Oligopoly with Mixed Bundling*

Volker Nocke† Susumu Sato‡

October 24, 2026

Abstract

We develop a framework for studying oligopoly with mixed bundling. We consider multinomial logit demand over product combinations and, under this demand system, allow for an arbitrary number of firms and for arbitrary product portfolios, consumer valuations, and marginal costs. Despite the high dimensionality of action spaces and strategic interactions, firms' optimal pricing has a simple structure: each firm optimally sets product prices and bundle discounts to equate the markup it earns from each trade. The primitives governing strategic interaction can be summarized by a type network that captures the value of same-firm purchases, cross-firm purchases, and outside options. We establish the existence and uniqueness of equilibrium. We then apply the framework to study the effects of changes in product costs and qualities, entry, and conglomerate mergers, and show how the welfare effects of mixed bundling depend on market coverage and firm asymmetries.

JEL codes: D43, L13, L41, L42.
Keywords: mixed bundling, oligopoly, potential games, mergers, lift.

1 Introduction

Mixed bundling entails a firm charging a distinct price for each product in its portfolio and, in addition, combining some of these products into a bundle and offering it at a discount. This business practice is widely used in many industries, including telecommunications, software, media, insurance, and energy.[1] Importantly, mixed bundling is particularly prevalent in oligopolistic markets dominated by a small number of large, multiproduct firms. Understanding firms' pricing decisions in such markets—and how these decisions respond to changes in the

*We thank audiences at the 2026 Decentralization Conference (Tokyo), the 2026 Digital Markets Workshop (Mannheim), Hitotsubashi, U Mannheim, and the 2025 World Congress of the Econometric Society (Seoul). We gratefully acknowledge support by the German Research Foundation (DFG) through CRC TR 224 (Project B03).

†University of Mannheim, Department of Economics and MaCCI; also affiliated with CEPR. Email: volker.nocke@gmail.com

‡Hitotsubashi University, Institute of Economic Research. Email: susumusato.econ@gmail.com

[1] In the U.K. telecommunications industry, Ofcom documents substantial savings from purchasing services in a bundle rather than separately (Ofcom, 2026). U.S. insurance regulators describe discounts for purchasing multiple policies, such as home and auto insurance, from the same insurer as a commonly used practice (National Association of Insurance Commissioners, 2023). Mixed bundling can also span distinct markets such as energy and broadband (see, e.g., New Zealand Commerce Commission, 2023).

market environment, including entry and mergers—therefore seems of first-order importance (as do the welfare consequences of these decisions).

Despite this, the literature on mixed bundling in oligopoly is still underdeveloped. Arguably, this is in large part due to technical difficulties. First, bundle discounts mechanically increase the dimensionality of firms' strategy space. Second, any product in a firm's portfolio may be a substitute or complement to other products in that firm's portfolio and to products offered by rival firms. Moreover, introducing bundle discounts can turn two other products, not contained in the bundle, from substitutes to complements. As a result, the payoff consequences of changing one product's price are highly complex, and so are strategic interactions. In light of this, the extant literature has focused on simple market structures, such as duopoly or symmetric oligopoly, and on covered markets (see Matutes and Regibeau, 1992, Armstrong and Vickers, 2010, and Zhou, 2021). Such simple market structures leave limited scope for applied work such as merger analysis, where changes in ownership and portfolio structure naturally arise, and firm asymmetries cannot be assumed away.

Against this background, we develop a framework for oligopoly with mixed bundling that allows for an arbitrary number of firms as well as arbitrary heterogeneity in product portfolios, valuations of product combinations, and marginal costs. In particular, we assume that there are two markets, with consumers demanding one product from each market. In each market, there is an arbitrary number of products as well as an outside option, so markets are not (necessarily) fully covered. The demand for each "product pair" (consisting of one product from each market) is assumed to have a multinomial logit (MNL) structure. Each firm may be active in a single market or in both, and offer an arbitrary number of products in each market. A firm sets a distinct price for each product in its portfolio and a bundle discount for each product pair in its portfolio.

In contrast to "standard" multiproduct-firm oligopoly with MNL demand (Nocke and Schutz, 2018), the pricing game is not aggregative. Despite this, we show that the dimensionality of the pricing game can be reduced by noting that each firm optimally sets a common absolute markup on each product (and product pair) in its portfolio. In the reduced pricing game (in which each firm sets a uni-dimensional markup on all its products), the sign of strategic interaction between two firms is governed by their *lift*—the probability that their products are purchased together relative to what independent purchase would imply. A key result consists in showing that this reduced pricing game admits a strictly concave *ordinal potential* (Monderer and Shapley, 1996) and that this potential has a global maximizer, thereby proving existence and uniqueness of a pure-strategy Nash equilibrium.

We also show that there is a low-dimensional representation of the primitives governing the strategic interactions between firms. Product-pair values, marginal costs, and product ownership can be aggregated into a *type network* consisting of an outside-option type, a standalone type for each firm, and a cross-firm type for each pair of firms. A firm's standalone type reflects the social value from consuming only products from that firm, while a cross-firm type represents the social value from consuming one product from each of the two corresponding firms. The type network representation is useful for deriving comparative statics, as changes in product values, marginal costs, and the set of firms can all be represented as changes in the type network. We show that an increase in a standalone type raises consumer surplus, whereas an increase in a cross-firm type may lower it. Despite these countervailing forces, an equal absolute reduction in the marginal costs of all products supplied by one firm raises consumer surplus, and so does firm entry. Entry can nevertheless raise some incumbent

markups and make some consumers worse off.

We apply our framework to mergers between firms that may sell products in the same market(s) or in different markets. A merger expands the strategy space of the merger partners by allowing the merged firm to offer bundle discounts on a larger set of product pairs, as the merged firm offers product pairs consisting of products from different merger partners. We provide a three-step decomposition of the effects of mergers on pricing and consumer welfare. At the first step, holding fixed the prices of all nonmerging outsiders as well as the set of product pairs on which the merger partners offer bundle discounts, the merger partners adjust their markups to internalize the pricing externality they impose on each other. At the second step, still holding fixed the prices of all outsiders, the merged firm introduces bundle discounts on the new product pairs it offers, which leads to further adjustment of its markups. At the third step, all firms reoptimize, resulting in the post-merger pricing equilibrium.

The three-step decomposition allows us to provide a rich characterization of the effects of mergers in the absence of merger-specific efficiencies. We show that, at the end of the first two steps, the merged firm sets a markup exceeding both merging firms' pre-merger markups. Consequently, any product pair that contains products of only one merging firm becomes more expensive. By contrast, the new bundle discounts eliminate double marginalization on product pairs containing products supplied by two different merger partners and may therefore lower their total prices. Hence, at the end of the second step, the effect on consumer surplus depends on whether the elimination of double marginalization is large enough to offset the increase in markups. We show that, at this point, consumer surplus rises when the merging firms' pre-merger lift exceeds a certain cutoff. Furthermore, the first two steps provide a one-sided screen for the merger's overall effect on consumer surplus: if a merger does not harm consumers when holding outsider prices fixed, the same remains true when taking into account all equilibrium effects. Lift remains informative about the merger's full-equilibrium effects: holding fixed the merging firms' pre-merger demands and the outsiders' purchase patterns, a higher lift makes the merger more likely to benefit consumers. We also derive sufficient conditions on market shares for a merger to benefit or harm consumers.

We finally study the effects of mixed bundling by comparing it with separate pricing. We focus on independent markets to isolate the effects of the cross-market interdependence created by bundle discounts. Separate pricing leaves the two markets independent, whereas mixed bundling links them by favoring same-firm purchases. Holding fixed consumers' market-level choice probabilities, this bundling-induced interdependence creates a deadweight loss. At the same time, bundle discounts can move choices within each market closer to the efficient allocation and thereby reduce the market-level deadweight losses. The overall effect of mixed bundling on aggregate surplus therefore depends on whether improvements in the market-level allocations are outweighed by the deadweight loss from cross-market interdependence. In a symmetric covered oligopoly, the market-level losses are absent, so mixed bundling lowers aggregate surplus. It also lowers each firm's profit. Consumer surplus rises when there are only two firms but falls when the number of firms is sufficiently large. When outside options are sufficiently attractive or firms are sufficiently asymmetric in both markets, reductions in the market-level losses dominate instead. In that case, mixed bundling raises aggregate surplus and industry profit but harms consumers.

Related literature. There is a large literature studying mixed bundling as a screening device, primarily under monopoly (Adams and Yellen, 1976; McAfee, McMillan, and Whinston,

1989; Armstrong, 1996, 1999; Rochet and Choné, 1998; Manelli and Vincent, 2006). The literature on mixed bundling under oligopoly is much smaller and focuses mostly on duopoly (e.g., Matutes and Regibeau, 1992; Anderson and Leruth, 1993; Thanassoulis, 2007; Armstrong and Vickers, 2010). One exception is Zhou (2021), who studies a perfectly symmetric oligopoly in which each firm offers one product in each of two markets under the assumption of full market coverage (no outside option). He allows for general distributions of the taste shocks, but is able to establish equilibrium existence only in special cases. By contrast, our framework imposes the assumption of multinomial logit taste shocks but allows for an arbitrary number of firms, heterogeneous product portfolios, and outside options. Moreover, we provide a proof of equilibrium existence and uniqueness under general conditions.

