Diversion Networks and Horizontal Acquisitions*

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

October 2026

Abstract

We develop a network-based screening framework to evaluate the effects of horizontal acquisitions. We show that the effect of a change in ownership can be locally decomposed into an upward-pricing-pressure (UPP) component and a network-centrality component that captures how pricing pressure propagates through the network implied by products' diversion ratios. For full mergers, the change in network centrality summarizes the unilateral effect, yielding a screening statistic with modest computational requirements. Using data from the European automobile industry, we show that changes in centrality provide substantial incremental explanatory power for predicted price effects beyond both UPP indices and changes in concentration. In screening exercises, we quantify the gains from adopting centrality-based rules, including one that can be computed without price data.

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

JEL Classification: D43, G34, L13, L41

*We are grateful to Lars Peter Hansen, Ali Hortaçsu, Scott Nelson, and Andrew Sweeting for several insightful discussions. We also thank Gaurab Aryal, Francesco Decarolis, Bruno Pellegrino, Julien Sauvagnat, Charles Taragin, and Tom Wollmann, as well as conference and seminar participants at Boston University and the 2026 EARIE Meeting for helpful comments.

†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 models in the tradition of Berry et al. (1995) offer a richer, microfounded 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 (Miller and Sheu, 2021; Valletti and Zenger, 2021).

Between concentration screens and full merger simulations lies a class of first-order approximation (FOA) approaches, such as upward pricing pressure (UPP) and the gross upward pricing pressure index (GUPPI). These pricing-pressure indices summarize the direct pricing incentive created by internalizing diversion between the merging products, without solving for the full post-merger equilibrium. They are attractive because they do not require an explicit market definition and can often be implemented using information on the merging products' margins and diversion ratios. However, these indices are local by design: they capture the direct incentive to raise prices created by common control, but not the strategic responses of rival products in the market.

This paper develops a screening framework for horizontal acquisitions that bridges FOAs and merger simulations while remaining tractable and scalable. We represent product competition as a network whose links are determined by diversion ratios.[2] As in traditional pricing-pressure screens, the predicted price effects capture the direct upward pricing pressure generated when merging firms internalize diversion between their products. As in merger simulations, however, the framework goes beyond direct pricing-pressure indices by accounting for how competitors' pricing responses propagate this initial pricing pressure through the network of differentiated products.

Our analysis proceeds in three steps. First, we develop the theoretical framework that microfounds our centrality-based screening statistic. We start by showing how the price effect of a minority stake acquisition can be locally decomposed into a UPP-style component and a network-centrality term (Proposition 1). The UPP-style component captures the direct diverted-margin incentive created by the acquisition, while the network-centrality term captures how this incentive propagates through competitors'

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

pricing responses. For full mergers, we show that unilateral price effects are locally characterized by changes in each product's network centrality (Proposition 3). A product's network centrality summarizes the competitive pressure it faces: the more central a product is in the diversion network, the lower its markup.

These results imply that the change in centrality induced by a horizontal acquisition can be used as a screening statistic. Unlike the HHI and ΔHHI, whose theoretical justification for merger screening relies on aggregative environments where substitution patterns are primarily driven by market shares (Nocke and Whinston, 2022), our centrality captures the closeness of competition between differentiated products and allows for flexible substitution patterns. This flexibility is important because substitution patterns can vary substantially across products within the same industry.[3]

One additional advantage of our approach is that it accommodates arbitrary changes in the ownership structure, including those implied by incumbents' acquisitions of minority stakes in rival firms or institutional investors' common ownership. 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.[4]

In the second step of the analysis, we bring our theoretical framework to the data using a parsimonious demand model. Following the specification in Pellegrino (2025), we derive demand from a representative consumer with linear-quadratic preferences (Singh and Vives, 1984; Vives, 1985; Anton et al., 2017), while allowing substitution patterns to depend on product characteristics, as in hedonic demand models (Lancaster, 1966; Rosen, 1974; Epple, 1987; Feenstra and Levinsohn, 1995). This specification is useful for our screening exercise for two reasons. First, because demand is linear, the local demand approximation underlying our centrality-based screening statistic holds exactly. Changes in centrality therefore capture the full equilibrium price effects of a merger. Second, diversion ratios can be readily computed from product characteristics: within this linear-quadratic hedonic specification, cross-price demand elasticities and diversion ratios depend on weighted inner products between vectors of product characteristics (Proposition 5). This structure captures explicitly the insight that substitution patterns should reflect products' relative positions in characteristic space (Gandhi and Houde, 2023). Thus, this specification provides a computationally tractable way to incorporate both the closeness of competition and competitors' pricing responses into merger screening. In many industries, information

[3]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.

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

on products' characteristics is available to agencies or can be collected at scale from public sources, such as websites and other unstructured sources. Advances in large language models (LLMs) facilitate the collection and organization of these data, reducing the information-gathering burden of implementing characteristics-based merger screens.

To assess whether this parsimonious specification captures substitution patterns relevant for merger screening, we compare its implied diversion ratios and merger effects with those obtained from standard discrete-choice demand models. We find that the network model on average generates diversion ratios similar to those obtained from a random coefficients nested logit (RCNL) model. When we simulate all possible mergers across European car manufacturers, the predicted price and consumer-surplus effects are highly correlated with those obtained from the RCNL model. These comparisons suggest that the characteristics-based diversion network captures much of the variation in substitution patterns and pricing responses relevant for merger screening.

In the third and final step, we evaluate the role of centrality as a screening statistic. In the automobile industry, changes in centrality provide substantial incremental explanatory power for predicted price effects beyond both GUPPI and changes in HHI. This indicates that our index captures information about merger harm that is not summarized by concentration measures or by traditional first-order screens based on direct diverted margins alone. In particular, when price effects are simulated using the RCNL benchmark, adding ΔHHI to GUPPI raises the R^2 from 0.765 to 0.814, whereas adding ΔCentrality raises it to 0.853, which is nearly the same explanatory power as the specification that includes all three screening statistics. This evidence is consistent with the idea that the network-propagation component captured by our framework is relevant for predicting the equilibrium price effects of horizontal mergers.

We then test the performance of centrality-based screening in a policy decision framework in which the regulator seeks to approve mergers that increase consumer surplus. We simulate all possible mergers under alternative marginal-cost efficiency assumptions of 2%, 5%, and 10%, using consumer-surplus effects from the RCNL model as the evaluation benchmark. Our centrality statistic delivers a screening rule whose threshold follows from the model's welfare criterion. Assumed cost efficiencies enter the statistic directly, allowing the rule to account jointly for the loss of competition, the benefits of lower marginal costs, and the propagation of both effects through the diversion network.

The centrality rule generates higher aggregate consumer-surplus gains from approved mergers than the 2023 concentration-based presumption in all three efficiency regimes. Among the screening approaches we consider, it is the only one to improve on this benchmark throughout. The alternative rules perform unevenly as assumed

efficiencies change: relative to the 2023 concentration-based presumption, a rule based on GUPPI is too permissive when efficiencies are small, whereas one based only on the merger-induced change in the HHI becomes overly restrictive when efficiencies are larger. In the low-efficiency regime, centrality-based screening also exhibits the lowest Type II error rate—the share of consumer-surplus-reducing mergers that are approved. As efficiencies increase, the decision rule accommodates more beneficial mergers by balancing the harm from reduced competition against the benefits of lower costs.

Lastly, we examine a simpler characteristics-based statistic, $\Delta\widehat{\text{Centrality}}$, that requires no demand-system estimation and uses product characteristics, ownership information, and pre-merger sales to aggregate product-level centralities. We assess its performance by comparing the mergers it flags with those flagged by the 2023 concentration-based presumption, holding the number of scrutinized deals fixed at the level implied by the concentration-based presumption. We find that using $\Delta\widehat{\text{Centrality}}$ yields higher aggregate consumer surplus from approved mergers in the baseline and high-efficiency regimes.

Overall, our framework complements existing approaches to merger screening with a tractable statistic based on changes in product-network centrality. By accounting for both the closeness of competition between differentiated products and competitors' pricing responses, this measure provides practitioners and policymakers with an economically grounded tool for identifying potentially anticompetitive acquisitions.

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 rationale for this practice by showing that, in a homogeneous-good Cournot model, a larger ΔHHI implies a larger compensating marginal cost reduction (CMCR) required to offset the merger's upward pricing incentives.[5] 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.[6]

Merger simulation approaches (Nevo, 2000) incorporate detailed demand and cost systems and are typically more predictive of actual price changes than other tools

[5] Bhattacharya et al. (2026b) show that, in the context of mergers of U.S. consumer packaged goods manufacturers, price changes both across and within mergers are correlated with ΔHHI.

[6] 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.

(Bhattacharya et al., 2026a). However, they 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).[7] Monte Carlo evidence shows that even simple UPP measures can predict merger price effects well under standard log-concave demand (Miller et al., 2017), and, with sufficiently accurate pass-through information, first-order approximations can also outperform misspecified merger simulations (Miller et al., 2016). These findings illustrate the predictive value of first-order approaches while also highlighting the importance of accessing reliable information on their inputs, particularly margins and diversion ratios.

Although the usefulness of diversion ratios is well understood and emphasized in the 2010 U.S. Merger Guidelines, measuring them accurately remains an empirical challenge (Conlon and Mortimer, 2021). A natural approach is to recover diversion ratios from a demand system that allows for flexible substitution, but estimating such a system is difficult to do at scale.[8] 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, network-based framework for the analysis of horizontal acquisitions. Like pricing-pressure indices, our framework captures the merging parties' incentive to raise prices as they internalize diversion between their products. Like merger simulations, however, it also accounts for how this initial pricing pressure propagates through the broader system of differentiated products. Whereas traditional FOAs linearize the equilibrium pricing conditions, we work with a local linear approximation to demand. We show that the propagation of price effects through the network is summarized by its Bonacich centrality, and that, for full mergers, the change in centrality provides an effective screening statistic. In our empirical implementation, products' positions in characteristic space determine the diversion network. This yields a simpler version of the screen based on product characteristics and ownership information, with pre-merger sales used to aggregate product-level centralities. Product characteristics are often available to agencies and are becoming increasingly accessible as advances in machine learning facilitate their extraction from unstructured data and their use in demand estimation (Magnolfi et al.,

[7] Jaffe and Weyl (2013) show that converting pricing pressures into quantitative estimates of the price change induced by a merger requires information on the pass-through matrix, which, under certain conditions, is approximated by the pre-merger cost pass-through.

[8] 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.

2022; Bajari et al., 2023; Compiani et al., 2025).[9]

Our work also contributes 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 à la Cournot or à la Bertrand, and product attributes induce a similarity network, providing empirical estimates of efficiency losses due to market power.[10] 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:

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

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.

[9] Consistent with the idea that concentration indices can be complemented with observable characteristics to help capture substitution patterns, Benson et al. (2024) show that, in the banking industry, introducing a threshold based on the proximity of merging firms' branch networks could improve merger screening.

[10] Ederer and Pellegrino (2025) extends the framework to incorporate common institutional ownership. Moreover, Antón et al. (2023) employ the same general equilibrium model to show that the economy-wide effects of mergers are substantial even when firms are assumed to internalize common ownership concerns before the mergers.

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):

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

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

We work with a local linearization of demand around the original equilibrium $(\boldsymbol{p}^*, \boldsymbol{q}^*)$. Specifically, we let S denote the Jacobian of $\boldsymbol{q}(\boldsymbol{p})$ evaluated at $\boldsymbol{p}^*$, define $\boldsymbol{\alpha} := -S^{-1}\boldsymbol{q}^* + \boldsymbol{p}^*$, and consider the affine demand system

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

which matches both demand levels and the demand Jacobian at the original equilibrium. This approach is related to the local analysis in Galeotti et al. (2024), who show that equilibrium comparative statics can be characterized using the status-quo Slutsky matrix when each product's own-price demand slope is locally constant. Our affine benchmark imposes the stronger restriction that the entire demand Jacobian S is held fixed. This additional structure delivers closed-form equilibrium prices and allows us to characterize welfare implications of ownership changes through the diversion network.[11]

With this linearization, we are ready to show that equilibrium markups admit an intuitive network-based decomposition.[12] 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

[11] Unlike FOAs that linearize firms' pricing first-order conditions, we linearize demand itself. For nonlinear demand, these procedures need not yield the same local equilibrium price response: although the affine approximation reproduces demand levels and all own- and cross-price demand derivatives at the pre-merger equilibrium, differentiating the pricing conditions introduces demand-curvature terms multiplied by equilibrium markups, which the affine approximation omits (see, e.g., Miller et al., 2016; Jaffe and Weyl, 2013). Our framework therefore provides a microfoundation for the centrality statistic rather than exact local comparative statics for arbitrary nonlinear demand. In the empirical application, we assess its usefulness by evaluating its performance against the nonlinear RCNL benchmark.

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

$\boldsymbol{b}(A, \delta, \boldsymbol{u})$ 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.[13]

The following equilibrium characterization result ensues.

