(0.335)		
EU Problematic	-0.312		
	(0.373)		
US 1982 High		0.530*	0.406
		(0.296)	(0.379)
US 1982 Moderate		0.604	0.333
		(0.487)	(0.460)
No. of Affected Markets	-0.00114	-0.00337	-0.00361
	(0.00545)	(0.00542)	(0.00551)
Constant	-0.464	-1.143***	-1.006***
	(0.302)	(0.190)	(0.168)
N	151	176	176

Standard errors in parentheses
* p<0.10, ** p<0.05, *** p<0.01

Table 5: EC and FTC/DOJ Classifications of mergers and consumer outcomes.

6 Merger Effects on Employment

While our analysis has so far focused on the CS effects of mergers, a common concern is that mergers lead to a reduction in employment. With our data and empirical strategy, we can directly test this hypothesis. Following the same approach as in Section 4.1, we estimate the SDID treatment effects and their event-study decomposition for the employment variable extracted from ORBIS, which captures the number of workers in the company.

Table 6 shows that the results are mixed. Around 80% of mergers have an insignificant impact on employment. While the amount of positively significant mergers is with 5% at the level to be expected by chance, the 14% of mergers with a negatively significant impact on employment are more systematic but small

Sample	TE	SE	% Impact	% Pos. Sign.	% Neg. Sign.	N
All	-0.075	0.361	-0.483	5.88	13.97	136
Phase 1	-0.063	0.373	-0.291	5.79	11.57	121
Phase 2	-0.175	0.263	-2.029	6.67	33.33	15

Table 6: Summary table of estimated treatment effects on employment. TE reports the average treatment effect $\tau = \frac{1}{M}\sum_{m=1}^{M} \tau^m$ and SE the corresponding standard errors computed using the placebo approach. % Impact translates the treatment effect into percentage changes in pre-merger employment. % Pos. Sign. and % Neg. Sign. indicate the share of cases in which we find positively and negatively significant treatment effects at the five percent level, respectively.

in magnitude. Figure 5 shows the event-study decomposition that confirms the estimated treatment effects. After an initial slight bump the employment effects remain close to zero for both Phase 1 and Phase 2 mergers.

Breaking the estimated treatment effects up by quartiles, we see in Figure 6 that the averages close to zero effects hide some heterogeneity, although the effects remain moderate. The highest quartile has reveals a positive, though insignificant, effect with a slight upward trend. The lowest quartile has a negative effect that remains constant over the post-merger periods.

7 Conclusion

This paper develops a scalable, sufficient-statistics approach to ex-post evaluation of the consumer-surplus effects of mergers. Building on recent theory on multiproduct-firm oligopoly, we show that under CES demand the sign of a merger's consumer-

