Product Networks and the Competitive Effects of Acquisitions*

Athos Carvalho[†] Shanglyu Deng[‡] Mario Leccese[§] Marco Loseto[¶]

March 2026

Abstract

We propose a network-based merger screening statistic to assess the competitive effects of acquisitions. We micro-found this measure in a standard differentiated oligopoly model where diversion ratios capture competitive links across products and equilibrium markups decompose into monopolistic and network centrality components. The centrality term summarizes the intensity of competition by capturing product proximity in the characteristic space and determines the unilateral effects of horizontal mergers. Using price and quantity data from the European automobile industry, we find that the network model matches substitution patterns and simulated merger effects from a random coefficients nested logit benchmark. Centrality based screening improves on the 2023 concentration-based presumption and performs comparably to consumer optimal ΔHHI thresholds.

Keywords: Diversion Ratio, Network Centrality, Antitrust, Merger Screening.

JEL Classification: D43, G34, L13, L41

*We are grateful to Lars Hansen, Ali Hortaçsu, Scott Nelson, and Andrew Sweeting for several insightful discussions.

[†]Bocconi University, athos.cavalcante@phd.unibocconi.it.

[‡]University of Macau, sdeng@um.edu.mo.

[§]Boston University, Questrom School of Business, leccese@bu.edu.

[¶]Bocconi University and IGIER, marco.loseto@unibocconi.it.

1 Introduction

The analysis of horizontal mergers requires tools that are both tractable and economically grounded. Traditional concentration indices, such as the Herfindahl-Hirschman Index (HHI) and its change (ΔHHI), are easy to compute and widely used in practice. Yet they do not capture the closeness of competition when products are differentiated and require an ex ante market definition, a step that is often difficult and frequently contested.[1] Structural demand models in the tradition of Berry et al. (1995) offer a richer, micro-founded way to recover substitution patterns and simulate merger effects. However, they require substantial data and are computationally intensive, which limits their scalability and practical use. Moreover, substitution patterns remain tied to modeling choices: estimates depend on the specified choice set and, possibly, on the chosen nesting structure, whose misspecification can lead to biased substitution patterns and merger predictions (Fosgerau et al., 2024).

A related challenge arises in the analysis of minority acquisitions, whereby firms or institutional investors acquire stakes in competitors without obtaining control. Recent research shows that even small shareholdings may soften competition (Azar et al., 2018; López and Vives, 2019), and enforcement agencies have begun scrutinizing such deals more closely, as reflected in the 2023 U.S. Merger Guidelines.[2] Existing tools, however, offer limited guidance: modified concentration indices still rely on an ex ante market definition, and structural demand models become difficult to interpret and implement when ownership links proliferate.

This paper develops a tractable and scalable framework that bridges concentration-based screening tools and structural merger simulation. Leveraging results from network theory, we show that equilibrium markups in oligopoly models admit a (locally) additive decomposition into a monopolistic component and a network centrality component, which captures the intensity of competition and is computed from the matrix of diversion ratios.[3] We leverage this network structure to characterize the competitive effects of acquisitions. Our results show that the effects of ownership changes are summarized by the induced changes in the network centrality of the merging products.

For mergers, we derive closed-form expressions for price effects, consumer surplus changes, and compensating marginal cost reductions (CMCRs), defined as the marginal cost reduction required to keep consumer surplus at its pre-merger level.

[1] As stated by Prof. Jonathan Baker, "the outcome of more cases has surely turned on market definition than on any other substantive issue." See `https://project-disco.org/competition/090518-antitrust-in-60-seconds-market-definition/`. See also Kaplow (2015).

[2] Guideline #11 states: "When an Acquisition Involves Partial Ownership or Minority Interests, the Agencies Examine Its Impact on Competition."

[3] The diversion ratio for products j and k is the fraction of consumers who leave product j after a price increase and switch to product k.

The impact of a merger can be decomposed into two forces. First, there is a direct effect reflecting the diversion between the merging products. Second, there is an indirect effect arising from equilibrium responses: when the merging firms adjust prices, rival firms respond, and these reactions feed back through substitution patterns across the broader set of products in the network. Because ownership links enter directly into the markup system, the framework extends naturally to arbitrary changes in the ownership structure, including those implied by incumbents' acquisitions of minority stakes in rival firms or institutional investors' common ownership.

These results imply that changes in centrality provide a screening statistic that summarizes unilateral effects. Unlike the HHI and ΔHHI, whose micro-foundations for merger screening rely on aggregative environments where substitution patterns are primarily driven by market shares (Nocke and Whinston, 2022), our centrality approach naturally admits a closeness measure of competition between products and thus allows for flexible substitution patterns. This flexibility is of great importance because substitution patterns can vary substantially across products within the same industry.[4] We further demonstrate how ownership changes alter pricing incentives via changes in centrality, driven by how firms internalize diversion across products.

To characterize the diversion network more explicitly, we consider a linear–quadratic demand model in which substitution patterns arise from a network representation of product characteristics. Each product is a node in a characteristic space, and its vector of attributes determines its position in that network. In this environment, the cross-price elasticity between two products, and hence their diversion ratio, is proportional to a weighted inner product of their vectors of attributes. This structure implies that the diversion network, and therefore changes in centrality induced by acquisitions, can be approximated using only observable product characteristics. In many industries, such information is publicly available and can be extracted at scale from websites or other unstructured sources. Recent advances in large language models (LLM) have substantially reduced the cost of acquiring and organizing such information, making it feasible to implement centrality-based merger screening without requiring agencies to seek additional information from firms or third-party market participants.

Lastly, we apply the framework to the European automobile industry and estimate our network model using price–quantity data. We show that the demand parameters are identified and can be estimated using a simple linear IV strategy, under the assumption that unobserved characteristics shift utility additively. The resulting estimates imply a diversion matrix determined by the network representation of product

[4]For example, in the beer market, Fan and Yang (2025) documents limited substitution between craft and non-craft beers despite their presence in the same product category. As a result, ownership changes across these segments may generate small price effects even when concentration measures change substantially.

characteristics, which in turn maps directly into the centrality-based markup system.

Comparing the demand estimates obtained from this framework with those we obtain from traditional discrete-choice models, we find that the network model captures economically meaningful substitution patterns in the data. In particular, it generates diversion ratios similar to those obtained from a Random Coefficients Nested Logit (RCNL) model, capturing stronger substitution within car segments without imposing a nesting structure ex ante. When we simulate all possible mergers across European car manufacturers, we find that predicted price and consumer surplus effects are highly correlated and relatively similar when we rely on estimates from the network model and the RCNL, suggesting that our approach captures the key substitution forces driving unilateral effects.

We also evaluate the role of centrality as a screening statistic. In the automobile industry, changes in centrality provide incremental explanatory power for predicted price effects beyond changes in HHI, indicating that our index captures information about substitution patterns that concentration indices alone neglect. We further assess the performance of centrality-based screening in a policy decision framework in which the regulator seeks to maximize expected consumer surplus from approved mergers. Using demand estimates from the RCNL model, we simulate all possible mergers in the sample under alternative efficiency assumptions (2%, 5%, and 10%) and compare decision rules based on changes in centrality and concentration. We find that centrality-based screening outperforms the concentration-based presumption rules adopted in the 2023 Merger Guidelines. Quantitatively, the welfare gains from centrality screening closely track those achieved by optimally chosen ΔHHI thresholds. However, centrality-based screening does not outperform these optimal ΔHHI thresholds. This is consistent with the view that the automobile industry has relatively well-defined boundaries and segments, with large firms operating across several of them, so that market shares capture much of the relevant substitution structure.

Overall, this exercise allows us to quantify the benefits of incorporating centrality-based rules to complement concentration-based thresholds into merger screening.

Related Literature. A large literature has developed tools for evaluating the unilateral effects of horizontal mergers. Traditional concentration-based measures such as the HHI and ΔHHI remain widely used because they rely only on pre-merger market shares and are straightforward to compute and interpret. Nocke and Whinston (2022) provide a micro-foundation for this practice by showing that, in a homogeneous-good Cournot model, a larger ΔHHI implies a larger CMCR required to offset the merger's upward pricing incentives. However, in differentiated-product oligopoly models the CMCR generally depends not only on concentration changes, but also on the merging parties' shares and demand-side primitives such as elasticities and margins which concentration indices alone cannot capture.[5]

Merger simulation approaches (Nevo, 2000) incorporate detailed demand and cost systems but are computationally demanding and sensitive to functional-form assumptions, such as the curvature of demand (Crooke et al., 1999). This has led to interest in local measures derived from first-order conditions, such as upward pricing pressure indices, which summarize merging firms' incentives to raise prices (Farrell and Shapiro, 2009, 2010; Jaffe and Weyl, 2013).[6] These indices rely on estimating diversion ratios. Although their usefulness is well understood and emphasized in the 2010 U.S. merger guidelines, their practical application remains limited due to the challenges in measuring them empirically (Conlon and Mortimer, 2021): the most natural approach is to estimate a rich demand system, which allows flexible substitution, and use it to compute the implied diversion ratios, but this is difficult to do at scale.[7] Recognizing this, a burgeoning literature has developed data-driven proxies for diversion ratios using second-choice survey data (Conlon et al., 2023), consumer churn (Qiu et al., 2024), and customer-overlap measures (Einav et al., 2026).

We contribute to this literature by developing a tractable framework with closed-form expressions in which the unilateral effects of acquisitions and mergers are determined by the Bonacich centrality of the merging parties in a product network based on diversion ratios. We further show how the diversion ratio matrix can be obtained using product characteristics only. While our approach requires a set of characteristics to construct the product network, recent advances in machine learning make this requirement far less restrictive, as product characteristics can now be extracted directly from unstructured data and used as input in demand estimation (Bajari et al., 2023; Compiani et al., 2025). Relatedly, Magnolfi et al. (2022) show that embeddings derived from online surveys of perceived product distances can be incorporated into random-coefficient logit estimation and yield elasticity estimates similar to those obtained using observable characteristics. This supports the broader idea that what matters for substitution is the relative positioning of products in latent characteristic space. Consistent with this insight, in our network model, any representation of product attributes that preserves the same similarity structure generates identical substitution patterns and competitive effects.

[5] For aggregative demand systems such as logit and CES, Nocke and Whinston (2022) show that the CMCR depends on a function of the merger partners' shares other than ΔHHI itself, and illustrate numerically that ΔHHI can still be informative about merger impacts. More generally, Nocke and Schutz (2025) provide a comprehensive analysis for price-setting models with logit and CES demand, deriving conditions under which a merger raises consumer surplus.

[6] Miller et al. (2016) and Miller et al. (2017) examine how well first-order condition measures predict the price effects that are found in merger simulations, finding that they work well when measures of pass-through are precise.

[7] Conlon and Mortimer (2021) further demonstrate that there is no single measure of diversion, just as there is no single measure of elasticity. Estimates of both diversion and elasticity depend on where one evaluates the demand curve.

A related strand of work studies how partial ownership or common ownership alters firms' pricing incentives. Bresnahan and Salop (1986) show how minority shareholdings can induce partial internalization of rivals' profits, and devise a modified HHI (MHHI) to quantify competitive incentives, although this characterization is specific to Cournot competition. A subsequent literature has focused on the potentially anti-competitive effects of common ownership, fueling an active policy debate (Azar et al., 2018, 2022; Antón et al., 2023).[8] Instead, Backus et al. (2021a) estimate a differentiated-product demand system and use non-nested tests to distinguish standard own-profit maximization from common ownership pricing. Their results show that ownership links can, in principle, generate price effects comparable to a merger, but the conduct test rejects such coordination in their setting. Rather than testing conduct, our framework offers a scalable and complementary tool for characterizing the pricing incentives that arise under partial or common ownership.

We also contribute to the empirical industrial organization (IO) literature on oligopolistic competition with differentiated products. Most empirical IO models build on discrete-choice demand with price competition a la Bertrand on the supply side (Berry, 1994; Berry et al., 1995).[9] These approaches flexibly incorporate unobserved product characteristics and enable counterfactual merger analysis, but realistic substitution patterns typically require random coefficients and non-convex estimation, and the resulting relationships between characteristics, substitution, and markups are only implicitly characterized. Moreover, practitioners must often specify a nesting structure ex ante to capture appropriate substitution patterns, which may not always be straightforward and, if misspecified, can bias estimates (Fosgerau et al., 2024).[10]

This paper is closer in spirit to Feenstra and Levinsohn (1995), who link substitution patterns to distances in characteristic space, but differs by providing an *analytical* characterization of markups rather than an implicit one. By framing differentiated-product competition as a network game, we show that equilibrium markups can be expressed through Bonacich centrality and that cross-price elasticities correspond to weighted inner products of product attributes. This structure accommodates both substitutability and complementarity,[11] allows for unobserved characteristics, delivers intuitive substitution patterns, and can be estimated with straightforward linear IV

[8]For example, the Federal Trade Commission featured a hearing on common ownership in 2018 (available at: `https://www.ftc.gov/system/files/documents/public_events/1422929/ftc_hearings_session_8_transcript_12-6-18_0.pdf`). See also Schmalz (2018) for a review.

[9]Bresnahan (1987) was among the first to estimate a discrete choice model of oligopolistic competition with products that are differentiated along one dimension (i.e., a la Hotelling). For an exception with continuous demand see Dubois et al. (2014).

[10]An emerging literature is studying new approaches to infer nests directly from consumer behavior. For example, Atalay et al. (2023) use co-purchasing patterns.

[11]See Gentzkow (2007) for a discrete-choice model with complementarities; see also Thomassen et al. (2017) and Lee and Allenby (2009) for quadratic-preference models across product categories.

methods.[12]

Lastly, our work relates to a growing literature that applies network theory to study oligopolistic competition. Following Ballester et al. (2006), several papers model differentiated-product competition as a network game. Ushchev and Zenou (2018) analyze price competition when varieties are located on a network and show that equilibrium prices are proportional to a sign-alternating Bonacich centrality. Galeotti et al. (2024) exploit the network structure of Slutsky matrices to characterize optimal tax–subsidy policies. Pellegrino (2025) develops a general equilibrium model in which firms compete a la Cournot or a la Bertrand, and product attributes induce a similarity network, providing empirical estimates of efficiency losses due to market power.[13] We complement this line of work by showing how network representations of competition can be used to characterize the competitive effects of acquisitions and to construct tractable merger screening tools.

2 A Network Approach to Equilibrium Analysis

We consider a market with J products available. The marginal (and average) cost of production is constant, and is denoted by c_j. The demand for product $j \in \{1, \dots, J\}$ at price vector $\boldsymbol{p} = (p_1, \dots, p_J)'$ is $q_j(\boldsymbol{p})$. We assume that the demand function is generated by utility maximization of a representative consumer with indirect utility function $V(\boldsymbol{p}, y)$, where y represents the income of this consumer. Therefore, the Jacobian of the demand function $S(\boldsymbol{p}) := d\boldsymbol{q}(\boldsymbol{p})/d\boldsymbol{p}'$ is symmetric.

In our subsequent analysis, we highlight the role of a network built upon the diversion ratio matrix, defined here:

$$D(\boldsymbol{p}) := I_J - [\text{diag}(S(\boldsymbol{p}))]^{-1} S(\boldsymbol{p}),$$

in determining the equilibrium prices and merger impacts. To see that this is the diversion ratio matrix, notice that $D_{ij} = 0$ for $i = j$ and $D_{ij} = -\frac{\partial q_j/\partial p_i}{\partial q_i/\partial p_i}$ for $i \neq j$ (using symmetry of $S(\boldsymbol{p})$), which is exactly the ratio of how much more of product j would be purchased to how much less i would be consumed when p_i is raised.

[12] Product proximity also matters for estimation. Gandhi and Houde (2023) show that optimal price instruments should reflect distances in product characteristics. In our framework, this intuition follows directly from the closed-form markup system based on product similarity.

[13] See also Ederer and Pellegrino (2025) for an extension incorporating common institutional ownership.

2.1 Independent Ownership

We start with the familiar independent ownership case, where each product is independently owned by a profit-maximizing firm, to illustrate the network approach to equilibrium analysis.

We consider Bertrand competition among the firms, where each firm independently and simultaneously sets the price for the product it owns. It is straightforward to derive the first-order condition (FOC):

$$\boldsymbol{q}(\boldsymbol{p}) + \mathrm{diag}(S(\boldsymbol{p}))(\boldsymbol{p} - \boldsymbol{c}) = 0.$$

We assume that the first-order condition is also sufficient, and a unique solution $\boldsymbol{p}^*$ exists.

Since our interests are in comparative statics and merger analysis, we take the same approach as Galeotti et al. (2024) and work with a locally linearized (or locally linear) demand function around the equilibrium $(\boldsymbol{p}^*, \boldsymbol{q}^*)$. We abuse notation slightly and let S denote the Jacobian of $\boldsymbol{q}(\boldsymbol{p})$ evaluated at $\boldsymbol{p}^*$. Similarly, let D denote the diversion ratio evaluated at $\boldsymbol{p}^*$. Let $\boldsymbol{\alpha} := -S^{-1}\boldsymbol{q}^* + \boldsymbol{p}^*$, then the linearized demand around the original equilibrium is $\boldsymbol{q} = -S(\boldsymbol{\alpha} - \boldsymbol{p})$.

With this linearization, we are ready to show that equilibrium markups admit an intuitive network-based decomposition.[14] For that purpose, we introduce the Bonacich centrality measure, which is defined below following Jackson (2008).

Definition 1 Consider a network with J nodes, adjacency matrix A, node weights $\boldsymbol{u}$, and attenuation factor $\delta > 0$. The J-vector of (weighted) Bonacich centralities $\boldsymbol{b}(A, \delta, \boldsymbol{u})$ is given by

$$\boldsymbol{b}(A, \delta, \boldsymbol{u}) := (I_J - \delta A)^{-1}\delta A\boldsymbol{u} = \sum_{k=1}^{\infty} \delta^k A^k \boldsymbol{u}. \tag{1}$$

The j-th element of $\boldsymbol{b}(A, \delta, \boldsymbol{u})$ summarizes how central node j is in the network. This measure of centrality is widely used in social networks because it captures a node's importance in terms of how connected this node is to others and how connected the nodes it is connected to. According to the definition of Bonacich centrality, a node's importance is a weighted sum of the walks that emanate from it. Moreover, if $\delta \in (0, 1)$, walks of shorter length are weighted more.[15]

The following equilibrium characterization result ensues.

[14] This decomposition is exact when consumers have linear-quadratic preferences and can be derived also in oligopoly models where firms compete on quantities (Loseto, 2023).

[15] This interpretation is motivated by the fact that when A is a binary $\{0, 1\}$ it k-th power A^k counts how many walks of length k are between any two nodes.