Lemma 1 *Under the linearized demand system, the equilibrium markup is*

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

Notably, $\frac{\alpha_j - c_j}{2}$ is the markup charged by a monopolist facing linear inverse-demand $p_j = \alpha_j - S_{jj}^{-1} q_j$. This decomposition is closely related to the ones derived by Chen et al. (2018), Ushchev and Zenou (2018), and Pellegrino (2025). Our approach helps isolate the role of changes in the centrality measure in predicting the effects of mergers.

Lemma 1 highlights 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:

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

The term $D^k \left(\frac{\boldsymbol{\alpha} - \boldsymbol{c}}{2}\right)$ 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-

[13] 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.

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. (2021) 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

$$H_{j\ell}=\frac{\sum_s \gamma_{js}\beta_{\ell s}}{\sum_s \gamma_{js}\beta_{js}} \qquad (j\neq\ell).^{14}$$

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.

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

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

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

We continue to work with the linearized demand system $\boldsymbol{q} = -S(\boldsymbol{\alpha} - \boldsymbol{p})$. 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 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 equilibrium markup under the linearized demand system is*

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

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

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

whose ℓth element is the discounted total length of all walks (of any length) from node ℓ to node j. 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 acquisition'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 the largest sellers 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 $M_{12} := I_J + E_{12} + E_{21}$. Under the linearized demand, we obtain the following result.

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, under the linearized demand system, 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

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 system $\boldsymbol{q}(\boldsymbol{p})$, 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?

To make progress, we focus on the case of linear-quadratic preferences, where our local approximation is exact. The linear-quadratic framework has a long-standing tradition in differentiated oligopoly models (e.g., Singh and Vives (1984); Vives (1985); Anton et al. (2017)) because it implies a linear demand system with a diversion ratio that does not depend on quantities. To relate diversion ratios and product attributes, we further specialize to the hedonic-quadratic demand of Pellegrino (2025) where the quadratic utility term is defined as the outer product of the matrix of product attributes. Within this framework, the diversion ratio matrix can be obtained directly from observed product characteristics, without estimating a full demand system.[15]

[15] Antón et al. (2023) use diversion ratios from Pellegrino (2025) to show that mergers increasingly

Combined with the results of Sections 2 and 3, this means that the data requirements 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. Following Pellegrino (2025), we consider a representative consumer that 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 $q_0 + \boldsymbol{p}'\boldsymbol{q} \leq y$. Preferences are

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

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 $\theta_{jj} \in (0, 1)$.[16]

occur between firms that sell similar products.

[16]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 $j \neq \ell$.

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.

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 $\lambda_k^{X'X} = \sum_{i=1}^{J}(\tilde{x}_{ik})^2$, 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 Effects

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.[17] Since equilibrium markups are determined by the Bonacich centrality of the diversion network (Lemmas 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:

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

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

[17] 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.

Low data requirements and compatibility with modern data sources. Conditional on a value of η, the diversion matrix D is constructed entirely from the attribute matrix X (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 diversion network can therefore be constructed with limited data collection. As we show below, combining this network with information on product ownership and pre-merger sales yields a merger-screening index that does not require price data.

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.

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. Additionally, the framework's modest computational requirements make it practical to include a broad product universe whenever the necessary characteristic data are available.

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 obtained by estimating demand using discrete-choice models. Lastly, we simulate the effects of mergers under different demand models and assess the effectiveness of merger screening using centrality-based rules.

5.1 Data

We use publicly available data from Brenkers and Verboven (2006).[18] 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

[18]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: Summary Statistics on the European Automobile Industry

	Variable	Mean	SD	Bel	Fra	Ger	Ita	U.K.
0	Sales (units)	19813.24	37719.92	3925.42	23305.81	31002.55	24292.14	19784.34
1	Horsepower (kW)	57.14	23.88	56.53	56.17	57.45	57.22	58.42
2	Fuel inefficiency (liter/100 km)	8.18	1.72	8.22	8.12	8.24	8.08	8.22
3	Width (cm)	164.38	9.62	164.16	164.28	164.57	164.06	164.82
4	Height (cm)	140.43	4.62	140.34	140.46	140.49	140.63	140.27
5	Foreign (0-1)	0.81	0.39	1.00	0.75	0.73	0.76	0.77
6	Price Eur	8352.52	5540.92	7674.24	8315.98	8119.62	8378.81	9384.34
7	Price Eur (post-tax)	10089.23	6606.51	9439.35	10476.46	9283.47	9946.35	11388.33
8	Price Eur (Deflated)	10793.17	5962.59	9029.31	10944.30	8394.99	13292.90	12870.79
9	Subcompact (0-1)	0.28	0.45	0.28	0.28	0.28	0.31	0.26
10	Compact (0-1)	0.23	0.42	0.23	0.24	0.23	0.22	0.24
11	Intermediate (0-1)	0.22	0.41	0.22	0.22	0.22	0.21	0.23
12	Standard (0-1)	0.19	0.39	0.19	0.19	0.18	0.17	0.19
13	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.

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).

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

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

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

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

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

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 $\log\left(y - \boldsymbol{p}'\boldsymbol{q}\right)$.[19] We also calibrate η to 0.136 following Pellegrino (2025).[20] 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 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.

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

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

[20] In Figure D.2 we illustrate the sensitivity of our estimated average substitution patterns to this assumption.

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.,[21]

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

To consistently estimate (28), we need to instrument p_{jt}/y_t because prices will be 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 using the linear specification in (28). We instrument p_{jt}/y_t with the demand instruments z_{jt} constructed following Brenkers and Verboven (2006). These include the product's own observed characteristics—horsepower, fuel efficiency, width, height, and a foreign-product indicator—and BLP-style functions of other products' characteristics (Berry et al., 1995). Specifically, we include product counts and sums of observed characteristics within three groups: other products of the same firm in the same segment; rival-firm products in the same segment; and rival-firm products in the same segment–foreign-status pair.

Equation (30) implies product characteristics affect firms' pricing decisions through the centrality term b_{jt} and the observable part of α_{jt}, implying that our instrument

[21] 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).

Table 2: Estimated Substitution Patterns across Demand Models

Segment	Within-segment diversion (%)		Between-segment diversion (%)	
	Median	Mean	Median	Mean
Logit				
Subcompact	0.05	0.05	0.03	0.04
Compact	0.05	0.05	0.04	0.04
Intermediate	0.04	0.04	0.04	0.04
Standard	0.02	0.02	0.04	0.05
Luxury	0.02	0.03	0.04	0.04
Nested Logit				
Subcompact	1.27	1.30	0.03	0.03
Compact	1.45	1.51	0.03	0.03
Intermediate	0.04	0.04	0.04	0.04
Standard	0.15	0.15	0.04	0.05
Luxury	4.89	6.12	0.03	0.03
RC Logit				
Subcompact	1.07	1.12	0.65	0.67
Compact	1.00	1.06	0.76	0.79
Intermediate	0.83	0.86	0.85	0.88
Standard	0.56	0.63	0.94	0.97
Luxury	0.65	1.03	0.91	0.93
RC Nested Logit				
Subcompact	2.28	2.33	0.39	0.41
Compact	2.11	2.20	0.55	0.56
Intermediate	0.79	0.82	0.79	0.83
Standard	0.72	0.78	0.86	0.89
Luxury	4.93	6.19	0.67	0.71
Network				
Subcompact	2.30	2.31	0.70	0.72
Compact	2.31	2.35	0.80	0.84
Intermediate	2.30	2.38	0.83	0.88
Standard	2.47	2.52	0.88	0.91
Luxury	3.09	3.18	0.98	1.03

Notes: The table summarizes, for each of the five demand models considered, the average diversion ratios across all markets by segment.

should satisfy the relevance condition. 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 $\mathbb{E}[\xi_{jt}|z_{jt}] = 0$.

Conditional on the product characteristics and the parameter η, the matrix Θ determines the relative pattern of substitution across products, i.e., diversion ratios can be recovered directly from Θ. The bottom panel of Table 2 reports average and median diversion ratios for product pairs within and across segments. For example,

following a marginal increase in the price of a compact car, the average diversion to any given compact model is 2.35%, compared with 0.84% to any given model in another segment.

Although the framework does not require a discrete market definition, its implementation depends on the availability and selection of product characteristics. To assess sensitivity to this choice, Table D.2 in the Appendix replaces each of the five baseline characteristics in turn with noise at three variance levels. Average diversion ratios remain broadly stable across these exercises, and within-segment diversion consistently exceeds across-segment diversion.

Comparison with discrete-choice models. To assess whether our model delivers reasonable substitution patterns, 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 characteristics. In each case, price is treated as endogenous and instrumented using the BLP-style instruments described above. The implementation details are discussed in Appendix E.

As expected, across logit-type models, substitution patterns become progressively more flexible. Logit imposes IIA and homogenous preferences; Nested Logit allows stronger substitution within segments; RCL introduces preference heterogeneity; and RCNL combines preference heterogeneity with within-segment correlation. These differences clearly emerge from Table 2. 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 strengthens substitution among products in the same nest.[22] The network model replicates these nesting patterns and generates diversion ratios that, on average, are remarkably close to those of the RCNL specification.[23]

We also emphasize that the network model allows consumers to exhibit a taste for variety and to consume more than one good. The automobile industry—where households generally purchase only one vehicle at a time—offers limited scope to assess this feature of the framework. Nevertheless, the network model's documented performance in this setting suggests that it may also prove useful in industries characterized by more extensive multi-product consumption.

[22] Table D.1 compares estimates across the RCNL and the Network model.

[23] Figure D.1 further shows that the diversion ratios implied by the Network model are positively correlated with the ones implied by the RCNL.

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 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.

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.[24] 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. The plots, however, also reveal some systematic differences. The network model tends to predict larger price increases for merging parties. Non-merging products also experience larger price adjustments compared to RCNL, consistent with competitive effects propagating more broadly through the similarity network rather than remaining concentrated within predefined nests. However, the two models exhibit relatively similar implications in terms of predicted change in consumer surplus.

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 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

[24] Figure D.3 shows the distribution of marginal costs across the network and the RCNL models, focusing on the most recent year of our sample (1999).

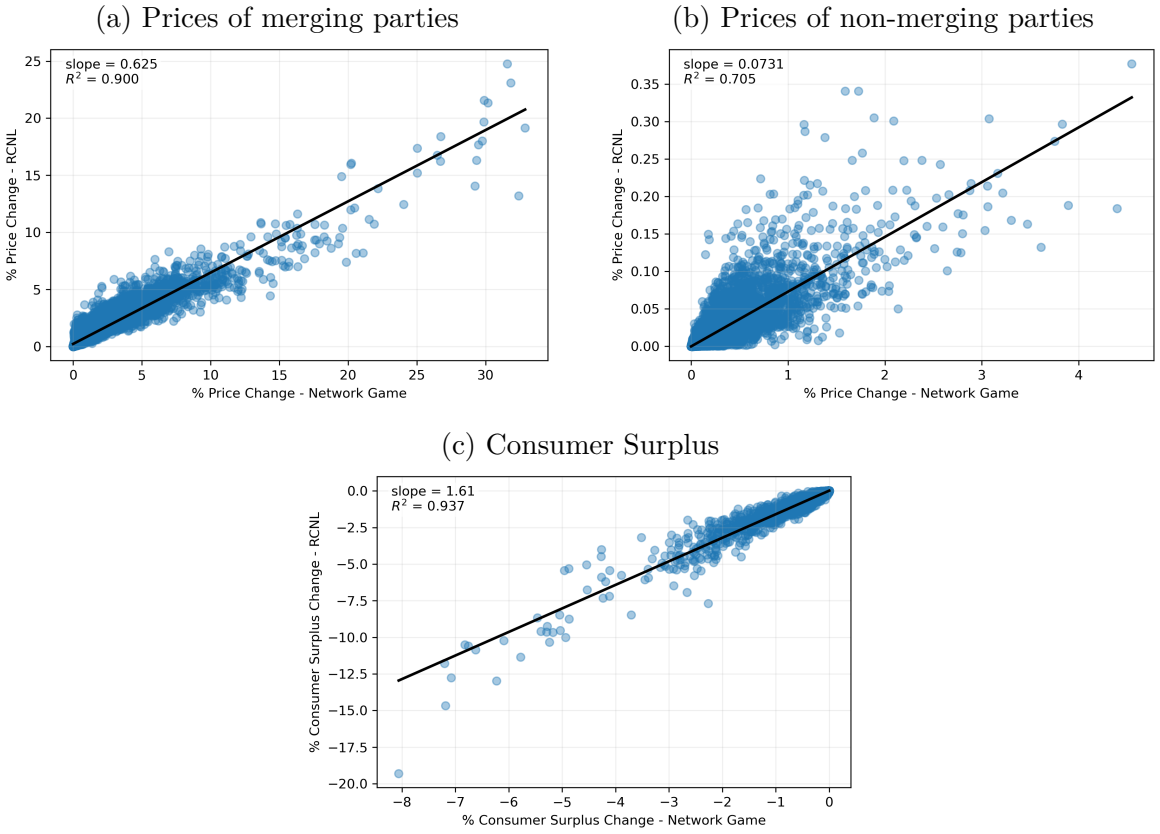


Figure 1: Simulation of the effects of all possible mergers

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.

points, leads the merged entity to increase prices by 20.2% under the network model and by 15.9% under RCNL. The associated consumer-surplus loss amounts to 4.1% in the network model and 5.4% in RCNL. By contrast, a merger between BMW and MCC in Germany, where the firms have limited overlap across segments, induces a small change in HHI (less than 19 points) and generates much smaller effects under both models, with price increases among merging parties between 0.3–0.7% and CS losses between 0.02–0.03%.