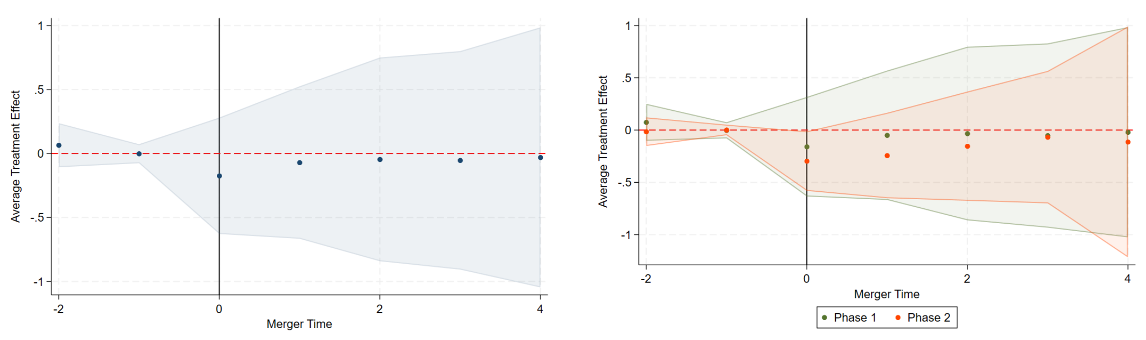

(a) Full sample (b) Sample Split by Procedure

Figure 5: Event-Study Estimation of Merger Effects on Employment

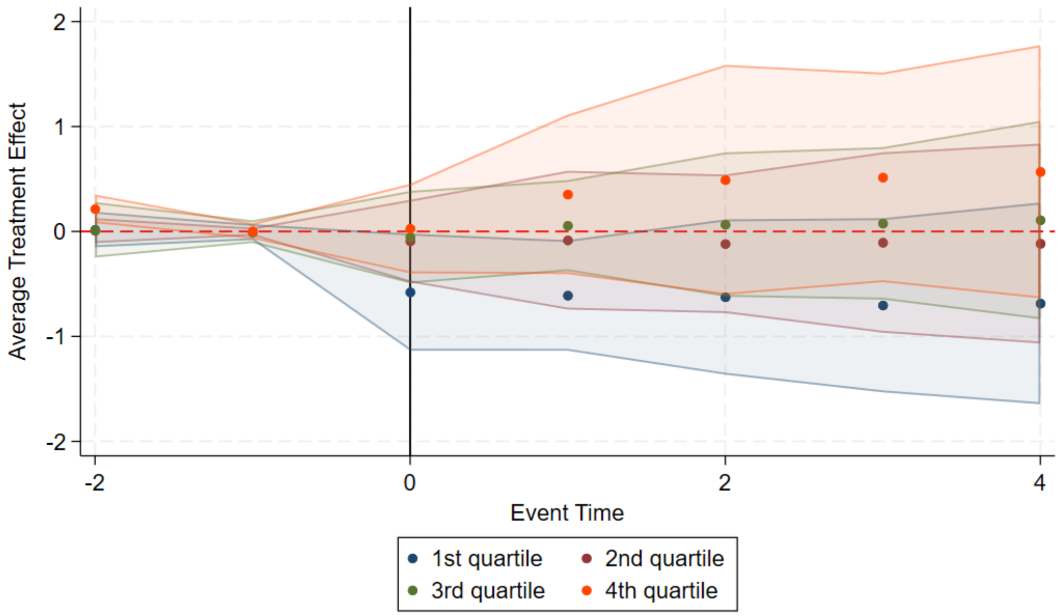

Figure 6: Event-Study Estimation of Employment Effects by Quartiles

surplus effect is identified by a single observable: the change in the merging parties' combined revenues. In this environment, consumer-surplus improvements correspond to a reallocation of expenditure toward inside goods, and, because the merger leaves outsiders' fundamentals unchanged, such a reallocation necessarily manifests itself in a higher revenue share for the merged entity. The same logic extends beyond the baseline horizontal setting: we provide an upstream-input-market extension in which the revenue-based sign test continues to apply.

We implement this diagnostic in a novel dataset that links European Commission merger decisions to firm-level accounts. We code the relevant market structure and procedural information from EC merger reports, and then identify merging parties in the ORBIS accounting database. While ORBIS coverage prevents us from following every reviewed merger, we recover the revenue panels needed for 202 mergers between 2002 and 2021—42.5% of the cases we coded, spanning 58% of two-digit NACE industries, and we show that matched and unmatched cases are broadly comparable on observables available in the reports.

We also compute the ex-ante compensating cost synergies required for a merger to be consumer-surplus neutral. The implied magnitudes are economically meaningful even for Phase 1 cases: on average, Phase 1 mergers require roughly 6–11% cost reductions to break even on consumer surplus, while Phase 2 mergers require about 14–24%. In the ex-post analysis, we estimate merger-induced revenue changes using a synthetic difference-in-differences design. We find that the average merger is close to consumer-surplus neutral, but heterogeneity is substantial: around 20% of approved mergers exhibit significantly negative consumer-surplus effects, with harm more common among Phase 2 cases (six out of 15 in the baseline sample). While there is a non-negligible amount of negative consumer-surplus effects, it is worth noting that around 80% of mergers generate sufficient synergies to overcome the anticompetitive market-power effects.

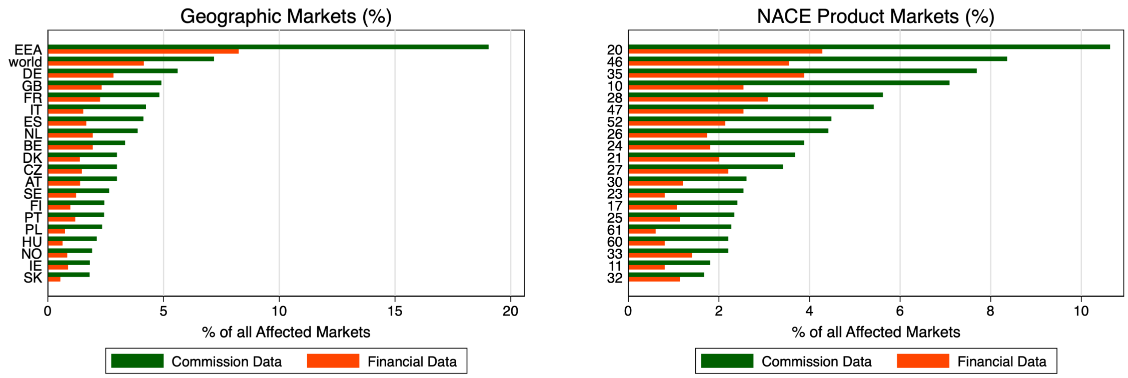

(a) Coverage by Geographic Market

(b) Coverage by Industry

Figure 7: Coverage of Merger Cases with ORBIS Data.

APPENDICES

A Data Coverage with ORBIS

In Figure 7a, we show our coverage of merger cases across geographic markets. The green bars refer to cases in our data set based only on the EC reports, while the orange bars indicate the share of those markets in our final sample, including the financial data. While the coverage varies, our sample does not appear to be substantially biased across geographical regions.

In Figure 7b, we display the share of the 15 most frequently affected two-digit industry NACE codes from our data set based on the EC reports with green bars. Orange bars indicate the analogous coverage in our final data set. While the coverage varies across industries, our final data set appears to cover the majority of industries well and not to be skewed heavily towards particular industries.

Figure 8 shows that our coverage of merger cases with financial data relative to the total number of merger cases collected is relatively stable over time. On average, we obtain financial data for more than half of the mergers in any given year. Note that our sample coverage drops drastically after the application year 2015. We require several years of data post-merger and data in ORBIS only fills up over time.