Lemma 1 *Around the original equilibrium, the new equilibrium markup after an ε-perturbation (on, e.g., $\boldsymbol{c}$ or $\boldsymbol{\alpha}$) is*

$$\boldsymbol{p}^{\varepsilon} - \boldsymbol{c}^{\varepsilon} = \frac{\boldsymbol{\alpha}^{\varepsilon} - \boldsymbol{c}^{\varepsilon}}{2} - \boldsymbol{b}\left(D, \frac{1}{2}, \frac{\boldsymbol{\alpha}^{\varepsilon} - \boldsymbol{c}^{\varepsilon}}{2}\right).$$

Notably, $\frac{\alpha_j^{\varepsilon} - c_j^{\varepsilon}}{2}$ is the markup charged by a monopolist facing linear inverse-demand $p_j = \alpha_j - S_{jj}^{-1} q_j$. The proposition reveals that the market power of each firm is shaped by its centrality in the competition network of all products: products that are more central face greater competitive pressure and therefore charge lower markups. Notably, the competition network takes the diversion ratio as the adjacency matrix, and the monopolist's markup (which is a measure of the standalone competition strength of each product) as the node weights. For product j, the competition pressure is higher if it is closer to (i.e., it has a high diversion ratio to) a strong competitor (with a high choke price or low cost), and therefore it must charge a lower markup.

The Bonacich centrality also captures multiple-step effects, which can be clearly seen from the walk expansion:

$$\boldsymbol{b}\Big(D, \tfrac{1}{2}, \frac{\boldsymbol{\alpha}^{\varepsilon} - \boldsymbol{c}^{\varepsilon}}{2}\Big) = \sum_{k=1}^{\infty} \left(\tfrac{1}{2}\right)^k D^k \left(\frac{\boldsymbol{\alpha}^{\varepsilon} - \boldsymbol{c}^{\varepsilon}}{2}\right).$$

The term $D^k\left(\frac{\boldsymbol{\alpha}^{\varepsilon} - \boldsymbol{c}^{\varepsilon}}{2}\right)$ captures the influence of products that are k diversion steps away from a given product: a strong competitor disciplines its direct rivals' markups, which then disciplines their rivals, and so on. Hence, competitive pressure propagates along walks in the diversion network. The geometric decay factor $(1/2)^k$ reflects the linear-demand marginal-revenue tradeoff: a unit of competitive pressure transmitted one more step must pass through another best-response adjustment, and under linear demand each such adjustment is dampened by the same $1/2$ factor (as in the single-product monopoly rule $p - c = (\alpha - c)/2$. Therefore, more distant competitors matter less, with influence declining at rate $1/2$ per additional diversion step.

2.2 General Control or Ownership Structure

Our analysis can be easily extended to the case with a general ownership or control structure of the products. We introduce the decision weight matrix H to capture the extent to which the ownership structure makes the pricing decision for one product internalize the profit of another. Specifically, let

$$H_{jj} = 1, \qquad H_{j\ell} \in [0, 1) \;\; (\ell \neq j). \tag{2}$$

Product j internalizes a fraction $H_{j\ell}$ of product ℓ's profit. That is, product j chooses p_j to maximize a weighted sum of profits:

$$\max_{p_j} \sum_{\ell=1}^{J} H_{j\ell}(p_\ell - c_\ell)q_\ell(\boldsymbol{p}). \tag{3}$$

This nests the standard independent ownership case when $H = I_J$. From the independent ownership, a merger between, say, firm 1 and firm 2, can be represented by $H = I + E_{12} + E_{21}$, where $E_{j\ell}$ is the matrix with all 0 entries except for an entry of 1 at row j, column ℓ.

This modeling approach allows flexibility in the analysis of common (partial) ownership and is also used in Backus et al. (2021b) and Ederer and Pellegrino (2025). The decision matrix H can be micro-founded from common ownership as follows. Let β_{fs} denote shareholder (or fund) s's cash-flow share in product j, and let γ_{js} be the corresponding governance weight that maps ownership into influence on price decisions. The ownership structure then implies a matrix of profit weights H with $H_{jj} = 1$ and

$$H_{j\ell} = \frac{\sum_s \gamma_{js}\beta_{\ell s}}{\sum_s \gamma_{js}\beta_{js}} \qquad (j \neq \ell).$$ [16]

which leads directly to the control matrix H, with $H_{j\ell} = \kappa_{j\ell}$. For our purposes, it suffices to work directly with the decision weight matrix H.

Given the decision weight matrix H, the FOC for the price competition game is

$$0 = \boldsymbol{q}(\boldsymbol{p}) + (H \odot S(\boldsymbol{p}))(\boldsymbol{p} - \boldsymbol{c}).$$

Again, we consider a local linearization of the equilibrium condition, with $\boldsymbol{q} = -S(\boldsymbol{\alpha} - \boldsymbol{p})$. We assume that $H \odot S$ is invertible and define the *internalized* diversion ratio matrix:

$$D_H := I_J - (H \odot S)^{-1}S = -(H \odot S)^{-1}[(E_J - H) \odot S], \tag{4}$$

where E_J is the $J \times J$ matrix with all 1s. Clearly, for the independent ownership case $H = I$, D_H is the standard diversion ratio matrix. At the other extreme when $H = E_J$ (i.e., every product fully internalizes the profit of all other products), we have $D_H = 0$. In general, $[(E_J - H) \odot S]$ is the non-internalized part of the demand diversion, and $-(H \odot S)^{-1}$ normalizes diversion with the internalized quantity effect (i.e., how much the quantity changes for internalized products), hence yielding a ratio. To further illustrate this definition, consider a two-product example. The diversion

[16] Under the commonly used proportional-control assumption $\gamma_{js} = \beta_{js}$, these weights reduce to a normalized ownership overlap, $H_{j\ell} = (\beta_j'\beta_\ell)/(\beta_j'\beta_j)$.

matrix, without considering profit internalization, is

$$D = \begin{pmatrix} 0 & D_{12} \\ D_{21} & 0 \end{pmatrix}.$$

Now suppose product 1 internalizes a fraction $h \in (0,1)$ of product 2's profit (this can arise if the owner of product 1 acquires a *passive* financial stake h in product 2 without obtaining control: product 2 still chooses p_2 to maximize its own profit, while product 1 sets p_1 taking into account the profit it earns from its stake in product 2):

$$H = \begin{pmatrix} 1 & h \\ 0 & 1 \end{pmatrix}.$$

Then the internalized diversion matrix in (4) becomes

$$D_H = \begin{pmatrix} h\, D_{12} D_{21} & (1-h) D_{12} \\ D_{21} & 0 \end{pmatrix}.$$

Relative to D, the effective diversion link $1 \to 2$ is scaled down by $(1-h)$, because the remaining fraction h is internalized. This internalization, in turn, generates a self-loop $1 \to 2 \to 1$ of size $h\, D_{12} D_{21}$ at node 1: When a higher p_1 shifts some demand to product 2, product 1 partially values the resulting gain at product 2 (that is, hD_{12}). Since this internalized gain is measured in product-2 units, D_{21} translates it into the corresponding strength of feedback for product 1's own pricing incentive, yielding $h\, D_{12} D_{21}$.

The following result shows how the diversion network approach to equilibrium analysis extends to a general control structure.

Lemma 1′ *Under a general ownership represented by the decision weight matrix H, the new equilibrium markup after an ε-perturbation (on, e.g., $\boldsymbol{c}$, $\boldsymbol{\alpha}$, or H) around the original equilibrium is*

$$\boldsymbol{p}^{\varepsilon} - \boldsymbol{c}^{\varepsilon} = \frac{\boldsymbol{\alpha}^{\varepsilon} - \boldsymbol{c}^{\varepsilon}}{2} - \boldsymbol{b}\left(D_{H^{\varepsilon}}, \frac{1}{2}, \frac{\boldsymbol{\alpha}^{\varepsilon} - \boldsymbol{c}^{\varepsilon}}{2}\right).$$

3 Network-Based Analysis of Acquisitions

In this section, we show how the diversion network approach provides new insight into the analysis of mergers and acquisitions. We start with the analysis of a small share acquisition and then move to the full merger case. For expositional ease, we first consider changes to the market structure from independent ownership, and then generalize the results to any control structure.

3.1 Minority Stake Acquisitions

In this part, we consider the impact of a local change to the control matrix $H = I_J$. Without loss of generality, we focus on a change to H_{12}, so that H becomes $I_J + \varepsilon E_{12}$. This can be understood as an acquisition of *passive* financial stake of product 2 by the owner of product 1. It is worth pointing out that, if a minority share acquisition is accompanied by proportional control, multiple elements of H would change at the same time. Nonetheless, our results easily extend to that case (or any other governance mapping), since the overall price or welfare impact—i.e., the total derivative with respect to multiple elements of H—is simply the sum of partial derivatives that we obtain below.

For later use, we introduce the Katz-Bonacich centrality for the J-node network with adjacency matrix A, node weights $\boldsymbol{u}$, and attenuation factor δ:

$$\boldsymbol{\kappa}(A, \delta, \boldsymbol{u}) \; \equiv \; (I_J - \delta A)^{-1}\boldsymbol{u} = \sum_{k=0}^{\infty} \delta^k A^k \boldsymbol{u}, \tag{5}$$

which is simply the Bonacich centrality plus the length-0 (i.e., self) term. Further, we define

$$\boldsymbol{\kappa}(j; D) := \boldsymbol{\kappa}\left(D, \tfrac{1}{2}, \boldsymbol{e}_j\right),$$

which is the discounted total length of all walks (of any length) from node j to each node. This is a proximity measure of each node to j in the competition network: Node ℓ is "close" to node j if there are many or strong short paths from j to ℓ.

The following result gives the impact on prices and consumer surplus (CS) of the minority share acquisition.

Proposition 1 *When a minority share acquisition changes H from I_J to $I_J + \varepsilon E_{12}$, the price changes are:*

$$\left.\frac{\partial p_j}{\partial \varepsilon}\right|_{\varepsilon=0} = \frac{p_2 - c_2}{2} D_{12}\left[\mathbf{1}\{j=1\} + b_j\Big(D, \tfrac{1}{2}, \boldsymbol{e}_1\Big)\right] = \frac{p_2 - c_2}{2} D_{12}\kappa_j(1; D). \tag{6}$$

By Roy's identity, the CS impact is:

$$\left.\frac{\partial CS}{\partial \varepsilon}\right|_{\varepsilon=0} = -\frac{p_2 - c_2}{2} D_{12}\boldsymbol{q}'\boldsymbol{\kappa}(1; D). \tag{7}$$

The proposition reveals an intuitive network-based decomposition of the merger's impact. For the price effects, the prefactor $\frac{p_2 - c_2}{2}D_{12}$ is the *direct upward pricing pressure*: firm 1 now partially captures the margin $p_2 - c_2$ on each unit diverted from product 1 to product 2 at rate D_{12}, with the factor $1/2$ reflecting the usual linear-demand best-response pass-through. In particular, this formalizes the intuition that the acquisition of a close substitute in terms of diversion ratio has a significant impact. The indicator $\mathbf{1}\{j = 1\}$ is the direct effect of a price increase for the acquirer. The Bonacich centrality term $b_j(D, \frac{1}{2}, \boldsymbol{e}_1)$ captures the competitive ripple effect that is the hallmark of the network approach: firm 1's price increase propagates outward along walks in the diversion network—its direct rivals raise their prices, which in turn relaxes competition for *their* rivals, and so on—so that *every* product in the market is affected, with the impact on product j governed by its network proximity to product 1. By Roy's identity, the impact on consumer surplus is the quantity-weighted sum of price impacts. Therefore, the total CS loss is larger when the acquirer is central to bestsellers in the diversion network.

Next, we consider the cost synergy required to offset the CS loss.

Proposition 2 *When a minority share acquisition changes H from I_J to $I_J + \varepsilon E_{12}$, assuming that the acquirer and the target gain the same amount of cost reduction (i.e., c_j becomes $c_j - c$ for $j = 1, 2$), the cost synergy required for the minority share acquisition to be CS-neutral is:*

$$\frac{dc}{d\varepsilon} = (p_2 - c_2)D_{12}\frac{\boldsymbol{q}'\boldsymbol{\kappa}(1; D)}{\boldsymbol{q}'\boldsymbol{\kappa}(1; D) + \boldsymbol{q}'\boldsymbol{\kappa}(2; D)}. \tag{8}$$

The required synergy inherits a transparent network interpretation. The numerator $\boldsymbol{q}'\boldsymbol{\kappa}(1; D)$ is the CS harm from the acquisition (Proposition 1), capturing the *upward pricing pressure* that originates at the acquirer and propagates through the diversion network. The denominator $\boldsymbol{q}'\boldsymbol{\kappa}(1; D) + \boldsymbol{q}'\boldsymbol{\kappa}(2; D)$ is the *downward pricing pressure* from a unit cost reduction shared equally by products 1 and 2, and it too propagates through the network via the same walk-based mechanism: a cost saving at product j lowers p_j, its rivals face intensified competition and reduce their prices, which then disciplines *their* rivals, and so on. The term $\boldsymbol{q}'\boldsymbol{\kappa}(j; D)$ measures the quantity-weighted reach of this pro-competitive cascade from node j, so the denominator reflects how broadly cost savings propagate based on the network proximity of each product to *both* the acquirer and the target. The ratio is therefore a number between zero and one that captures how much of the combined pass-through benefit is needed to neutralize the anticompetitive harm. Notably, this ratio depends *in opposite directions* on the network centrality of the acquirer and the target: if the target (product 2) is highly central—connected, via diversion, to many high-volume products—then $\boldsymbol{q}'\boldsymbol{\kappa}(2; D)$ is large, the denominator grows, and the synergy bar falls, because cost reductions at a central target propagate widely and generate large consumer benefits. Conversely, when the acquirer is central relative to the target, the ratio approaches one, and the required synergy is close to the full direct upward pricing pressure $(p_2 - c_2)D_{12}$.

3.2 Full Mergers

We now consider a full merger between firms 1 and 2. This means that the control matrix changes from I_J to $M_{12} := I_J + E_{12} + E_{21}$. We again consider the case of linear or linearized demand, assuming that the price impact is not so large that the linearization is invalid.

Proposition 3 *A merger between firms 1 and 2 changes prices by:*

$$\Delta \boldsymbol{p} = \boldsymbol{b}\left(D, \frac{1}{2}, \frac{\boldsymbol{\alpha} - \boldsymbol{c}}{2}\right) - \boldsymbol{b}\left(D_{M_{12}}, \frac{1}{2}, \frac{\boldsymbol{\alpha} - \boldsymbol{c}}{2}\right). \tag{9}$$

For the quantity-weighted price change $\boldsymbol{q}'\Delta\boldsymbol{p}$ (which approximates CS loss when price changes are small) to be 0, the required cost synergy (assumed to be the same for both firms) is

$$\Delta c = \frac{2\boldsymbol{q}'\left[\boldsymbol{b}\left(D, \frac{1}{2}, \frac{\boldsymbol{\alpha}-\boldsymbol{c}}{2}\right) - \boldsymbol{b}\left(D_{M_{12}}, \frac{1}{2}, \frac{\boldsymbol{\alpha}-\boldsymbol{c}}{2}\right)\right]}{\boldsymbol{q}'\boldsymbol{\kappa}(1; D_{M_{12}}) + \boldsymbol{q}'\boldsymbol{\kappa}(2; D_{M_{12}})}. \tag{10}$$

Just like in the minority share acquisition case, centralities in the diversion-based competition network play a key role in determining the merger impacts and the bar of compensating synergies. Equation (9) comes directly from the markup decomposition result (Lemma 1′), highlighting that a merger affects prices by altering the competition network and thus the centralities of each product. Equation (10)—its denominator in particular—again reveals the insight that the downward pricing pressure from cost synergies propagates throughout the network from the merging parties.

3.3 Extensions and Discussions

The analysis of merger impacts through the lens of diversion network could enrich the antitrust authority's toolbox and entails two major benefits. First, the centrality measures extend naturally to a general control structure. In contrast, the conventional screening tool, HHI index, is not easily defined when there is partial common ownership. Second, the computation of centrality measures does not require a full-fledged estimation of the demand system, and can be done with product characteristics data, which is relatively easy to obtain. We focus on the first point in this part, and develop the second point in detail in Section 4.

To show that the diversion network-based merger screening measures are applicable under a general initial market structure, we extend previous results accordingly. For a market structure change that perturbs the control matrix to $H^{\varepsilon} = H + \varepsilon \Delta H$, where ΔH is an arbitrary $J \times J$ matrix, it is useful to define

$$\boldsymbol{\omega}_{\Delta H}(H) \;:=\; -(H \odot S)^{-1}\Big[(\Delta H \odot S)(\boldsymbol{p} - \boldsymbol{c})\Big], \tag{11}$$

which measures the immediate price adjustment pressure (typically upward when $\Delta H \geq 0$) after the change in the internalization pattern H. The term $(\Delta H \odot S)(\boldsymbol{p}-\boldsymbol{c})$ aggregates, for each product j, the newly internalized "recapture" incentives created by the perturbation: a positive $\Delta H_{j\ell}$ makes the decision maker for p_j place more weight on product ℓ's margin $(p_\ell - c_\ell)$, and this incentive is stronger when the local substitution slope $S_{j\ell}$ indicates that changing p_j shifts demand toward ℓ. Just like in the definition of the internalized diversion ratio matrix D_H in (4), the prefactor $-(H \odot S)^{-1}$ *normalizes* these raw incentive shifts by the *internalized quantity effect* under the baseline structure H—that is, by how a price change for product j translates into quantity changes for the portfolio of products whose profits j internalizes. In this way, $\boldsymbol{\omega}_{\Delta H}(H)$ expresses the direct incentive effect of ΔH in the same internalization-adjusted units that govern competitive interactions, so that it can be propagated through the diversion network in the subsequent equilibrium adjustment.

For minority share acquisitions that change H locally, we obtain the following.