In Table 3, the network model predicts larger average price increases for the merging firms in 18 of the 19 cases, and larger price increases for non-merging firms in every case. For merging firms, the prediction gap tends to be larger for mergers generating greater increases in concentration, although the relationship is not monotonic. In Germany, for example, the gap is 0.40 percentage points for BMW–MCC, compared with 3.41 percentage points for Opel–VW and 3.76 percentage points for Ford–VW.

These patterns motivate examining whether the prediction gap varies systematically with the magnitude of the merger's effects. A related concern arises for first-order approximations, which can become less accurate when mergers induce larger equilib-

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%	20.1922%	0.3048%	1.8882%	-5.4262%	-4.1103%
GBR	BMW–Ford	348.1630	3.7497%	5.0744%	0.0418%	0.5327%	-1.1106%	-0.8028%
GBR	BMW–Opel	255.0992	2.9597%	4.0671%	0.0455%	0.4036%	-0.8305%	-0.5887%
GBR	Ford–Opel	651.7778	5.3994%	6.6246%	0.0782%	0.7383%	-2.7334%	-1.5969%
GER	BMW–Ford	151.0954	1.9720%	2.9422%	0.0343%	0.3797%	-0.3921%	-0.2957%
GER	BMW–MCC	18.7394	0.3090%	0.7060%	0.0026%	0.0371%	-0.0236%	-0.0279%
GER	BMW–Mercedes	155.4387	3.0930%	2.8300%	0.0768%	0.3079%	-0.9659%	-0.3597%
GER	BMW–Opel	234.0387	2.3207%	3.6973%	0.0438%	0.4529%	-0.5545%	-0.4279%
GER	BMW–VW	494.3758	3.9663%	5.3277%	0.0655%	0.9781%	-1.4288%	-0.9534%
GER	Ford–MCC	24.0961	0.9525%	1.5064%	0.0089%	0.0752%	-0.0762%	-0.0593%
GER	Ford–Mercedes	199.8714	2.9257%	4.2558%	0.0523%	0.4001%	-0.5967%	-0.4150%
GER	Ford–Opel	300.9395	3.6512%	6.0142%	0.0794%	0.6115%	-0.9313%	-0.6704%
GER	Ford–VW	635.6948	5.9811%	9.7381%	0.0768%	1.4913%	-1.9570%	-1.5193%
GER	MCC–Mercedes	24.7888	0.6152%	1.7067%	0.0033%	0.0431%	-0.0303%	-0.0392%
GER	MCC–Opel	37.3236	1.6758%	2.2470%	0.0139%	0.0898%	-0.1265%	-0.0813%
GER	MCC–VW	78.8412	1.5702%	2.4809%	0.0120%	0.2082%	-0.2077%	-0.1765%
GER	Mercedes–Opel	309.5901	3.1861%	4.5839%	0.0684%	0.4110%	-0.8399%	-0.5827%
GER	Mercedes–VW	653.9680	4.0745%	5.6634%	0.0948%	0.8824%	-2.0943%	-1.2837%
GER	Opel–VW	984.6571	6.3387%	9.7461%	0.1091%	1.5010%	-3.0836%	-2.0885%

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.

rium changes (Jaffe and Weyl, 2013). Although our network model computes the full post-merger equilibrium under its linear demand specification, differences in demand curvature may likewise become more consequential for its agreement with RCNL as price changes grow. We therefore extend the comparison to the full simulation sample, and, in Figure D.4, we plot the difference between the models' predicted average percentage price changes for merging firms' products against both ΔHHI and RCNL's predicted price increase. The figure shows that the network model tends to predict larger price increases relative to RCNL as the size of the merger effects grows, with the largest discrepancies occur toward the upper tails of concentration changes and predicted price effects.[25]

[25]The left panel implies that a 100-point increase in ΔHHI is associated with an approximately 0.40 percentage-point increase in the gap. In the right panel, a one-percentage-point increase in the RCNL-predicted price increase for merging firms' products is associated with an approximately 0.44 percentage-point increase in the gap.

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."[26] 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-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′) provide a rationale for 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.[27] 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.[28] Further, we define the change in centrality induced by a merger as

$$\Delta\text{Centrality} := \boldsymbol{q}_1'\boldsymbol{b}\left(D_H^1, \frac{1}{2}, \frac{\boldsymbol{\alpha}-\boldsymbol{c}_0}{2}\right) - \boldsymbol{q}_0'\boldsymbol{b}\left(D_H^0, \frac{1}{2}, \frac{\boldsymbol{\alpha}-\boldsymbol{c}_0}{2}\right), \tag{31}$$

where $\boldsymbol{q}_0$ is the vector of pre-merger quantities, $\boldsymbol{q}_1$ is the vector of predicted post-

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

[27] Appendix C.1 provides additional details on the computation of this index.

[28] 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 $\boldsymbol{\alpha} = y\widetilde{\boldsymbol{\alpha}}$.

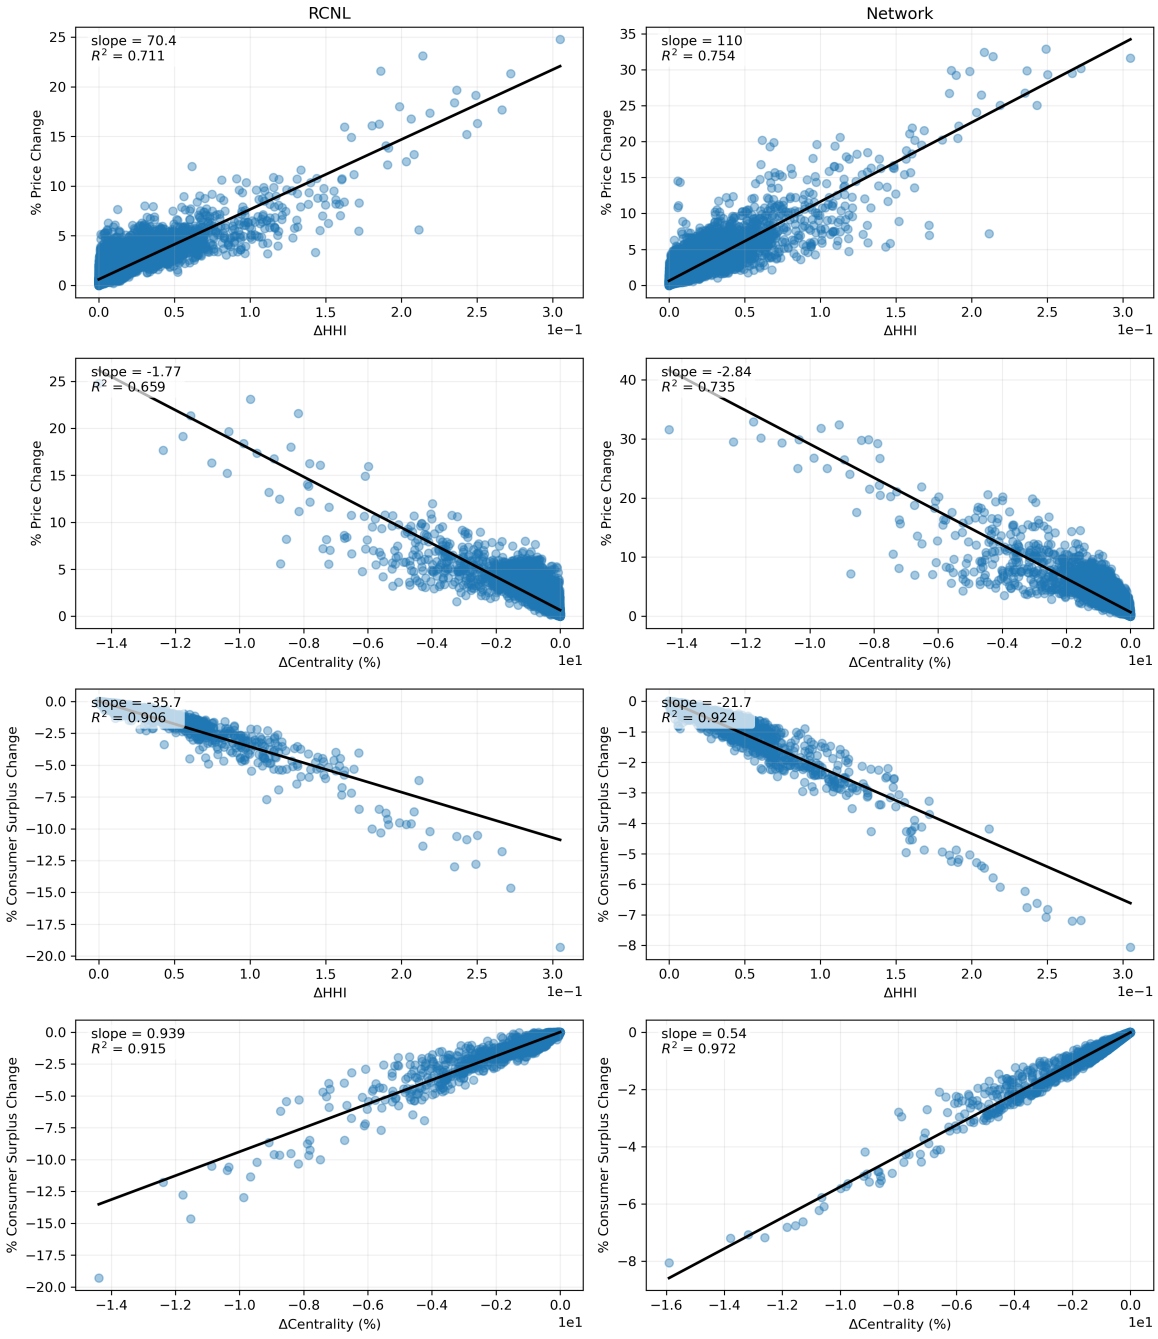


Figure 2: Correlation between the Effects of Simulated Mergers and Screening Tools

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

merger quantities, and D_H^0 and D_H^1 are the pre- and post-merger internalized diversion matrices, respectively. By construction, holding marginal costs constant at the pre-merger level $\boldsymbol{c}_0$, a merger always decreases centrality.

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

Table 4: Simulated Price Effects and Screening Tools

	(1) $\%\Delta P$	(2) $\%\Delta P$	(3) $\%\Delta P$	(4) $\%\Delta P$	(5) $\%\Delta P$	(6) $\%\Delta P$	(7) $\%\Delta P$
Panel A: RCNL							
ΔHHI	0.702***			0.432***	0.343***		0.058***
	(0.014)			(0.028)	(0.016)		(0.018)
ΔCentrality		-2.650***		-1.207***		-1.364***	-1.246***
		(0.138)		(0.150)		(0.094)	(0.106)
GUPPI			1.032***		0.613***	0.662***	0.623***
			(0.016)		(0.017)	(0.021)	(0.018)
N	21786	21786	21786	21786	21786	21786	21786
R^2	0.738	0.705	0.765	0.775	0.814	0.853	0.854
Panel B: Network game							
ΔHHI	1.097***			0.684***	0.676***		0.243***
	(0.025)			(0.048)	(0.029)		(0.034)
ΔCentrality		-4.130***		-1.845***		-2.388***	-1.892***
		(0.226)		(0.252)		(0.189)	(0.202)
GUPPI			1.545***		0.718***	0.896***	0.734***
			(0.033)		(0.029)	(0.045)	(0.035)
N	21786	21786	21786	21786	21786	21786	21786
R^2	0.788	0.748	0.749	0.825	0.833	0.867	0.873

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. ΔCentrality and GUPPI are in percentage terms. All regressions include market (country-year) and year fixed effects. Standard errors in parentheses.

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 the 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.

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. In the network model, the relationship between CS losses and the percentage change in centrality is exact, as established in Section 3, and is therefore reflected in the bottom-right panel of Figure 2. Additionally, our analysis confirms the strong predictive power of ΔHHI, in line with evidence from prior studies in other industries (e.g., Bhattacharya et al. (2026b)): 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 its own. A natural question, however, is whether, by capturing substitution patterns overlooked in concentration indices, ΔCentrality adds incremental information beyond ΔHHI. We also compare it with the Gross Upward Pricing Pressure Index (GUPPI), a standard first-order pricing-pressure measure. This comparison is particularly relevant because FOAs were developed precisely to incorporate diversion-based information about closeness of competition that concentration measures may miss. In contrast to GUPPI, however, ΔCentrality captures not only incentives arising from the diversion among merging products, but also how those incentives propagate through the broader network of competing non-merging products.[29]

To assess whether ΔCentrality adds information beyond ΔHHI and GUPPI, we regress the simulated percentage price change of the merging parties across all possible mergers on ΔHHI, ΔCentrality, and GUPPI, first separately and then jointly in alternative combinations, 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 on the conventional 0–10,000 HHI scale, is associated with approximately a 0.7 percentage-point increase in post-merger prices. Column (2) indicates that a one-percent increase in ΔCentrality predicts a 2.6 percentage-point decrease in prices, while Column (3) shows that a one-percent increase in GUPPI is associated with a 1.0 percentage-point increase in prices. When ΔCentrality is included jointly with ΔHHI in Column (4), both coefficients remain statistically and economically significant, and the R^2 rises to 0.775 from 0.738 in the univariate specification.