Table 7 illustrates that the mergers for which we obtain financial data are comparable to the full sample of merger cases in our dataset. The column *Total* refers to all cases on which we obtained information, the column *Subset* refers to

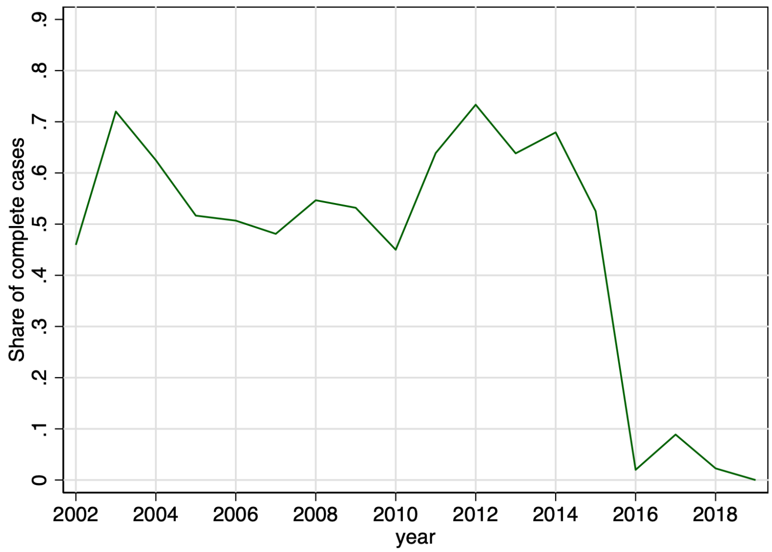

Figure 8: Coverage of Merger Cases by Application Year.

	Phase 1				Phase 2			
	Total		Subset		Total		Subset	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD
N Affected Markets								
Geographic Markets	8.28	8.82	6.68	6.71	21.62	10.48	15.17	8.63
Product Markets	9.86	11.74	10.20	13.77	64.93	55.01	18.23	14.44
Combined	30.34	40.36	22.73	24.05	252.77	201.49	85.65	71.12
Market Shares								
Acquirers	0.16	0.16	0.17	0.16	0.24	0.22	0.24	0.19
Targets	0.11	0.16	0.11	0.15	0.24	0.23	0.26	0.25
All	0.28	0.19	0.28	0.18	0.46	0.27	0.50	0.26
HHI								
Pre-Merger HHI	0.43	0.26	0.44	0.25	0.38	0.19	0.34	0.18
Post-Merger HHI	0.45	0.26	0.46	0.25	0.45	0.20	0.43	0.21
Delta HHI	0.02	0.04	0.02	0.03	0.07	0.08	0.09	0.09

Table 7: Summary Statistics of Merger-Specific Information from EC Merger Decisions. The table compares the information available from the EC decision reports for our full sample of merger cases and the subset of merger cases for which we obtain financial data.

the cases on which we obtained financial data on ORBIS. The market shares and HHIs are very similar in both samples. The number of affected markets is also comparable for Phase 1 mergers in both samples. For the affected markets in Phase 2 mergers, however, the subset of cases that we can use in our empirical analysis is restricted to mergers with relatively few markets, compared to the full sample of Phase 2 mergers.

B HHI Computation from Market Shares

As the EC decision reports contain only information on market share intervals of the main competitors, we need to impose additional structure when computing HHIs from these. We first provide several alternative assumptions before discussing their empirical implications.

B.1 Theoretical Specifications

We compare three different approaches for the computation of pre- and post-merger HHIs. First, note that typically the reported market shares, even when choosing the upper bound of all firms' market share intervals, the sum of those market shares does not add up to one. In these cases, we need to decide how to allocate the remaining market shares across firms. Due to the interval-reporting, we can compute an interval of potential missing market shares. To decide on scalar market shares, we follow Mini (2018)'s approach. We simulate a distribution of market shares by uniformly sampling from the reported intervals, and keeping only feasible realizations, i.e., those with a total market share of less than or equal to one. For all realizations strictly below one, we record the missing market share. From those feasible realizations, we compute each firm's expected market share using the empirical distributions.

Second, each time that there is a missing market share, we need to decide how to allocate it to potential non-reported competitors. We suggest three procedures as outlined in Table 8. In specification HHI_1, we allocate the remaining market shares proportionally to the firms mentioned in the competition report. In specification HHI_2, we allocate the remaining market share to one additional (possibly large)

Specification	Description
HHI_1	merging parties + reported competitors
HHI_2	merging parties + reported competitors + one large competitor
HHI_3	merging parties + reported competitors + competitive fringe

Table 8: HHI Computation: Missing Market Shares.

competitor. In specification HHI_3, we allocate the remaining market share to a competitive fringe with 1% of the missing market share each.

Hence, we obtain the following HHI formulas for each specification, where M denotes the number of merging parties and C the number of main competitors mentioned in the EC report.

$$HHI_1 = \sum_{m=1}^{M} \alpha_m^2 + \underbrace{\sum_{c=1}^{C} \alpha_c^2}_{\text{reported competitors}} \tag{13}$$

$$HHI_2 = \sum_{m=1}^{M} \alpha_m^2 + \sum_{c=1}^{C} \alpha_c^2 + \underbrace{\left(1 - \sum_{m=1}^{M} \alpha_m - \sum_{c=1}^{C} \alpha_c\right)^2}_{\text{one large competitor}} \tag{14}$$

$$HHI_3 = \sum_{m=1}^{M} \alpha_m^2 + \sum_{c=1}^{C} \alpha_c^2 + \underbrace{100 \cdot \left(\left(1 - \sum_{m=1}^{M} \alpha_m - \sum_{c=1}^{C} \alpha_c\right) \cdot 0.01\right)^2}_{\text{competitive fringe}} \tag{15}$$

Our preferred specification that we apply in the main text is specification HHI_3, as it ensures that the resulting market shares of the firms mentioned in the EC decision remain in the intervals reported (which is not guaranteed under specification HHI_1). Moreover, it ensures that there is no large competitor that appears to be "overlooked" by the EC (which may happen under specification HHI_2). In Figure 9 we compare the resulting distributions of HHIs. Note that specification HHI_2 provides an upper bound for the true underlying market share and HHI_3 the corresponding lower bound.