Proposition 1′ *Consider a market structure change that perturbs the control matrix to $H^\varepsilon = H + \varepsilon \Delta H$, where ΔH is an arbitrary $J \times J$ matrix. Then the equilibrium price changes satisfy*

$$\left.\frac{\partial \boldsymbol{p}}{\partial \varepsilon}\right|_{\varepsilon=0} = \frac{1}{2}\, \boldsymbol{\kappa}\Big(D_H, \tfrac{1}{2}, \boldsymbol{\omega}_{\Delta H}(H)\Big)\,. \tag{12}$$

By Roy's identity, the consumer-surplus effect is

$$\left.\frac{\partial CS}{\partial \varepsilon}\right|_{\varepsilon=0} = -\frac{1}{2}\, \boldsymbol{q}'\boldsymbol{\kappa}\Big(D_H, \tfrac{1}{2}, \boldsymbol{\omega}_{\Delta H}(H)\Big)\,. \tag{13}$$

Proposition 2′ *Consider a market structure change that perturbs the control matrix to $H^\varepsilon = H + \varepsilon \Delta H$. The cost synergies required for the market structure change to be CS-neutral must satisfy:*

$$\sum_{j=1}^{J} \boldsymbol{q}'\boldsymbol{\kappa}(j; D_H)\left(-\frac{dc_j}{d\varepsilon}\right)\Bigg|_{\varepsilon=0} = \boldsymbol{q}'\boldsymbol{\kappa}\Big(D_H, \tfrac{1}{2}, \boldsymbol{\omega}_{\Delta H}(H)\Big)\,. \tag{14}$$

For more significant market structure changes, such as full mergers, assuming that the induced price changes are still within the range of linear approximation, we have the following.

Proposition 3′ *Consider an arbitrary change to the control matrix from H to $\widehat{H}$. The price changes (without any cost synergy) are:*

$$\Delta \boldsymbol{p} = \boldsymbol{b}\left(D_H, \frac{1}{2}, \frac{\boldsymbol{\alpha}-\boldsymbol{c}}{2}\right) - \boldsymbol{b}\left(D_{\widehat{H}}, \frac{1}{2}, \frac{\boldsymbol{\alpha}-\boldsymbol{c}}{2}\right). \tag{15}$$

Suppose there are synergies that reduce costs from $\boldsymbol{c}$ to $\boldsymbol{c} - \Delta \boldsymbol{c}$ at the same time *of the control matrix change. For the quantity-weighted price change $\boldsymbol{q}'\Delta\boldsymbol{p}$ (which approximates CS loss when price changes are small) to be 0, the required cost synergies must satisfy:*

$$\frac{1}{2}\sum_{j=1}^{J}\boldsymbol{q}'\boldsymbol{\kappa}(j;D_{\widehat{H}})\Delta c_j = \boldsymbol{q}'\left[\boldsymbol{b}\left(D_H,\frac{1}{2},\frac{\boldsymbol{\alpha}-\boldsymbol{c}}{2}\right)-\boldsymbol{b}\left(D_{\widehat{H}},\frac{1}{2},\frac{\boldsymbol{\alpha}-\boldsymbol{c}}{2}\right)\right]. \tag{16}$$

4 Diversion Network and Product Attributes: A Micro-foundation

Sections 2 and 3 established that equilibrium markups and the competitive effects of acquisitions are governed by the diversion ratio matrix D and its centrality structure. In those sections, we consider the local linearization of a general demand function $\boldsymbol{q}(\boldsymbol{p})$, and take its Jacobian S (and thus the diversion ratio matrix D) as given. The practical question for a regulator or analyst is: where does D come from, and what data are needed to compute it?

In this section, we provide a micro-foundation based on the linear-quadratic demand model that has long been a workhorse in industrial organization for differentiated oligopoly (e.g., Singh and Vives (1984); Vives (1985)) and has more recently been used to link product attributes to competition networks (Pellegrino, 2025). We start from a representative consumer with quadratic preferences over the vector of product characteristics.[17] We show that this model generates a Jacobian S—and hence the diversion ratio matrix—that is determined by weighted inner products of product characteristics. Within this framework, the diversion ratio matrix can be obtained directly from observed product characteristics, without estimating a full demand system. Combined with the results of Sections 2 and 3, this means that the data requirement and computational burden of the centrality-based tools for merger screening are minimal.

4.1 Preferences and Demand

Products. There are J products available. Each product j is characterized by a vector of K attributes $\boldsymbol{x}_j = (x_{j1},\ldots,x_{jK})'$, collected in the $J\times K$ matrix X. The entry x_{jk} measures how much of attribute k a consumer obtains per unit of product j.

Utility. A representative consumer takes prices $\boldsymbol{p} = (p_1,\ldots,p_J)'$ as given and chooses quantities $\boldsymbol{q} = (q_1,\ldots,q_J)'$ and an outside good q_0 subject to $q_0+\boldsymbol{p}'\boldsymbol{q}\le y$. Preferences are

$$u(\boldsymbol{q},X) \;=\; q_0+\boldsymbol{q}'\boldsymbol{\alpha}-\frac{\beta}{2}\,\boldsymbol{q}'(I_J+\eta XX')\,\boldsymbol{q}, \tag{17}$$

[17] One can also obtain the demand system as aggregate demand from a mass of consumers, each with possibly heterogeneous linear-quadratic preferences. For more details see Loseto (2023).

where $\eta > 0$ governs the importance of product differentiation, $\boldsymbol{\alpha}$ (a J-vector) captures the consumer willingness to pay for each product, and $\beta > 0$ captures the consumer's preference for variety. The quadratic term $\boldsymbol{q}'(I_J + \eta XX')\boldsymbol{q}$ ensures that products with more similar attribute vectors are closer substitutes: the cross-product penalty for consuming products j and ℓ jointly is proportional to $\eta\, \boldsymbol{x}_j'\boldsymbol{x}_\ell$.

Demand. The consumer maximizes (17) subject to $q_0 + \boldsymbol{p}'\boldsymbol{q} \leq y$. Substituting the budget constraint, the first-order condition yields

$$\boldsymbol{q}(\boldsymbol{p}) \;=\; \frac{1}{\beta}(I_J + \eta XX')^{-1}\,(\boldsymbol{\alpha} - \boldsymbol{p}), \tag{18}$$

which is always well defined because $I_J + \eta XX'$ is positive definite.

The substitution matrix. Define the $J \times J$ matrix $\Theta := (\theta_{j\ell})_{j,\ell=1}^{J}$ with

$$\theta_{j\ell} \;:=\; \boldsymbol{x}_j'\,\Omega^{-1}\,\boldsymbol{x}_\ell, \qquad \Omega \;:=\; \frac{1}{\eta}I_K + X'X. \tag{19}$$

It can be verified that Ω is a positive-definite weighting matrix, and $\theta_{j\ell} \in (-1, 1)$ for $j \neq \ell$ and $\theta_{jj} \in (0, 1)$.[18]

Proposition 4 *The demand system can be written in matrix form as*

$$\boldsymbol{q}(\boldsymbol{p}) \;=\; \frac{1}{\beta}(I_J - \Theta)\,(\boldsymbol{\alpha} - \boldsymbol{p}). \tag{20}$$

Equivalently, for each product j,

$$q_j(\boldsymbol{p}) = a_j - \frac{1}{\beta}(1 - \theta_{jj})\,p_j + \frac{1}{\beta}\sum_{\ell \neq j}\theta_{j\ell}\,p_\ell,$$

where a_j is a product-specific demand intercept given by $\boldsymbol{a} = \frac{1}{\beta}(I_J - \Theta)\boldsymbol{\alpha}$.

Because preferences are quadratic, demand is linear in prices. All substitution patterns are encoded in the $J \times J$ matrix Θ with entries $\theta_{j\ell}$. The Slutsky matrix (Jacobian of demand) takes the form $S = -(I_J - \Theta)/\beta$. Products j and ℓ are substitutes when $\theta_{j\ell} > 0$ and complements when $\theta_{j\ell} < 0$. Crucially, $\theta_{j\ell}$ is a *weighted inner product* of the attribute vectors $\boldsymbol{x}_j$ and $\boldsymbol{x}_\ell$, which means that the closer j and ℓ are in this inner-product space, the more substitutable they are.

[18] Because $\eta > 0$, the matrix $\Omega = \frac{1}{\eta}I_K + X'X$ is positive definite: for any nonzero $\boldsymbol{u} \in \mathbb{R}^K$, $\boldsymbol{u}'\Omega\boldsymbol{u} = \frac{1}{\eta}\|\boldsymbol{u}\|^2 + \|X\boldsymbol{u}\|^2 > 0$. Moreover, $\Theta = X\Omega^{-1}X'$ is the ridge *hat matrix*. Writing an SVD $X = U\Sigma V'$, we have $\Theta = U\,\Sigma(\Sigma^2 + \frac{1}{\eta}I)^{-1}\Sigma\,U'$, so its eigenvalues are $\sigma_r^2/(\sigma_r^2 + 1/\eta) \in [0, 1)$, which implies $\theta_{jj} = \boldsymbol{e}_j'\Theta\boldsymbol{e}_j \in (0, 1)$ (provided $\boldsymbol{x}_j \neq \boldsymbol{0}$). Finally, since $\Omega^{-1} \succ 0$, letting $\boldsymbol{z}_j := \Omega^{-1/2}\boldsymbol{x}_j$ gives $\theta_{j\ell} = \boldsymbol{z}_j'\boldsymbol{z}_\ell$, and thus by Cauchy–Schwarz, $|\theta_{j\ell}| \leq \|\boldsymbol{z}_j\|\,\|\boldsymbol{z}_\ell\| = \sqrt{\theta_{jj}\theta_{\ell\ell}} < 1$ for $j \neq \ell$.

The next result shows that the weights in the inner product $\theta_{j\ell}$ have a transparent interpretation in terms of the principal components of the characteristics matrix X.

Proposition 5 *Let U be the $K \times K$ matrix whose columns are the principal-component directions of X, and let $\tilde{\boldsymbol{x}}_j = U'\boldsymbol{x}_j$ be the projection of $\boldsymbol{x}_j$ onto these directions. Then*

$$\theta_{j\ell} = \sum_{k=1}^{K} \frac{\tilde{x}_{jk}\,\tilde{x}_{\ell k}}{1/\eta + \lambda_k^{X'X}}, \tag{21}$$

where $\lambda_k^{X'X}$ is the k-th eigenvalue of $X'X$.

Expression (21) delivers two insights. First, because U is orthogonal, the rotation from $\boldsymbol{x}_j$ to $\tilde{\boldsymbol{x}}_j$ does not alter the inner-product structure ($\tilde{X}\tilde{X}' = XX'$), so the projection is without loss.

Second, the weight on the k-th principal component is inversely related to the variance of that component across all products. To see this, note that $\lambda_k^{X'X} = \tilde{\boldsymbol{x}}_k'\tilde{\boldsymbol{x}}_k$, so characteristics along which products are more homogeneous receive a higher weight. The intuition is that attributes shared by nearly all products in the market are especially informative about pairwise substitutability: consumers cannot easily switch to a differentiated alternative along a dimension on which all options are similar. Conversely, attributes that vary widely across products matter less for any given pair.

To sum up, the substitutability between products j and ℓ depends not only on how similar their attribute vectors are, but also on how the similarity is distributed across dimensions relative to market-wide variation.

4.2 Product Attributes, Diversion Network, and Merger Impacts

We now connect the substitution matrix Θ to the diversion ratio matrix D introduced in Section 2. Recall that D is defined as $D = I_J - [\text{diag}(S)]^{-1}S$, so its off-diagonal entries are

$$D_{j\ell} = -\frac{S_{j\ell}}{S_{jj}} = \frac{\theta_{j\ell}}{1 - \theta_{jj}}, \qquad j \neq \ell.$$

Under the linear-quadratic demand model, the diversion ratio from product j to product ℓ is therefore determined entirely by the attribute-based similarity $\theta_{j\ell}$, normalized by the own-price sensitivity term $1 - \theta_{jj}$. In particular:

- $D_{j\ell}$ is large when $\boldsymbol{x}_j$ and $\boldsymbol{x}_\ell$ are similar in the Ω^{-1}-weighted inner product, i.e., when the two products are close substitutes in characteristic space.
- Like the linear and logit models discussed in Conlon and Mortimer (2021), the network model delivers diversion ratios that do not vary with price. However,

the economic source of this constancy is entirely different. In our model, diversion is constant because substitution patterns are generated by the geometry of product attributes under quadratic preferences, not because of IIA or because treatment effects are homogeneous. This structure allows us to preserve realistic heterogeneity in demand intercepts and marginal utilities while obtaining closed-form markups and centrality-based merger effects.

The entire diversion matrix D can be constructed from product characteristics alone, without estimating a demand system. This has an immediate implication for the results of Sections 2–3.[19] Since equilibrium markups are determined by the Bonacich centrality of the diversion network (Lemma 1 and 1′), and since merger impacts propagate through walks in this same network (Propositions 1–2 and their primed extensions), the entire centrality-based apparatus can be computed from product attributes.

The results of this section, combined with those of Sections 2–3, deliver a practical procedure for antitrust screening:

$$\text{Product characteristics } X \to \text{Diversion network } D \to \text{Centrality-based merger effects.}$$

We close this section by reiterating several properties that make this procedure attractive for policy.

Low data requirements and compatibility with modern data sources. The diversion matrix D is constructed from the attribute matrix X and the single parameter η (equation (19)). No price or quantity data, no estimation of a structural demand system, and no market-share information are needed to build the competition network. In settings where product characteristics are publicly available—as they are in automobiles, consumer electronics, pharmaceuticals, and many other industries—the entire screening exercise can be carried out with minimal data collection. At the same time, the framework places essentially no restriction on *how* characteristics are measured: recent advances in machine learning show that embeddings derived from text descriptions, images, or other unstructured product data can capture demand-relevant positioning and substitute for hand-coded attributes (Bajari et al., 2023; Compiani et al., 2025). Relatedly, Magnolfi et al. (2022) demonstrate that embeddings from perceived-distance surveys yield elasticity estimates comparable to those from observable attributes. Because our construction depends on X only through inner products, any representation of product attributes that preserves the same inner-product structure generates the same substitution patterns and thus the same diversion network and centrality-based screening measures.

[19] We assume that a unique equilibrium, where all prices p_j lie in the interior of $[c_j, \alpha_j]$, exists for the linear-quadratic utility model throughout. A sufficient condition on the primitives, as shown in Appendix B, is $(1 - \theta_{jj})(\alpha_j - c_j) > \sum_{k \neq j} |\theta_{jk}|(\alpha_k - c_k)$ for all j.

No ex ante market definition. As emphasized by Hoberg and Phillips (2016), a key advantage of characteristic-based approaches is that they bypass the need to define a relevant market. In our framework, the network structure *is* the market: products that are close in attribute space share strong links and exert competitive pressure on each other, whereas distant products are effectively independent. Even a very broad product universe does not distort the analysis, because the inner-product weights in (21) automatically downweight characteristics that vary widely and hence are less informative about pairwise substitutability.

Micro-founded in a setting with differentiated products. Related to the previous point, the centrality measure naturally emerges in the equilibrium of a price-setting game between arbitrarily differentiated products. By contrast, while functions of merging parties' shares related to ΔHHI are shown to be indicative of merger impacts in differentiated price competition for aggregative games (Nocke and Whinston, 2022), the exact micro-foundation for ΔHHI is obtained within a homogeneous-good Cournot competition framework.

Coherent under arbitrary ownership structure. A further practical advantage is that the centrality statistic used for screening is well defined under *any* ownership or control structure. Ownership enters the model as a primitive of the firms' objective functions through the decision-weight matrix H. As shown in Sections 2–3, partial acquisitions and common ownership modify incentives directly via the internalized diversion network D_H, which governs equilibrium pricing. The screening statistic is computed from this object, so changes in control—whether full mergers or minority stake acquisitions by incumbents or financial investors—affect the measure through the same structure that determines markups.

Scale-free substitution patterns. Under Bertrand competition, the elements $\theta_{j\ell}$ lie in $(-1, 1)$ regardless of the units in which characteristics are measured. This means that the diversion network—and hence the centrality measures and merger effects—are invariant to rescaling of the attribute vectors. In practice, no normalization of product characteristics is required.

5 Application to the Automobile Industry

In this section, we estimate the model using data on the European automobile industry. The same data have been used extensively in the empirical industrial organization literature (e.g., Brenkers and Verboven (2006)). We then introduce our empirical framework and use it to estimate demand, comparing the results with those we obtain by estimating demand using discrete-choice models. Lastly, we simulate the effects of

Table 1: Summary Statistics on the European Automobile Industry

Variable	Mean	SD	Bel	Fra	Ger	Ita	U.K.
Sales (units)	19813.24	37719.92	3925.42	23305.81	31002.55	24292.14	19784.34
Horsepower (kW)	57.14	23.88	56.53	56.17	57.45	57.22	58.42
Fuel inefficiency (liter/100 km)	8.18	1.72	8.22	8.12	8.24	8.08	8.22
Width (cm)	164.38	9.62	164.16	164.28	164.57	164.06	164.82
Height (cm)	140.43	4.62	140.34	140.46	140.49	140.63	140.27
Foreign (0-1)	0.81	0.39	1.00	0.75	0.73	0.76	0.77
Price Eur	8352.52	5540.92	7674.24	8315.98	8119.62	8378.81	9384.34
Price Eur (post-tax)	10089.23	6606.51	9439.35	10476.46	9283.47	9946.35	11388.33
Price Eur (Deflated)	10793.17	5962.59	9029.31	10944.30	8394.99	13292.90	12870.79
Subcompact (0-1)	0.28	0.45	0.28	0.28	0.28	0.31	0.26
Compact (0-1)	0.23	0.42	0.23	0.24	0.23	0.22	0.24
Intermediate (0-1)	0.22	0.41	0.22	0.22	0.22	0.21	0.23
Standard (0-1)	0.19	0.39	0.19	0.19	0.18	0.17	0.19
Luxury (0-1)	0.08	0.27	0.08	0.07	0.09	0.09	0.08

Notes: The number of observations is 11549; SD means standard deviation.

mergers under different assuming different demand models and assess the effectiveness of screening mergers using centrality-based rules.

5.1 Data

We use publicly available data from Brenkers and Verboven (2006).[20] The dataset covers prices, quantities, and detailed product characteristics for virtually all passenger car models sold in five major European markets—Belgium, France, Germany, Italy, and the United Kingdom—between 1970 and 1999. Over this period, roughly 350 distinct models appear in the data, though many represent successive generations of existing nameplates (e.g., Volkswagen Golf, Toyota Corolla, BMW 5-Series). In total, the panel contains 11,549 model–market–year observations, corresponding to an average of about 80 models available per country in a given year.

For each model, we observe the list price for the base version offered in each market, as reported in contemporary consumer catalogs. Sales are measured using new vehicle registrations at the model range level. The dataset further includes a rich set of physical and performance characteristics drawn from the same catalogs. These variables capture vehicle size (e.g., weight, length, width, height), engine specifications (horsepower and displacement), and performance metrics (top speed, acceleration, and fuel economy). Finally, we observe identifiers for the model, brand, and manufacturer, along with the country of origin or production location and the market segment classification (subcompact, compact, intermediate, standard, luxury).

[20] The database is available at: https://sites.google.com/site/frankverbo/data-and-software/data-set-on-the-european-car-market?authuser=0 and has also been used in Goldberg and Verboven (2001) and Goldberg and Verboven (2005).