Importantly for our framework, Column (6) shows that ΔCentrality also adds substantial predictive information beyond GUPPI: the R^2 increases from 0.765 with GUPPI alone to 0.853 when the two measures are included together. Additionally, comparing Columns (5) and (6) indicates that complementing GUPPI with ΔCentrality rather than with ΔHHI substantially increases explanatory power, raising the R^2 from 0.814 to 0.853. This is consistent with changes in centrality implied by mergers, capturing the propagation of price effects across the network instead of simply merging parties' incentives to raise prices. In the fully saturated specification in Column (7), all three measures remain statistically significant, and, relative to the specification with only ΔHHI and GUPPI, the addition of ΔCentrality increases the R^2 from 0.814 to 0.854. By contrast, the additional contribution of ΔHHI beyond GUPPI and ΔCentrality in explaining the variation in prices implied by the merger is negligible.

[29] Since GUPPI is typically defined at the product level, in Section C.2 of the Appendix, we describe how we construct our aggregate merger-level GUPPI measure building on Conlon and Mortimer (2013).

Panel B presents the same exercise using price effects simulated from the network model. In this case, the economic significance of the coefficient on ΔCentrality is larger than in Panel A. This is unsurprising, since centrality is theoretically grounded in the network structure of demand that generates the simulated merger effects. Still, the joint specifications show similar qualitative patterns, with all three screens remaining statistically significant in the fully saturated specification of Column (7).

Overall, these results indicate that changes in centrality provide complementary predictive power that is not subsumed by either concentration measures or traditional first-order pricing-pressure tools. This distinction is central to our framework. While ΔHHI summarizes changes in market structure and GUPPI captures direct diverted-margin incentives, ΔCentrality reflects how merger-induced pricing pressure propagates through the broader network of differentiated products. The fact that centrality remains highly informative even in the RCNL simulations suggests that this network component captures economically relevant variation in merger effects beyond the information contained in standard screening statistics.

Unweighted centrality. One caveat on the analyses presented is that, in practice, merger screening requires tools that are easy to construct with limited data. Our theoretical measure of ΔCentrality in Equation (31) requires estimating the full network model to recover $(\boldsymbol{\alpha} - \boldsymbol{c}_0)/2$, which may not always be feasible for enforcement agencies operating under data constraints. To address this concern, we construct a simpler measure that replaces $(\boldsymbol{\alpha} - \boldsymbol{c}_0)/2$ with a vector of ones 1_J.[30] The resulting statistic corresponds to an unweighted Bonacich centrality given by:

$$\widehat{\boldsymbol{b}}\left(D_H, \frac{1}{2}, 1_J\right) = \left(I_J - \frac{1}{2}D_H\right)^{-1} \frac{1}{2} D_H 1_J, \tag{32}$$

where Θ is constructed using the same characteristics that enter indirect utility in Brenkers and Verboven (2006). This measure still encapsulates diversion information but only requires product characteristics, which are often accessible for both merging and non-merging parties. Moreover, advances in data collection and machine learning tools can allow agencies to quickly extract such information from firm websites and industry sources.

For any potential merger, we then compute $\Delta\widehat{\text{Centrality}}$ as:

$$\Delta\widehat{\text{Centrality}} := \boldsymbol{q}_0' \widehat{\boldsymbol{b}}\left(D_H^1, \frac{1}{2}, 1_J\right) - \boldsymbol{q}_0' \widehat{\boldsymbol{b}}\left(D_H^0, \frac{1}{2}, 1_J\right). \tag{33}$$

[30] Alternatively, as we show in Appendix A.11, if price sensitivity is already known, the Bonacich centrality can be obtained from product attributes and sales information: $\boldsymbol{b}\left(D_H, \frac{1}{2}, \frac{\boldsymbol{\alpha}-\boldsymbol{c}}{2}\right) = \frac{\beta}{2}\left[(I_J - \Theta)^{-1} - \left(H \odot (I_J - \Theta)\right)^{-1}\right]\boldsymbol{q}$.

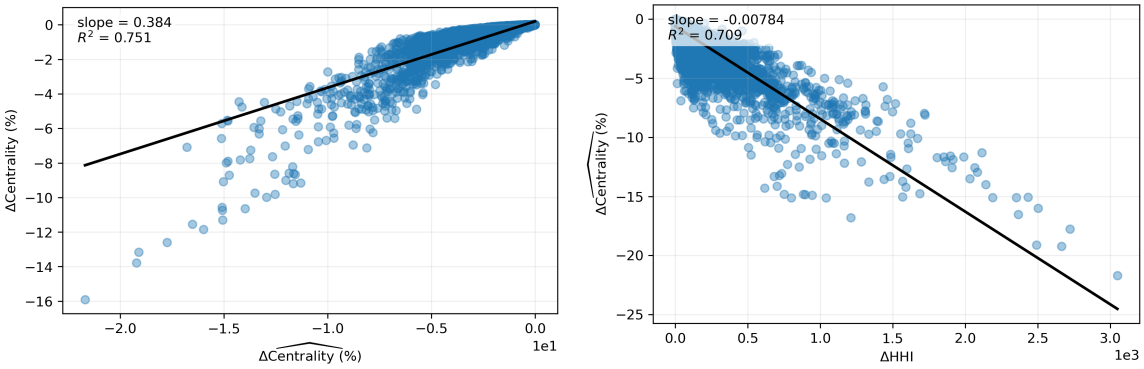


Figure 3: Unweighted Change in Centrality

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

Thus, constructing this index only requires product characteristics and pre-merger sales. Figure 3 shows the strong correlation between the theoretical ΔCentrality and its unweighted counterpart, demonstrating that this simpler measure captures economically relevant variation in our structural index. Moreover, $\Delta\widehat{\text{Centrality}}$ is highly correlated with ΔHHI. Replacing ΔCentrality with $\Delta\widehat{\text{Centrality}}$ in the regression of the simulated percentage price change of the merging parties across all possible mergers on the various screens considered yields a negative and statistically significant coefficient even after controlling for both ΔHHI and GUPPI, although the incremental explanatory power is smaller than for theoretical centrality (Table D.3). Thus, the version of the change in centrality constructed without price data retains information about merger-induced price changes beyond conventional screening measures.

Performance in Merger Screening. We evaluate screening rules based on ΔCentrality against concentration-based presumptions and a rule that approves mergers with a GUPPI below 5%. The concentration-based rules include the structural presumption introduced by the 2023 Merger Guidelines, which we refer to as the concentration-based presumption (CBP), and a rule that approves mergers with ΔHHI below 100 regardless of the post-merger concentration level. We implement these criteria as binary approval rules for this exercise.[31] We consider three regimes with different assumed merger-specific marginal cost reductions: strict (2%), baseline (5%), and lax (10%).

Table 5 reports approval rates and consumer surplus outcomes, all evaluated using the demand estimates from the RCNL model. Within each regime, the opti-

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

Table 5: Merger Screening with GUPPI, Concentration- and Centrality-based Rules

	Proposed mergers	Allowed mergers (%)	Optimal CS ($000,000)	ΔCS for allowed mergers	Type I Errors (%)	Type I loss ($000,000)	Type II Errors (%)	Type II loss ($000,000)
Panel A: Strict Regime (2% assumed efficiency)								
CBP	21,779	92.4	170948.036	144123.986	1.7	3008.262	62.1	23815.788
$\Delta\widehat{\text{Centrality}}$	21,779	92.4	170948.036	134887.188	1.4	1976.589	60.1	34084.260
GUPPI > 5%	21,779	99.2	170948.036	63831.092	0.0	0.000	95.2	107116.944
ΔHHI > 100	21,779	86.9	170948.036	157052.842	2.9	6006.811	33.3	7888.384
ΔCentrality, efficiency-adjusted	21,779	81.1	170948.036	147684.077	7.2	20374.938	20.3	2889.021
ΔCentrality, RCNL costs	21,779	88.4	170948.036	156480.697	2.1	5429.820	38.6	9037.519
Panel B: Baseline Regime (5% assumed efficiency)								
CBP	21,786	92.4	620882.616	576463.744	4.7	42293.410	32.4	2125.462
$\Delta\widehat{\text{Centrality}}$	21,786	92.4	620882.616	587899.592	4.4	28502.271	26.5	4480.753
GUPPI > 5%	21,786	99.2	620882.616	590596.751	0.0	0.000	83.4	30285.865
ΔHHI > 100	21,786	86.8	620882.616	543209.767	9.3	77008.783	8.2	664.066
ΔCentrality, efficiency-adjusted	21,786	91.3	620882.616	579862.307	4.9	40204.033	13.4	816.276
ΔCentrality, RCNL costs	21,786	96.7	620882.616	610806.977	0.7	4342.140	43.3	5733.499
Panel C: Lax Regime (10% assumed efficiency)								
CBP	21,786	92.4	1538830.263	1356424.796	6.7	182374.925	2.5	30.542
$\Delta\widehat{\text{Centrality}}$	21,786	92.4	1538830.263	1403742.332	6.7	135036.418	2.0	51.513
GUPPI > 5%	21,786	99.2	1538830.263	1536468.890	0.1	319.828	25.0	2041.545
ΔHHI > 100	21,786	86.8	1538830.263	1240871.417	12.4	297954.880	0.5	3.966
ΔCentrality, efficiency-adjusted	21,786	95.8	1538830.263	1460878.912	3.4	77760.766	6.0	190.585
ΔCentrality, RCNL costs	21,786	99.1	1538830.263	1535373.367	0.2	1949.852	24.5	1507.044

Notes: The table shows the number of proposed mergers (i.e., those that are profitable), the percentage of deals allowed, the maximum possible CS gains and the total 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), as well as the CS loss associated with each of them. We present 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). Six screening approaches using concentration-based, pricing-pressure, and centrality-based indices.

mal rule approves all proposed mergers that increase consumer surplus and blocks those that reduce it. We also report the probability of blocking a consumer-surplus-enhancing merger (Type I error), the probability of approving a consumer-surplus-reducing merger (Type II error), and the associated consumer surplus losses. These losses distinguish the frequency of mistakes from their economic importance: the shortfall from the optimal rule equals the gains forgone from blocking beneficial mergers plus the harm caused by approving detrimental mergers.

We assess each screening rule by the fraction of these maximum attainable gains that it delivers. The relative performance of the conventional screens varies substantially across regimes. CBP delivers 84.3%, 92.8%, and 88.1% of optimal consumer surplus gains in the strict, baseline, and lax regimes, respectively. The corresponding shares are 91.87%, 87.49%, and 80.64% for the ΔHHI rule, and 37.3%, 95.1%, and 99.8% for GUPPI. Thus, neither alternative outperforms CBP in all three regimes.

The ΔHHI rule reduces Type II errors relative to CBP in every regime, but blocks more beneficial mergers; the resulting forgone gains account for nearly all its welfare shortfall in the baseline and lax regimes. Conversely, GUPPI approves approximately 99.2% of proposed mergers. This permissiveness is costly under small efficiencies, where it allows 95.2% of harmful mergers, but delivers nearly optimal consumer surplus under lax efficiencies. Its Type II error rate remains 25.0% in the lax regime, illustrating that a high conditional error rate need not imply a large aggregate welfare loss.

We next consider an efficiency-adjusted centrality-based rule that approves mergers satisfying

$$\boldsymbol{q}_1'\boldsymbol{b}\left(D_H^1, \frac{1}{2}, \frac{\boldsymbol{\alpha}-\boldsymbol{c}_1}{2}\right) - \boldsymbol{q}_0'\boldsymbol{b}\left(D_H^0, \frac{1}{2}, \frac{\boldsymbol{\alpha}-\boldsymbol{c}_0}{2}\right) > 0. \tag{34}$$

Here, subscripts and superscripts 0 and 1 denote pre- and post-merger outcomes, respectively, and $\boldsymbol{q}_1$ denotes the predicted post-merger equilibrium quantities. Pre-merger marginal costs $\boldsymbol{c}_0$ are recovered from the network demand estimates, and $\boldsymbol{c}_1$ incorporates the cost reductions assumed under the relevant regime. We refer to the expression on the left-hand side as efficiency-adjusted ΔCentrality. Relative to the alternative screens considered above, the network model supplies a rule with a theory-grounded threshold that incorporates merger efficiencies directly into the criterion. Its implementation uses the parsimonious network demand estimates together with observed product characteristics, prices, and quantities.