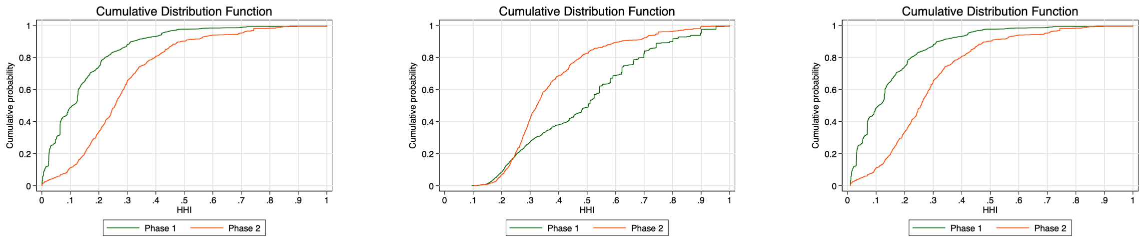

(a) Specification HHI_1 (b) Specification HHI_2 (c) Specification HHI_3

Figure 9: Comparison of HHI Computation.

C Compensating Synergies

In this section, we also report the compensating *type* synergies rather than the cost synergies, required to ensure that a merger is consumer-surplus neutral, as derived in Nocke and Whinston (2022). The formula for general type synergies (without imposing assumptions on which synergies obtain), is

$$TS = \frac{s_{\mathcal{M}}\left(\sigma + \frac{s_{\mathcal{M}}}{1-s_{\mathcal{M}}}\right)^{\sigma-1}}{\sum_{f\in\mathcal{M}}\left(s_f\left(\sigma + \frac{s_f}{1-s_f}\right)^{\sigma-1}\right)}.$$

For our merger sample, we obtain the compensating type synergies as reported in Table 9. While the type synergies are difficult to interpret economically, it is apparent that substantial synergies are necessary to ensure CS neutrality. However, the table shows that there is substantial heterogeneity across mergers.

D Robustness Checks

In this section, we provide several robustness checks to illustrate that our results are not driven by the specific choices of our baseline specification.

D.1 Different Time Horizon

In principle, by choosing different pre-merger windows, the construction of the counterfactual revenue evolution changes. Furthermore, by choosing different

	Phase 1		Phase 2	
	Mean	SD	Mean	SD
Merger Level				
$\sigma = 4$	5.09	24.18	3.09	3.10
$\sigma = 5$	10.94	65.82	4.69	6.78
$\sigma = 6$	23.06	154.17	7.40	13.93
$\sigma = 7$	45.60	320.58	11.81	26.49
Market Level				
$\sigma = 4$	22.03	245.00	2.96	7.57
$\sigma = 5$	57.81	681.48	4.60	17.92
$\sigma = 6$	133.43	1608.90	7.52	38.53
$\sigma = 7$	275.81	3359.28	12.43	75.31

Table 9: Compensating Type Synergies

$[-t, t]$	N	TE	SE	% Impact	% Pos Sign.	% Neg. Sign.
[-2, 2]	204	-0.186	0.262	-0.798	4.90	18.14
[-2, 3]	173	-0.165	0.275	-0.705	2.89	18.50
[-2, 4]	152	-0.166	0.297	-0.712	3.29	20.39
[-2, 5]	138	-0.157	0.321	-0.667	6.52	18.12
[-3, 3]	156	-0.175	0.274	-0.766	3.21	21.15
[-3, 4]	135	-0.174	0.293	-0.759	2.22	17.78
[-3, 5]	121	-0.173	0.314	-0.749	4.96	15.70
[-4, 4]	122	-0.174	0.289	-0.755	5.74	17.21
[-4, 5]	108	-0.173	0.335	-0.747	4.63	15.74
[-5, 5]	90	-0.218	0.345	-0.921	7.78	21.11

Table 10: Estimation Results for Different Time Windows

post-merger windows, the time horizon over which merger effects can realize or disappear changes.

Table 10 highlights that our estimation results are robust to different specifications of both the pre- and post-merger time windows. As the number of mergers considered differs in Table 10, we illustrate the density of the estimated treatment effects *on the same sample* in Figure 10 for different time windows. As is evident from the figures, the estimated treatment effects are very similar, no matter the time window.