Our merger analysis relates to work on cross-market mergers that enable merger partners to offer bundle discounts. The literature typically studies stylized two-market settings—in which two single-product firms operate in each of two separate markets—and examines the effects of cross-market mergers that turn otherwise independent products into complements (e.g., Choi, 2008; Flores-Fillol and Moner-Colonques, 2011; Thanassoulis, 2011). Chen and Rey (2026) study a merger between two single-product firms operating in two different markets. In their model, a merger makes products complementary because of (heterogeneous) joint-purchase benefits on the consumer side, even in the absence of any bundling. Compared to these studies, we analyze the effects of mergers in much richer environments.

Methodologically, our equilibrium analysis is also related to the literature on multiproduct-firm oligopoly. Under multinomial logit demand, a multiproduct firm optimally charges the same absolute markup on all its products (Konovalov and Sándor, 2010). Nocke and Schutz (2018) introduce a broader class of IIA demand systems under which firms charge a common adjusted markup and the pricing game is aggregative. Their aggregative-games approach does not apply here because a firm's payoff is a high-dimensional function of rival firms' prices. We nevertheless establish that the common markup property carries over to our setting, which greatly reduces the dimensionality of firms' pricing problem. More recently, Nocke and Schutz (2025) pioneer a potential games approach to multiproduct pricing games, including environments in which products are purchased jointly.[2] Our original mixed-bundling game generally does not admit an ordinal potential, but once the game is reduced to one in which each firm sets a uni-dimensional markup, it does admit a (strictly concave) ordinal potential.

2 Model

Consider two markets, A and B, each with a finite set of products and an outside option. Let $\mathcal{A}$ and $\mathcal{B}$ denote the corresponding choice sets, with generic elements a and b, respectively. In both sets, 0 denotes the outside option, representing no purchase in the corresponding market.

Consumers. The representative consumer has quasi-linear preferences. Given the vector of prices for product pairs $ab \in \mathcal{A} \times \mathcal{B}$,

$$\mathbf{p} = (p_{ab})_{ab \in \mathcal{A} \times \mathcal{B}},$$

indirect utility is

$$V(\mathbf{p}, y) = \lambda \log H + y,$$

[2]The concept of (ordinal) potential games was formalized by Monderer and Shapley (1996).

where

$$H \equiv \sum_{a\in\mathcal{A}}\sum_{b\in\mathcal{B}} h_{ab},$$
$$h_{ab} \equiv \mathrm{e}^{\frac{s_{ab}-p_{ab}}{\lambda}},$$

$s_{ab} \in \mathbb{R}\cup\{-\infty\}$ is the deterministic gross utility from product pair ab, $\lambda > 0$ is a price-sensitivity parameter, and y is income. Consumer surplus therefore equals $\lambda \log H$. In the following, we will sometimes refer to H as the *CS-index*. We say that product pair ab is *available* if $s_{ab} > -\infty$; if $s_{ab} = -\infty$, it is *unavailable*.

Applying Roy's identity, we obtain the demand for product pair ab:

$$q_{ab}(\mathbf{p}) = -\frac{\partial V(\mathbf{p},y)/\partial p_{ab}}{\partial V(\mathbf{p},y)/\partial y} = \frac{h_{ab}}{H}.$$

The associated marginal demands are

$$q_a \equiv \sum_{b\in\mathcal{B}} q_{ab} \quad \text{and} \quad q_b \equiv \sum_{a\in\mathcal{A}} q_{ab}.$$

This demand system can equivalently be derived from either simultaneous or sequential discrete choice by a population of heterogeneous consumers.

Simultaneous discrete choice with taste shocks at the product-pair level. Consumer l obtains utility

$$u_{abl} = \log h_{ab} + \varepsilon_{abl}$$

from product pair ab, where the shocks $(\varepsilon_{abl})_{ab\in\mathcal{A}\times\mathcal{B}}$ are i.i.d. type-I extreme value with unit scale and locations normalized to give mean zero. The probability that the consumer chooses ab is

$$\Pr\left(u_{abl} \geq u_{a'b'l} \text{ for every } a'b' \in \mathcal{A}\times\mathcal{B}\right) = \frac{h_{ab}}{H}.$$

Sequential discrete choice with taste shocks at the product level. Consumer l obtains utility

$$u_{abl} = \log h_{ab} + \varepsilon_{al} + \varepsilon_{bl}$$

from product pair ab, where $(\varepsilon_{al})_{a\in\mathcal{A}}$ and $(\varepsilon_{bl})_{b\in\mathcal{B}}$ are independent collections of i.i.d. type-I extreme-value shocks with unit scale and locations normalized to give mean zero. The consumer first observes $(\varepsilon_{al})_{a\in\mathcal{A}}$ and chooses a product in market A. She then observes $(\varepsilon_{bl})_{b\in\mathcal{B}}$ and chooses a product in market B.

Conditional on the choice of a, the probability of choosing b is

$$q_{b|a} = \frac{h_{ab}}{h_a}, \quad h_a \equiv \sum_{b'\in\mathcal{B}} h_{ab'}.$$

The expected utility from choosing a is

$$\mathbb{E}\left[\max_{b\in\mathcal{B}} u_{abl}\right] = \log h_a + \varepsilon_{al}.$$

The probability of choosing a is therefore

$$q_a = \frac{h_a}{\sum_{a' \in \mathcal{A}} h_{a'}} = \frac{h_a}{H}.$$

It follows that the probability of choosing product pair ab is

$$q_a q_{b|a} = \frac{h_{ab}}{H}.$$

Firms. Let $\mathcal{N}$ be the (finite) set of firms. For each $i \in \mathcal{N}$, let $\mathcal{A}^i \subseteq \mathcal{A} \setminus \{0\}$ and $\mathcal{B}^i \subseteq \mathcal{B} \setminus \{0\}$ denote the products supplied by firm i in the two markets. For any $i, j \in \mathcal{N}$, $i \neq j$, we have $\mathcal{A}^i \cap \mathcal{A}^j = \emptyset$ and $\mathcal{B}^i \cap \mathcal{B}^j = \emptyset$; moreover,

$$\mathcal{A} \setminus \{0\} = \bigcup_{i \in \mathcal{N}} \mathcal{A}^i, \quad \mathcal{B} \setminus \{0\} = \bigcup_{i \in \mathcal{N}} \mathcal{B}^i.$$

That is, each product is offered by exactly one firm. A firm may offer multiple products in each market, or only products in one market.

Firms produce at constant marginal costs that can differ arbitrarily across products and firms: c_a for $a \in \mathcal{A} \setminus \{0\}$ and c_b for $b \in \mathcal{B} \setminus \{0\}$. We adopt the convention that the outside option is costless: $c_0 = 0$ for both $0 \in \mathcal{A}$ and $0 \in \mathcal{B}$. Each firm i sets price p_a for product $a \in \mathcal{A}^i$ and p_b for product $b \in \mathcal{B}^i$. Firm i may also offer discount δ_{ab} for $ab \in \mathcal{A}^i \times \mathcal{B}^i$, i.e., whenever a consumer purchases an ab-bundle from firm i. The total price of product pair ab, p_{ab}, is therefore given by

$$p_{ab} = p_a + p_b - \delta_{ab},$$

with $\delta_{ab} = 0$ if products a and b are not offered by the same firm. The price of the outside option in each market is normalized to zero, so that $p_{a0} = p_a$, $p_{0b} = p_b$, and $p_{00} = 0$. Firm i's profit is given by

$$\Pi^i = \sum_{a \in \mathcal{A}^i} q_a (p_a - c_a) + \sum_{b \in \mathcal{B}^i} q_b (p_b - c_b) - \sum_{ab \in \mathcal{A}^i \times \mathcal{B}^i} q_{ab} \delta_{ab}.$$

Pricing game. In the mixed-bundling pricing game, firms simultaneously set prices and bundle discounts $\{(p_a)_{a \in \mathcal{A}^i}, (p_b)_{b \in \mathcal{B}^i}, (\delta_{ab})_{ab \in \mathcal{A}^i \times \mathcal{B}^i}\}_{i \in \mathcal{N}}$. The solution concept is pure-strategy Nash equilibrium, which we call *pricing equilibrium*.

3 Equilibrium Analysis

In this section, we begin by showing that each firm's optimal pricing problem can be reduced to choosing a common markup across all product combinations in its portfolio. Next, we show that the dimensionality of the primitives governing the strategic interaction between firms—the gross utilities of all product combinations, the marginal costs of all products, and the ownership structure of the different products—can be greatly reduced and represented by a type network. Finally, we show that the pricing game with mixed bundling has an equilibrium and that it is unique.

3.1 Common markup property

We first consider a firm's optimal pricing problem, holding fixed all rival prices (including bundle discounts). In general, this problem is complicated because of its high dimensionality and because any product may be a substitute to some products but a complement to others. Fortunately, despite the game not being aggregative, the pricing problem can be reduced to choosing a common markup for all products (or product combinations) in the firm's portfolio.