The efficiency-adjusted ΔCentrality delivers 86.4%, 93.4%, and 94.9% of optimal consumer surplus gains, exceeding CBP in every regime. In the strict regime, efficiency-adjusted centrality raises the share of optimal consumer surplus gains by 2.08 percentage points relative to CBP, equivalent to an additional $3.56 billion across the simulated merger counterfactuals, or approximately $0.16 million per proposed merger.[32] The source of the improvement differs across regimes. Under strict efficiencies, the rule reduces the Type II error rate from 62.1% to 20.3%, at the cost of increasing the Type I error rate from 1.7% to 7.2%. The reduction in losses from harmful approvals more than offsets the additional forgone gains. In the baseline regime, the improvement not only reflects lower losses from approving harmful mergers, but also from blocking beneficial mergers, although the Type I error rate increases slightly. In the lax regime, the gains instead arise from blocking fewer beneficial mergers: the Type I error rate falls from 6.7% to 3.4%, while the Type II error rate rises from 2.5% to 6.0%. The additional harm from these approvals is small relative to the gains preserved.

[32] Aggregate consumer surplus amounts sum over separate merger counterfactuals and should not be interpreted as the welfare effect of mergers occurring jointly. The per-merger figure divides the aggregate gain relative to CBP by the 21,779 proposed mergers.

Individual merger counterfactuals illustrate the magnitude of the harm that concentration screens can miss. Table C.1 considers a hypothetical Fiat–Ford merger in France in 1984. The deal falls below the ΔHHI threshold of 100, yet the RCNL benchmark predicts annual consumer surplus losses of $18.9 million under the 2% efficiency assumption. The efficiency-adjusted ΔCentrality blocks the transaction when efficiencies are 2%, whereas it approves it with a 5% efficiency, when it would result in a $3.8 million gain in consumer surplus.[33] This illustrates the importance of incorporating efficiencies into the screening criterion. Table D.5 further examines the complementarity between centrality and concentration-based screening by documenting where the rules disagree and the associated consumer surplus effects. We find that centrality identifies harmful mergers missed by concentration-based screening when efficiencies are small and helps avoid blocking beneficial mergers when efficiencies are larger. We find that centrality identifies harmful mergers cleared by CBP when efficiencies are small, while CBP helps identify beneficial mergers blocked by centrality when efficiencies are larger.

To examine the role of marginal cost estimates, we also implement the centrality criterion using costs recovered from the RCNL model and the corresponding regime-specific cost reductions. This version delivers 91.5%, 98.4%, and 99.8% of optimal consumer surplus gains, improving on the network-cost implementation in every regime. These results follow from the choice of the RCNL demand model as the evaluation benchmark. If consumer surplus were instead evaluated using network demand, the original efficiency-adjusted centrality rule would attain the optimum by construction. Its shortfall therefore reflects differences between the network and RCNL models' welfare predictions, and which are mitigated by the use of RCNL costs.

Next, we consider an approach relying on $\Delta\widehat{\text{Centrality}}$, which uses product characteristics and pre-merger sales. This screen requires neither demand nor marginal cost estimation and can be implemented without product-level prices. We rank proposed mergers by this index and choose a cutoff that matches CBP's approval rate of 92.4% in all regimes. Holding fixed the share of mergers flagged for further scrutiny allows us to assess whether centrality improves the allocation of agencies' review capacity across deals at the same screening intensity. The characteristics-based rule delivers 78.9%, 94.7%, and 91.2% of optimal consumer surplus gains in the strict, baseline, and lax regimes, respectively. At the same approval rate, the characteristics-based rule reduces both Type I and Type II error rates in all three regimes. Centrality therefore identifies

[33] Table C.1 also illustrates the opposite scenario. Given a 2% efficiency, a hypothetical BMW–Renault merger in the same market would lead to an increase in CS and would be approved by the efficiency-adjusted ΔCentrality. However, the deal exceeds the ΔHHI threshold of 100. As a benchmark, Nevo (2000) predicts an annual consumer surplus loss of approximately $14 million from the approved Post–Nabisco merger under the assumption of no marginal cost reductions.

harmful mergers that the concentration-based presumption misses while allowing beneficial mergers that it flags, supporting its use in refining concentration-based scrutiny. These improvements in classification translate into higher consumer surplus in the baseline and lax regimes, although not in the strict regime, where the larger losses associated with the remaining harmful approvals outweigh the gains from fewer errors.

We perform two additional sets of analyses in Appendix D.1. First, the differences in predicted prices for large mergers across the network and the RCNL model documented in Section 5.4 motivate us to assess screening outcomes in highly concentrated markets—specifically mergers increasing the HHI to 2,500 or more. Figure D.5 shows that centrality-based screens rarely approve mergers associated with substantial consumer-surplus losses, suggesting that imperfect quantitative price predictions need not undermine their screening value. Second, we assess the robustness of our main findings in a scenario in which only firm-level—rather than product-level—information on non-merging firms' pre-merger sales is available. Table D.6 shows that splitting pre-merger sales equally across a firm's products does not affect our conclusions about the performance of $\Delta\widehat{\text{Centrality}}$.

Taken together, these results suggest that centrality-based measures can usefully complement traditional merger screens. A fixed Δ threshold performs well in the strict regime, but becomes overly restrictive as assumed efficiencies increase.[34] By contrast, centrality-based screening delivers more balanced performance across regimes, significantly improving on the CBP in the baseline and lax cases and outperforming GUPPI whenever approving harmful mergers is a central concern.

Finally, we note that there are also reasons to view the comparison with concentration-based rules in this setting as conservative for centrality-based screening. First, the automobile industry has relatively clear market segments, so product substitution patterns may already be reasonably well summarized by market shares. This limits the scope for ΔCentrality to improve on concentration-based measures. Because centrality incorporates product proximity in characteristic space, it may be more informative in industries where market boundaries are blurrier and substitution patterns are more heterogeneous. Second, the reference model used for the screening analysis in Table 5 is the RCNL. In logit-type models, substitution patterns are closely tied to market shares, which gives HHI-based screening measures a natural advantage. Centrality measures, by contrast, rely on distances in product characteristics. Under alternative demand systems in which substitution is less tightly linked to shares, centrality-based screens may offer larger improvements over traditional concentration thresholds.

[34] Table D.4 shows that, when the threshold of ΔCentrality and ΔHHI is optimally chosen, the two screens perform comparably. However, this requires using a threshold of 391 (896) in the baseline (lax) regime for the ΔHHI.

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. In this framework, 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 provides a unified characterization of the unilateral effects of horizontal acquisitions. We show that the price effects of a change in ownership can be decomposed into two forces: a direct upward-pricing-pressure term, familiar from first-order approaches, and a network-propagation term that captures how this initial pricing pressure spills over through the system of differentiated products. The first term depends on internalized diverted margins between the merging parties; the second depends on the structure of the broader diversion network. Hence, the framework extends first-order approaches by accounting for the indirect price responses that arise as pricing pressure propagates across competing products.

For full mergers, these direct and indirect effects aggregate into a change in centrality, which summarizes the unilateral effects of the acquisition and can therefore serve as a screening statistic for harmful mergers. The same structure also delivers closed-form expressions for consumer surplus changes and for the compensating marginal cost reductions needed to offset the anti-competitive effects of acquisitions.

Under linear-quadratic demand, diversion ratios are weighted inner products of product attributes, linking the diversion network directly to observable characteristics. This structure offers both computational and informational advantages: equilibrium merger effects can be evaluated through matrix operations, while a characteristics-based screen requires only product attributes and pre-merger sales. Product characteristics are often readily observable and increasingly accessible at scale, including from unstructured sources using LLMs. Like traditional first-order approaches, the framework avoids the explicit market delineation required by concentration-based screens and accommodates partial and common ownership.

In our application to the European automobile industry, the network model generates simulated merger effects comparable to those of an RCNL model. Changes in centrality are strongly associated with predicted price and consumer surplus changes and remains informative alongside concentration-based and first-order measures. In screening exercises, efficiency-adjusted ΔCentrality delivers higher consumer surplus

than the 2023 concentration-based presumptions. Moreover, at the same approval rate, a characteristics-based centrality reduces both types of classification errors in every regime and raises consumer surplus under baseline and lax efficiencies relative to the concentration-based presumptions. These findings indicate that using centrality-based rules can help refine merger scrutiny by incorporating information about substitution patterns and merger efficiencies in the review process.

References

Antón, M., F. Ederer, M. Giné, and B. Pellegrino (2023): "Mergers and acquisitions under common ownership," in *AEA Papers and Proceedings*, American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, vol. 113, 294–298.

Anton, M., F. Ederer, M. Gine, and M. Schmalz (2017): "Common Ownership, Competition, and Top Management Incentives," *Journal of Political Economy*, 81, 669–728.

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 (2021): "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.

Benson, D., S. Blattner, S. Grundl, Y. S. Kim, and K. Onishi (2024): "Concentration and Geographic Proximity in Antitrust Policy: Evidence from Bank Mergers," *American Economic Journal: Microeconomics*, 16, 107–33.

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.

Bhattacharya, V., G. Illanes, A. A. Kreps, J. D. Salas, and D. Stillerman (2026a): "A Large-Scale Evaluation of Merger Simulations," Tech. rep., National Bureau of Economic Research.

Bhattacharya, V., G. Illanes, and D. Stillerman (2026b): "Merger effects and antitrust enforcement: Evidence from us consumer packaged goods," *American Economic Review*, 116, 2996–3035.

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

Chen, Y., Y. Zenou, and J. Zhou (2018): "Competitive pricing strategies in social networks," *RAND Journal of Economics*, 49, 672–705.

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.

Conlon, C. T. and J. H. Mortimer (2013): "An experimental approach to merger evaluation," Tech. rep., National Bureau of Economic Research.

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.

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.

Epple, D. (1987): "Hedonic Prices and Implicit Markets: Estimating Demand and Supply Functions for Differentiated Products," *Journal of Political Economy*, 95, 59–80.

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.

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.

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.

Lancaster, K. J. (1966): "A New Approach to Consumer Theory," *Journal of Political Economy*, 74, 132–157.

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.

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.

Miller, N. H. and G. Sheu (2021): "Quantitative Methods for Evaluating the Unilateral Effects of Mergers," *Review of Industrial Organization*, 58, 143–177.

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.

Rosen, S. (1974): "Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition," *Journal of Political Economy*, 82, 34–55.

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

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

Valletti, T. and H. Zenger (2021): "Mergers with Differentiated Products: Where Do We Stand?" *Review of Industrial Organization*, 58, 179–212.

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

Appendices to "Diversion Networks and Horizontal Acquisitions"

by Athos Carvalho, Shanglyu Deng, Mario Leccese, and Marco Loseto

A Proofs

A.1 Proof of Lemma 1

Proof. This is the special case of Lemma 1′ with $H = I_J$, so $D_H = D$ and the stated expression follows immediately. ■

A.2 Proof of Lemma 1′

Proof. Fix the lineaized demand system $\boldsymbol{q} = -S(\boldsymbol{\alpha} - \boldsymbol{p})$. 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

$$q_j(\boldsymbol{p}) \;+\; \sum_{\ell=1}^{J} H_{j\ell}\,S_{\ell j}(\boldsymbol{p})\,m_\ell \;=\; 0, \qquad j = 1, \ldots, J. \tag{35}$$

By symmetry of S, the stacked FOCs are

$$\boldsymbol{q}(\boldsymbol{p}) + (H \odot S(\boldsymbol{p}))(\boldsymbol{p} - \boldsymbol{c}) = 0. \tag{36}$$

Substituting $\boldsymbol{q} = -S(\boldsymbol{\alpha} - \boldsymbol{p})$ into (36) 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

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

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

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

we have

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

Therefore

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

which is the claimed expression. ■

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

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

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,

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

Therefore total price change relative to the baseline H is

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

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

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

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

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

which is (16). ■

A.9 Proof of Proposition 4

Proof. From (18),

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

By the Woodbury matrix identity:

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

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). ■

A.11 Proposition 6 and Its Proof

Proposition 6 *For the linear-quadratic model, the following holds in equilibrium:*

$$\boldsymbol{b}\left(D_H, \frac{1}{2}, \frac{\boldsymbol{\alpha} - \boldsymbol{c}}{2}\right) = \frac{\beta}{2}\left[(I_J - \Theta)^{-1} - \Big(H \odot (I_J - \Theta)\Big)^{-1}\right] \boldsymbol{q}. \tag{37}$$

This implies that the centrality measure can be obtained from product attributes and sales information, provided that the price sensitivity β is known.