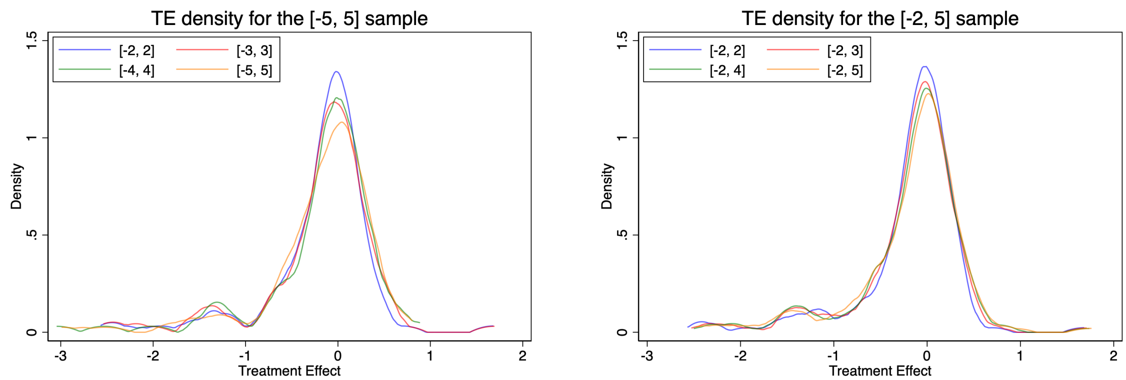

(a) Symmetric time windows (90 cases) (b) Varying post-merger time (139 cases)

Figure 10: Comparison of Treatment Effects Across Time Windows. Each specification of the time window is computed on the same sample that is available for all specifications of the time window.

D.2 Consolidated Accounts

In our baseline analysis, we include both cases involving unconsolidated and consolidated accounts. Recall that consolidated accounts include the mother company, including its subsidiaries, while unconsolidated accounts contain the merging entity alone. We see that the estimated treatment effect is more negative for consolidated cases. The sign is negative for both account types, however, and our main interest lies in the sign of the treatment effect. The main difference lies in the share of treatment effects that is significantly negative, which is lower for unconsolidated accounts. The sample of unconsolidated accounts is less than half the one for consolidated accounts, however. In all, while the qualitative picture across the account types is similar, we do not know to which extent the differences arise from structural differences in consolidated and unconsolidated accounts, or a selection of the merger cases for which we obtain the respective account types.

D.3 Single-Entity Analysis

In our baseline specification, we treat the merging firms as independent entities in the pre-merger period, and let the SDID procedure recover a separate counterfactual revenue path for each firm. The merger effect is then defined as the average

	Consolidated				Unconsolidated			
	Affected		Problematic		Affected		Problematic	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD
N Markets								
Geographic Markets	4.27	5.39	0.79	3.44	2.68	3.93	0.06	0.24
Product Markets	6.62	8.91	0.71	2.60	5.51	6.93	0.19	0.94
Combined	13.91	23.73	0.95	4.25	8.04	11.54	0.07	0.43
Market Shares								
Acquirers	0.20	0.18	0.22	0.16	0.16	0.17	0.59	0.25
Targets	0.20	0.22	0.47	0.28	0.14	0.21	0.17	0.16
All	0.40	0.24	0.70	0.24	0.29	0.23	0.76	0.13
HHI								
Pre-Merger HHI	0.22	0.17	0.38	0.18	0.16	0.21	0.44	0.27
Post-Merger HHI	0.27	0.21	0.52	0.21	0.18	0.22	0.58	0.19
Delta HHI	0.05	0.07	0.14	0.10	0.02	0.04	0.14	0.14
Revenues (Billion Euros)								
Operating Revenues	26.10	32.16	53.94	63.99	9.32	16.29	18.72	30.78

Table 11: Summary Statistics of Merger Cases. We report descriptive statistics for the markets associated with the mergers in our sample, separately for cases with consolidated and unconsolidated accounts, distinguishing between all affected markets and those classified as problematic. For each subsample, we report the mean and standard deviation of the number of geographic and product markets (and their sum), the market shares of acquirers and targets (and their combined share), as well as the pre-merger, post-merger, and change in the Herfindahl-Hirschman Index. The bottom row shows the operating revenues (in billion euros) of the merging parties.

Sample	TE	SE	% Impact	% Pos Sgn.	% Neg. Sign.	N
All	-0.166	0.297	-0.712	3.29	20.39	152
Cons	-0.219	0.277	-0.947	0.95	22.86	105
Uncons	-0.048	0.340	-0.187	8.51	14.89	47

Table 12: Estimated Treatment Effects by Account Type.

treatment effect across these firms. As an alternative, one could instead aggregate the acquirer and target into a single hypothetical merged firm by summing their yearly revenues, and apply SDID to this aggregate. In that case, the estimator delivers a merger-specific treatment effect on the sum of the parties' revenues rather than an average effect at the firm level.

Sample	TE	SE	% Impact	% Pos Sgn.	% Neg. Sign.	N
All	-0.179	0.383	-0.757	3.90	12.34	154
Phase 1	-0.167	0.385	-0.711	4.32	12.23	139
Phase 2	-0.292	0.364	-1.181	0.00	13.33	15

Table 13: Estimated Treatment Effects for Revenue Sum.

Filter	TE	SE	% Impact	% Pos. Sign.	% Neg. Sign.	N
Employment	-0.190	0.306	-0.814	3.95	16.45	152
3-Digit NACE Code	-0.279	0.497	-1.249	3.29	14.47	152
Fixed Assets	-0.252	0.366	-1.158	3.27	17.65	153

Table 14: Estimated Treatment Effects for Different Control Filters.