Table 1 provides summary statistics for all the key characteristics of the cars in our sample across market and years, and separately for each country across years. The last five rows show the penetration of each segment: not surprisingly, the luxury segment only accounts for about 8% of the sales. By contrast, the subcompact segment tends to be the largest (28% of the sales), followed closely by the compact and intermediate segments, with 23% and 22% of the sales, respectively.

5.2 The Empirical Nash-Bertrand Network Model

We operationalize the network model described in Section 4, showing how to adapt it to the specific context of the automobile industry.

An implicit assumption of our model is the absence of income effects due to the quasi-linearity of consumer preferences in the outside good. In the context of car purchases, this assumption might be unrealistic. For this reason, rather than parsimoniously extending the demand model to accommodate income effects, we follow Berry, Levinsohn, and Pakes (1995) and model consumer preference for the inside and outside goods in a Cobb-Douglas fashion:

$$U\left(y - \boldsymbol{p}'\boldsymbol{q}, \boldsymbol{q}; X\right) = \left(y - \boldsymbol{p}'\boldsymbol{q}\right)^{\gamma} \left[G\left(\boldsymbol{q}, X\right)\right]^{\phi} \tag{22}$$

where the first term in U already substitutes for the budget constraint and

$$G(\boldsymbol{q}, X) \equiv \exp\left\{\boldsymbol{q}'\boldsymbol{\alpha} - \frac{\beta}{2}\boldsymbol{q}'\left(I_J + \eta X X'\right)\boldsymbol{q}\right\}.$$

Next, substituting G into (22) and taking logs, the representative consumer's utility can be written as

$$u(\boldsymbol{q}, X) \equiv \log\left(U\right) = \gamma \log(y - \boldsymbol{p}'\boldsymbol{q}) + \boldsymbol{q}'\boldsymbol{\alpha} - \frac{\beta}{2}\boldsymbol{q}'\left(I_J + \eta X X'\right)\boldsymbol{q} \tag{23}$$

$$\approx \boldsymbol{q}'\left(\boldsymbol{\alpha} - \frac{\boldsymbol{p}}{y}\right) - \frac{\beta}{2}\boldsymbol{q}'\left(I_J + \eta X X'\right)\boldsymbol{q} \tag{24}$$

where we normalize ϕ and γ to 1 because they cannot be separately identified from α and β, and we use a first-order Taylor expansion to approximate $\log\left(y - \boldsymbol{p}'\boldsymbol{q}\right)$.[21] We also calibrate η to 0.12 following Pellegrino (2025). The preferences in (24) are identical to the ones described in Section 4 except for the fact that prices are now measured relative to income. The representative consumer's demand system can be

[21] The same approximation has been used in Berry, Levinshon, and Pakes (1999).

derived as before:

$$\boldsymbol{q}(\boldsymbol{p}) = \frac{1}{\beta}\left(I_J + \eta XX'\right)^{-1}\left(\boldsymbol{\alpha} - \frac{\boldsymbol{p}}{y}\right). \tag{25}$$

To map this framework to automobile data, let t denote the market (i.e., a country-year pair) and note that equation (25) can be rearranged in vector form as

$$\tilde{\boldsymbol{q}}_t \equiv (I_{J_t} + \eta X_t X_t')\boldsymbol{q}_t = \frac{1}{\beta}\left(\boldsymbol{\alpha}_t - \frac{\boldsymbol{p}_t}{y_t}\right) \tag{26}$$

where J_t is the number of car models available in market t, X_t is a $J_t \times K$ matrix of product characteristics and the linear preference parameter vector α is allowed to vary over time. Next, consider the jth equation of the above system

$$\tilde{q}_{jt} = \frac{1}{\beta}\left(\alpha_{jt} - \frac{p_{jt}}{y_t}\right) \tag{27}$$

and note that, upon calibrating η and assuming that the matrix X_t contains only observables characteristics, the left-hand side in (25), denoted by $\tilde{q}_{jt}$, is directly measurable. Conversely, on the right-hand side of (25), only $\boldsymbol{p}_t$ and y_t are observable.[22]

More generally, equation (27) suggests that we can estimate demand using the following linear specification

$$\tilde{q}_{jt} = -\frac{1}{\beta}\frac{p_{jt}}{y_t} + w_{jt}'\zeta + \xi_{jt} \tag{28}$$

where w_{jt} is a vector of observable product and demographic characteristics which, in this context, includes both x_{jt} and y_t. On the other hand, ξ_{jt} includes characteristics that are unobservable to the econometrician but known by the agents.

To sum up, two assumptions allow us to estimate demand from (28). First, all the characteristics that enter consumer preferences in the quadratic term are observable. Second, any unobserved characteristic (ξ_{jt}) enters consumer preferences only through the linear parameter vector α, i.e.,[23]

$$\alpha_{jt} = \beta(w_{jt}'\zeta + \xi_{jt}) \tag{29}$$

To consistently estimate (28), we need to instrument p_{jt}/y_t because prices will be

[22] To measure income in a given year, we use the simulated draws that come with the pyBLP package developed in Conlon and Gortmaker (2020), and we average them using the weights provided. The draws come from a log-normal distribution of income whose location and scale parameters are estimated from the Current Population Survey (CPS) each year as described in Berry, Levinsohn, and Pakes (1995).

[23] An additional assumption is that β is constant across markets. However, this can be partially relaxed by interacting prices with any market level observable in equation (28).

correlated with the unobservable component ξ_{jt}. The reason is that firms internalize ξ_{jt} before setting prices simultaneously. To see this formally, recall that Nash equilibrium prices are given by

$$p_{jt} = c_{jt} + \frac{y_t \alpha_{jt} - c_{jt}}{2} - b_{jt} \tag{30}$$

and note that those prices are a function of α_{jt} (and in turn function of ξ_{jt}) both directly and indirectly through product j's Bonacich centrality b_{jt}. Under this setting, demand can be estimated with a simple linear instrumental variable strategy which is described next.

5.3 Demand Estimation

We begin by estimating demand from the linear specification in (28). To do so, we instrument the term p_{jt}/y_t using the set of demand instruments z_{jt} constructed in Brenkers and Verboven (2006). These include the product's own observed characteristics (horsepower, fuel efficiency, width, height, and a dummy for whether the product is a foreign one), as well as functions of competitors' characteristics (Berry et al., 1995). Specifically, we use: the number of products and the sums of characteristics of other products of the same firm belonging to the same segment, and the number of products and the sums of the characteristics of competing products belonging to the same segment.

Equation (30) implies product characteristics affect firms' pricing decisions through the centrality term b_{jt} and the observable part of α_{jt}, suggesting the relevance of our instrument. To ensure that demand remains constant while instruments shift the supply, the identifying assumption relies on the idea that firms choose characteristics before observing any demand shock ξ_{jt} which formally boils down to requiring that $\mathbb{E}[\xi_{jt}|z_{jt}] = 0$.

Because θ_{jl} is a function of observable product characteristics, estimating $1/\beta$ is sufficient to recover own-price elasticities, cross-price elasticities, and diversion ratios. The bottom part of Table 2 shows the estimates we obtain for the average diversion ratio, own- and cross-price elasticities, both within and across segments. For example, our estimates imply that following a marginal increase in the price of a compact car, on average 2.35% of the lost sales are diverted to other compact models, whereas 0.82% are diverted to models in other segments.

Comparison with discrete-choice models. We compare the network demand estimates with four standard discrete-choice specifications: Logit, Nested Logit, Random Coefficients Logit (RCL), and Random Coefficients Nested Logit (RCNL). All models are estimated using the same product-level data and the same set of observable char-

Table 2: Estimated Substitution Patterns across Demand Models

Segment	Own-price	Cross-price elasticity		Diversion ratio	
		Within segment	Across segments	Within segment	Across segments
Logit					
Subcompact	-1.32	0.00	0.00	0.05	0.05
Compact	-1.85	0.00	0.00	0.05	0.05
Intermediate	-2.19	0.00	0.00	0.04	0.04
Standard	-3.06	0.00	0.00	0.02	0.02
Luxury	-4.22	0.00	0.00	0.03	0.03
Nested Logit					
Subcompact	-2.25	0.03	0.00	1.30	0.05
Compact	-3.08	0.05	0.00	1.51	0.04
Intermediate	-2.77	0.00	0.00	0.04	0.03
Standard	-3.96	0.01	0.00	0.15	0.02
Luxury	-6.99	0.40	0.00	6.12	0.02
RC Logit					
Subcompact	-1.54	0.02	0.02	1.12	0.97
Compact	-2.15	0.02	0.02	1.06	1.06
Intermediate	-2.55	0.02	0.02	0.86	0.82
Standard	-3.59	0.02	0.02	0.63	0.47
Luxury	-4.93	0.04	0.02	1.03	0.59
RC Nested Logit					
Subcompact	-2.87	0.07	0.02	2.33	0.82
Compact	-3.47	0.07	0.02	2.20	0.84
Intermediate	-3.15	0.03	0.02	0.82	0.60
Standard	-4.55	0.03	0.02	0.78	0.36
Luxury	-7.72	0.44	0.02	6.19	0.48
Network					
Subcompact	-4.50	0.11	0.06	2.31	0.73
Compact	-5.46	0.13	0.05	2.35	0.82
Intermediate	-6.93	0.17	0.06	2.38	0.85
Standard	-9.53	0.24	0.06	2.53	0.87
Luxury	-9.33	0.29	0.05	3.29	0.96

Notes: The table summarizes, for each of the five demand models considered, the average elasticities and diversion ratios across all markets by segment. Values are trimmed at the 25th percentile for the Network model.

acteristics. In each case, price is treated as endogenous and instrumented using the BLP-style instruments described above.

We first estimate a homogeneous multinomial logit model. Consumer utility is specified as

$$u_{ijt} = x_{jt}\beta - \alpha \frac{p_{jt}}{y_t} + \xi_{jt} + \varepsilon_{ijt}, \tag{31}$$

where x_{jt} includes observable characteristics and ε_{ijt} follows the type-I extreme value distribution. Mean utility is recovered using the standard logit inversion

$$\delta_{jt} = \log s_{jt} - \log s_{0t}. \tag{32}$$

We then estimate

$$\delta_{jt} = x_{jt}\beta - \alpha \frac{p_{jt}}{y_t} + \xi_{jt} \tag{33}$$

via two-stage least squares. This specification imposes homogeneous price sensitivity and the Independence of Irrelevant Alternatives (IIA), and hence substitution patterns depend only on market shares.

To relax the IIA property, we estimate a nested logit model that groups vehicles into the five market segments available in our data: subcompact, compact, intermediate, standard, and luxury. This allows consumers to have correlated preferences for products in the same segment, which share a common set of features. Utility is specified as:

$$u_{ijt} = \delta_{jt} + \zeta_{ig} + (1 - \sigma)\varepsilon_{ijt}, \tag{34}$$

where g indexes the nest and $\sigma \in [0, 1)$ captures within-segment correlation. The estimating equation becomes

$$\log s_{jt} - \log s_{0t} = x_{jt}\beta - \alpha \frac{p_{jt}}{y_t} + \sigma \log s_{j|g,t} + \xi_{jt}, \tag{35}$$

where $s_{j|g,t}$ denotes the within-nest share. Both price and the within-nest share term are instrumented using BLP-style excluded instruments. Specifically, we interact these instruments with the set of segment dummies to allow the nesting parameter to differ across segments. This model allows stronger substitution within segments than across segments while maintaining homogeneous price sensitivity across consumers.

We also estimate an RCL model in the spirit of Berry, Levinsohn, and Pakes (1995). Utility is specified as

$$u_{ijt} = x_{jt}\beta - \alpha \frac{p_{jt}}{y_t} + \mu_{ijt} + \xi_{jt} + \varepsilon_{ijt}, \tag{36}$$

where μ_{ijt} captures unobserved taste heterogeneity through random coefficients on selected product characteristics and a random intercept. Specifically, coefficients on observable characteristics are assumed to vary across consumers according to a multivariate normal distribution.

Importantly, the price coefficient remains homogeneous across consumers. We do not model income as a demographic interaction with price in these discrete-choice specifications because we do not observe the income distribution at the market-year level for the European automobile markets in our sample. The model is estimated by simulated GMM using the `pyBLP` package, integrating over the distribution of random coefficients (Conlon and Gortmaker, 2020). Price is instrumented using the same excluded instruments as in the logit specification.

Finally, we estimate an RCNL model, which combines taste heterogeneity with within-segment correlation (Grigolon and Verboven, 2014). Utility takes the form:

$$u_{ijt} = x_{jt}\beta - \alpha \frac{p_{jt}}{y_t} + \mu_{ijt} + \sigma\zeta_{ig} + \xi_{jt} + \varepsilon_{ijt}, \tag{37}$$

where μ_{ijt} represents random coefficients on product characteristics (including a random intercept) and σ captures within-nest correlation. As in the RCL specification, price sensitivity is not heterogeneous across consumers. The model is estimated via simulated GMM using `pyBLP`, with the same instrument set used in the previous specifications.

Across models, substitution patterns become progressively more flexible. Logit imposes IIA and homogeneous price sensitivity; Nested Logit allows stronger substitution within segments; RCL introduces unobserved taste heterogeneity; and RCNL combines taste heterogeneity with within-segment correlation. These differences clearly emerge from Table 2.[24] In Logit and RCL, average diversion ratios are similar within and across segments, reflecting the absence of explicit segment-level correlation. In contrast, Nested Logit and RCNL generate higher within-segment diversion, as preference correlation mechanically strengthens substitution among products in the same nest. Interestingly, the network model is able to replicate this pattern without exogenously imposing a nesting structure, generating average diversion ratios that are remarkably close to those of the RCNL specification. This is particularly valuable because, in many settings, it is difficult to determine the appropriate nesting structure ex ante, and misspecifying the nests can lead to biased estimates of substitution patterns (Fosgerau et al., 2024).

We also emphasize that the network model allows consumers to exhibit a taste for variety and to consume more than one good. In this respect, the automobile industry—where households typically own no more than two cars and, in most cases, purchase only one vehicle at a time—does not represent the ideal environment to apply our framework. Therefore, the fact that the model performs well in this setting provides a conservative assessment of its empirical relevance. Its performance is likely to be even stronger in industries characterized by multi-product consumption and richer substitution patterns.

[24] Table C.1 in the appendix shows the results we obtain when the values of the own-price elasticities are trimmed at the 10th percentile before averaging.

5.4 Merger simulations

Next, we consider all possible mergers—across all countries and years—between the car manufacturers in our sample. In what follows, we will base our analyses on demand estimates obtained from the network and RCNL models.

The first step entails recovering firms' marginal costs. To that end, we estimate $\hat{\alpha}_{jt}$ for each product j and market t by simply plugging our demand estimates $(\hat{\beta}, \hat{\zeta})$ and the estimated regression residuals $(\hat{\xi}_{jt})$ into equation (29). Then, from the pricing equation in (30), we can recover marginal costs c_{jt}, price-cost margins $p_{jt} - c_{jt}$ and decompose the markup into monopolistic price-cost margins $(\alpha_{jt} - c_{jt})/2$ and network centrality b_{jt}. This decomposition implies that, in the network model, network centrality is the primary determinant of the effects of acquisitions.

Figure C.1 shows the distribution of marginal costs across the network and the RCNL models, focusing on the most recent year of our sample (1999). The median marginal cost across models and markets is \$7,636.17 (\$6,493.51), and 75% of car models have a marginal cost lower than \$12,555.56 (\$10,839.04) in the network (RCNL) framework. Overall, the estimated marginal costs are reasonably close across the two demand models considered.

Figure 1 compares merger simulations obtained from the network model and the RCNL specification, holding marginal costs fixed at pre-merger levels. The two models generate strongly correlated predictions across all outcomes, indicating substantial agreement in the ranking of mergers by their competitive impact. This is consistent with the broadly similar diversion patterns implied by the two demand systems. The plots, however, also reveal some systematic differences. The network model tends to predict smaller price increases and smaller consumer-surplus losses, consistent with its higher (in absolute value) own-price elasticities. At the same time, merger outsiders appear more responsive under the network specification: non-merging products experience relatively larger price adjustments compared to RCNL, suggesting that competitive effects propagate more broadly through the similarity network rather than remaining concentrated within predefined nests.

To gain a better understanding of these differences in magnitudes, we zoom in on domestic mergers, which are arguably more likely to raise competitive concerns, and compare the simulation results obtained by relying on demand estimates obtained using the RCNL and the network models. To tag domestic mergers, we use the information on the country of origin or production location of each manufacturer. Since some manufacturers produce in multiple countries (e.g., BMW and Ford), simulated mergers involving them can span multiple countries. Table 3 summarizes the results for the potential domestic mergers that could have happened in 1999.

We find that the two models generate reasonably close price and CS effects across all

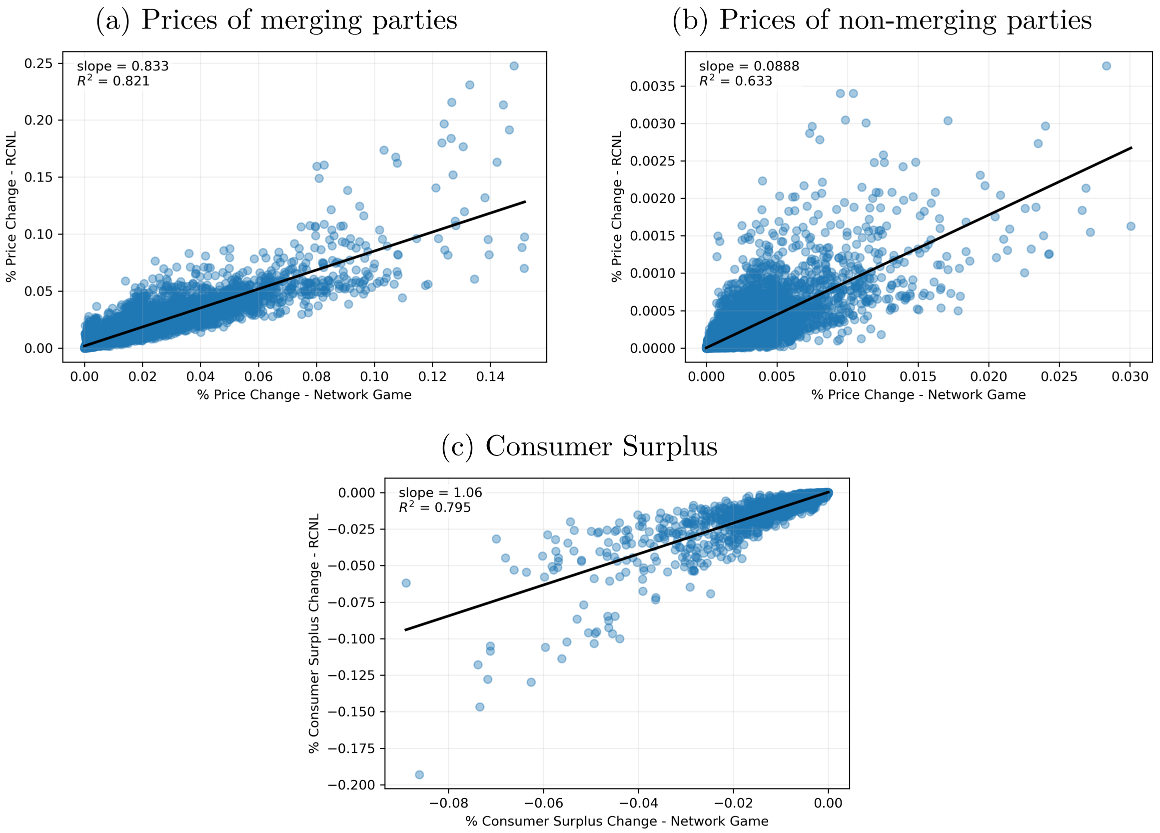

Figure 1: Simulation of the effects of all possible mergers

Notes: Each figure contrasts the effect of a potential merger estimated using the network model with that estimated using the RCNL model. Each figure focuses on the percentage change in a different outcome: the prices of the merged parties, the prices of the non-merged parties, and consumer surplus.

simulated domestic mergers. For example, a merger between the two largest domestic manufacturers in France (Peugeot and Renault), which raises the HHI by almost 1,624 points, leads the merged entity to increase prices by 8.0% under the network model and by 16% under RCNL. The associated consumer-surplus loss amounts to 3.4% in the network model and 5.4% in RCNL. By contrast, a merger between BMW and MCC in Germany, which have limited overlap across segments and induces a small change in HHI (less than 19 points), generates much smaller effects under both models, with price increases between 0.3–0.49% and CS losses between 0.02–0.03%.