To see this, consider first a relaxed problem in which for any of its products, firm i can charge a price that depends on the other product with which this product is consumed. Letting $\mathcal{Z}^i \equiv \{ab \in \mathcal{A} \times \mathcal{B} | a \in \mathcal{A}^i \text{ or } b \in \mathcal{B}^i\}$ denote the set of product pairs involving at least one of firm i's products, the demand for each $z \in \mathcal{Z}^i$ is

$$q_z^i = \frac{\mathrm{e}^{\frac{\tilde{s}_z^i - p_z^i}{\lambda}}}{H},$$

where

$$\tilde{s}_z^i = \begin{cases} s_{ab} - p_b & \text{if } a \in \mathcal{A}^i,\ b \notin \mathcal{B}^i, \\ s_{ab} - p_a & \text{if } a \notin \mathcal{A}^i,\ b \in \mathcal{B}^i, \\ s_{ab} & \text{if } a \in \mathcal{A}^i,\ b \in \mathcal{B}^i. \end{cases}$$

Whenever product pair $z \in \mathcal{Z}^i$ combines one of firm i's products and one product offered by a rival, the firm only partially controls the total price for that pair. However, for the purpose of firm i's optimal pricing problem, what matters is only $\tilde{s}_z^i$—the difference between the gross surplus of the product pair z and the price charged by the rival firm, which firm i takes as given. Holding fixed rival prices, firm i's relaxed problem can thus be written as

$$\max_{(p_z^i)_{z \in \mathcal{Z}^i}} \sum_{z \in \mathcal{Z}^i} q_z^i (p_z^i - \tilde{c}_z^i),$$

where $\tilde{c}_z^i$ is firm i's marginal cost of offering product pair z, i.e.,

$$\tilde{c}_z^i = \begin{cases} c_a & \text{if } a \in \mathcal{A}^i,\ b \notin \mathcal{B}^i, \\ c_b & \text{if } a \notin \mathcal{A}^i,\ b \in \mathcal{B}^i, \\ c_a + c_b & \text{if } a \in \mathcal{A}^i,\ b \in \mathcal{B}^i. \end{cases}$$

Note that firm i's relaxed problem is isomorphic to the multiproduct-firm pricing problem under standard multinomial logit demand. The solution to that simpler problem entails firm i charging the same absolute markup μ^i on all of its products (Nocke and Schutz, 2018): $\mu^i = (p_z^i - \tilde{c}_z^i)/\lambda > 1$ for all $z \in \mathcal{Z}^i$. Indeed, the first-order condition implicitly gives the value of μ^i by

$$\mu^i (1 - Q^i) = 1, \tag{1}$$

where

$$Q^i \equiv \sum_{z \in \mathcal{Z}^i} q_z^i$$

is the demand for product pairs containing at least one product supplied by firm i.

Importantly, the solution to the relaxed problem is also feasible in firm i's original pricing problem (where the price of a product offered by firm i cannot depend on which other product

it is consumed with). Indeed, if product pair $z \in \mathcal{Z}^i$ contains only one of firm i's products, by the common markup property the price firm i optimally charges for that product in the relaxed problem does not depend on which other product it is consumed with:

$$p_z^i = \tilde{c}_z^i + \lambda\mu^i = \begin{cases} p_a = c_a + \lambda\mu^i & \text{if } a \in \mathcal{A}^i,\ b \notin \mathcal{B}^i, \\ p_b = c_b + \lambda\mu^i & \text{if } a \notin \mathcal{A}^i,\ b \in \mathcal{B}^i, \\ p_a + p_b - \delta_{ab} = c_a + c_b + \lambda\mu^i & \text{if } a \in \mathcal{A}^i,\ b \in \mathcal{B}^i. \end{cases}$$

Hence, the solution to the relaxed problem is also the solution to the original (more constrained) problem.

This proves the following result:

Lemma 1 (Common markup property)**.** *The solution to firm i's optimal pricing problem entails a common markup: there exists a scalar $\mu^i > 1$ such that, for any $a \in \mathcal{A}^i$ and $b \in \mathcal{B}^i$,*

$$p_a - c_a = p_b - c_b = \delta_{ab} = \lambda\mu^i.$$

Lemma 1 establishes a stark property: holding fixed rival prices, each firm optimally charges the same absolute markup on all of its products, including product bundles offered by that firm. This means that whenever a consumer decides to purchase two products from the same firm, it is as if the firm sells the second product at cost. As will become apparent in what follows, this common markup property greatly reduces the dimensionality of the pricing game, making its analysis tractable.

Lemma 1 implies that the pricing game can be reduced to one in which each firm i chooses a single scalar μ^i—its markup. For a given markup vector $\boldsymbol{\mu} \equiv (\mu^i)_{i\in\mathcal{N}}$, firm i's profit can therefore be written as

$$\Pi^i(\boldsymbol{\mu}) = \lambda\mu^i Q^i(\boldsymbol{\mu}). \tag{2}$$

3.2 Type-network representation

Under common-markup pricing, the demand for firm i's products is given by

$$\begin{aligned} Q^i(\boldsymbol{\mu}) &= \frac{1}{H(\boldsymbol{\mu})} \sum_{z\in\mathcal{Z}^i} \mathrm{e}^{\frac{\tilde{s}_z^i - \tilde{c}_z^i}{\lambda} - \mu^i} \\ &= \frac{\mathrm{e}^{-\mu^i}}{H(\boldsymbol{\mu})} \left[T^i + \sum_{j\in\mathcal{N}\setminus\{i\}} T^{ij} \mathrm{e}^{-\mu^j} \right], \end{aligned} \tag{3}$$

where

$$\begin{aligned} T^i &\equiv \sum_{a\in\mathcal{A}^i}\sum_{b\in\mathcal{B}^i} \mathrm{e}^{\frac{s_{ab}-c_a-c_b}{\lambda}} + \sum_{a\in\mathcal{A}^i} \mathrm{e}^{\frac{s_{a0}-c_a}{\lambda}} + \sum_{b\in\mathcal{B}^i} \mathrm{e}^{\frac{s_{0b}-c_b}{\lambda}} \quad \text{and} \\ T^{ij} &\equiv \sum_{a\in\mathcal{A}^i}\sum_{b\in\mathcal{B}^j} \mathrm{e}^{\frac{s_{ab}-c_a-c_b}{\lambda}} + \sum_{a\in\mathcal{A}^j}\sum_{b\in\mathcal{B}^i} \mathrm{e}^{\frac{s_{ab}-c_a-c_b}{\lambda}} \equiv T^{ji}. \end{aligned} \tag{4}$$

In words, $T^{ij} = T^{ji}$ captures the social value from consuming a product of firm i in one market (either A or B) jointly with a product of firm j in the other market. Similarly, T^i represents the social value from consuming either two products of firm i or combining a product of that

firm in one market with the outside option in the other market. Finally, $T^0 \equiv \mathrm{e}^{s_{00}/\lambda}$ reflects the social value from consuming the outside option in both markets. In this notation, the CS-index can be written as

$$H(\boldsymbol{\mu}) = T^0 + \sum_{i\in\mathcal{N}} T^i \mathrm{e}^{-\mu^i} + \frac{1}{2}\sum_{i\in\mathcal{N}}\sum_{j\in\mathcal{N}\setminus\{i\}} T^{ij}\mathrm{e}^{-\mu^i-\mu^j}.$$

In the following, we call

$$\mathbf{T} \equiv \left\{T^0, (T^i)_{i\in\mathcal{N}}, (T^{ij})_{i,j\in\mathcal{N};i\neq j}\right\}$$

the *type network*, as it summarizes (in conjunction with the vector of endogenous markups) all the necessary information about firms' product portfolios and consumer preferences over these products. In the language of network economics, T^{ij} represents the *weight* or *strength* of the link between nodes i and j (see, e.g., Elliott and Golub, 2019; Galeotti, Golub, and Goyal, 2020), while T^i can be interpreted as the *standalone weight* of node i. We put a minimal restriction on the type network: either $T^0 > 0$, or for every $i \in \mathcal{N}$ there is $j \in \mathcal{N}\setminus\{i\}$ with $T^j > 0$ or there are distinct $j, k \in \mathcal{N}\setminus\{i\}$ with $T^{jk} > 0$. That is, we require that no firm has unlimited monopoly power, in the sense that consumers always have a viable alternative to completely bypass any given firm.

3.3 Equilibrium existence and uniqueness

Despite these substantial reductions in dimensionality, establishing equilibrium existence is not straightforward. Markups have no upper bound, so Brouwer's fixed-point theorem cannot be applied directly. Standard results for aggregative games do not directly apply either. To see this, consider the expression for Q^i in equation (3): firm i's demand would depend on other firms' markups exclusively through H only if $T^{ij} = 0$ for all $j \neq i$, i.e., if no consumer would want to consume a product from firm i jointly with a product supplied by another firm.