The results for this alternative specification are shown in Table 13. They are broadly comparable to those from our baseline specification. The main difference is that the share of significantly negative treatment effects is somewhat smaller in the single-entity analysis. It is not clear, however, that this alternative counterfactual is economically more relevant, as it is based on a hypothetical pre-merger aggregation of the merging parties.

D.4 Filtering of Control Group

While in our baseline specification, we apply all three filters to reduce the set of potential control firms, we show in Table 14 that our results are robust to applying each of the filters separately.

D.5 Difference-in-Differences and Synthetic Control

We compare our results obtained from the synthetic difference-in-differences estimation that follows Arkhangelsky et al. (2021) with the traditional difference-in-differences (DID) and synthetic control (SC) approach. The comparison of results is in Table 15 for different time windows. It appears that both SDID and DID obtain similar results both in terms of sign, magnitude, and standard errors. The

		SDID		*DID*		*SC*	
$[-t,t]$	N	TE	SE	TE	SE	TE	SE
$[-2,2]$	203	-0.180	0.256	-0.198	0.308	0.125	0.35
$[-2,3]$	174	-0.172	0.272	-0.181	0.310	0.137	0.35
$[-2,4]$	153	-0.173	0.289	-0.186	0.326	0.228	0.34
$[-2,5]$	138	-0.157	0.325	-0.182	0.367	0.270	0.39
$[-3,3]$	156	-0.175	0.286	-0.194	0.323	0.216	0.36
$[-3,4]$	135	-0.174	0.291	-0.210	0.359	0.313	0.38
$[-3,5]$	121	-0.172	0.320	-0.217	0.389	0.334	0.41
$[-4,4]$	122	-0.174	0.296	-0.213	0.379	0.342	0.40
$[-4,5]$	108	-0.174	0.330	-0.222	0.420	0.368	0.43
$[-5,5]$	90	-0.218	0.337	-0.242	0.421	0.318	0.43

Table 15: Comparison of ATT Estimates of Different Methods

SC yield very different results, however; in particular, positive treatment effects. We interpret this as evidence in favor of our SDID approach. We study, at least predominantly, the largest companies in the European Union, as we are focusing ourselves on merging firms that are investigated by the EC following the merger notification. For the SC approach, the control firms need to be of comparable size as the matching is based on revenue levels. In constrast to that, SDID and DID include unit fixed effects that net out the levels and instead match the changes over time. Now, if the firms in the control group are systematically smaller than the treated units, the SC counterfactual will lie below the treated series both before and after the merger, and if this pre-merger gap is not taken into account, it can *underestimate* the revenues counterfactual post-merger. Hence, mechanically it will tend to provide a positive treatment effect.

E Upstream Revenue as Sufficient Statistic

In this appendix, we formally demonstrate that the revenue sufficient statistic obtained in Proposition 1 extends to upstream horizontal mergers where the merging firms sell intermediate inputs to downstream producers. We consider a vertical structure where upstream firms supply inputs to downstream firms, who in turn sell to final consumers.

E.1 Framework, Upstream Aggregator and Consumer Surplus

Final Demand. Considerkg a representative consumer with quasilinear preferences over a set of differentiated final goods $\mathcal{N}^D$. The indirect utility is $V(y,p) = y + \frac{1}{\eta}\ln H_D(\mathbf{p})$, where H_D is the CES aggregator for the downstream sector

$$H_D(\mathbf{p}) \equiv \sum_{j\in\mathcal{N}^D} b_j p_j^{1-\sigma},\ \sigma > 1,\ b_j > 0. \tag{16}$$

As in the baseline model, total expenditure on final goods is exogenous $M \equiv (\sigma-1)/\eta$. The Marshallian demand for product j is $D_j(\mathbf{p}) = \frac{b_j p_j^{-\sigma}}{H_D(\mathbf{p})}M$.

Downstream Production. Downstream firms $g \in \mathcal{F}^D$ produce final goods using a composite of differentiated upstream inputs $k \in \mathcal{N}^U$. We impose a *common CES input bundle* assumption. The unit cost of this common bundle is equal to the upstream price index

$$P_U(\mathbf{w}) = \left(\sum_{k\in\mathcal{N}^U} a_k w_k^{1-\rho}\right)^{\frac{1}{1-\rho}},\ \rho > 1,\ a_k > 0, \tag{17}$$

where w_k is the wholesale price of input k. The marginal cost of producing the final good j is $c_j(\mathbf{w}) = P_U(\mathbf{w})/\phi_j$, where ϕ_j is a product-specific productivity parameter.

Analogous to the baseline model, we define the upstream aggregator as $H_U(\mathbf{w}) \equiv \sum_{k\in\mathcal{N}^U} a_k w_k^{1-\rho}$. Note that $P_U(\mathbf{w}) = H_U(\mathbf{w})^{-\frac{1}{\rho-1}}$. Since $\rho > 1$, a higher upstream aggregator H_U corresponds to a lower input price index P_U.

Vertical Transmission. We first establish the link between the upstream aggregator and consumer welfare under an assumption on the relation between downstream prices and the input costs. Lateron, we microfound this relation.