Proof. By Proposition 4,

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

Therefore,

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

Moreover, since $S = -(I_J - \Theta)/\beta$, the general-ownership first-order condition

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

implies

$$\boldsymbol{p} - \boldsymbol{c} = \beta\Big(H \odot (I_J - \Theta)\Big)^{-1}\boldsymbol{q}.$$

Hence

$$\frac{\boldsymbol{\alpha}-\boldsymbol{c}}{2}=\frac{\beta}{2}\left[(I_J-\Theta)^{-1}+\Big(H\odot(I_J-\Theta)\Big)^{-1}\right]\boldsymbol{q}.$$

By Lemma 1′,

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

Rearranging and substituting the two expressions above gives

$$\begin{aligned}\boldsymbol{b}\left(D_H,\frac{1}{2},\frac{\boldsymbol{\alpha}-\boldsymbol{c}}{2}\right)&=\frac{\boldsymbol{\alpha}-\boldsymbol{c}}{2}-(\boldsymbol{p}-\boldsymbol{c})\\&=\frac{\beta}{2}\left[(I_J-\Theta)^{-1}-\Big(H\odot(I_J-\Theta)\Big)^{-1}\right]\boldsymbol{q},\end{aligned}$$

which proves the claim. ■

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

$$(1 - \theta_{jj})(\alpha_j - c_j) > \sum_{k \neq j} |\theta_{jk}| \, (\alpha_k - c_k) \qquad \text{for all } j, \tag{38}$$

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

$$T_j(\boldsymbol{p}) := \frac{\alpha_j + c_j}{2} - \frac{1}{2(1-\theta_{jj})} \sum_{k \neq j} \theta_{jk} \Big[(\alpha_k - c_k) - (1 + H_{jk})(p_k - c_k) \Big]. \tag{39}$$

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 (39), 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 (38). 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 $\boldsymbol{p}, \tilde{\boldsymbol{p}} \in K$, (39) 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

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

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

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

where the strict inequality again follows from (38). Therefore there exists $\rho \in (0, 1)$ such that

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

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 Details on the Construction of Aggregate Screens

C.1 Aggregate Centrality

To compute the Bonacich centrality of a merger, we first proceed by computing the vector of centralities for all products, as presented in Equation 1:

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

We then aggregate the individual centralities to our market level measure using the weighted sum of all of them, $\boldsymbol{q}'\boldsymbol{b}\left(D_H, \frac{1}{2}, \frac{\boldsymbol{\alpha}-\boldsymbol{c}}{2}\right)$.

To simulate a merger, we compute new centralities taking into account the new ownership structure, which then allows us to update prices, using Lemma 1′. Then, we update quantities and market shares.

For a practical example, we consider the market in France, 1984, with two different proposed mergers, one between BMW and Renault, and the other between Fiat and Ford. Table C.1 presents some statistics regarding the mergers.

Table C.1: Characteristics and Implications of Selected Mergers

	Merger 1		Merger 2	
	BMW	Renault	Fiat	Ford
Merger characteristics				
Pre-merger HHI	2351.7		2351.7	
Post-merger HHI	2461.0		2446.2	
ΔHHI	109.3		94.5	
Δ%Centrality	-0.163%		-0.427%	
ΔP%	0.126%		0.547%	
ΔP% - Merged Firms	1.139%		2.733%	
ΔCS - 2% Efficiency ($000,000)	13.671		-18.852	
ΔCS - 5% Efficiency ($000,000)	66.186		3.778	
ΔCS - 10% Efficiency ($000,000)	157.187		43.345	
Number of Products	2	6	11	5
Quantity Sold	28,734	537,145	96,618	138,194
Market Share	1.71%	31.96%	5.75%	8.22%
Diversion statistics				
Average diversion to all products	0.79%	2.98%	0.42%	1.32%
Average diversion from all products	0.76%	0.79%	0.78%	0.76%
Average diversion from merger pair	2.09%	2.86%	0.87%	0.69%
Average diversion to merger pair	0.43%	3.42%	0.51%	1.49%
Average Cross-diversion to partner	2.33%	0.39%	1.57%	0.48%

The two mergers illustrate how the efficiency-adjusted centrality screen can yield different conclusions from the ΔHHI threshold of 100. Fiat–Ford falls below this threshold (ΔHHI = 94.5), yet the RCNL benchmark predicts an annual consumer-surplus loss of $18.9 million under the 2% efficiency assumption. The efficiency-adjusted centrality screen blocks this merger at 2% efficiency and approves it at 5%, when the benchmark predicts a $3.8 million gain. BMW–Renault illustrates the opposite case: its ΔHHI of 109.3 exceeds the threshold, but at 2% efficiency the benchmark predicts a $13.7 million gain and the efficiency-adjusted centrality screen approves the merger. In these examples, the centrality-based decisions therefore agree with the benchmark consumer-surplus effects.

The diversion statistics in the bottom panel describe substitution between the merger partners and their connections to other products. The centrality measure incorporates the full diversion matrix, while the reported averages summarize selected aspects of these substitution patterns. Fiat–Ford involves substantially more cross-firm product pairs than BMW–Renault. Its larger decline in centrality is consistent with internalizing diversion across this broader set of product-level links, even though average cross-diversion to the merger partner is not uniformly higher.

C.2 Aggregate GUPPI

Because the Gross Upward Pricing Pressure Index (GUPPI) is typically defined at the product-level and in our analyses we consider mergers of multi-product firms, this section describes how we construct an aggregate version of the GUPPI. We start from the GUPPI for product j with respect to product ℓ, which is given by:

$$\text{GUPPI}_{j\ell} = \frac{p_\ell^{(0)} - c_\ell}{p_j^{(0)}} D_{j\ell}\left(\mathbf{p}^{(0)}\right), \tag{40}$$

where $p_r^{(0)}$ denotes the pre-merger price of product r, c_r denotes its marginal cost, and $D_{j\ell}\left(\mathbf{p}^{(0)}\right)$ is the diversion ratio from product j to product ℓ, evaluated at pre-merger prices.

Since we consider mergers between multi-product firms, we aggregate product-level GUPPIs to obtain firm-level and merger-level measures. Consider two merging firms indexed by $f \in \{g, h\}$, with product sets

$$\mathcal{J}_f = \{1, \ldots, J_f\}.$$

For a product $j \in \mathcal{J}_g$, its GUPPI with respect to all products of firm h is defined as

the sum of the individual GUPPIs:

$$\text{GUPPI}_{j\to\mathcal{J}_h} = \sum_{\ell\in\mathcal{J}_h} \frac{p_\ell^{(0)} - c_\ell}{p_j^{(0)}} D_{j\ell}\left(\mathbf{p}^{(0)}\right). \tag{41}$$

This measures the upward pricing pressure on product j generated by the products owned by the merger partner.

We then aggregate across all products owned by firm g using a sales-weighted average:

$$\text{GUPPI}_g = \sum_{j\in\mathcal{J}_g} \omega_j^g \text{GUPPI}_{j\to\mathcal{J}_h}, \tag{42}$$

where

$$\omega_j^g = \frac{\text{sales}_j}{\sum_{r\in\mathcal{J}_g}\text{sales}_r} \tag{43}$$

is product j's share of firm g's total sales.

Analogously, for firm h, we compute

$$\text{GUPPI}_h = \sum_{k\in\mathcal{J}_h} \omega_k^h \text{GUPPI}_{k\to\mathcal{J}_g}, \tag{44}$$

where ω_k^h is product k's share of firm h's total sales.

Finally, we aggregate the two firm-level GUPPIs into a merger-level GUPPI using firm-level sales weights:

$$\text{GUPPI}_{\text{merger}} = \sum_{f\in\{g,h\}} \Omega_f \text{GUPPI}_f, \tag{45}$$

where

$$\Omega_f = \frac{\sum_{j\in\mathcal{J}_f}\text{sales}_j}{\sum_{r\in\mathcal{J}_g}\text{sales}_r + \sum_{r\in\mathcal{J}_h}\text{sales}_r} \tag{46}$$

is firm f's share of the merging parties' total sales. We use this final aggregation, $\text{GUPPI}_{\text{merger}}$, as our screening tool

D Additional Empirical Analyses

Table D.1: RCNL vs Network demand estimates.

Variable	RCNL	Network
Price	-3.7720***	-0.0012***
	(0.6959)	(0.0001)
Horsepower	0.0109	0.0023***
	(0.0102)	(0.0003)
Fuel Efficiency	-0.8009	-0.0024*
	(0.5307)	(0.0014)
Width	0.0520***	0.0066***
	(0.0063)	(0.0006)
Height	0.0089	0.0035***
	(0.0241)	(0.0007)
Foreign	-1.8058***	-0.2485***
	(0.1682)	(0.0040)

Notes: Table reports the estimated coefficients of the main product characteristics for both the RCNL and Network models. The dependent variable of the Network model is $\tilde{q}_t \equiv (I + \eta X X')q_t$, which is winsorized at 5%. The coefficient of the Network model are divided by the sample average of $\tilde{q}_{jt}$. Following Brenkers and Verboven (2006), the variable q_{jt} is the defined as product j's sales in market t divided by the number of households in market t.

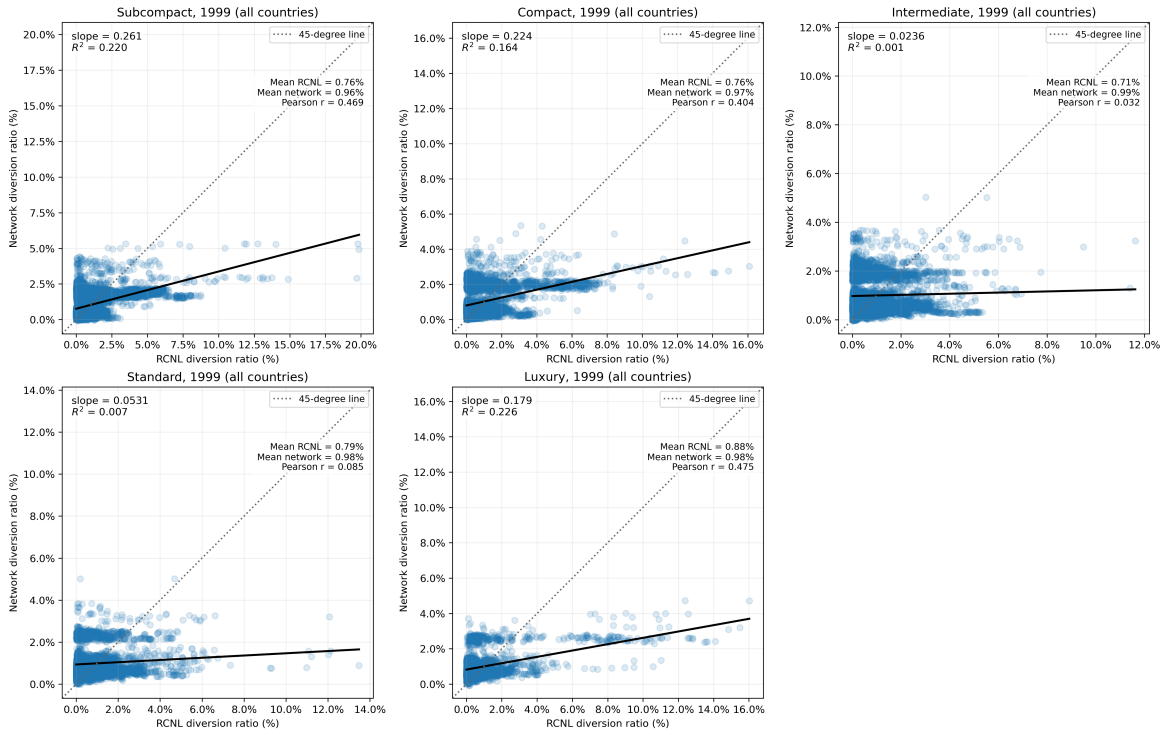


Figure D.1: Comparison between diversion ratios in the RCNL and the Network model.