While, on average, the network model tends to predict smaller price effects for the merging parties, this pattern is not uniform across all transactions. For instance, in the BMW–VW merger in Germany, the network model predicts larger price increases for both merged and non-merged products relative to RCNL, yet the associated CS loss is smaller. This reflects differences in demand curvature and substitution patterns across the two specifications. Although the network model implies stronger diversion between BMW and VW, leading to larger unilateral price effects, it also generates broader substitution toward alternative inside products and relatively less reallocation

Table 3: Domestic Mergers in 1999: RCNL vs. Network

Country	Merger	ΔHHI	ΔP% - Merged		ΔP% - Non-Merged		ΔCS%	
			RCNL	Network	RCNL	Network	RCNL	Network
FRA	Peugeot–Renault	1623.5437	15.9407%	8.0120%	0.3048%	0.9834%	-5.4262%	-3.3852%
GBR	BMW–Ford	348.1630	3.7497%	2.8420%	0.0418%	0.3665%	-1.1106%	-0.8891%
GBR	BMW–Opel	255.0992	2.9597%	2.2277%	0.0455%	0.2677%	-0.8305%	-0.6477%
GBR	Ford–Opel	651.7778	5.3994%	3.5448%	0.0782%	0.4839%	-2.7334%	-1.5025%
GER	BMW–Ford	151.0954	1.9720%	2.1863%	0.0343%	0.2974%	-0.3921%	-0.3608%
GER	BMW–MCC	18.7394	0.3090%	0.4900%	0.0026%	0.0260%	-0.0236%	-0.0297%
GER	BMW–Mercedes	155.4387	3.0930%	2.0623%	0.0768%	0.2260%	-0.9659%	-0.3604%
GER	BMW–Opel	234.0387	2.3207%	2.6840%	0.0438%	0.3396%	-0.5545%	-0.4789%
GER	BMW–VW	494.3758	3.9663%	4.4607%	0.0655%	0.8168%	-1.4288%	-1.1102%
GER	Ford–MCC	24.0961	0.9525%	0.7483%	0.0089%	0.0408%	-0.0762%	-0.0516%
GER	Ford–Mercedes	199.8714	2.9257%	2.3549%	0.0523%	0.2588%	-0.5967%	-0.4130%
GER	Ford–Opel	300.9395	3.6512%	3.3459%	0.0794%	0.4110%	-0.9313%	-0.6367%
GER	Ford–VW	635.6948	5.9811%	5.8639%	0.0768%	1.0172%	-1.9570%	-1.4532%
GER	MCC–Mercedes	24.7888	0.6152%	0.9370%	0.0033%	0.0233%	-0.0303%	-0.0359%
GER	MCC–Opel	37.3236	1.6758%	1.1744%	0.0139%	0.0474%	-0.1265%	-0.0711%
GER	MCC–VW	78.8412	1.5702%	1.3616%	0.0120%	0.1153%	-0.2077%	-0.1548%
GER	Mercedes–Opel	309.5901	3.1861%	2.3994%	0.0684%	0.2526%	-0.8399%	-0.5392%
GER	Mercedes–VW	653.9680	4.0745%	3.6063%	0.0948%	0.6038%	-2.0943%	-1.2339%
GER	Opel–VW	984.6571	6.3387%	6.1152%	0.1091%	1.0259%	-3.0836%	-1.9051%

Notes: The table reports simulated effects of hypothetical domestic mergers among automobile manufacturers in 1999. ΔHHI denotes the change in the HHI resulting from the merger, computed using pre-merger market shares and expressed on the 0–10,000 scale. Columns 4–5 report the average percentage price change for the products of the merging firms, while Columns 6–7 report the average percentage price change for the products of non-merging firms. Columns 8–9 report the percentage change in consumer surplus relative to the pre-merger equilibrium. Results are computed using the RCNL model and the network demand model described in the paper. All price and welfare changes are computed from the post-merger equilibrium implied by each model, holding marginal costs fixed at the pre-merger level.

toward the outside option. As a result, market contraction is attenuated, and the overall welfare loss is smaller despite larger price increases.

5.5 Changes in Bonacich Centrality as a Screening Tool

Competition authorities have long relied on concentration indices (post-merger HHI and ΔHHI) to assess the competitive effects of horizontal acquisitions. The 2010 Merger Guidelines establish that markets with an HHI below 1,500 are "unconcentrated" whereas markets with an HHI above 2,500 are "highly concentrated."[25] In "unconcentrated" markets, deals are unlikely to be challenged. In highly concentrated markets, mergers are unlikely to be challenged only when the ΔHHI is below 100. In contrast, mergers are likely to be challenged when the ΔHHI is greater than 200. The other mergers happening in highly concentrated markets are "more likely than not" to be challenged. In the remaining "moderately concentrated" markets, with post-

[25] The 1982 Meger Guidelines already establish similar thresholds. The 2010 Guidelines maintained the structure but made the thresholds less stringent.

merger HHI between 1,500 and 2,500, a merger is considered more likely than not to be challenged if ΔHHI is above 100; otherwise, agencies are unlikely to challenge it.

These thresholds have recently been updated by the 2023 Merger Guidelines, which have re-established, as per the 1982 Merger Guidelines, markets with an HHI greater than 1,800 as highly concentrated. Additionally, the new guidelines introduced a structural presumption: deals leading to a change in HHI greater than 100 in highly concentrated markets or involving parties with a combined market share above 30% are presumed to substantially lessen competition or tend to create a monopoly. We refer to this criterion as the 2023 concentration-based presumption (CBP).

Nocke and Whinston (2022) provides a foundation for using solely ΔHHI in screening mergers for whether their unilateral price effects will harm consumers. In our framework, the decomposition of the price-cost margin into a monopolist and a centrality component (Lemma 1′) and the theoretical merger analysis (Proposition 3′) micro-found the Bonacich Centrality as a novel measure to be used for merger screening. Specifically, we define a market- or industry-level index, which we call *centrality*, as $\boldsymbol{q}'\boldsymbol{b}\left(D_H, \frac{1}{2}, \frac{\boldsymbol{\alpha}-\boldsymbol{c}}{2}\right)$—i.e., the quantity-weighted sum of all products' Bonacich centrality. As mentioned in Section 2, for non-linear demands, the internalized diversion ratio matrix D_H and demand intercept $\boldsymbol{\alpha}$ are obtained through local linearization around the equilibrium prices.[26] Further, we define

$$\Delta\text{Centrality} := \boldsymbol{q}'\boldsymbol{b}\left(D_{H^{\text{pre-merger}}}, \frac{1}{2}, \frac{\boldsymbol{\alpha}-\boldsymbol{c}}{2}\right) - \boldsymbol{q}'\boldsymbol{b}\left(D_{H^{\text{post-merger}}}, \frac{1}{2}, \frac{\boldsymbol{\alpha}-\boldsymbol{c}}{2}\right).$$

Next, we assess the performance of ΔCentrality in predicting the unilateral effects of mergers. We compare the performance with that of ΔHHI, and investigate to what extent the novel measure of ΔCentrality may complement traditional tools adopted in merger screening. We emphasize that ΔCentrality is an industry-level index, much like conventional concentration measures. This is because a merger affects the centrality scores of all products, not only those of the merging parties. As a result, focusing only on the centrality change of the merging products would overlook the part of merger incentives generated by network propagation.

Figure 2 shows the predicted percentage changes across all possible mergers in the prices of merging parties or CS against either ΔCentrality or ΔHHI.[27] The four left figures rely on estimates from the RCNL model, whereas the four right ones use the ones from the network model. Overall, across demand models, mergers leading to a larger decrease in centrality generate greater price increases and reductions in CS.

[26] When demand is estimated using prices normalized by income, the linearized system is $\boldsymbol{q} = -\widetilde{S}(\widetilde{\boldsymbol{\alpha}} - \boldsymbol{p}/y)$; to compute Bonacich centralities in the original price units, we then set $S = \widetilde{S}/y$ and $\boldsymbol{\alpha} = y\widetilde{\boldsymbol{\alpha}}$.

[27] Figure C.2 in the appendix reproduces Figure 2 using the percentage change in centrality.

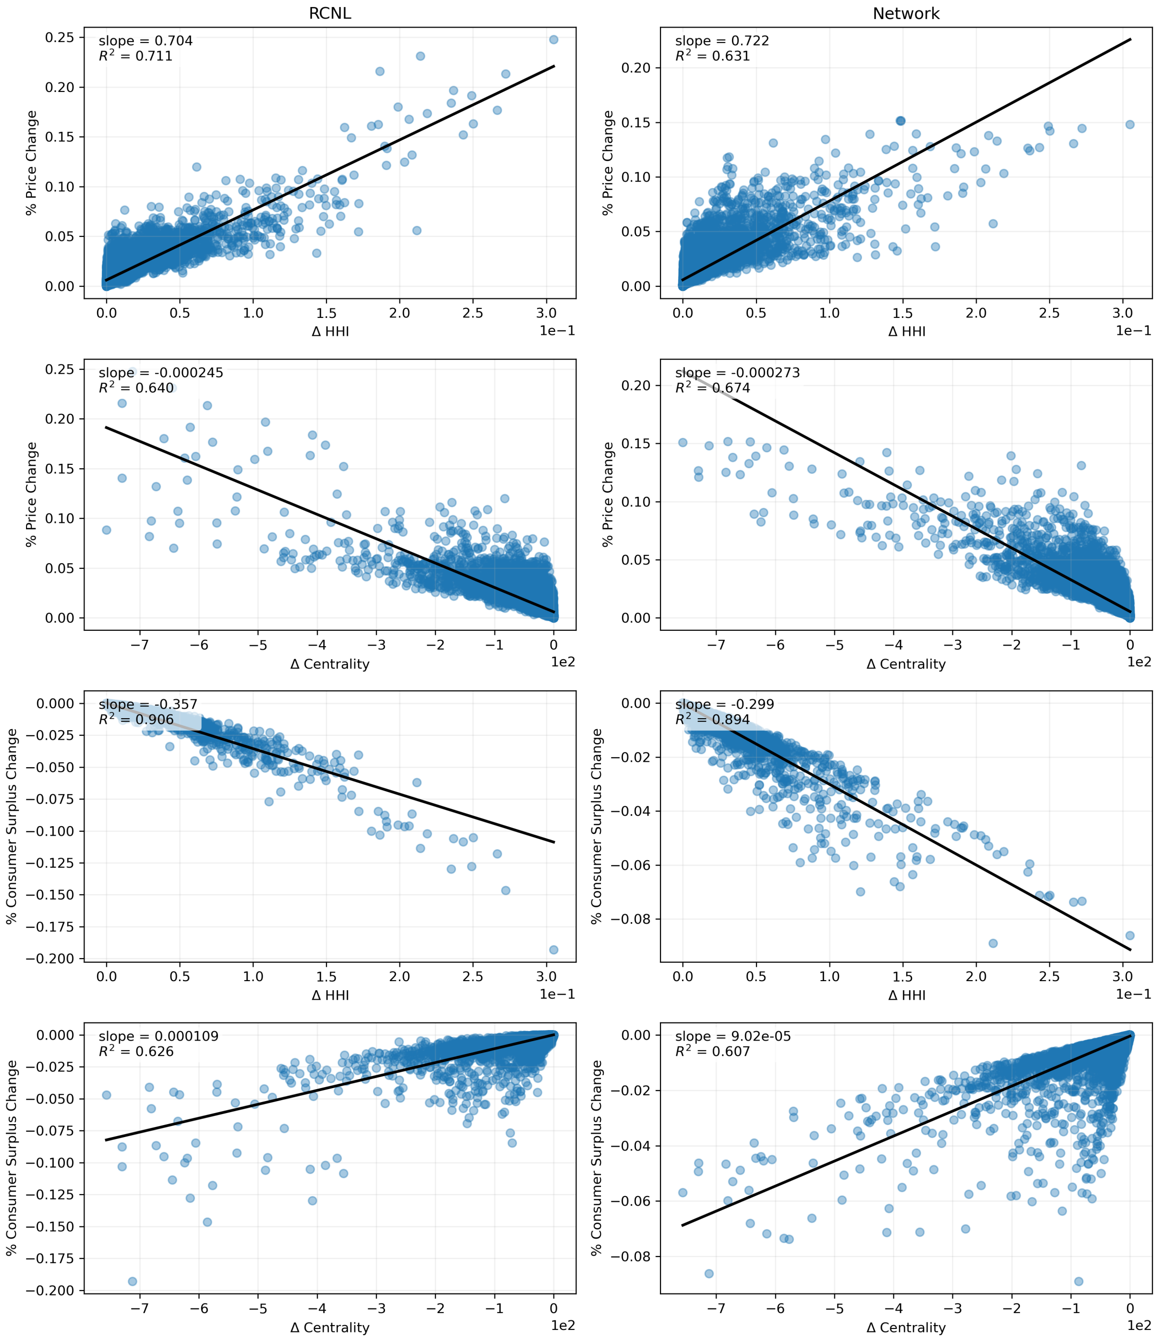

Figure 2: Correlation between the Effects of Simulated Mergers and Screening Tools

Notes: The figure shows the % price or consumer surplus changes against ΔHHI or ΔCentrality (in levels), for all possible mergers across all markets (country-year). The figures on the left (right) use demand estimates from the RCNL (network) model.

Additionally, our analyses confirm the high predictive power of ΔHHI: across both demand models, a higher increase in market concentration is associated with greater price increases by merging parties, as well as larger losses in CS.

The evidence presented so far suggests that ΔCentrality can serve as a useful screening statistic for mergers, while also confirming that ΔHHI performs well on

Table 4: Simulated Price Effects and the Changes in Centrality and Concentration

	(1) $\%\Delta P$	(2) $\%\Delta P$	(3) $\%\Delta P$
Panel A: RCNL			
ΔHHI	0.702***		0.537***
	(0.014)		(0.043)
Δ%Centrality		-2.190***	-0.555***
		(0.096)	(0.160)
N	21786	21786	21786
R^2	0.738	0.701	0.742
Panel B: Network game			
ΔHHI	0.717***		0.168**
	(0.016)		(0.080)
Δ%Centrality		-2.359***	-1.847***
		(0.094)	(0.306)
N	21786	21786	21786
R^2	0.670	0.708	0.712

Notes: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. The dependent variable is the estimated percentage price change for the merging parties following the merger. All regressions include market (country-year) and year fixed effects. Standard errors in parentheses.

its own. A natural question, however, is whether, by capturing substitution patterns overlooked in concentration indices, ΔCentrality adds incremental information beyond ΔHHI. To explore this, we regress the simulated percentage price change of the merging parties (across all possible mergers) first on ΔHHI alone, then on Δ%Centrality alone, and finally on both jointly, controlling for market and year fixed effects.

Panel A reports results using the RCNL demand estimates. Column (1) shows that a 0.01 increase in ΔHHI (corresponding to a 100-point increase in the conventional 0–10,000 HHI scale) is associated with approximately a 0.7 percent increase in post-merger prices. Column (2) instead indicates that a one-percent increase in Δ%Centrality predicts a 2.2 percentage-point decrease in prices. When both indices are included jointly in Column (3), both coefficients remain statistically and economically significant, and the R^2 increases from 0.738 and 0.701 in the univariate specifications to 0.742 in the joint regression.

Panel B presents the same exercise using price effects simulated from the network model. In this case, %ΔCentrality alone explains more variation than ΔHHI (with $R^2 = 0.708$ versus 0.670), and the coefficient on ΔHHI declines substantially once both measures are included jointly. While the joint specification yields only a modest additional increase in fit ($R^2 = 0.712$), both measures remain statistically significant. It is not surprising that ΔCentrality performs particularly well in this setting, as it is micro-founded in the underlying network structure of demand. Importantly, however, ΔCentrality also retains substantial predictive power within the RCNL model. Taken together, these results indicate that changes in centrality provide complementary predictive power relative to traditional concentration measures, highlighting the value of incorporating substitution network structure into merger screening.