Lemma 2. *Suppose downstream equilibrium prices are proportional to marginal costs, $p_j(\boldsymbol{w}) = \mu_j c_j(\boldsymbol{w})$, where the markup vector $\boldsymbol{\mu}$ is invariant to the upstream merger. Then, final consumer surplus is strictly increasing in the upstream aggregator H_U.*

Proof. First, we substitute the pricing rule and the marginal cost function into the downstream aggregator to obtain

$$H_D(\mathbf{p}(\mathbf{w})) = \sum_{j\in\mathcal{N}^D} b_j \left(\mu_j \frac{P_U(\mathbf{w})}{\phi_j}\right)^{1-\sigma}$$
$$= P_U(\mathbf{w})^{1-\sigma} \underbrace{\sum_{j\in\mathcal{N}^D} b_j \mu_j^{1-\sigma} \phi_j^{\sigma-1}}_{\equiv C>0}. \tag{18}$$

Substituting $P_U(\mathbf{w}) = H_U(\mathbf{w})^{-\frac{1}{\rho-1}}$, we obtain:

$$H_D(\mathbf{p}(\mathbf{w})) = C \cdot H_U(\mathbf{w})^{\frac{\sigma-1}{\rho-1}}. \tag{19}$$

Since C is positive and because $\sigma > 1$ and $\rho > 1$, the exponent is positive. Thus, H_D is strictly increasing in H_U. Since consumer surplus is strictly increasing in H_D by definition of the indirect utility function, the result follows. □

E.2 Upstream Revenue as Sufficient Statistic

Let $\mathcal{M} \subset \mathcal{F}^U$ be the set of merging firms in the upstream market. Total demand for input k is derived via Shephard's Lemma

$$X_k(\mathbf{w}) = \frac{a_k w_k^{-\rho}}{H_U(\mathbf{w})} M_U(\mathbf{w}), \tag{20}$$

where $M_U(\mathbf{w}) \equiv \sum_{k\in\mathcal{N}_U} w_k X_k(\mathbf{w})$ is the total expenditure by the downstream sector on upstream inputs. The upstream game is a standard CES pricing game conditional on the market size $M_U(\mathbf{w})$. Denote by $\overline{R}_\mathcal{M}$ the post-merger revenue of the merged entity and by $R^*_\mathcal{M}$ the sum of pre-merger revenues of the merging firms.

The next proposition generalizes Proposition 1 to the vertical supply chain, assuming that the upstream market potential M_U does not vary with the merger. Under this assumption, the upstream pricing game is isomorphic to a standard CES oligopoly as in our baseline model.

Proposition 2. *If the total downstream expenditure on inputs, M_U, is invariant to the* *merger, then*

$$\overline{R}_{\mathcal{M}} > R^*_{\mathcal{M}} \iff \Delta CS > 0. \tag{21}$$

Proof. The CES oligopoly with fixed market potential M_U is identical to the one in our baseline model and the equilibrium market share of any non-merging firm $g \notin \mathcal{M}$ is given by $s_g^U(H_U) = S_U(T_g^U/H_U)$. Since the fitting-in function $S_U(\cdot)$ is strictly increasing (by Lemma 1), the outsider market share s_g^U is strictly decreasing in the level of the aggregator H_U.

As usual, the sum of market shares must equal unity

$$s_{\mathcal{M}}^U(H_U) + \sum_{g \notin \mathcal{M}} s_g^U(H_U) = 1. \tag{22}$$

Since the sum of outsider shares is strictly decreasing in H_U, the combined market share of the merging firms, $s_{\mathcal{M}}^U$, must be strictly increasing in H_U.

Total revenue for the merging parties is $R_{\mathcal{M}} = s_{\mathcal{M}}^U M_U$. Since M_U is constant across the pre- and post-merger equilibria, we obtain

$$\overline{R}_{\mathcal{M}} > R^*_{\mathcal{M}} \iff \overline{s}_{\mathcal{M}}^U > s_{\mathcal{M}}^{U,*} \iff \overline{H}_U > H_U^*. \tag{23}$$

By Lemma 2, $\overline{H}_U > H_U^*$ implies $\Delta CS > 0$. □

E.3 Microfoundations for Fixed Upstream Expenditure

Proposition 2 relies on the condition that M_U is invariant to the upstream merger. We identify two downstream market structures where this condition holds.

Case A: Downstream Monopolistic Competition. If downstream firms ignore their effect on the aggregate downstream price index, they charge constant markups $\mu^{MC} = \sigma/(\sigma - 1)$. Thus, the total input expenditure, that is, the variable cost, is a constant fraction of consumers' final goods expenditure:

$$M_U = \frac{M}{\mu^{MC}} = \frac{\sigma - 1}{\sigma} M. \tag{24}$$

Case B: Downstream CES Oligopoly. We assume that the downstream market is a multiproduct price-setting oligopoly. An upstream merger alters the input price index P_U. Under our assumption of a common input bundle, this acts as a uniform multiplicative cost shock. All downstream marginal costs are scaled by some factor $\gamma = \overline{P}_U / P_U^*$ (where the numerator is the post-merger input price index and the denominator the pre-merger input price index).