We address this challenge by showing that the reduced pricing game, in which each firm sets a single scalar (its markup), admits an ordinal potential. Equilibrium existence then reduces to showing that the potential attains a maximum. Define

$$\Phi(\boldsymbol{\mu}) \equiv \sum_{i\in\mathcal{N}} \left(\log\mu^i - \mu^i\right) - \log H(\boldsymbol{\mu}). \tag{5}$$

To see that Φ is an ordinal potential, note from equations (2) and (3) that the logarithm of firm i's profit is

$$\log\Pi^i(\boldsymbol{\mu}) = \log\mu^i - \mu^i - \log H(\boldsymbol{\mu}) + \log\lambda + \log\left[T^i + \sum_{j\in\mathcal{N}\setminus\{i\}} T^{ij}\mathrm{e}^{-\mu^j}\right].$$

Hence, for every positive markup vector

$$\frac{\partial\Phi(\boldsymbol{\mu})}{\partial\mu^i} = \frac{1}{\mu^i} - 1 - \frac{\partial\log H(\boldsymbol{\mu})}{\partial\mu^i} = \frac{\partial\log\Pi^i(\boldsymbol{\mu})}{\partial\mu^i}. \tag{6}$$

Thus Φ is an exact potential of the game in which each firm maximizes $\log\Pi^i$ and therefore is an ordinal potential of the reduced pricing game.

The question is whether Φ attains a maximum and, if so, whether its maximizer is unique. The proof of the following result establishes that Φ does indeed have a unique maximizer, where the uniqueness follows from the strict concavity of the ordinal potential Φ.

Theorem 1 (Equilibrium existence and uniqueness)**.** *The reduced pricing game has a unique equilibrium $\boldsymbol{\mu}^*$. All equilibria in the original pricing game are outcome equivalent.*

Proof. See Appendix A.1.1. □

To see why the potential Φ in equation (5) is strictly concave, note that the potential has a firm-specific component $(\log \mu^i - \mu^i)$, which is strictly concave in μ^i, and the logarithm of the CS-index $(-\log H(\boldsymbol{\mu}))$. As the CS-index $H(\boldsymbol{\mu})$ is the weighted sum of exponentials whose exponents are linear in firms' markups, the convexity of the log-sum-exp function (Boyd and Vandenberghe, 2004, Sections 3.1.5 and 3.2.2) implies that $\log H(\boldsymbol{\mu})$ is convex, so $-\log H(\boldsymbol{\mu})$ is concave. Together, they constitute a strictly concave potential.

Although every best response has the common-markup form, the original pricing game generally does not admit an ordinal potential. Appendix A.1.2 constructs a cycle of unilateral changes in firms' price vectors, each of which strictly raises the deviating firm's profit. The common-markup reduction therefore retains every best response while yielding a potential game.

4 Comparative Statics

In this section, we analyze the effects of changes in the type network $\mathbf{T}$ on the equilibrium outcome. In particular, we consider changes in costs or product qualities, firm entry, and mergers.

Recall from equation (1) that firm i's optimal markup satisfies $\mu^i = 1/(1 - Q^i)$, where Q^i represents firm-level demand. Hence, any rival action that shifts Q^i down lowers firm i's optimal markup, whereas any rival action that shifts Q^i up has the opposite effect.

Consider an increase in the markup charged by firm $j \neq i$, μ^j. We have:

$$\begin{aligned}\frac{\partial Q^i}{\partial \mu^j} &= -\frac{\partial H/\partial \mu^j}{H} Q^i - \frac{T^{ij} \mathrm{e}^{-\mu^i - \mu^j}}{H} \\ &= Q^i Q^j - Q^{ij} = Q^i Q^j (1 - \ell^{ij}),\end{aligned}$$

where

$$Q^{ij} \equiv \frac{T^{ij} \mathrm{e}^{-\mu^i - \mu^j}}{H}$$

is the demand for product pairs involving one product from firm i and one from firm j, and

$$\ell^{ij} \equiv \frac{Q^{ij}}{Q^i Q^j}$$

is the *lift* of firms i and j. The lift—a concept from the literature on market basket analysis (Hastie, Tibshirani, and Friedman, 2009)—measures the extent to which the products of two firms are consumed together (see also Nocke and Schutz, 2025). Hence, in the reduced pricing game, if $\ell^{ij} < 1$, the offerings of firms i and j are substitutes and their markups are strategic complements. By contrast, if $\ell^{ij} > 1$, the offerings are complements and the markups are

strategic substitutes. Intuitively, the products of firms i and j are more likely to be consumed together, and so the lift of the two firms is large, if T^{ij} is large. This intuition is made formal in the following result:

Proposition 1 (Strategic complementarity and substitutability)**.** *If the lift of firms i and j is larger than one, $\ell^{ij} > 1$, the offerings of firms i and j are (local) complements—and their markups (local) strategic substitutes. By contrast, if the lift of firms i and j is smaller than one, $\ell^{ij} < 1$, the offerings of firms i and j are (local) substitutes—and their markups (local) strategic complements. If every available product pair involves firm i or firm j, then $\ell^{ij} \leq 1$; if firm i sells only together with firm j, or vice versa, then $\ell^{ij} \geq 1$. Otherwise, $\ell^{ij} < 1$ if T^{ij} is sufficiently small, and $\ell^{ij} > 1$ if T^{ij} is sufficiently large.*

Proof. See Appendix A.2.2. □

We note that whether the lift exceeds one does not depend on the markups of firms i and j; it depends only on the type network and the other firms' markups. When $\mathcal{N} = \{i, j\}$, it therefore depends only on the type network: if $T^0 > 0$, there is a unique cutoff $\widehat{T}^{ij} \equiv T^i T^j / T^0$ such that $\ell^{ij} > 1$ if $T^{ij} > \widehat{T}^{ij}$ and $\ell^{ij} < 1$ if $T^{ij} < \widehat{T}^{ij}$; if $T^0 = 0$, the lift is always below one. See Appendix A.2.2.

4.1 Cost and quality changes

We first consider changes in individual elements of the type network, holding all other elements fixed.

The first result considers an increase in firm i's standalone type T^i, which may be the result of an increase in the gross surplus s_{ab}, s_{a0} or s_{0b} for $a \in \mathcal{A}^i$ and $b \in \mathcal{B}^i$.

Proposition 2 (Standalone types and consumer surplus)**.** *An increase in firm i's standalone type T^i strictly raises equilibrium consumer surplus.*

Proof. See Appendix A.2.3. □

Proposition 2 shows that an increase in a firm's standalone type unambiguously increases consumer surplus. As an improvement in a firm's standalone type directly benefits consumers and shifts demand toward one firm, that firm tends to raise its markup, while all its rivals tend to lower theirs. Our demand structure bounds the improving firm's markup increase so that it cannot overturn the direct benefit and the induced markup reductions by its rivals.

The conclusion need not hold for the cross-firm type T^{ij}, defined in equation (4), which reflects the social values of product pairs containing one product from firm i and one from firm j. An increase in T^{ij} can induce both firms to raise their markups, which may overturn the original benefit of an improvement of the cross-firm type. The following example illustrates this case (see Appendix A.2.4 for details).

Example 1. *Suppose that $\mathcal{N} = \{1, 2\}$ and $(T^0, T^1, T^2, T^{12}) = (0, 1, 1, x)$. The equilibrium is symmetric with $\mu^{1*} = \mu^{2*} = \mu^*$ and satisfies*

$$\mu^* = 2 + x\mathrm{e}^{-\mu^*}, \quad H^* = 2\mathrm{e}^{-\mu^*} + x\mathrm{e}^{-2\mu^*} = \mu^*\mathrm{e}^{-\mu^*}. \tag{7}$$

Then, μ^* is strictly increasing in x, and so H^* is strictly decreasing in x.

We next consider changes in product costs. Consider a reduction, by the same absolute amount, in the marginal cost of every product firm i supplies, holding everything else fixed. This means that the marginal cost of an ab-pair supplied by firm i falls by twice as much as the marginal cost of a single product.

Proposition 3 (A uniform cost reduction within a firm)**.** *Suppose that the marginal costs of all products supplied by firm i decrease by the same absolute amount. Then, equilibrium consumer surplus strictly increases.*

Proof. See Appendix A.2.5. □

The uniform cost reduction raises not only firm i's standalone type T^i (which, on its own, would raise consumer surplus) but also all cross-firm types T^{ij}, $j \neq i$ (which, as we have just seen, may reduce consumer surplus). Despite this tension, Proposition 3 establishes that the uniform cost reduction raises equilibrium consumer surplus.

When a cost reduction applies only to a subset of a firm's products, it need not benefit consumers, as the following example illustrates (see Appendix A.2.6 for details).

Example 2. *Suppose that $\mathcal{N} = \{1, 2\}$, $\mathcal{A}^1 = \{a, a'\}$, $\mathcal{B}^2 = \{b\}$, and $\mathcal{B}^1 = \mathcal{A}^2 = \emptyset$. Suppose also that $s_{a0} = s_{a'b} = s_{00} = -\infty$, that is, the only product pairs generating surplus are ab, $a'0$, and $0b$. Suppose also that the gross surpluses and marginal costs are such that the resulting type network is $(T^0, T^1, T^2, T^{12}) = (0, 1, 1, x)$, as in Example 1. Then, a decrease in c_a only raises x and thus reduces consumer surplus, whereas a decrease in $c_{a'}$ only raises T^1 and thus increases consumer surplus.*

4.2 Entry

We now turn to analyzing the effects of firm entry—i.e., the introduction of a new node and new links in the type network—on the equilibrium outcome. Consider entrant E, and let T^E denote its standalone type and T^{iE}, $i \in \mathcal{N}$, the corresponding cross-firm types. The post-entry markup of firm $i \in \mathcal{N} \cup \{E\}$ is denoted $\overline{\mu}^{i*}$.