Performance in Merger Screening. We evaluate the performance of screening rules based on ΔCentrality. Let $m = 1, \ldots, M$ index potential mergers and $r \in \{\text{strict}, \text{baseline}, \text{lax}\}$ denote enforcement regimes, which differ in their assumed post-merger efficiency realizations. For each merger m and regime r, we compute (i) the change in consumer surplus $\Delta CS_{m,r}$ from a full merger simulation, (ii) the change in centrality ΔC_m, and (iii) whether the merger would be privately profitable and therefore proposed, denoted by $1\{\Pi_{m,r} > 0\}$.[28]

We restrict attention to monotone threshold rules based on ΔC_m. For a given cutoff τ, the regulator challenges the merger m if $\Delta C_m > \tau$ and allows it otherwise. For simplicity, we assume that challenged deals are blocked, though we recognize that, in practice, such mergers may still be allowed after further analysis. Let $d_m(\tau) = 1\{\Delta C_m > \tau\}$ denote the blocking decision. Conditional on the set of proposed mergers under regime r, $\mathcal{P}_r = \{m : \Pi_{m,r} > 0\}$, the regulator's objective is to maximize expected consumer surplus from allowed mergers:

$$W_r(\tau) = \sum_{m \in \mathcal{P}_r} (1 - d_m(\tau))\, \Delta CS_{m,r}. \tag{38}$$

The welfare-optimal ΔCentrality threshold in regime r is therefore

$$\tau_r^* = \arg\max_{\tau} W_r(\tau). \tag{39}$$

Because decisions change only when τ crosses an observed value of ΔC_m, the maximization can be performed by evaluating $W_r(\tau)$ over the finite set of candidate thresholds. For screening rules using concentration indices, we consider the structural presumption introduced by the 2023 Merger Guidelines. When we analyze hybrid regimes combining ΔCentrality with the structural presumption, we hold the concentration threshold defined by the presumption fixed and optimize only over τ. Optimization is conducted separately for each enforcement regime and each decision rule. The analysis relies on demand estimates from the RCNL model.

Table 5 summarizes our results.[29] Across regimes, ΔCentrality delivers meaning-

[28] As before, the unit of analysis is a pair of firms in a given market (country-year). Thus, the same merger appears as a separate observation across markets.

[29] We abstract from the dynamic effects of enforcement regimes. In practice, stricter regimes may affect consumer surplus by deterring merger activity, potentially discouraging both harmful and beneficial transactions. Moreover, we implicitly assume that merging parties will pass through realized

Table 5: Merger Screening with ΔCentrality

	Proposed mergers	Allowed mergers (%)	ΔCS for allowed mergers ($000)	Type I Errors (%)	Type II Errors (%)
Panel A: Strict Regime (2% assumed efficiency)					
ΔHHI^*	21,779	86.9	3.125	2.8	33.4
CBP	21,779	93.4	2.406	1.5	66.5
ΔCentrality	21,779	86.0	2.992	4.6	36.7
CBP or ΔCentrality	21,779	85.8	2.997	4.7	36.3
CBP & ΔCentrality	21,779	93.5	2.401	1.4	66.9
Panel B: Baseline Regime (5% assumed efficiency)					
ΔHHI^*	21,786	96.6	12.305	0.7	41.1
CBP	21,786	93.3	11.748	3.9	36.5
ΔCentrality	21,786	96.4	12.186	1.2	47.0
CBP or ΔCentrality	21,786	93.1	11.749	4.0	34.1
CBP & ΔCentrality	21,786	96.7	12.186	1.1	49.4
Panel C: Lax Regime (10% assumed efficiency)					
ΔHHI^*	21,786	99.4	31.150	0.1	44.5
CBP	21,786	93.3	28.585	5.8	4.5
ΔCentrality	21,786	99.0	31.083	0.5	38.5
CBP or ΔCentrality	21,786	93.3	28.585	5.8	4.5
CBP & ΔCentrality	21,786	99.0	31.083	0.5	38.5

Notes: The table shows the number of proposed mergers (i.e., those that are profitable), the percentage of deals allowed, the average change in CS for allowed mergers, the probability of blocking a welfare-enhancing deal (type I error) and that of approving a welfare-reducing deal (type II error) under three different regimes: strict, which assumes a 2% marginal cost efficiency (Panel A); baseline, which assumes a 5% marginal cost efficiency (Panel B); and lax, which assumes a 10% marginal cost efficiency (Panel C). Four screening approaches using the CBP and the ΔCentrality are considered. The thresholds for the ΔCentrality are optimally set for each regime-decision rule pair to maximize the total CS after the merger. As a benchmark, we also show the results obtained by screening mergers by optimally choosing the ΔHHI threshold (ΔHHI^*).

ful improvements relative to the current CBP thresholds. In the strict regime, where assumed cost efficiencies reduce marginal costs by 2%, centrality-based screening increases average consumer surplus by nearly 25% relative to CBP, from an average of about $2,400 per-merger to roughly $3,000. In the lax regime, with a 10% assumed cost efficiency, the increase in consumer surplus induced by the ΔCentrality thresholds

efficiencies, though this may not necessarily be the case in practice (Harrington, 2021; Leccese et al., 2025).

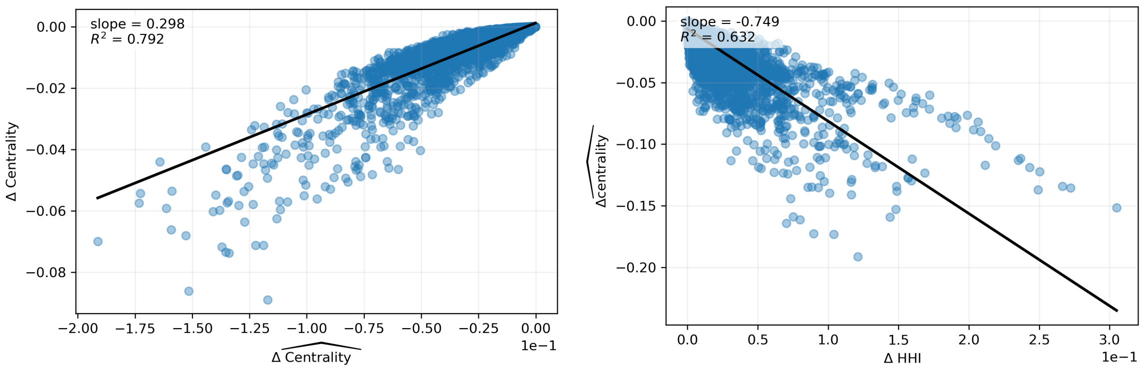

Figure 3: Unweighted Change in Centrality

Notes: The left (right) figure illustrates the relationship between the unweighted change in centrality and ΔCentrality (ΔHHI) across all mergers.

is about 9% relative to CBP.

Although centrality-based screening improves on CBP in all regimes, the sources of these gains differ. In the strict regime, the ΔCentrality rule blocks more mergers than CBP. As a result, it is more successful at preventing harmful mergers and reduces what we refer to as Type II error by nearly half, from 66.5% to 36.7%. In the lax regime, the pattern reverses. Centrality-based screening is more successful at approving welfare-enhancing mergers, reducing the Type I error.

Quantitatively, the gains from ΔCentrality screening track closely those generated by the welfare optimal ΔHHI^* thresholds, although ΔCentrality does not outperform them. There are at least two reasons. First, the auto industry has relatively clear market segments, so product substitution patterns are likely already well summarized by market shares. This limits the scope for ΔCentrality to improve on ΔHHI. Because our centrality measure incorporates product proximity in characteristic space, it may be more informative in industries where market boundaries are less well defined and substitution patterns more heterogeneous. Second, the reference model used for the screening analysis in Table 5 is the RCNL. In logit-type models, substitution patterns are typically well summarized by market shares, which gives a mechanical advantage to HHI-based screening measures. By contrast, centrality measures rely on distances in product characteristics. Under alternative demand models, centrality screening may instead have an advantage.

Unweighted centrality. One caveat of the analyses presented is that, in practice, merger screening requires tools that are easy to construct with limited data. Our theoretical measure of ΔCentrality requires estimating the full network model to recover $(\boldsymbol{\alpha} - \boldsymbol{c})/2$, which may not always be feasible for enforcement agencies operating under time and data constraints. To address this concern, we construct a simpler measure that replaces $(\boldsymbol{\alpha} - \boldsymbol{c})/2$ with a vector of ones 1_J. The resulting statistic corresponds to an unweighted Bonacich centrality given by:

$$\widehat{\boldsymbol{b}}\left(D_H, \frac{1}{2}, 1_J\right) = \left(I_J - \frac{1}{2}D_H\right)^{-1} \frac{1}{2}D_H 1_J, \tag{40}$$

which can be computed directly from product characteristics, and still encapsulates diversion information. In particular, we construct Θ using the same characteristics that enter indirect utility in Brenkers and Verboven (2006).[30] For any potential merger, we then compute $\Delta\widehat{\text{Centrality}}$ as the change induced in (40). Figure 3 shows the strong correlation between the theoretical ΔCentrality and its unweighted counterpart, demonstrating that this simpler measure captures most of the economically relevant variation in our structural index. Moreover, we also show that $\Delta\widehat{\text{Centrality}}$ is highly correlated with the ΔHHI.

Importantly, this approximation only requires product characteristics, which are typically publicly available for both merging and non-merging parties. Advances in data collection and machine learning tools allow agencies to quickly extract such information from firm websites and industry sources, enabling low-cost merger screening that we have shown can complement traditional concentration-based measures.

6 Conclusion

This paper develops a tractable framework for analyzing horizontal mergers and acquisitions in differentiated-product industries using product networks. The key object is the diversion network: products are connected according to how demand is reallocated after a price change, and firms' pricing incentives depend on their position in the network. We show that equilibrium markups admit an additive decomposition into a monopolistic component and a Bonacich-centrality component. Mergers matter because they change how firms internalize diversion and therefore reshape the network that governs competitive pressure.

This perspective yields transparent formulas for the unilateral effects of ownership changes. For full mergers, minority stake acquisitions, and more general control structure changes, the framework delivers closed-form expressions for price effects, consumer-surplus changes, and the compensating marginal cost reductions needed to offset the anti-competitive effects of ownership concentration. A central implication is that merger harm is determined not only by direct diversion between the parties, but also by the way pricing incentives propagate through the broader product network. This makes changes in centrality a natural screening statistic for unilateral effects.

We also provide a micro-foundation that links the diversion network directly to

[30] Recall that η is calibrated to 0.12.

product characteristics. Under linear-quadratic demand, diversion ratios are weighted inner products of product attributes. This gives the approach practical appeal: unlike detailed price–quantity data, product characteristics are often readily observable and can increasingly be collected at scale, including from unstructured sources with the help of LLMs. Moreover, our approach does not require an ex ante market definition or a nest structure, and it extends naturally to settings with partial ownership and common ownership where concentration-based measures are less straightforward to interpret.

In our application to the European automobile industry, the network model performs well despite not imposing any substitution nests ex ante. It generates substitution patterns and simulated merger impacts that are highly comparable to those from an RCNL model. We also find that changes in centrality alone are very informative about merger harm, and this measure improves on the concentration-based presumption introduced by the 2023 Merger Guidelines. At the same time, the gains relative to HHI-based benchmarks are moderate in the automobile industry, partly reflecting its well-defined industry boundaries and segments.

References

Antón, M., F. Ederer, M. Giné, and M. Schmalz (2023): "Common ownership, competition, and top management incentives," *Journal of Political Economy*, 131, 1294–1355.

Atalay, E., E. Frost, A. T. Sorensen, C. J. Sullivan, and W. Zhu (2023): "Scalable demand and markups," Tech. rep., National Bureau of Economic Research.

Azar, J., S. Raina, and M. Schmalz (2022): "Ultimate ownership and bank competition," *Financial Management*, 51, 227–269.

Azar, J., M. C. Schmalz, and I. Tecu (2018): "Anticompetitive effects of common ownership," *The Journal of Finance*, 73, 1513–1565.

Backus, M., C. Conlon, and M. Sinkinson (2021a): "Common Ownership and Competition in the Ready-To-Eat Cereal Industry," *NBER working paper*.

——— (2021b): "Common ownership in America: 1980–2017," *American Economic Journal: Microeconomics*, 13, 273–308.

Bajari, P., Z. Cen, V. Chernozhukov, M. Manukonda, S. Vijaykumar, J. Wang, R. Huerta, J. Li, L. Leng, G. Monokroussos, et al. (2023): "Hedonic prices and quality adjusted price indices powered by AI," *arXiv preprint arXiv:2305.00044*.

Ballester, C., A. Calvó-Armenagol, and Y. Zenou (2006): "Who's Who in Networks. Wanted: The Key Player," *Econometrica*, 74, 1403–1417.

Berry, S. (1994): "Estimating Discrete-Choice Models of Product Differentiation," *The RAND Journal of Economics*, 25, 242–262.

Berry, S., J. Levinshon, and A. Pakes (1999): "Voluntary Export Restraints on Automobiles: Evaluating a Trade Policy," *American Economic Review*, 89, 400–430.

Berry, S., J. Levinsohn, and A. Pakes (1995): "Automobile Prices in Market Equilibrium," *Econometrica*, 63, 841–890.

Brenkers, R. and F. Verboven (2006): "Liberalizing a distribution system: The European car market," *Journal of the European Economic Association*, 4, 216–251.

Bresnahan, T. F. (1987): "Competition and collusion in the American automobile industry: The 1955 price war," *The Journal of Industrial Economics*, 457–482.

Bresnahan, T. F. and S. C. Salop (1986): "Quantifying the competitive effects of production joint ventures," *International Journal of Industrial Organization*, 4, 155–175.

Compiani, G., I. Morozov, and S. Seiler (2025): "Demand estimation with text and image data," *arXiv preprint arXiv:2503.20711.*

Conlon, C. and J. Gortmaker (2020): "Best practices for differentiated products demand estimation with PyBLP," *RAND Journal of Economics*, 51, 1108–1161.

Conlon, C., J. Mortimer, and P. Sarkis (2023): "Estimating preferences and substitution patterns from second choice data alone," *Preliminary and incomplete.*

Conlon, C. and J. H. Mortimer (2021): "Empirical properties of diversion ratios," *The RAND Journal of Economics*, 52, 693–726.

Crooke, P., L. Froeb, S. Tschantz, and G. J. Werden (1999): "Effects of assumed demand form on simulated postmerger equilibria," *Review of Industrial Organization*, 15, 205–217.

Dubois, P., R. Griffith, and A. Nevo (2014): "Do Prices and Attributes Explain International Differences in Food Purchases?" *American Economic Review*, 104, 832–867.

Ederer, F. and B. Pellegrino (2025): "A Tale of Two Networks: Common Ownership and Product Market Rivalry," *Review of Economic Studies*, rdaf057.

Einav, L., M. D. Guido, and P. J. Klenow (2026): "Customer Overlap and Diversion Ratios," Tech. rep., National Bureau of Economic Research.

Fan, Y. and C. Yang (2025): "Estimating discrete games with many firms and many decisions: An application to merger and product variety," *Journal of Political Economy*, 133, 1886–1931.

Farrell, J. and C. Shapiro (2009): "Recapture, pass-through, and market definition," *Antitrust LJ*, 76, 585.

——— (2010): "Antitrust evaluation of horizontal mergers: An economic alternative to market definition," *BE Journal of Theoretical Economics*, 10, 1–41.

Feenstra, R. and J. Levinsohn (1995): "Estimating Markups and Market Conduct with Multidimensional Product Attributes," *Review of Economic Studies*, 62, 19–52.

Fosgerau, M., J. Monardo, and A. de Palma (2024): "The Inverse Product Differentiation Logit Model," *American Economic Journal: Microeconomics*, 16, 329–70.

Galeotti, A., B. Golub, S. Goyal, E. Talamas, and O. Tamuz (2024): "Robust market interventions," *arXiv preprint arXiv:2411.03026*.

Gandhi, A. and J. Houde (2023): "Measuring Substitution Patterns in Differentiated-Products Industries," NBER working paper.

Gentzkow, M. (2007): "Valuing New Goods in a Model with Complementarity: Online Newspapers," *American Economic Review*, 97, 713–744.

Goldberg, P. K. and F. Verboven (2001): "The evolution of price dispersion in the European car market," *The Review of Economic Studies*, 68, 811–848.

——— (2005): "Market integration and convergence to the Law of One Price: evidence from the European car market," *Journal of International Economics*, 65, 49–73.

Grigolon, L. and F. Verboven (2014): "Nested logit or random coefficients logit? A comparison of alternative discrete choice models of product differentiation," *Review of Economics and Statistics*, 96, 916–935.

Harrington, J. E. (2021): "There may be no pass through of a merger-related cost efficiency," *Economics Letters*, 208, 110050.

Hoberg, G. and G. Phillips (2016): "Text-Based Network Industries and Endogenous Product Differentiation," *Journal of Political Economy*, 124, 1423–1465.

Jackson, M. (2008): *Social and Economic Networks*, Princeton University Press.

Jaffe, S. and E. G. Weyl (2013): "The First-Order Approach to Merger Analysis," *American Economic Journal: Microeconomics*, 5, 188–218.

Kaplow, L. (2015): "Market definition, market power," *International Journal of Industrial Organization*, 43, 148–161.

Leccese, M., A. Sweeting, and X. Tao (2025): "Should We Expect Merger Synergies to be Passed Through to Consumers?" *The Journal of Industrial Economics*, 73, 1–30.

Lee, S. and G. Allenby (2009): "A Direct Utility Model for Market Basket Data," *working paper*.

López, Á. L. and X. Vives (2019): "Overlapping ownership, R&D spillovers, and antitrust policy," *Journal of Political Economy*, 127, 2394–2437.

LOSETO, M. (2023): "Network Games of Imperfect Competition: An Empirical Framework," SSRN working paper.

MAGNOLFI, L., J. MCCLURE, AND A. SORENSEN (2022): "Triplet Embeddings for Demand Estimation," SSRN working paper.

MILLER, N. H., M. REMER, C. RYAN, AND G. SHEU (2016): "Pass-through and the prediction of merger price effects," *The Journal of Industrial Economics*, 64, 683–709.

——— (2017): "Upward pricing pressure as a predictor of merger price effects," *International Journal of Industrial Organization*, 52, 216–247.

NEVO, A. (2000): "Mergers with differentiated products: The case of the ready-to-eat cereal industry," *The Rand journal of economics*, 395–421.

NOCKE, V. AND N. SCHUTZ (2025): "An Aggregative Games Approach to Merger Analysis in Multiproduct-Firm Oligopoly," *The RAND Journal of Economics*, 56, 233–250.

NOCKE, V. AND M. D. WHINSTON (2022): "Concentration Thresholds for Horizontal Mergers," *American Economic Review*, 112, 1915–48.

PELLEGRINO, B. (2025): "Product differentiation and oligopoly: A network approach," *American Economic Review*, 115, 1170–1225.

QIU, Y. J., M. SAWADA, AND G. SHEU (2024): "Win/loss data and consumer switching costs: Measuring diversion ratios and the impact of mergers," *The Journal of Industrial Economics*, 72, 327–355.

SCHMALZ, M. C. (2018): "Common-ownership concentration and corporate conduct," *Annual Review of Financial Economics*, 10, 413–448.

SINGH, N. AND X. VIVES (1984): "Price and quantity competition in a differentiated duopoly," *RAND Journal of Economics*, 15, 546–554.

THOMASSEN, O., H. SMITH, S. SEILER, AND P. SCHIRALDI (2017): "Multi-Category Competition and Market Power: A Model of Supermarket Pricing," *American Economic Review*, 107, 2308–2351.