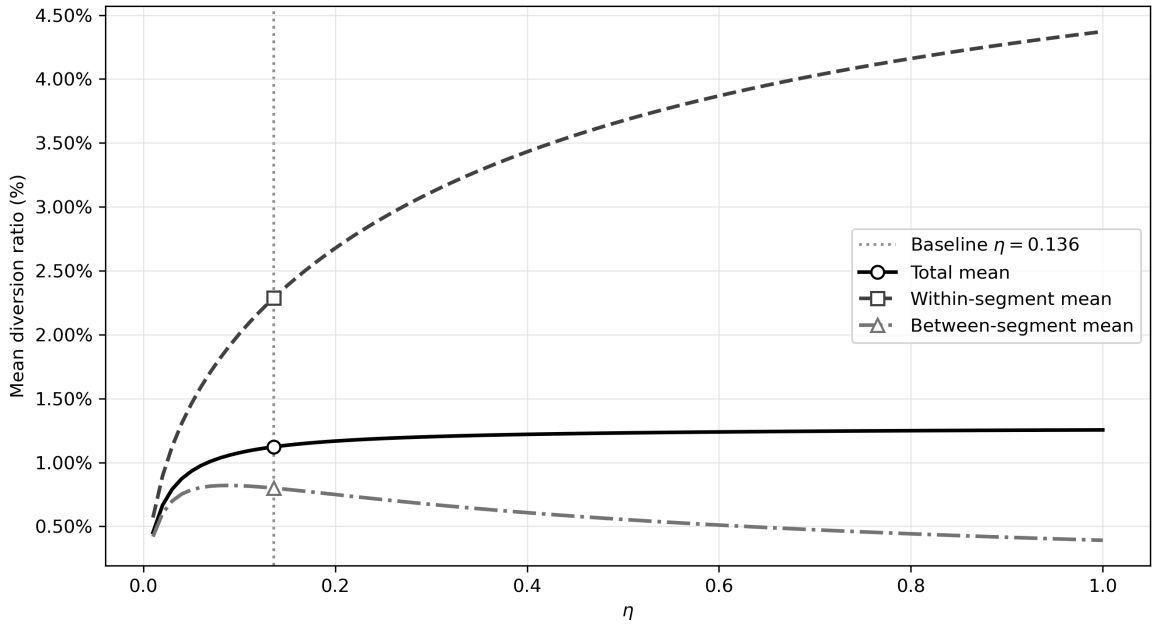


Figure D.2: Average Diversion Ratio for Different Values of η

Table D.2: Diversion ratios under characteristic replacement

	Subcompact		Compact		Intermediate		Standard		Luxury	
	Within segment	Across segments	Within segment	Across segments	Within segment	Across segments	Within segment	Across segments	Within segment	Across segments
Panel A: Replacement variance = 0.5										
Baseline	2.31	0.72	2.35	0.84	2.38	0.88	2.52	0.91	3.18	1.03
Drop horsepower	2.23	0.74	2.36	0.81	2.42	0.84	2.53	0.87	3.25	0.99
Drop fuel efficiency	2.31	0.68	2.38	0.79	2.43	0.83	2.59	0.86	3.38	0.99
Drop width	2.31	0.67	2.37	0.79	2.42	0.84	2.58	0.87	3.32	1.01
Drop height	2.30	0.66	2.38	0.79	2.42	0.84	2.57	0.88	3.29	1.02
Drop foreign	2.31	0.69	2.37	0.81	2.41	0.86	2.54	0.89	3.20	1.03
Panel B: Replacement variance = 1										
Baseline	2.31	0.72	2.35	0.84	2.38	0.88	2.52	0.91	3.18	1.03
Drop horsepower	2.16	0.74	2.27	0.81	2.36	0.84	2.45	0.86	3.12	0.99
Drop fuel efficiency	2.24	0.68	2.29	0.79	2.37	0.83	2.52	0.86	3.27	0.99
Drop width	2.23	0.67	2.29	0.79	2.37	0.84	2.51	0.87	3.21	1.01
Drop height	2.22	0.67	2.29	0.79	2.37	0.84	2.50	0.88	3.19	1.01
Drop foreign	2.23	0.70	2.28	0.81	2.36	0.86	2.48	0.89	3.11	1.02
Panel C: Replacement variance = 2										
Baseline	2.31	0.72	2.35	0.84	2.38	0.88	2.52	0.91	3.18	1.03
Drop horsepower	2.05	0.74	2.17	0.81	2.22	0.83	2.34	0.86	2.91	0.97
Drop fuel efficiency	2.12	0.68	2.19	0.79	2.23	0.83	2.42	0.86	3.08	0.98
Drop width	2.12	0.67	2.18	0.79	2.23	0.83	2.41	0.87	3.04	1.00
Drop height	2.10	0.67	2.18	0.79	2.23	0.83	2.41	0.88	3.02	1.00
Drop foreign	2.12	0.70	2.18	0.81	2.22	0.85	2.38	0.89	2.95	1.02

Notes: Values are percentages. Baseline retains all five characteristics without noise. Panels report the results for the three replacement variances. Diversion is averaged within product rows, then within segments, then equally across markets with valid values.

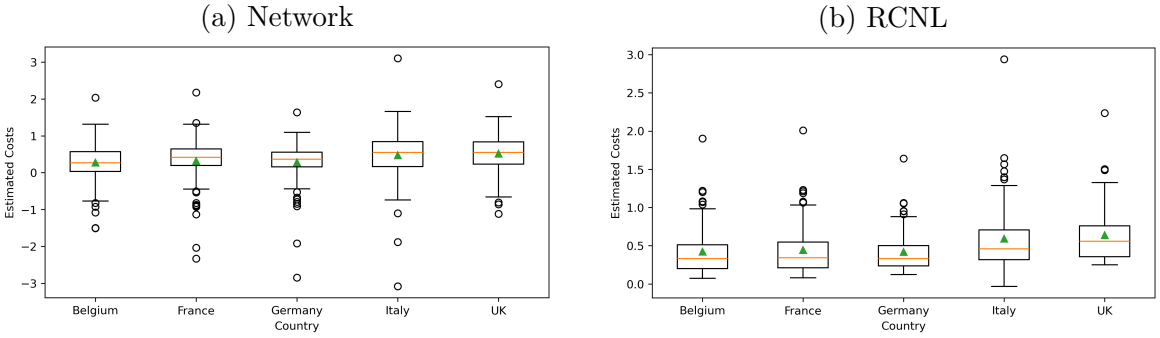


Figure D.3: Estimated marginal costs by country for 1999

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.

Table D.3: Simulated Price Effects and Screening Tools

	(1) $\%\Delta P$	(2) $\%\Delta P$	(3) $\%\Delta P$	(4) $\%\Delta P$	(5) $\%\Delta P$	(6) $\%\Delta P$	(7) $\%\Delta P$
Panel A: RCNL							
ΔHHI	0.702***			0.615***	0.343***		0.349***
	(0.014)			(0.015)	(0.016)		(0.017)
$\Delta\widehat{\text{Centrality}}$		-0.692***		-0.256***		-0.109***	-0.125***
		(0.021)		(0.012)		(0.015)	(0.011)
GUPPI			1.032***		0.613***	0.966***	0.530***
			(0.016)		(0.017)	(0.020)	(0.020)
N	21786	21786	21786	21786	21786	21786	21786
R^2	0.738	0.345	0.765	0.774	0.814	0.770	0.821
Panel B: Network game							
ΔHHI	1.097***			0.975***	0.676***		0.687***
	(0.025)			(0.027)	(0.029)		(0.031)
$\Delta\widehat{\text{Centrality}}$		-1.051***		-0.360***		-0.187***	-0.219***
		(0.035)		(0.018)		(0.027)	(0.020)
GUPPI			1.545***		0.718***	1.432***	0.573***
			(0.033)		(0.029)	(0.040)	(0.037)
N	21786	21786	21786	21786	21786	21786	21786
R^2	0.788	0.340	0.749	0.819	0.833	0.756	0.843

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. $\Delta\widehat{\text{Centrality}}$ denotes the change in the centrality measure constructed without price data. GUPPI is expressed in percentage terms. All regressions include market (country-year) and year fixed effects. Standard errors in parentheses.

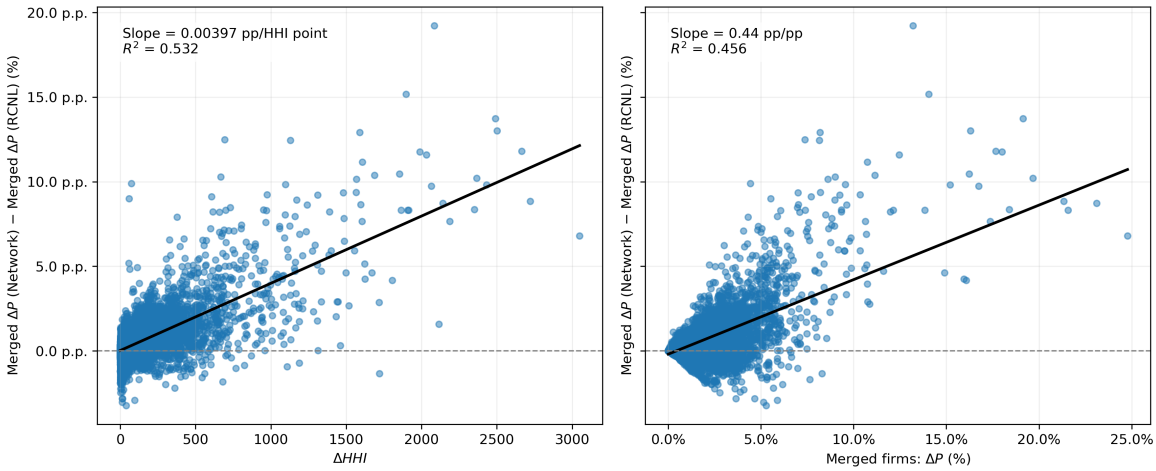


Figure D.4: The Simulated Price Effects of Large Acquisitions

Notes: The figure illustrates the difference between Network and RCNL predicted percentage price changes for merging firms plotted against the ΔHHI (left panel) and the percentage price changes among merging parties simulated using the RCNL demand estimates (right panel). Every data point is a simulated merger.

D.1 Additional Evidence on Centrality-based Rules Performance in Merger Screening

One advantage of centrality-based screening is that it delivers a theory-grounded threshold that accounts for the efficiency assumed by the agency. In what follows, we abstract away from this and compare centrality-based rules against one relying on ΔHHI, setting for all rules the threshold to maximize the change in CS implied by the merger.

To construct the rules, we proceed as follows. 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 opposite of the change in centrality ΔC_m, and (iii) whether the merger would be privately profitable and therefore proposed, denoted by $1\{\Pi_{m,r} > 0\}$.[35] As we do in the main text, the optimal screening rule only approves mergers that increase consumer surplus and is defined relative to consumer surplus effects simulated under the RCNL benchmark.

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:

$$W_r(\tau) = \sum_{m \in \mathcal{P}_r} (1 - d_m(\tau)) \, \Delta CS_{m,r}. \tag{47}$$

The welfare-optimal ΔCentrality threshold in regime r is therefore

$$\tau_r^* = \arg\max_{\tau} W_r(\tau). \tag{48}$$

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. Optimization is conducted separately for each enforcement regime and each decision rule.

Table D.4 summarizes the results. In the strict regime, optimized centrality captures 92.9% of optimal consumer surplus gains, compared with 92.1% for optimized

[35] 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.

Table D.4: Optimally Chosen Screening Thresholds

	Proposed mergers	Allowed mergers (%)	Optimal CS ($000,000)	ΔCS for allowed mergers	Type I Errors (%)	Type I loss ($)	Type II Errors (%)	Type II loss ($000,000)
Panel A: Strict Regime (2% assumed efficiency)								
ΔHHI*	21,779	87.4	170948.196	157388.710	2.6	5224.401	34.9	8335.084
ΔCentrality* (%)	21,779	87.5	170948.196	158765.035	2.8	5126.906	36.8	7056.254
$\Delta\widehat{\text{Centrality}}^*$ (%)	21,779	85.7	170948.196	149683.481	4.8	10205.690	36.2	11059.024
Panel B: Baseline Regime (5% assumed efficiency)								
ΔHHI*	21,786	97.1	620883.053	612206.673	0.5	2090.040	46.2	6586.340
ΔCentrality* (%)	21,786	96.5	620883.053	611519.261	0.9	4748.877	43.3	4614.916
$\Delta\widehat{\text{Centrality}}^*$ (%)	21,786	96.0	620883.053	600747.536	1.6	9228.279	47.4	10907.238
Panel C: Lax Regime (10% assumed efficiency)								
ΔHHI*	21,786	99.4	1538831.180	1533014.646	0.1	1075.926	44.5	4740.608
ΔCentrality* (%)	21,786	98.9	1538831.180	1532438.641	0.5	4571.077	26.0	1821.462
$\Delta\widehat{\text{Centrality}}^*$ (%)	21,786	98.9	1538831.180	1527426.926	0.5	6058.423	40.0	5345.831

Notes: The table shows the number of proposed mergers (i.e., those that are profitable), the percentage of deals allowed, the maximum possible CS gains and the total 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), as well as the CS loss associated with each of them. We present 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). The thresholds for all screens are optimally set for each regime-decision rule pair to maximize the total CS after the merger.

ΔHHI. In the baseline regime, these shares are 98.5% and 98.6%, respectively, while both reach approximately 99.6% in the lax regime. Achieving this performance with ΔHHI, however, requires raising its threshold from 100 to 391 in the baseline regime and 896 in the lax regime, allowing 97.1% and 99.4% of proposed mergers, respectively. Relative to optimized centrality, these more permissive concentration screens block fewer beneficial mergers but approve more harmful ones. These results provide an in-sample benchmark for each index when its cutoff can be selected using benchmark welfare outcomes, rather than prescribed independently of those outcomes.

Complementarity between CBP and ΔCentrality. To examine whether concentration and centrality provide distinct information for screening, we cross-classify the proposed mergers in the main screening exercise according to whether they are flagged by CBP and by efficiency-adjusted ΔCentrality. Table D.5 summarizes the results. The four cells in distinguish mergers flagged by both rules, by either rule alone, or by neither. For each regime, we report each cell's share of proposed mergers, the fraction that reduce consumer surplus under the RCNL benchmark, and their aggregate signed consumer-surplus change, including both gains and losses.