The following result shows that consumers on average benefit from entry:

Proposition 4 (Entry and consumer surplus)**.** *The entry of firm E strictly raises equilibrium consumer surplus.*

Proof. See Appendix A.2.7. □

Proposition 4 follows from Proposition 3 through a limiting argument. Starting from the post-entry economy, raise all of the entrant's marginal costs by the same amount. As this amount tends to infinity, the equilibrium converges to the pre-entry equilibrium; see Appendix A.2.7.

Note that entry of firm E may increase the markups of some rival firms that predominantly offer complements to E's products. While consumer surplus increases, in the discrete-choice microfoundations, some consumers may therefore be hurt by entry. More formally, entry makes a positive measure of consumers worse off if and only if at least one product pair ab, with $a \notin \mathcal{A}^E$ and $b \notin \mathcal{B}^E$, becomes more expensive. This is illustrated in the following example (see Appendix A.2.8 for details).

Example 3. *Suppose $\mathcal{N} = \{1\}$. Then the entry of firm E raises firm 1's markup if and only if the lift of firms 1 and E strictly exceeds 1, $\ell^{1E} > 1$. This occurs if and only if $T^{1E} > \widehat{T}^{1E} \equiv T^1 T^E / T^0$. In that case, entry hurts all those consumers who, in the post-entry equilibrium, only purchase firm 1's products (possibly combined with the outside option).*

4.3 Mergers

We finally consider mergers between two firms. We focus on no-synergy mergers, which change product ownership but leave product values and marginal costs unchanged; synergies can be incorporated by invoking the analysis of Section 4.1. Formally, a no-synergy merger among firms in a set $\mathcal{M} \subset \mathcal{N}$, with $|\mathcal{M}| = 2$, places all products previously supplied by the two merger partners under the ownership of a single merged firm M. To ease notation, we label the merger partners as firms f and g, so that $\mathcal{M} = \{f, g\}$. The merger transforms the set of firms $\mathcal{N}$ to $\overline{\mathcal{N}} \equiv (\mathcal{N} \setminus \{f, g\}) \cup \{M\}$ and type network $\mathbf{T}$ to a new type network $\overline{\mathbf{T}}$ such that

$$\overline{T}^M \equiv T^f + T^g + T^{fg},$$

and, for each nonmerging firm $k \in \mathcal{O} \equiv \mathcal{N} \setminus \{f, g\}$,

$$\overline{T}^{Mk} \equiv T^{fk} + T^{gk}.$$

The other types remain unchanged: $\overline{T}^0 = T^0$, $\overline{T}^k = T^k$ for every $k \in \mathcal{O}$, and $\overline{T}^{kl} = T^{kl}$ for all $k, l \in \mathcal{O}$, $k \neq l$. In the language of the network literature, a no-synergy merger is a node contraction: the merging firms' nodes are replaced by a single node, and their connections to each remaining node are added together (e.g., Pemmaraju and Skiena, 2003; Balbuzanov and Kotowski, 2024). We assume that the minimal requirement on the type network—no firm has unlimited monopoly power—still holds after the merger. Let $\overline{\boldsymbol{\mu}}^* = (\overline{\mu}^{i*})_{i \in \overline{\mathcal{N}}}$ and $\overline{H}^*$ denote the resulting post-merger markups and the CS-index.

To understand how mergers affect prices, we decompose mergers into the following three conceptual steps:

Step 1 (Profit internalization). Holding all outsiders' markups at their pre-merger levels, let the merging firms choose their markups to maximize joint profit while leaving the pre-merger pricing structure unchanged: that is, the products previously supplied by each firm carry one common markup, and merging firms in this step cannot offer discounts on the purchase of product pairs containing products of both merging firms. Let $(\breve{\mu}^f, \breve{\mu}^g)$ denote a profile of markups that maximizes joint profit in this step, and $\breve{H}$ the associated CS-index.[3]

Step 2 (Introduction of new bundle discounts). Still holding the outsiders' markups fixed, let the merged firm offer bundle discounts on the purchase of product pairs containing products of both merging firms and adjust the markups accordingly. Let $\widetilde{\mu}^M$ and $\widetilde{H}$ denote the resulting markup and CS-index.

Step 3 (Full equilibrium adjustment). Finally allow every firm to reoptimize, yielding the post-merger equilibrium $\overline{\boldsymbol{\mu}}^*$ and $\overline{H}^*$.

[3]The maximizer need not be unique because joint profit may not be concave in the merging firms' markups. Nonetheless, the results below hold for every maximizer.

For each merging firm $i \in \{f, g\}$, this decomposition gives

$$\overline{\mu}^{M*} - \mu^{i*} = \underbrace{\breve{\mu}^i - \mu^{i*}}_{\text{Step 1}} + \underbrace{\widetilde{\mu}^M - \breve{\mu}^i}_{\text{Step 2}} + \underbrace{\overline{\mu}^{M*} - \widetilde{\mu}^M}_{\text{Step 3}}.$$

Steps 1 and 2 formalize how the merging firms' pricing incentives change after mergers—Step 1 isolates joint profit maximization under the pre-merger pricing structure, whereas Step 2 introduces the discounts available only after the merger. Step 3 allows all firms to reoptimize and yields the full post-merger equilibrium. In what follows, we first analyze the price changes at each step and then characterize their combined effect on consumer surplus.

Step 1: Internalizing cross-firm substitutability/complementarity. At the pre-merger equilibrium, a marginal increase in the markup of firm $i \in \{f, g\}$ changes the merging firms' combined profit according to

$$\left.\frac{\partial(\Pi^f + \Pi^g)}{\partial \mu^i}\right|_{\boldsymbol{\mu}=\boldsymbol{\mu}^*} = \lambda Q^{f*} Q^{g*} \mu^{j*}(1 - \ell^{fg*}), \quad \{i, j\} = \{f, g\}. \tag{8}$$

Equation (8) captures a merging firm's incentive to change its markup while holding the other's fixed. This incentive is positive when the lift of the merging firms is below one and negative when it is above one. If the lift is below one, both markups rise in Step 1. If it is above one, at least one markup falls: the firm with lower pre-merger demand lowers its markup, whereas the firm with higher pre-merger demand may raise its markup. A lower markup on the smaller firm's products expands demand for the larger firm's products, which initially carry the higher markup. The resulting profit gain can outweigh the larger firm's direct incentive to lower its own markup, widening the markup gap (see Lemma A.2 in Appendix A.3.3).

Lemma 2 (Step 1 and consumer surplus)**.** *Every Step-1 optimum satisfies*

$$\operatorname{sign}\big(\breve{H} - H^*\big) = \operatorname{sign}\big(\ell^{fg*} - 1\big). \tag{9}$$

Thus consumer surplus rises in Step 1 if and only if the pre-merger lift exceeds one.

Proof. See Appendix A.3.2. □

Step 2: Allowing the merged firm to offer new bundle discounts. At Step 2, the merged firm is allowed to offer bundle discounts on the purchase of product pairs containing products of both merging firms. This leads the merged firm to set a single scalar $\widetilde{\mu}^M$. The new bundle discounts reduce the total prices of these product pairs. As this expands firm M's demand and reduces the margin that firm M charges from the combined purchases, it induces the merged firm to raise the product-level prices. The following lemma formally establishes this observation.

Lemma 3 (The merged firm's markup at Step 2)**.** *The Step-2 markup satisfies*

$$\widetilde{\mu}^M > \max\{\mu^{f*}, \mu^{g*}\}.$$

Proof. See Appendix A.3.4. □

Lemma 3 shows a striking result: although the merging firms may lower their markups at Step 1, the introduction of new bundle discounts at Step 2 overturns any such reduction and induces the merged firm to raise its markup above every pre-merger markup of the merger partners. Before the new discounts become available, the merging firms can respond to complementarity between their products only through their individual markups: lowering either markup reduces the price of purchases involving both firms, but also lowers the markup on purchases involving that firm alone. Step 2 separates these two pricing decisions. The merged firm can raise its product markups while using the new discounts to prevent purchases involving both firms from bearing two such markups.

At the end of Step 2, the prices of product pairs involving the merged firm's products may rise or fall. A pair involving products from only one merging firm before the merger becomes more expensive because the merged firm's markup exceeds that firm's pre-merger markup. By contrast, a pair involving products previously supplied by both merging firms changes in price by $\lambda(\widetilde{\mu}^M - \mu^{f*} - \mu^{g*})$, which may be positive or negative. The merged firm therefore charges a higher markup on every product, while the new discounts eliminate double marginalization on pairs involving both merging firms.

Hence, the consumer-surplus effect at the end of Step 2 depends on which dominates: the higher markup or the elimination of double marginalization. The pre-merger demands and lift determine whether consumer surplus rises or falls at the end of Step 2. Lift enters this tradeoff in two ways. First, when lift is above one, joint profit maximization in Step 1 creates downward pressure on markups. Second, holding pre-merger demands fixed, a higher lift means that a larger share of purchases combines the firms' products. These are precisely the purchases on which the merger eliminates double marginalization. Hence, a high lift tends to raise consumer surplus at the end of Step 2. The following lemma establishes the exact cutoff.