USHCHEV, P. AND Y. ZENOU (2018): "Price competition in product variety networks," *Games and Economic Behavior*, 110, 226–247.

VIVES, X. (1985): "On the efficiency of Bertrand and Cournot equilibria with product differentiation," *Journal of Economic Theory*, 36, 166–175.

A Proofs

A.1 Proof of Lemma 1

Proof. This is the special case of Lemma 1′ with $H^\varepsilon = I_J$, so $D_{H^\varepsilon} = D$ and the stated expression follows immediately. ■

A.2 Proof of Lemma 1′

Proof. Fix the local linearization and suppress the superscript ε for notational convenience. The locally linear demand system can be written as $\boldsymbol{q} = -S(\boldsymbol{\alpha} - \boldsymbol{p})$. Let $\boldsymbol{m} := \boldsymbol{p} - \boldsymbol{c}$.

Under the decision-weight matrix H, the decision maker for product j chooses p_j to maximize

$$\max_{p_j} \sum_{\ell=1}^{J} H_{j\ell}\,(p_\ell - c_\ell)\,q_\ell(\boldsymbol{p}),$$

taking p_{-j} as given. Differentiating the objective with respect to p_j gives

$$\begin{aligned}
0 &= \frac{\partial}{\partial p_j}\sum_{\ell=1}^{J} H_{j\ell}\,(p_\ell - c_\ell)\,q_\ell(\boldsymbol{p}) \\
&= \sum_{\ell=1}^{J} H_{j\ell}\left[\underbrace{\frac{\partial(p_\ell - c_\ell)}{\partial p_j}}_{=\,\mathbf{1}\{\ell=j\}}\; q_\ell(\boldsymbol{p}) \;+\; (p_\ell - c_\ell)\,\underbrace{\frac{\partial q_\ell(\boldsymbol{p})}{\partial p_j}}_{=:\,S_{\ell j}(\boldsymbol{p})}\right] \\
&= H_{jj}\,q_j(\boldsymbol{p}) \;+\; \sum_{\ell=1}^{J} H_{j\ell}\,(p_\ell - c_\ell)\,S_{\ell j}(\boldsymbol{p}).
\end{aligned}$$

Using $H_{jj} = 1$, the j-th first-order condition can be written as

$$q_j(\boldsymbol{p}) \;+\; \sum_{\ell=1}^{J} H_{j\ell}\,S_{\ell j}(\boldsymbol{p})\,m_\ell \;=\; 0, \qquad j = 1, \ldots, J. \tag{41}$$

By symmetry of S, the stacked FOCs are

$$\boldsymbol{q}(\boldsymbol{p}) + (H \odot S(\boldsymbol{p}))(\boldsymbol{p} - \boldsymbol{c}) = 0. \tag{42}$$

Substituting $\boldsymbol{q} = -S(\boldsymbol{\alpha} - \boldsymbol{p})$ into (42) and writing $\boldsymbol{p} = \boldsymbol{c} + \boldsymbol{m}$ gives

$$-S(\boldsymbol{\alpha} - \boldsymbol{c} - \boldsymbol{m}) + (H \odot S)\boldsymbol{m} = 0 \iff -S(\boldsymbol{\alpha} - \boldsymbol{c}) + (S + H \odot S)\boldsymbol{m} = 0.$$

Premultiplying by $(H \odot S)^{-1}$ and using

$$D_H = I_J - (H \odot S)^{-1} S \quad \Longrightarrow \quad (H \odot S)^{-1} S = I_J - D_H,$$

we obtain

$$(2I_J - D_H)\boldsymbol{m} = (I_J - D_H)(\boldsymbol{\alpha} - \boldsymbol{c}).$$

Hence

$$\boldsymbol{m} = (2I_J - D_H)^{-1} (I_J - D_H)(\boldsymbol{\alpha} - \boldsymbol{c}) = \Big(I_J - \frac{1}{2} D_H\Big)^{-1} (I_J - D_H) \frac{\boldsymbol{\alpha} - \boldsymbol{c}}{2}.$$

Let $\boldsymbol{u} := (\boldsymbol{\alpha} - \boldsymbol{c})/2$. Since

$$\boldsymbol{b}\Big(D_H, \frac{1}{2}, \boldsymbol{u}\Big) = \Big(I_J - \frac{1}{2} D_H\Big)^{-1} \frac{1}{2} D_H \, \boldsymbol{u},$$

we have

$$\boldsymbol{u} - \boldsymbol{b}\Big(D_H, \frac{1}{2}, \boldsymbol{u}\Big) = \Big(I_J - \frac{1}{2} D_H\Big)^{-1} (I_J - D_H)\boldsymbol{u} = \boldsymbol{m}.$$

Therefore

$$\boldsymbol{p} - \boldsymbol{c} = \frac{\boldsymbol{\alpha} - \boldsymbol{c}}{2} - \boldsymbol{b}\Big(D_H, \frac{1}{2}, \frac{\boldsymbol{\alpha} - \boldsymbol{c}}{2}\Big),$$

which is the claimed expression (and replacing $(\boldsymbol{\alpha}, \boldsymbol{c}, H)$ by their ε-perturbed counterparts yields the stated ε version). ■

A.3 Proof of Proposition 1

Proof. This is Proposition 1′ specialized to $H = I_J$ and $\Delta H = E_{12}$, for which $D_H = D$ and $\boldsymbol{\omega}_{\Delta H}(H) = D_{12}(p_2 - c_2)\boldsymbol{e}_1$, and Roy's identity then gives the CS expression. ■

A.4 Proof of Proposition 1′

Proof. Consider $H^\varepsilon = H + \varepsilon \Delta H$, with $\boldsymbol{\alpha}, \boldsymbol{c}, S = \frac{d\boldsymbol{q}}{d\boldsymbol{p}'}$ fixed under local linearization. The locally linearized demand is

$$\boldsymbol{q}(\boldsymbol{p}^\varepsilon) = -\, S(\boldsymbol{\alpha} - \boldsymbol{p}^\varepsilon).$$

The FOC is

$$\boldsymbol{q}(\boldsymbol{p}^\varepsilon) + (H^\varepsilon \odot S)\,(\boldsymbol{p}^\varepsilon - \boldsymbol{c}) = 0,$$

or equivalently

$$-\, S(\boldsymbol{\alpha} - \boldsymbol{p}^\varepsilon) + (H^\varepsilon \odot S)\,(\boldsymbol{p}^\varepsilon - \boldsymbol{c}) = 0.$$

Differentiating at $\varepsilon = 0$ gives

$$S\dot{\boldsymbol{p}} + (\Delta H \odot S)(\boldsymbol{p} - \boldsymbol{c}) + (H \odot S)\dot{\boldsymbol{p}} = 0,$$

so

$$\big(S + H \odot S\big)\dot{\boldsymbol{p}} = -(\Delta H \odot S)(\boldsymbol{p} - \boldsymbol{c}).$$

Premultiplying by $(H \odot S)^{-1}$ and using $(H \odot S)^{-1}S = I_J - D_H$ yields

$$(2I_J - D_H)\dot{\boldsymbol{p}} = -\,(H \odot S)^{-1}\Big[(\Delta H \odot S)(\boldsymbol{p} - \boldsymbol{c})\Big] = \boldsymbol{\omega}_{\Delta H}(H).$$

Therefore

$$\dot{\boldsymbol{p}} = (2I_J - D_H)^{-1}\,\boldsymbol{\omega}_{\Delta H}(H) = \frac{1}{2}\left(I_J - \frac{1}{2}D_H\right)^{-1}\boldsymbol{\omega}_{\Delta H}(H) = \frac{1}{2}\,\boldsymbol{\kappa}\Big(D_H, \tfrac{1}{2}, \boldsymbol{\omega}_{\Delta H}(H)\Big)\,,$$

which is (12). By Roy's identity, evaluated at the baseline equilibrium,

$$\left.\frac{\partial CS}{\partial \varepsilon}\right|_{\varepsilon=0} = -\boldsymbol{q}'\dot{\boldsymbol{p}} = -\frac{1}{2}\,\boldsymbol{q}'\boldsymbol{\kappa}\Big(D_H, \tfrac{1}{2}, \boldsymbol{\omega}_{\Delta H}(H)\Big)\,,$$

which is (13). ■

A.5 Proof of Proposition 2

Proof. This is Proposition 2′ specialized to $H = I_J$, $\Delta H = E_{12}$, and $\left.\frac{dc_1}{d\varepsilon}\right|_0 = \left.\frac{dc_2}{d\varepsilon}\right|_0 = \frac{dc}{d\varepsilon}$ with all other cost reductions equal to zero, together with $\boldsymbol{\omega}_{E_{12}}(I_J) = D_{12}(p_2 - c_2)\boldsymbol{e}_1$ from Proposition 1. ■

A.6 Proof of Proposition 2′

Proof. Let $H^\varepsilon = H + \varepsilon\Delta H$ and write post-change costs as $\boldsymbol{c}^\varepsilon = \boldsymbol{c} - \varepsilon\boldsymbol{s}$, where s_j is a marginal cost reduction. The locally linearized demand is $\boldsymbol{q}(\boldsymbol{p}^\varepsilon) = -S(\boldsymbol{\alpha} - \boldsymbol{p}^\varepsilon)$, so the FOC is

$$-\,S(\boldsymbol{\alpha} - \boldsymbol{p}^\varepsilon) + (H^\varepsilon \odot S)\big(\boldsymbol{p}^\varepsilon - \boldsymbol{c}^\varepsilon\big) = 0.$$

Differentiating at $\varepsilon = 0$:

$$S\dot{\boldsymbol{p}} + (\Delta H \odot S)(\boldsymbol{p} - \boldsymbol{c}) + (H \odot S)(\dot{\boldsymbol{p}} + \boldsymbol{s}) = 0.$$

Hence

$$\big(S + H \odot S\big)\dot{\boldsymbol{p}} = -(\Delta H \odot S)(\boldsymbol{p} - \boldsymbol{c}) - (H \odot S)\boldsymbol{s}.$$

Premultiplying by $(H \odot S)^{-1}$:

$$(2I_J - D_H)\dot{\boldsymbol{p}} = \boldsymbol{\omega}_{\Delta H}(H) - \boldsymbol{s},$$

so

$$\dot{\boldsymbol{p}} = \frac{1}{2}\boldsymbol{\kappa}\Big(D_H, \tfrac{1}{2}, \boldsymbol{\omega}_{\Delta H}(H) - \boldsymbol{s}\Big).$$

CS-neutrality requires $0 = \boldsymbol{q}'\dot{\boldsymbol{p}}$, therefore

$$\boldsymbol{q}'\boldsymbol{\kappa}\Big(D_H, \tfrac{1}{2}, \boldsymbol{\omega}_{\Delta H}(H)\Big) = \boldsymbol{q}'\boldsymbol{\kappa}\Big(D_H, \tfrac{1}{2}, \boldsymbol{s}\Big).$$

By linearity in node weights and $\boldsymbol{s} = \sum_{j=1}^{J} s_j \boldsymbol{e}_j$,

$$\boldsymbol{q}'\boldsymbol{\kappa}\Big(D_H, \tfrac{1}{2}, \boldsymbol{s}\Big) = \sum_{j=1}^{J} s_j\, \boldsymbol{q}'\boldsymbol{\kappa}(j; D_H),$$

which gives the stated condition. ■

A.7 Proof of Proposition 3

Proof. This is Proposition $3'$ specialized to $H = I_J$, $\widehat{H} = M_{12}$, and (for (10)) a common cost reduction vector $\Delta\boldsymbol{c} = \Delta c\,(\boldsymbol{e}_1 + \boldsymbol{e}_2)$. ■

A.8 Proof of Proposition $3'$

Proof. From Lemma $1'$, with no cost synergy,

$$\boldsymbol{p}(H) - \boldsymbol{c} = \frac{\boldsymbol{\alpha} - \boldsymbol{c}}{2} - \boldsymbol{b}\Big(D_H, \frac{1}{2}, \frac{\boldsymbol{\alpha} - \boldsymbol{c}}{2}\Big),$$

$$\boldsymbol{p}(\widehat{H}) - \boldsymbol{c} = \frac{\boldsymbol{\alpha} - \boldsymbol{c}}{2} - \boldsymbol{b}\Big(D_{\widehat{H}}, \frac{1}{2}, \frac{\boldsymbol{\alpha} - \boldsymbol{c}}{2}\Big).$$

Subtracting gives

$$\Delta\boldsymbol{p} := \boldsymbol{p}(\widehat{H}) - \boldsymbol{p}(H) = \boldsymbol{b}\Big(D_H, \frac{1}{2}, \frac{\boldsymbol{\alpha} - \boldsymbol{c}}{2}\Big) - \boldsymbol{b}\Big(D_{\widehat{H}}, \frac{1}{2}, \frac{\boldsymbol{\alpha} - \boldsymbol{c}}{2}\Big),$$

which is (15).

Now add a cost-reduction vector $\Delta\boldsymbol{c}$ so post-change costs are $\boldsymbol{c} - \Delta\boldsymbol{c}$ under $\widehat{H}$. Again by Lemma $1'$,

$$\begin{aligned}
\boldsymbol{p}(\widehat{H}, \Delta\boldsymbol{c}) - (\boldsymbol{c} - \Delta\boldsymbol{c}) &= \frac{\boldsymbol{\alpha} - (\boldsymbol{c} - \Delta\boldsymbol{c})}{2} - \boldsymbol{b}\Big(D_{\widehat{H}}, \frac{1}{2}, \frac{\boldsymbol{\alpha} - (\boldsymbol{c} - \Delta\boldsymbol{c})}{2}\Big) \\
&= \frac{\boldsymbol{\alpha} - \boldsymbol{c} + \Delta\boldsymbol{c}}{2} - \boldsymbol{b}\Big(D_{\widehat{H}}, \frac{1}{2}, \frac{\boldsymbol{\alpha} - \boldsymbol{c} + \Delta\boldsymbol{c}}{2}\Big).
\end{aligned}$$

Using linearity of $\boldsymbol{b}$ in node weights and $\boldsymbol{p}(\widehat{H})$ above,

$$\boldsymbol{p}(\widehat{H}, \Delta\boldsymbol{c}) - \boldsymbol{p}(\widehat{H}) = -\frac{\Delta\boldsymbol{c}}{2} - \boldsymbol{b}\Big(D_{\widehat{H}}, \frac{1}{2}, \frac{\Delta\boldsymbol{c}}{2}\Big) = -\frac{1}{2}\boldsymbol{\kappa}\Big(D_{\widehat{H}}, \frac{1}{2}, \Delta\boldsymbol{c}\Big).$$

Therefore total price change relative to the baseline H is

$$\Delta\boldsymbol{p}_{\text{total}} = \boldsymbol{b}\Big(D_H, \frac{1}{2}, \frac{\boldsymbol{\alpha}-\boldsymbol{c}}{2}\Big) - \boldsymbol{b}\Big(D_{\widehat{H}}, \frac{1}{2}, \frac{\boldsymbol{\alpha}-\boldsymbol{c}}{2}\Big) - \frac{1}{2}\boldsymbol{\kappa}\Big(D_{\widehat{H}}, \frac{1}{2}, \Delta\boldsymbol{c}\Big).$$

Imposing CS-neutrality in the quantity-weighted approximation, $\boldsymbol{q}'\Delta\boldsymbol{p}_{\text{total}} = 0$, gives

$$\frac{1}{2}\,\boldsymbol{q}'\boldsymbol{\kappa}\Big(D_{\widehat{H}}, \frac{1}{2}, \Delta\boldsymbol{c}\Big) = \boldsymbol{q}'\Big[\boldsymbol{b}\Big(D_H, \frac{1}{2}, \frac{\boldsymbol{\alpha}-\boldsymbol{c}}{2}\Big) - \boldsymbol{b}\Big(D_{\widehat{H}}, \frac{1}{2}, \frac{\boldsymbol{\alpha}-\boldsymbol{c}}{2}\Big)\Big].$$

Finally, by linearity and $\Delta\boldsymbol{c} = \sum_{j=1}^{J} \Delta c_j \boldsymbol{e}_j$,

$$\boldsymbol{q}'\boldsymbol{\kappa}\Big(D_{\widehat{H}}, \frac{1}{2}, \Delta\boldsymbol{c}\Big) = \sum_{j=1}^{J} \boldsymbol{q}'\boldsymbol{\kappa}(j; D_{\widehat{H}})\Delta c_j,$$

which is (16). ■

A.9 Proof of Proposition 4

Proof. From (18),

$$\boldsymbol{q}(\boldsymbol{p}) = \frac{1}{\beta}\left(I_J + \eta XX'\right)^{-1}(\boldsymbol{\alpha} - \boldsymbol{p}).$$

By the Woodbury matrix identity:

$$\left(I_J + \eta XX'\right)^{-1} = I_J - X\left(\frac{1}{\eta}I_K + X'X\right)^{-1} X' = I_J - \Theta,$$

where $\Theta = X\Omega^{-1}X'$ and $\Omega = \frac{1}{\eta}I_K + X'X$. Therefore

$$\boldsymbol{q}(\boldsymbol{p}) = \frac{1}{\beta}(I_J - \Theta)(\boldsymbol{\alpha} - \boldsymbol{p}),$$

which is (20). Taking the j-th row yields

$$q_j(\boldsymbol{p}) = a_j - \frac{1}{\beta}(1 - \theta_{jj})p_j + \frac{1}{\beta}\sum_{\ell \neq j}\theta_{j\ell}p_\ell,$$

with $\boldsymbol{a} = \frac{1}{\beta}(I_J - \Theta)\boldsymbol{\alpha}$. ■

A.10 Proof of Proposition 5

Proof. By definition,

$$\theta_{j\ell} = \boldsymbol{x}_j' \Omega^{-1} \boldsymbol{x}_\ell, \qquad \Omega = \frac{1}{\eta} I_K + X'X.$$

Let $X'X = U\Lambda^{X'X}U'$, where $U = [u_1, \ldots, u_K]$ is orthogonal and $\Lambda^{X'X} = \mathrm{diag}(\lambda_1^{X'X}, \ldots, \lambda_K^{X'X})$. Then

$$\Omega^{-1} = U \left(\frac{1}{\eta} I_K + \Lambda^{X'X} \right)^{-1} U'.$$

Therefore

$$\theta_{j\ell} = \boldsymbol{x}_j' U \left(\frac{1}{\eta} I_K + \Lambda^{X'X} \right)^{-1} U' \boldsymbol{x}_\ell.$$

Define $\tilde{\boldsymbol{x}}_j := U'\boldsymbol{x}_j$, so $\tilde{x}_{jk} = u_k' \boldsymbol{x}_j$. Since $\left(\frac{1}{\eta} I_K + \Lambda^{X'X} \right)^{-1}$ is diagonal, we get

$$\theta_{j\ell} = \sum_{k=1}^{K} \frac{\tilde{x}_{jk} \tilde{x}_{\ell k}}{1/\eta + \lambda_k^{X'X}},$$

which is (21). ■

B Equilibrium Existence and Uniqueness

In this appendix, we provide a sufficient condition for existence and uniqueness of an interior equilibrium for the linear-quadratic utility model under a general control matrix H.