Lemma 3. *In a CES oligopoly with type aggregation, a uniform scaling of all firms' marginal costs by a factor $\gamma > 0$ leaves equilibrium market shares and markups unchanged.*

Proof. Let T_g^D be downstream firm g's type defined as $T_g^D = \sum_{j \in g} b_j c_j^{1-\sigma}$. Under uniform scaling from c_j to γc_j, the new types are

$$\bar{T}_g^D = \sum_{j \in g} b_j (\gamma c_j)^{1-\sigma} = \gamma^{1-\sigma} T_g^D. \tag{25}$$

The equilibrium condition requires $\sum_g S_D(\overline{T}_g^D / \overline{H}_D) = 1$. Denote the post-upstream-merger downstream-aggregator be $\overline{H}_D = \gamma^{1-\sigma} H_D^*$. Then, the argument of the fitting-in function becomes

$$\frac{\overline{T}_g^D}{\overline{H}_D} = \frac{\gamma^{1-\sigma} T_g^D}{\gamma^{1-\sigma} H_D^*} = \frac{T_g^D}{H_D^*}. \tag{26}$$

Since the argument is unchanged relative to the pre-merger case, the equilibrium market shares s_g^D and markups μ_g are unchanged. □

As downstream firm g's expenditure on inputs is equal to its total variable cost $VC_g = R_g^D / \mu_g$, the total upstream market size is

$$M_U = \sum_{g \in \mathcal{F}^D} \frac{s_g^D M}{\mu_g}. \tag{27}$$

Since both s_g^D and μ_g are invariant to the uniform cost shock caused by the upstream merger, due to Lemma 3, total input expenditure M_U is constant. Thus, the revenue sufficient statistic is valid.

References

Abadie, Alberto, Diamond, Alexis, and Hainmueller, Jens (2010). "Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program". In: *Journal of the American statistical Association* 105.490, pp. 493–505.

Abadie, Alberto, Diamond, Alexis, and Hainmueller, Jens (2015). "Comparative politics and the synthetic control method". In: *American Journal of Political Science* 59.2, pp. 495–510.

Affeldt, Pauline (2019). "EU Merger Policy Predictability Using Random Forests". In: *DIW Berlin Discussion Paper*.

Arkhangelsky, Dmitry et al. (2021). "Synthetic difference-in-differences". In: *American Economic Review* 111.12, pp. 4088–4118.

Ashenfelter, Orley C, Hosken, Daniel S, and Weinberg, Matthew C (2015). "Efficiencies Brewed: Pricing and Consolidation in the US Beer Industry". In: *The RAND Journal of Economics* 46.2, pp. 328–361.

Asker, John and Nocke, Volker (2021). "Collusion, mergers, and related antitrust issues". In: *Handbook of industrial organization*. Vol. 5. 1. Elsevier, pp. 177–279.

Bhattacharya, Vivek, Illanes, Gastón, and Stillerman, David (2023). *Merger effects and antitrust enforcement: Evidence from us consumer packaged goods*. Tech. rep. National Bureau of Economic Research.

Ciccia, Diego (2024). "A short note on event-study synthetic difference-in-differences estimators". In: *arXiv preprint arXiv:2407.09565*.

Das, Somnath (2019). "Effect of merger on market price and product quality: American and US airways". In: *Review of Industrial Organization* 55.3, pp. 339–374.

De Loecker, Jan and Warzynski, Frederic (2012). "Markups and Firm-Level Export Status". In: *American Economic Review* 102.6, pp. 2437–71.

European Commission (2004). *Guidelines on the assessment of horizontal mergers under the Council Regulation on the control of concentrations between undertakings*. Communication from the Commission (2004/C 31/03), Official Journal of the European Union C 31, 05.02.2004. Feb. 2004.

Focarelli, Dario and Panetta, Fabio (2003). “Are mergers beneficial to consumers? Evidence from the market for bank deposits”. In: *American Economic Review* 93.4, pp. 1152–1172.

Kim, E Han and Singal, Vijay (1993). “Mergers and market power: Evidence from the airline industry”. In: *The American Economic Review*, pp. 549–569.

Luo, Dan (2014). “The price effects of the Delta/Northwest airline merger”. In: *Review of Industrial Organization* 44.1, pp. 27–48.

Miller, Nathan H and Weinberg, Matthew C (2017). “Understanding the price effects of the MillerCoors joint venture”. In: *Econometrica* 85.6, pp. 1763–1791.

Mini, Federico (2018). “Fifty is the New Forty: EU Merger Policy Permits Higher Market Shares After the 2004 Reform”. In: *Review of Industrial Organization* 53.3, pp. 535–561.

Nocke, Volker and Schutz, Nicolas (2018). “Multiproduct-firm Oligopoly: An Aggregative Games Approach”. In: *Econometrica* 86.2, pp. 523–557.

Nocke, Volker and Schutz, Nicolas (2025). “An Aggregative Games Approach to Merger Analysis in Multiproduct-Firm Oligopoly”. In: *The RAND Journal of Economics* 56.3, pp. 233–250.

Nocke, Volker and Whinston, Michael D (2022). “Concentration thresholds for horizontal mergers”. In: *American Economic Review* 112.6, pp. 1915–1948.

Peters, Craig (2006). “Evaluating the performance of merger simulation: Evidence from the US airline industry”. In: *The Journal of law and economics* 49.2, pp. 627–649.

Stiebale, Joel and Szücs, Florian (2022). “Mergers and market power: evidence from rivals’ responses in European markets”. In: *The RAND Journal of Economics* 53.4, pp. 678–702.

U.S. Department of Justice and Federal Trade Commission (2023). *Merger Guidelines*. Issued December 18, 2023. Dec. 2023.

Werden, Gregory J (2017). “Cross-market Balancing of Competitive Effects: What Is the Law, and What Should It Be”. In: *J. Corp. L.* 43, p. 119.