Lemma 4 (The lift cutoff for Steps 1–2)**.** *Fix the pre-merger demands of the merging firms,* (Q^{f*}, Q^{g*})*. There exists a unique cutoff* $\widehat{\ell}(Q^{f*}, Q^{g*}) \in \big(1,\, 1/\max\{Q^{f*}, Q^{g*}\}\big)$ *such that consumer surplus rises through Steps 1–2 if* $\ell^{fg*} > \widehat{\ell}(Q^{f*}, Q^{g*})$ *and falls if* $\ell^{fg*} < \widehat{\ell}(Q^{f*}, Q^{g*})$*.*

Proof. See Appendix A.3.5. □

Importantly, Lemma 4 also shows that the lift cutoff exceeds one, so a lift above one—the complementarity between merging firms' offerings—is necessary but insufficient for consumer surplus to rise through Steps 1–2. By contrast, the cutoff for Step 1 alone is exactly one (Lemma 2). Hence, if $1 < \ell^{fg*} < \widehat{\ell}(Q^{f*}, Q^{g*})$, then $\widetilde{H} < H^* < \breve{H}$: consumer surplus rises in Step 1, but the bundle discounts introduced in Step 2 more than offset this gain.

Step 3: Letting outsiders and the merged firm reoptimize. Finally, consider the full-equilibrium effects. At the end of Step 2, the merged firm's markup exceeds both merging firms' pre-merger markups. The outsiders' responses may amplify or dampen this increase. What generally holds is that, even with such equilibrium adjustment, the post-merger markup of the merged firm exceeds at least one merging firm's pre-merger markup.

Proposition 5 (The merged firm's markup and individual consumer harm)**.** *Suppose* $T^f, T^g > 0$*. Then, the post-merger markup of the merged firm satisfies* $\overline{\mu}^{M*} > \min\{\mu^{f*}, \mu^{g*}\}$*. Consequently, the merger makes a positive measure of consumers strictly worse off.*

Proof. See Appendix A.3.6. □

Proposition 5 implies that at least some consumers are worse off if both merging firms have positive standalone types.[4] On average, consumers may still benefit from the merger. The next proposition shows that outsiders' responses preserve, and strengthen, any consumer-surplus gain at the end of Step 2.

Proposition 6 (A consumer-surplus gain through Steps 1–2)**.** *If $\widetilde{H} > H^*$, then $\overline{H}^* \geq \widetilde{H} > H^*$, where the first inequality is strict if $\mathcal{O} \neq \emptyset$.*

Proof. See Appendix A.3.7. □

Proposition 6 shows that Step 2 provides a useful screen for evaluating merger effects—whenever a merger raises consumer surplus at the end of Step 2, the equilibrium response of outsiders reinforces it but never overturns it. To see why this holds, recall that at the end of Step 2, the merged firm raises the markup for all products, which reduces the value of the product pairs involving one product from a merging firm and one from outsiders. If the CS-index H at the end of Step 2 is higher than before the merger, these changes unambiguously reduce outsiders' demands. Hence, outsiders have an incentive to lower their markups. Under the two-market logit structure, this initial response eventually leads to a further increase in the CS-index, thereby benefiting consumers in the aggregate.

By Lemma 4 and Proposition 6, the merger raises consumer surplus in the full post-merger equilibrium whenever its pre-merger lift exceeds the cutoff defined in Lemma 4. If outsiders are present, consumer surplus rises further in Step 3. When the lift lies below the cutoff, consumer surplus falls through Steps 1–2. With no outsiders, the Step 2 outcome is also the full post-merger equilibrium outcome; with outsiders, Step 3 may reverse this loss.

As in Steps 1–2, the merger's full-equilibrium effects can be characterized using lift, once the merging firms' pre-merger demands and the outsiders' purchase patterns are held fixed. For each firm i, let q^{i*} denote the pre-merger probability of choosing a product pair supplied only by firm i, and let $\mathbf{q}^{\mathcal{O}*} \equiv (q^{k*}, Q^{ik*})_{k\in\mathcal{O}, i\neq k}$ denote the vector of pre-merger purchase probabilities involving outsiders. Holding (Q^{f*}, Q^{g*}) and $\mathbf{q}^{\mathcal{O}*}$ fixed, pre-merger purchase patterns are pinned down by the probability Q^{fg*} of buying from both merging firms, and hence by the lift ℓ^{fg*}. The following proposition shows that whether the merger raises consumer surplus in the full equilibrium depends on whether this lift exceeds a cutoff.

Proposition 7 (A full-equilibrium lift cutoff)**.** *Fix the pre-merger demands of the merging firms, (Q^{f*}, Q^{g*}), and the purchase probabilities involving outsiders, $\mathbf{q}^{\mathcal{O}*}$. There is a cutoff $\ell^F = \ell^F(Q^{f*}, Q^{g*}, \mathbf{q}^{\mathcal{O}*}) \leq \widehat{\ell}(Q^{f*}, Q^{g*})$ such that the merger raises consumer surplus if $\ell^{fg*} > \ell^F$ and lowers it if $\ell^{fg*} < \ell^F$.*

Proof. See Appendix A.3.8. □

Intuitively, holding demands fixed, a higher lift means that more consumers buy from both merging firms despite paying two markups, which reveals that they value these combinations highly. Since the merger removes the double markup precisely on these purchases, a higher lift makes the merger better for consumers.

The bound $\ell^F \leq \widehat{\ell}(Q^{f*}, Q^{g*})$ follows from Lemma 4 and Proposition 6: outsiders' responses can only lower the lift a merger needs to raise consumer surplus. The following example shows

[4]Without this condition, the conclusion may fail. For instance, if $\mathcal{N} = \{f, g\}$ and $T^f = T^g = 0$, the merging firms sell complementary components of a single bundle, and the merger lowers the bundle's price, so no consumer is worse off.

that the full-equilibrium cutoff can lie strictly below $\widehat{\ell}$, and even below one (see Appendix A.3.9 for details).

Example 4. *Suppose* $\mathcal{N} = \{1, 2, 3\}$ *and* $\{f, g\} = \{1, 2\}$*, with the type network*

$$\mathbf{T} = (T^0, T^1, T^2, T^3, T^{12}, T^{13}, T^{23}) = (1, 1, 1, x, 1, 0, x), \quad x > 0. \tag{10}$$

Then $\widetilde{H} < H^*$ *for every* $x > 0$*, while* $\overline{H}^* > H^*$ *for all sufficiently large* x*. Thus, although consumer surplus falls through Steps 1–2, the merger raises it when* x *is sufficiently large.*

The type network in Example 4 can be generated by two independent markets with additive product values. For instance, consider a case where firms 1 and 3 are single-product firms in market A, while firm 2 is a single-product firm in market B. In this case, no bundle discounts are available before the merger. Joint profit maximization leaves firms 1 and 2's markups unchanged in Step 1. Furthermore, independence implies $\ell^{12*} = 1$, which by Lemma 4 implies that consumer surplus falls through Steps 1–2. Nevertheless, Example 4 shows that firm 3's response more than offsets the loss. In this construction, x measures the marginal-cost-adjusted attractiveness of firm 3's product relative to firm 1's product in market A: $\lambda \log x$ is their quality–cost difference. Independence across markets makes the same factor x enter both T^3 and T^{23}. Thus, a large x makes firm 3 a strong outsider in market A.

Finally, we examine the role of the size of merging firms, measured by the pre-merger probability that consumers purchase at least one of their products. Define

$$Q^{M*} \equiv Q^{f*} + Q^{g*} - Q^{fg*}.$$

The following proposition shows that when merging firms are sufficiently small, the merger raises consumer surplus.

Proposition 8 (Consumer surplus when the merging firms are small)**.** *If* Q^{M*} *is sufficiently small while* Q^{fg*}/Q^{M*} *is bounded away from zero, then the merger raises consumer surplus.*

Proof. See Appendix A.3.10. □

A sufficient condition in terms of primitives is obtained by fixing all types other than T^0 and assuming that $T^{fg} > 0$. Proposition 8 then applies for all sufficiently large T^0. As T^0 becomes large, demand for every firm becomes small, and all pre- and post-merger markups converge to one—the monopolistically competitive level. Then elimination of double marginalization is the only first-order effect of the merger, so consumer surplus rises.

At the opposite limit, suppose that consumers have few attractive alternatives that avoid both merging firms, making the merging firms collectively large. Recall that q^{i*} denotes the probability of choosing a product pair supplied only by firm i.

Proposition 9 (Consumer surplus when the merging firms are large)**.** *If* Q^{M*} *is sufficiently close to one while* q^{f*} *and* q^{g*} *are bounded away from zero, then the merger lowers consumer surplus.*

Proof. See Appendix A.3.11. □

A sufficient condition in terms of primitives is obtained by fixing all types involving firm f or firm g. Proposition 9 then applies when T^0, all standalone types T^k of outsiders, and all cross-firm types T^{kl} between outsiders are jointly sufficiently small but not all zero. Before the merger, each merging firm's demand is bounded away from one, because purchases from only the other merging firm have probability bounded away from zero; this keeps both markups bounded. After the merger, the probability that the chosen pair contains a product of the merged firm approaches one, and its first-order condition drives its markup to infinity.