The cross-classification reveals a consistent ordering across regimes: harmful merg-

Table D.5: Complementarity between CBP and Efficiency-Adjusted ΔCentrality

Regime	CBP met	ΔCentrality < 0, eff.-adj.	Mergers (% of proposed)	Harmful mergers (% within cell)	Aggregate ΔCS ($000,000)
Strict	Yes	Yes	6.66	87.79	-157,824.95
	Yes	No	0.90	28.06	914.70
	No	Yes	12.20	57.08	-2,645.39
	No	No	80.25	3.76	146,769.37
Baseline	Yes	Yes	4.91	61.50	-76,998.22
	Yes	No	2.65	3.81	26,822.72
	No	Yes	3.74	26.13	23,424.15
	No	No	88.69	0.58	553,039.59
Lax	Yes	Yes	2.97	28.55	-456.83
	Yes	No	4.59	1.00	137,514.57
	No	Yes	1.25	1.10	33,060.45
	No	No	91.18	0.01	1,323,364.34

Notes: The table cross-classifies proposed mergers according to the CBP and efficiency-adjusted centrality screening rules. "Yes" indicates that the corresponding blocking condition is satisfied. The strict, baseline, and lax regimes assume marginal cost reductions of 2%, 5%, and 10%, respectively. For each cell, we report its share of all proposed mergers within the regime, the percentage of mergers in that cell with $\Delta CS < 0$, and the sum of their signed consumer surplus changes, including both beneficial and harmful mergers. Consumer surplus changes are evaluated using the RCNL benchmark. Aggregate amounts sum over separate merger counterfactuals and do not represent the effects of mergers occurring jointly.

ers are most prevalent among transactions flagged by both screens, less prevalent among those flagged by only one, and least prevalent among those flagged by neither. Each screen therefore provides information conditional on the other's decision. In the strict regime, for example, 87.8% of jointly flagged mergers are harmful, compared with 57.1% of those flagged only by ΔCentrality, 28.1% of those flagged only by CBP, and 3.8% of those flagged by neither.

The welfare consequences of disagreement depend on assumed efficiencies. Under strict efficiencies, mergers flagged by centrality but cleared by CBP account for 12.2% of proposals and generate aggregate consumer-surplus losses of $2.65 billion. ΔCentrality thus identifies harmful transactions missed by concentration-based screening. Under baseline and lax efficiencies, however, the corresponding groups generate aggregate gains of $23.42 billion and $33.06 billion, respectively. Clearing these groups would improve welfare relative to relying on ΔCentrality alone, indicating that CBP can help identify cases where centrality is overly restrictive. These results illustrate the existence of potential complementarities arising from using both approaches in merger review.

Screening with coarse information about non-merging firms. Information available in merger investigations may be more detailed for the merging parties than for their competitors. A practical attraction of first-order screening tools such as GUPPI

Table D.6: Merger Screening with Coarser Information about Merger Outsiders

	Proposed mergers	Allowed mergers (%)	Optimal CS ($000,000)	ΔCS for allowed mergers	Type I Errors (%)	Type I loss ($000,000)	Type II Errors (%)	Type II loss ($000,000)
Panel A: Strict Regime (2% assumed efficiency)								
CBP	21,779	92.4	170948.036	144123.986	1.7	3008.262	62.1	23815.788
$\Delta\widehat{\text{Centrality}}$	21,779	92.4	170948.036	134887.188	1.4	1976.589	60.1	34084.260
$\Delta\widehat{\text{Centrality}}$, coarse quantities	21,779	92.4	170948.036	134792.984	1.4	2029.047	60.1	34126.006
Panel B: Baseline Regime (5% assumed efficiency)								
CBP	21,786	92.4	620882.616	576463.744	4.7	42293.410	32.4	2125.462
$\Delta\widehat{\text{Centrality}}$	21,786	92.4	620882.616	587899.592	4.4	28502.271	26.5	4480.753
$\Delta\widehat{\text{Centrality}}$, coarse quantities	21,786	92.4	620882.616	587424.913	4.4	28974.884	26.6	4482.819
Panel C: Lax Regime (10% assumed efficiency)								
CBP	21,786	92.4	1538830.263	1356424.796	6.7	182374.925	2.5	30.542
$\Delta\widehat{\text{Centrality}}$	21,786	92.4	1538830.263	1403742.332	6.7	135036.418	2.0	51.513
$\Delta\widehat{\text{Centrality}}$, coarse quantities	21,786	92.4	1538830.263	1402480.912	6.7	136297.838	2.0	51.513

Notes: The table shows the number of proposed mergers (i.e., those that are profitable), the percentage of deals allowed, the maximum possible CS gains and the total 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), as well as the CS loss associated with each of them. We present 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). Three screening approaches comparing concentration-based and centrality-based indices.

is their focus on information about the merging parties and substitution between their products. We therefore examine whether characteristic-based centrality screening retains its performance when competitors' product-level sales are unavailable.

For each candidate merger, we retain the merging parties' observed product quantities. For each non-merging firm f, we replace individual product quantities with an equal allocation of its total sales, setting $q_j = Q_f/J_f$, where Q_f is the firm's total quantity and J_f its number of products in the relevant market-year. This preserves firm-level sales and market shares, leaving concentration-based screening unchanged. We recompute $\Delta\widehat{\text{Centrality}}$ using these imputed quantities and retain the original approval-rate target, matching CBP. The merger sample and RCNL welfare outcomes used to evaluate decisions remain unchanged.

Table D.6 shows that screening performance is largely preserved. Relative to the specification using observed product-level quantities, the share of optimal consumer-surplus gains achieved falls by only 0.06, 0.08, and 0.08 percentage points in the strict, baseline, and lax regimes, respectively. Error rates are virtually unchanged, and the screen continues to outperform CBP in the baseline and lax regimes. These results suggest that detailed product-level sales information for non-merging firms contributes little to the screen's performance in this application, provided their total sales and product characteristics remain available.

Screening outcomes in highly concentrated markets. The differences between the network and RCNL price predictions become more pronounced for mergers with larger concentration changes or predicted price effects. Such quantitative differences need not translate into poor screening decisions if the affected mergers are nevertheless identified as warranting scrutiny. As a complementary assessment, we examine the consumer-surplus consequences of mergers approved by each screening rule in highly concentrated markets.

We restrict the simulation sample to mergers with a post-merger HHI of at least 2,500, computed by combining the merging firms' pre-merger market shares. We retain the screening rules and thresholds used in the main analysis and evaluate consumer-surplus changes using the RCNL benchmark. Figure D.5 reports results separately for the strict, baseline, and lax regimes, corresponding to assumed efficiencies of 2%, 5%, and 10%.

For each rule, the left-hand panels report the share of approved mergers whose percentage consumer-surplus gain exceeds the threshold on the horizontal axis. The right-hand panels report the share whose percentage consumer-surplus loss exceeds the indicated magnitude. Specifically, letting A_r denote approval under rule r and g denote the RCNL percentage change in consumer surplus, the curves report

$$G_r(x) = \Pr(g > x \mid A_r, \mathrm{HHI}_{\mathrm{post}} \geq 2{,}500), \quad L_r(x) = \Pr(g < -x \mid A_r, \mathrm{HHI}_{\mathrm{post}} \geq 2{,}500).$$

Both use all approved mergers under the relevant rule in the restricted sample as the denominator. At zero, the curves therefore identify the shares of approvals generating gains and losses, respectively. Outcomes beyond the displayed 3% threshold remain included in the tail probabilities. The legends report each rule's approval rate within the restricted sample.

The gain curves largely overlap across screening rules within each regime, whereas the loss curves reveal clearer differences. In the strict regime, approximately 15% of mergers approved by GUPPI reduce consumer surplus, and approximately 4% generate losses exceeding 1%. Substantial losses are much less common among approvals under centrality-based rules: their loss tails are essentially zero at the 1% threshold. The characteristic-based centrality rule has a somewhat heavier loss tail than the other centrality rules in this regime, but these losses remain concentrated at smaller magnitudes.

Harmful approvals become less common as assumed efficiencies increase. In the baseline and lax regimes, efficiency-adjusted centrality approves 86.6% and 91.8% of mergers, respectively, compared with 82.8% under the concentration-based presumption, while retaining very small shares of harmful approvals. The centrality rule using RCNL costs approves 90.7% and 96.2%, respectively, also with limited harmful ap-

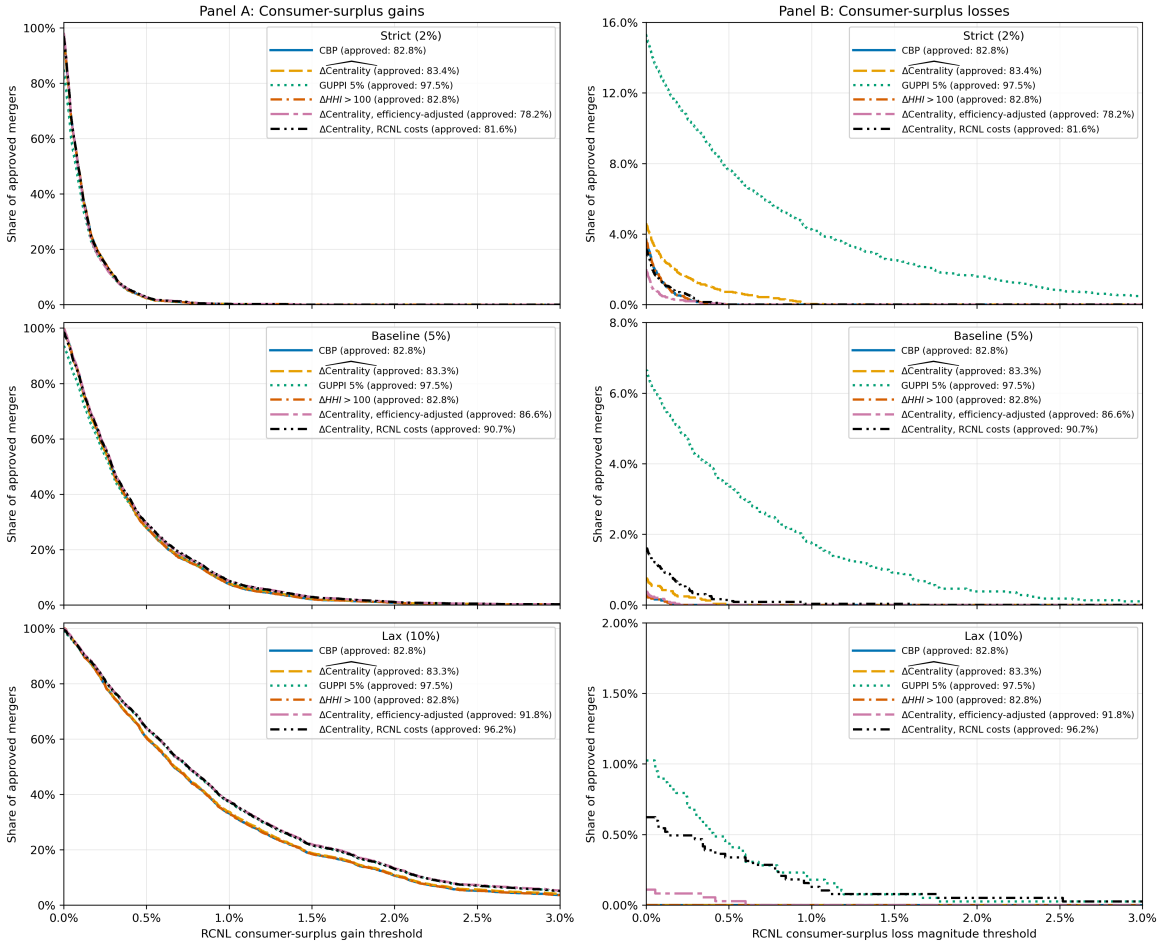


Figure D.5: Screens' Performance in Highly Concentrated Markets

Notes: The figure shows the share of approved mergers with post-merger $HHI > 2500$ according to each screening rule, for all scenarios. Panel A shows the the share of approved mergers with an overall consumer-surplus gain above a threshold X. Panel B does the same for consumer-surplus losses

provals. These results support the screening value of centrality in highly concentrated markets despite imperfect quantitative agreement between the models' price predictions.

E Comparison between Network and Logit models

To assess whether our network demand model delivers reasonable substitution patterns we compare it with four workhorse demand models: logit, nested logit, random-coefficient logit, and random-coefficient nested logit.

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

$$u_{ijt} = x_{jt}\beta - \alpha\frac{p_{jt}}{y_t} + \xi_{jt} + \varepsilon_{ijt}, \tag{49}$$

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

$$\delta_{jt} = \log s_{jt} - \log s_{0t}. \tag{50}$$

We then estimate

$$\delta_{jt} = x_{jt}\beta - \alpha\frac{p_{jt}}{y_t} + \xi_{jt} \tag{51}$$

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:

$$u_{ijt} = \delta_{jt} + \zeta_{ig} + (1-\sigma)\varepsilon_{ijt}, \tag{52}$$

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

$$\log s_{jt} - \log s_{0t} = x_{jt}\beta - \alpha\frac{p_{jt}}{y_t} + \sigma \log s_{j|g,t} + \xi_{jt}, \tag{53}$$

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

$$u_{ijt} = x_{jt}\beta - \alpha\frac{p_{jt}}{y_t} + \mu_{ijt} + \xi_{jt} + \varepsilon_{ijt}, \tag{54}$$

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:

$$u_{ijt} = x_{jt}\beta - \alpha\frac{p_{jt}}{y_t} + \mu_{ijt} + \sigma\zeta_{ig} + \xi_{jt} + \varepsilon_{ijt}, \tag{55}$$

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.