Proposition 6 *Suppose $\beta > 0$, $\alpha_j > c_j$ for all j, and fix a control matrix H with $H_{jj} = 1$ and $0 \le H_{jk} \le 1$ for all $j \neq k$. If*

$$(1 - \theta_{jj})(\alpha_j - c_j) > \sum_{k \neq j} |\theta_{jk}| (\alpha_k - c_k) \qquad \text{for all } j, \tag{43}$$

then there exists a unique Nash equilibrium $\boldsymbol{p}^*$, and it is interior—i.e., $p_j^* \in (c_j, \alpha_j)$ for all j.

Proof. Let

$$K := \prod_{j=1}^{J} [c_j, \alpha_j].$$

Since $\alpha_j > c_j$ for all j, the set K is nonempty and closed. Endow $\mathbb{R}^J$ with the weighted sup norm

$$\|\boldsymbol{x}\|_* := \max_j \frac{|x_j|}{\alpha_j - c_j}.$$

Because K is a closed subset of the finite-dimensional normed space $(\mathbb{R}^J, \|\cdot\|_*)$, it is complete.

Under the control matrix H, product j chooses p_j to maximize

$$\sum_{\ell=1}^{J} H_{j\ell}(p_\ell - c_\ell) q_\ell(\boldsymbol{p}).$$

The first-order condition with respect to p_j is

$$0 = q_j(\boldsymbol{p}) + \sum_{\ell=1}^{J} H_{j\ell}(p_\ell - c_\ell) \frac{\partial q_\ell}{\partial p_j}.$$

Because $I_J - \Theta$ is symmetric, we have $\theta_{j\ell} = \theta_{\ell j}$ and hence

$$\frac{\partial q_j}{\partial p_j} = -\frac{1}{\beta}(1 - \theta_{jj}), \qquad \frac{\partial q_\ell}{\partial p_j} = \frac{1}{\beta}\theta_{j\ell} \quad (\ell \neq j).$$

Substituting these expressions into the first-order condition and multiplying through by $\beta > 0$, we obtain

$$0 = (1 - \theta_{jj})(\alpha_j - p_j) - \sum_{k \neq j} \theta_{jk}(\alpha_k - p_k) - (1 - \theta_{jj})(p_j - c_j) + \sum_{k \neq j} H_{jk}\theta_{jk}(p_k - c_k).$$

Rearranging gives the best-reply map

$$T_j(\boldsymbol{p}) := \frac{\alpha_j + c_j}{2} - \frac{1}{2(1-\theta_{jj})} \sum_{k \neq j} \theta_{jk} \Big[(\alpha_k - c_k) - (1 + H_{jk})(p_k - c_k) \Big]. \tag{44}$$

Moreover, product j's objective is strictly concave in its own price, since

$$\frac{\partial^2}{\partial p_j^2} \sum_{\ell=1}^{J} H_{j\ell}(p_\ell - c_\ell) q_\ell(\boldsymbol{p}) = -\frac{2}{\beta}(1 - \theta_{jj}) < 0.$$

Thus $T_j(\boldsymbol{p})$ is the unique best reply of product j to $\boldsymbol{p}_{-j}$.

We first show that $T(K) \subseteq K$. Take any $\boldsymbol{p} \in K$. Then for every k,

$$0 \leq p_k - c_k \leq \alpha_k - c_k.$$

Since $0 \leq H_{jk} \leq 1$, it follows that

$$-H_{jk}(\alpha_k - c_k) \leq (\alpha_k - c_k) - (1 + H_{jk})(p_k - c_k) \leq \alpha_k - c_k,$$

and therefore

$$|(\alpha_k - c_k) - (1 + H_{jk})(p_k - c_k)| \leq \alpha_k - c_k.$$

Using (44), we obtain

$$\begin{aligned}
\left| T_j(\boldsymbol{p}) - \frac{\alpha_j + c_j}{2} \right| &\leq \frac{1}{2(1-\theta_{jj})} \sum_{k \neq j} |\theta_{jk}| \, |(\alpha_k - c_k) - (1 + H_{jk})(p_k - c_k)| \\
&\leq \frac{1}{2(1-\theta_{jj})} \sum_{k \neq j} |\theta_{jk}| \, (\alpha_k - c_k) \\
&< \frac{\alpha_j - c_j}{2},
\end{aligned}$$

where the last inequality follows from (43). Hence

$$c_j < T_j(\boldsymbol{p}) < \alpha_j \qquad \text{for all } j.$$

Therefore

$$T(K) \subseteq \prod_{j=1}^{J} (c_j, \alpha_j) \subseteq K.$$

We next show that T is a contraction on K. For any $\boldsymbol{p}, \tilde{\boldsymbol{p}} \in K$, (44) implies

$$T_j(\boldsymbol{p}) - T_j(\tilde{\boldsymbol{p}}) = \frac{1}{2(1-\theta_{jj})} \sum_{k \neq j} (1 + H_{jk}) \theta_{jk} (p_k - \tilde{p}_k).$$

Hence

$$\begin{aligned}\frac{|T_j(\boldsymbol{p}) - T_j(\tilde{\boldsymbol{p}})|}{\alpha_j - c_j} &\le \frac{1}{2(1-\theta_{jj})(\alpha_j - c_j)} \sum_{k\neq j} (1+H_{jk})|\theta_{jk}|\,|p_k - \tilde{p}_k| \\ &\le \frac{1}{2(1-\theta_{jj})(\alpha_j - c_j)} \sum_{k\neq j} (1+H_{jk})|\theta_{jk}|\,(\alpha_k - c_k)\,\|\boldsymbol{p} - \tilde{\boldsymbol{p}}\|_*.\end{aligned}$$

Since $0 \le H_{jk} \le 1$, we have $(1 + H_{jk})/2 \le 1$, so

$$\begin{aligned}\frac{|T_j(\boldsymbol{p}) - T_j(\tilde{\boldsymbol{p}})|}{\alpha_j - c_j} &\le \frac{\sum_{k\neq j} |\theta_{jk}|\,(\alpha_k - c_k)}{(1-\theta_{jj})(\alpha_j - c_j)} \|\boldsymbol{p} - \tilde{\boldsymbol{p}}\|_* \\ &< \|\boldsymbol{p} - \tilde{\boldsymbol{p}}\|_*,\end{aligned}$$

where the strict inequality again follows from (43). Therefore there exists $\rho \in (0, 1)$ such that

$$\|T(\boldsymbol{p}) - T(\tilde{\boldsymbol{p}})\|_* \le \rho \|\boldsymbol{p} - \tilde{\boldsymbol{p}}\|_* \qquad \text{for all } \boldsymbol{p}, \tilde{\boldsymbol{p}} \in K.$$

Thus T is a contraction on the complete metric space $(K, \|\cdot\|_*)$.

By Banach's fixed point theorem, T has a unique fixed point $\boldsymbol{p}^* \in K$. Since each T_j is the unique best reply of product j, this fixed point is the unique Nash equilibrium. Moreover, because $T(K) \subseteq \prod_{j=1}^{J}(c_j, \alpha_j)$, the fixed point satisfies $p_j^* \in (c_j, \alpha_j)$ for all j.
■

C Additional Details on Empirical Analyses

Segment	Own-price	Cross-price elasticity		Diversion ratio	
		Within segment	Across segments	Within segment	Across segments
Logit					
Subcompact	-1.32	0.00	0.00	0.05	0.05
Compact	-1.84	0.00	0.00	0.05	0.05
Intermediate	-2.17	0.00	0.00	0.04	0.04
Standard	-2.84	0.00	0.00	0.02	0.02
Luxury	-3.18	0.00	0.00	0.04	0.04
Nested Logit					
Subcompact	-2.25	0.03	0.00	1.30	0.05
Compact	-3.05	0.05	0.00	1.54	0.05
Intermediate	-2.77	0.00	0.00	0.04	0.03
Standard	-3.75	0.01	0.00	0.16	0.02
Luxury	-5.05	0.22	0.00	4.23	0.05
RC Logit					
Subcompact	-1.54	0.02	0.02	1.12	0.99
Compact	-2.15	0.02	0.02	1.06	1.05
Intermediate	-2.53	0.02	0.02	0.87	0.80
Standard	-3.32	0.02	0.02	0.72	0.53
Luxury	-3.69	0.05	0.02	1.33	0.96
RC Nested Logit					
Subcompact	-2.87	0.07	0.02	2.33	0.84
Compact	-3.43	0.08	0.02	2.24	0.85
Intermediate	-3.14	0.03	0.02	0.83	0.58
Standard	-4.30	0.03	0.02	0.85	0.38
Luxury	-5.34	0.29	0.02	4.97	1.57
Network					
Subcompact	-7.38	0.17	0.10	2.31	0.73
Compact	-9.36	0.22	0.09	2.35	0.83
Intermediate	-12.40	0.30	0.11	2.38	0.86
Standard	-20.40	0.51	0.13	2.51	0.89
Luxury	-19.85	0.62	0.11	3.19	1.00

Table C.1: Average elasticities and diversion ratios across all markets by segment. Own elasticity values are trimmed at the 10th percentile before averaging.

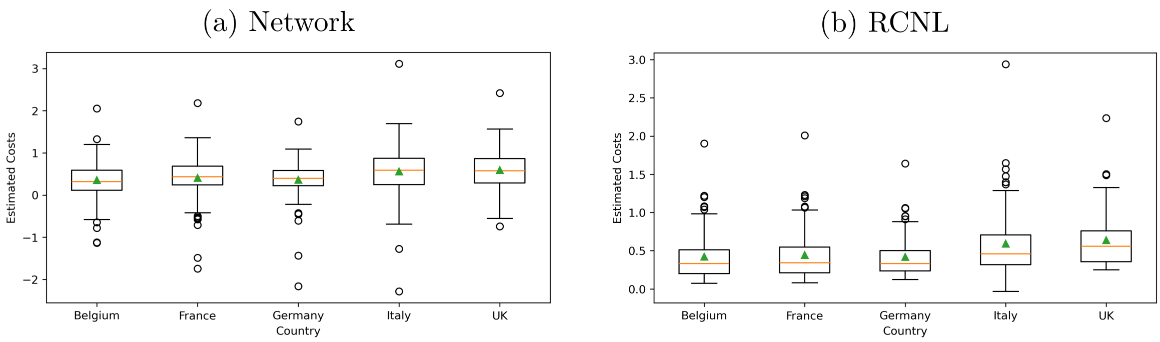

Figure C.1: Estimated marginal costs by country for 1999

Notes: The figure illustrates the distribution of the estimated marginal costs across countries for the network (left) and the RCNL (models) in 1999 using a box-plot.

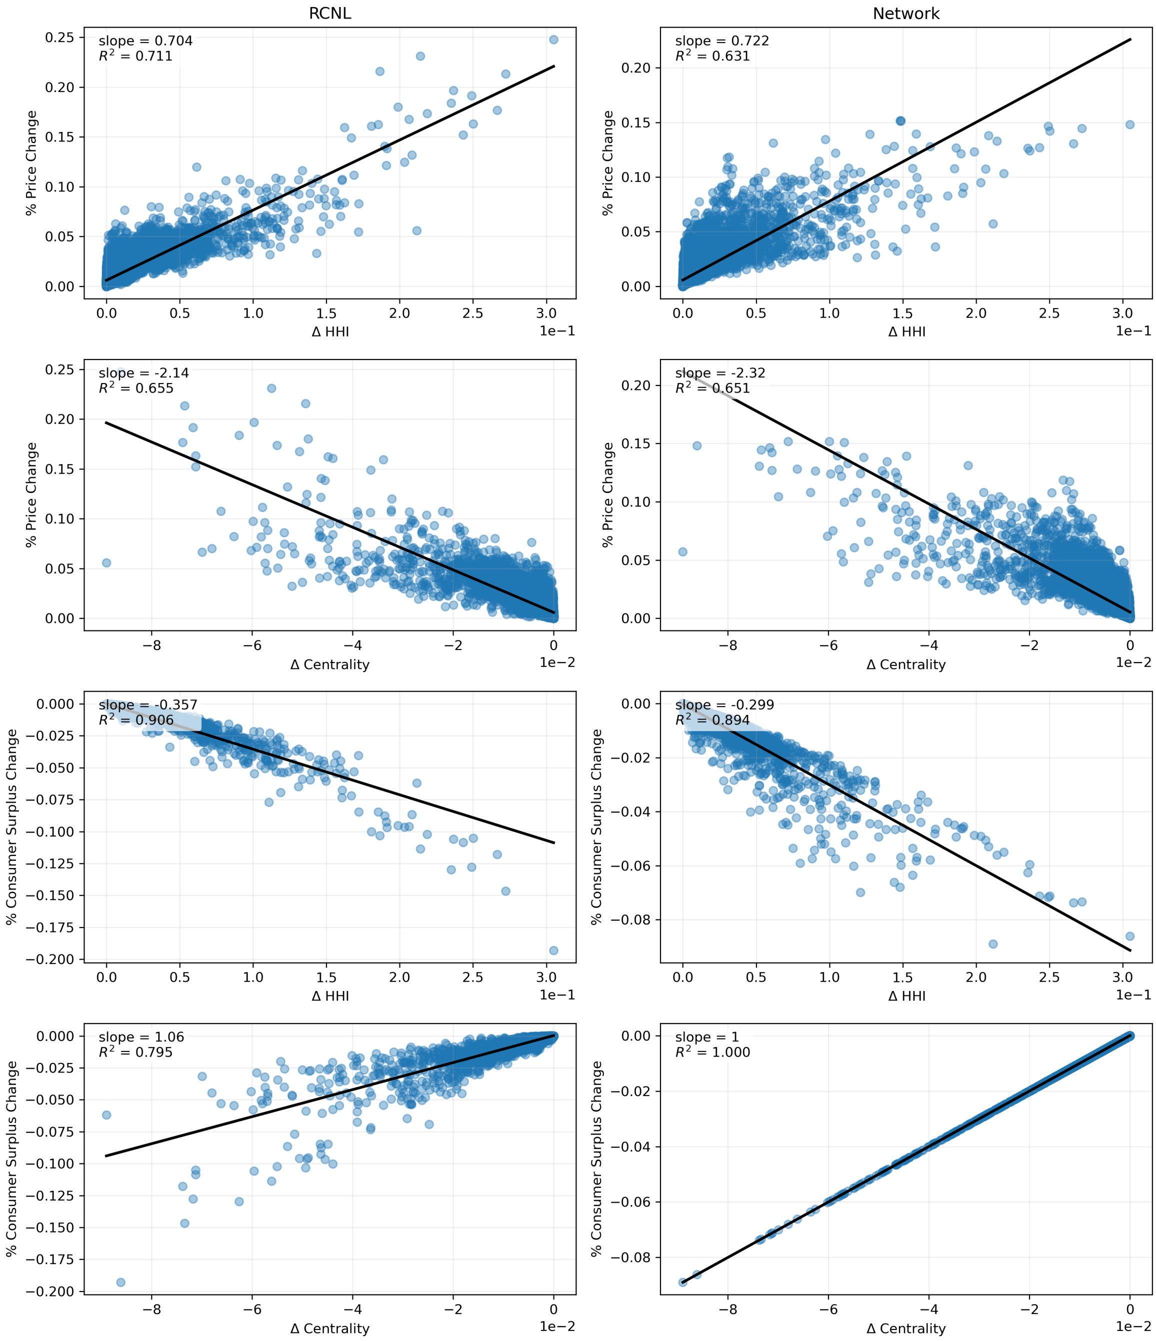

Figure C.2: % Price and Consumer Surplus changes vs ΔHHI and ΔCentrality, for all possible mergers across all markets (country-year). ΔCentrality is the percentage change in centrality after the merger.

Table C.2: Merger Screening with $\Delta\widehat{\text{Centrality}}$

	Proposed mergers	Allowed mergers (%)	ΔCS for allowed mergers ($000)	Type I Errors (%)	Type II Errors (%)
Panel A: Strict Regime (2% assumed efficiency)					
ΔHHI^*	21,779	86.9	3.125	2.8	33.4
CBP	21,779	93.4	2.406	1.5	66.5
$\Delta\widehat{\text{Centrality}}$	21,779	82.6	2.741	7.7	32.4
CBP or $\Delta\widehat{\text{Centrality}}$	21,779	82.4	2.757	7.9	31.5
CBP & $\Delta\widehat{\text{Centrality}}$	21,779	93.6	2.391	1.3	67.4
Panel B: Baseline Regime (5% assumed efficiency)					
ΔHHI^*	21,786	96.6	12.305	0.7	41.1
CBP	21,786	93.3	11.748	3.9	36.5
$\Delta\widehat{\text{Centrality}}$	21,786	96.4	11.885	1.8	57.6
CBP or $\Delta\widehat{\text{Centrality}}$	21,786	92.7	11.674	4.4	33.9
CBP & $\Delta\widehat{\text{Centrality}}$	21,786	97.0	11.960	1.3	60.1
Panel C: Lax Regime (10% assumed efficiency)					
ΔHHI^*	21,786	99.4	31.150	0.1	44.5
CBP	21,786	93.3	28.585	5.8	4.5
$\Delta\widehat{\text{Centrality}}$	21,786	99.2	30.827	0.5	67.0
CBP or $\Delta\widehat{\text{Centrality}}$	21,786	93.3	28.585	5.8	4.5
CBP & $\Delta\widehat{\text{Centrality}}$	21,786	99.2	30.827	0.5	67.0

Notes: The table shows the number of proposed mergers (i.e., those that are profitable), the percentage of deals allowed, the average change in CS for allowed mergers, the probability of blocking a welfare-enhancing deal (type I error) and that of approving a welfare-reducing deal (type II error) under three different regimes: strict, which assumes a 2% marginal cost efficiency (Panel A); baseline, which assumes a 5% marginal cost efficiency (Panel B); and lax, which assumes a 10% marginal cost efficiency (Panel C). Four screening approaches using the CBP and the $\Delta\widehat{\text{Centrality}}$ are considered. The thresholds for the $\Delta\widehat{\text{Centrality}}$ are optimally set for each regime-decision rule pair to maximize the total CS after the merger. As a benchmark, we also show the results obtained by screening mergers by optimally choosing the ΔHHI threshold (ΔHHI^*).