The preceding results also speak to the long-standing debate over conglomerate mergers exemplified by GE/Honeywell. The European Commission viewed the merged firm's ability to offer discounted packages of GE aircraft engines and Honeywell avionics as a potential source of competitive harm, whereas the U.S. Department of Justice emphasized the short-run gains from internalizing pricing externalities between complementary products. Our analysis places these two arguments within a common framework and characterizes when each force dominates. Complementarity in the sense of a lift above one is necessary but not sufficient for profit internalization and the newly available bundle discounts to raise consumer surplus: the lift must exceed a cutoff determined by the merging firms' pre-merger demands. Moreover, whenever these immediate effects raise consumer surplus, the subsequent equilibrium responses of outsiders reinforce the gain. Above the cutoff, any theory of harm to aggregate consumer surplus must therefore rely on an additional mechanism beyond static price adjustment, such as exit, reduced innovation, or foreclosure. Below the cutoff, by contrast, the merger can harm consumers even without such mechanisms.

5 Mixed Bundling versus Separate Pricing

We now examine the effects of mixed bundling as compared to separate pricing, under which firms choose product prices while all bundle discounts remain zero. Separate pricing may appear analytically simpler because it removes bundle discounts. In fact, it generally loses the sources of tractability available under mixed bundling. The common-markup reduction in Section 3 is no longer available because bundle discounts are needed to implement the pair-specific optimum. Nor does the transformed-potential approach of Nocke and Schutz (2025) apply in general.[5] Indeed, the improvement cycle in Appendix A.1.2 keeps all bundle discounts equal to zero and therefore establishes directly that the separate-pricing game does not admit an ordinal potential in general.

Given this difficulty, we focus on a simpler setting in which the two markets are independent in the sense that product values are additive across markets:

$$s_{ab} = s_a + s_b \quad \text{for every } a \in \mathcal{A},\ b \in \mathcal{B}. \tag{11}$$

Under separate pricing, additive values make each firm's profit separable across the two markets. The pricing game therefore decomposes into two independent single-market logit

[5]Their market-basket demand structure requires firm-level separability aligned with ownership. In particular, the relative demands for two products supplied by the same firm cannot depend on prices set by other firms. Under separate pricing in our model, for $a, a' \in \mathcal{A}^i$,

$$\frac{q_a}{q_{a'}} = \mathrm{e}^{-(p_a - p_{a'})/\lambda} \frac{\sum_{b\in\mathcal{B}} \mathrm{e}^{(s_{ab}-p_b)/\lambda}}{\sum_{b\in\mathcal{B}} \mathrm{e}^{(s_{a'b}-p_b)/\lambda}},$$

which generally depends on prices set by other firms.

pricing games: equilibrium prices in either market are unaffected by which products firms supply or the prices they charge in the other market. The unique separate-pricing equilibrium is obtained by combining the unique equilibria of the two single-market games, and consumers' choices in the two markets are independent (see Appendix A.4.1 for details).

By contrast, when a firm supplies products in both markets, mixed bundling can link the markets through its discounts on same-firm pairs. For each firm $i \in \mathcal{N}$, define its market-level types by

$$T_A^i \equiv \sum_{a \in \mathcal{A}^i} \mathrm{e}^{(s_a - c_a)/\lambda}, \quad T_B^i \equiv \sum_{b \in \mathcal{B}^i} \mathrm{e}^{(s_b - c_b)/\lambda}. \tag{12}$$

Let T_A^0 and T_B^0 denote the corresponding market-level types of the outside options. We retain the market subscript on market-level types throughout this section because T^i denotes firm i's standalone type. The market-level types map into the type network as

$$\begin{aligned} T^0 &= T_A^0 T_B^0, \\ T^i &= T_A^i T_B^i + T_A^i T_B^0 + T_A^0 T_B^i, \quad i \in \mathcal{N}, \\ T^{ij} &= T_A^i T_B^j + T_A^j T_B^i, \qquad\qquad i, j \in \mathcal{N},\ i \neq j. \end{aligned} \tag{13}$$

Hence, the market-level types determine demand under both pricing regimes. Under separate pricing, they enter the two market-specific demand systems separately. Under mixed bundling, they determine the standalone and cross-firm types in the reduced markup game according to equation (13).

Independent markets also let us obtain a useful decomposition of welfare using the language of information theory. For probability distributions f and $\widehat{f}$ on the same finite set, let

$$K(f \| \widehat{f}) \equiv \sum_{z : f_z > 0} f_z \log\left(\frac{f_z}{\widehat{f}_z}\right)$$

denote the Kullback–Leibler divergence, or relative entropy, of f from $\widehat{f}$. Let $c_{ab} \equiv c_a + c_b$ for $ab \in \mathcal{A} \times \mathcal{B}$ and $\mathbf{c} = (c_{ab})_{ab \in \mathcal{A} \times \mathcal{B}}$. For any vector of total prices $\mathbf{p}$, let $\mathbf{q}(\mathbf{p}) \equiv (q_{ab}(\mathbf{p}))_{ab \in \mathcal{A} \times \mathcal{B}}$ and let $\mathbf{q}_A(\mathbf{p}) \equiv (q_a(\mathbf{p}))_{a \in \mathcal{A}}$ and $\mathbf{q}_B(\mathbf{p}) \equiv (q_b(\mathbf{p}))_{b \in \mathcal{B}}$ denote its marginal choice distributions. Let $\mathbf{q} \equiv \mathbf{q}(\mathbf{p})$, $\mathbf{q}_A \equiv \mathbf{q}_A(\mathbf{p})$, $\mathbf{q}_B \equiv \mathbf{q}_B(\mathbf{p})$, $\widehat{\mathbf{q}}_A \equiv \mathbf{q}_A(\mathbf{c})$, and $\widehat{\mathbf{q}}_B \equiv \mathbf{q}_B(\mathbf{c})$. Additive values imply that the choice distribution under marginal-cost pricing is $\widehat{\mathbf{q}}_A \otimes \widehat{\mathbf{q}}_B$. Let $\mathrm{AS}(\mathbf{p})$ denote aggregate surplus—consumer surplus plus industry profit—when total prices are $\mathbf{p}$. The entropy representation of logit demand gives[6]

$$\begin{aligned} \mathrm{DWL}(\mathbf{p}) &\equiv \mathrm{AS}(\mathbf{c}) - \mathrm{AS}(\mathbf{p}) \\ &= \lambda K\big(\mathbf{q} \| \widehat{\mathbf{q}}_A \otimes \widehat{\mathbf{q}}_B\big). \end{aligned}$$

Let $\mathbf{q}_A \otimes \mathbf{q}_B \equiv (q_a q_b)_{ab \in \mathcal{A} \times \mathcal{B}}$ denote the product distribution that would arise if choices in the two markets were independent while their marginal distributions remained unchanged. Applying the chain rule for Kullback–Leibler divergence yields

$$K\big(\mathbf{q} \| \widehat{\mathbf{q}}_A \otimes \widehat{\mathbf{q}}_B\big) = K(\mathbf{q}_A \| \widehat{\mathbf{q}}_A) + K(\mathbf{q}_B \| \widehat{\mathbf{q}}_B) + K\big(\mathbf{q} \| \mathbf{q}_A \otimes \mathbf{q}_B\big). \tag{14}$$

[6]More generally, under multinomial logit demand and arbitrary product values, deadweight loss satisfies $\mathrm{DWL}(\mathbf{p}) = \lambda K(\mathbf{q}(\mathbf{p}) \| \mathbf{q}(\mathbf{c}))$. This identity is a straightforward consequence of the Shannon-entropy representation of logit demand (see, e.g., Anderson, de Palma, and Thisse, 1988). We are unaware of prior work that explicitly identifies market deadweight loss under multinomial logit demand with the Kullback–Leibler divergence of choice probabilities under a given price profile from the choice probabilities under marginal-cost pricing.

The first two terms measure within-market misallocation: whether the marginal choice probabilities in markets A and B coincide with those under marginal-cost pricing. The last term is the mutual information between the two market choices and measures cross-market dependence (see, e.g., Cover and Thomas, 2006, Chapter 2). It is nonnegative and equals zero exactly when the choices are independent.

Separate pricing makes the two choices independent and therefore sets the cross-market term to zero. By favoring same-firm pairs, mixed bundling can make choices across markets dependent; at the same time, the induced price changes can move the marginal choice distribution in either market toward or away from its marginal-cost level. The cases below differ in whether separate pricing already achieves the efficient within-market allocation and, if not, whether mixed bundling corrects enough of that distortion to offset the dependence it creates.

Symmetric covered market. Consider $N \geq 2$ firms in two covered markets, so $T_m^0 = 0$ for $m \in \{A, B\}$. Firms are symmetric within each market: $T_m^i = T_m > 0$ for every firm i and $m \in \{A, B\}$. The common market-level type T_m may differ across markets.

In symmetric covered markets, it is well known that mixed bundling lowers aggregate surplus by inducing excessive purchases from the same firm (see, e.g., Armstrong and Vickers, 2010; Zhou, 2021). The following result also characterizes the effects on firm profits and consumer surplus.

Proposition 10 (Mixed bundling in a symmetric covered market)**.** *In a symmetric covered market, mixed bundling strictly lowers aggregate surplus and each firm's profit. It strictly raises consumer surplus when $N = 2$ and strictly lowers consumer surplus for all sufficiently large N.*

Proof. See Appendix A.4.2. □

The symmetric covered market is the benchmark in which separate pricing already achieves the efficient marginal allocation in each market. Symmetry makes $\mathbf{q}_A = \widehat{\mathbf{q}}_A$ and $\mathbf{q}_B = \widehat{\mathbf{q}}_B$ under both pricing regimes, so there is no within-market distortion for mixed bundling to correct. Its allocative effect is instead to tilt purchases toward same-firm pairs. Equation (14) then shows immediately that mixed bundling lowers aggregate surplus: separate pricing leaves the two choices independent, whereas mixed bundling creates cross-market dependence and therefore a strictly positive mutual-information term. Hence $\mathrm{AS}^{B*} < \mathrm{AS}^{S*}$, where the superscripts B and S denote the mixed-bundling and separate-pricing equilibria.

The decline in each firm's profit instead results from reciprocal discounting. Each firm can profitably introduce a bundle discount at the separate-pricing equilibrium because doing so shifts demand toward its same-firm pairs. When every firm can do so, however, these business-stealing gains are more than offset by discounting, and each firm's equilibrium profit falls. The fall in consumer surplus for large N has a simple explanation in terms of prices. For large N, mixed bundling raises standalone prices relative to separate pricing, while the share of consumers who purchase both products from the same firm and receive the bundle discount converges to zero. Most consumers therefore pay the higher standalone prices without receiving the discount, so consumer surplus falls.[7]

[7]The $N = 2$ consumer-surplus result is established by Anderson and Leruth (1993). Proposition 10 generalizes their result on firm profit to an arbitrary number of firms and shows that the consumer-surplus result reverses when the number of firms is large.

Attractive outside option. The covered-market benchmark leaves no within-market distortion for mixed bundling to correct. With an outside option, separate pricing instead generates under-purchasing in each market because positive markups make the marginal purchase probabilities lower than their marginal-cost levels. Mixed bundling can therefore improve the within-market allocation even though it creates dependence across markets. Suppose there are $N \geq 2$ firms and $T_m^0 > 0$ for each $m \in \{A, B\}$. Define the aggregate type of the inside suppliers relative to the outside option and firm i's share of that aggregate type by

$$t_m^{\mathcal{N}} \equiv \frac{\sum_{i \in \mathcal{N}} T_m^i}{T_m^0}, \quad \sigma_m^i \equiv \frac{T_m^i}{\sum_{j \in \mathcal{N}} T_m^j}.$$

Thus, $T_m^i = t_m^{\mathcal{N}} \sigma_m^i T_m^0$ and $\sum_i \sigma_m^i = 1$. We hold the type shares $(\sigma_A^i)_{i \in \mathcal{N}}$ and $(\sigma_B^i)_{i \in \mathcal{N}}$ fixed and vary $t_A^{\mathcal{N}}$ and $t_B^{\mathcal{N}}$. The outside option in market m becomes increasingly attractive relative to the inside suppliers as $t_m^{\mathcal{N}}$ approaches zero.

Proposition 11 (Mixed bundling with an attractive outside option)**.** *Suppose that* $\sum_i \sigma_A^i \sigma_B^i > 0$*. For all sufficiently small* $t_A^{\mathcal{N}}, t_B^{\mathcal{N}} > 0$*, mixed bundling raises aggregate surplus and industry profit and lowers consumer surplus relative to separate pricing. It also raises firm* i*'s profit whenever* $\sigma_A^i \sigma_B^i > 0$*.*

Proof. See Appendix A.4.3. □

The condition $\sum_i \sigma_A^i \sigma_B^i > 0$ requires at least one firm to supply products in both markets. If it fails, no bundle discounts are available and the two pricing regimes coincide. The sum is the probability that two suppliers drawn independently according to the type shares in markets A and B have the same identity. The symmetric case in which $T_m^i = x T_m^0$ for every firm is obtained by setting $t_m^{\mathcal{N}} = Nx$ and $\sigma_m^i = 1/N$.

When $t_A^{\mathcal{N}}$ and $t_B^{\mathcal{N}}$ are small, purchasing in both markets is rare. Under separate pricing, a consumer who purchases from the same firm in both markets pays two product markups. Mixed bundling removes one of them and shifts probability away from purchasing in exactly one market and toward purchasing in both markets or in neither. The probability of purchasing in each individual market nevertheless rises. For sufficiently attractive outside options, these increases move both marginal choice distributions closer to their marginal-cost levels. The first two terms in equation (14) therefore fall by more than the mutual-information term rises, so total Kullback–Leibler divergence falls and aggregate surplus rises.

Every firm active in both markets gains from mixed bundling, and industry profit rises by more than aggregate surplus. Consumer surplus therefore falls. The probability of choosing the outside option in both markets also rises: the logit numerator for that choice is unchanged across the two pricing regimes, so its probability rises exactly when the CS-index falls.

Firm asymmetry. Consider a covered duopoly in which the firms differ in both markets. Suppose that $T_m^0 = 0$ and $T_m^1, T_m^2 > 0$ for $m \in \{A, B\}$, and let

$$\Delta_m \equiv \lambda \log\left(\frac{T_m^1}{T_m^2}\right) \geq 0, \quad m \in \{A, B\}.$$

Thus, Δ_m measures firm 1's marginal-cost-adjusted attractiveness advantage in market m. These advantages may arise from product values, marginal costs, or product variety.

Proposition 12 (Firm asymmetry in a covered duopoly)**.** *If Δ_A and Δ_B are both sufficiently large, mixed bundling raises firm 1's profit, industry profit, and aggregate surplus and lowers firm 2's profit and consumer surplus relative to separate pricing.*

Proof. See Appendix A.4.4. □

Firm asymmetry creates a different within-market distortion under separate pricing. When $\Delta_m > 0$, firm 1 has the larger market-level type in market m but also charges a higher markup. Firm 2 therefore receives more demand in that market than it would under marginal-cost pricing. Mixed bundling shifts demand away from cross-firm supplier pairs toward same-firm pairs. When both Δ_A and Δ_B are sufficiently large, most of the demand shifted away from the cross-firm pairs goes to the pair in which firm 1 supplies both products, rather than the pair in which firm 2 supplies both products. Firm 2's marginal share consequently moves closer to its marginal-cost level in both markets, reducing the within-market allocation distortions under separate pricing.

Mixed bundling also makes the two market choices more dependent. Conditional on firm 2 being chosen in one market, purchases involving firm 2 become increasingly concentrated on the pair in which it supplies both products. But firm 2 itself is chosen with vanishing probability as both advantages grow, so the mutual-information term converges to zero. Under separate pricing, each marginal divergence converges to one. Under mixed bundling, their sum converges to one, even when Δ_A and Δ_B grow at different rates. The reduction in the two marginal divergences therefore converges to one and dominates the increase in cross-market dependence.

The efficiency gain does not necessarily accrue to consumers. As both advantages become large, the aggregate-surplus gain converges to λ, whereas firm 1's profit gain grows without bound and firm 2's profit loss vanishes. The increase in industry profit therefore eventually exceeds the aggregate-surplus gain, so consumer surplus falls. Thus, the same asymmetry that allows mixed bundling to improve the allocation also allows the stronger firm to capture more than the resulting efficiency gain.

6 Conclusion

In this paper, we develop a framework of oligopoly with mixed bundling. The analysis of oligopoly with mixed bundling is technically difficult because firms face high-dimensional pricing problems and complex strategic interactions: firms choose separate prices for each product and discounts for each same-firm product pair, while consumers can mix and match products across firms, so a product can be a substitute for a rival's product in some purchase configurations and a complement in others. The literature has therefore made a number of simplifying assumptions about consumer preferences, the number of firms, costs, and product portfolios. By contrast, we allow for an arbitrary number of firms and arbitrary costs, product values, and product portfolios, while requiring consumer demand to follow a multinomial logit structure.

The logit demand structure allows us to obtain a strikingly sharp characterization of oligopoly with mixed bundling. On the pricing side, the common markup property implies that the pair-specific optimum in a relaxed pricing problem can be implemented through product prices and bundle discounts, reducing each firm's pricing problem to one dimension. The type-network representation provides a second reduction, summarizing all primitives

governing strategic interaction. Finally, the resulting reduced pricing game admits a strictly concave potential, which allows us to establish the existence and uniqueness of an equilibrium outcome. Taken together, these results make the otherwise complex problem of oligopolistic mixed bundling tractable.

This tractable equilibrium characterization enables us to derive comparative-statics results in substantially more general environments than those considered in existing analyses of mixed bundling. Changes in costs and product values, entry, and mergers can all be represented as changes in the associated type network and therefore analyzed in a unified manner. An increase in a standalone type raises consumer surplus, whereas an increase in a cross-firm type may raise or lower it. Thus, changes in costs and product values can have ambiguous effects depending on the changes they induce in standalone and cross-firm types. Nevertheless, an equal absolute reduction in the marginal costs of all products supplied by one firm raises consumer surplus, as does firm entry.

Mergers provide a particularly rich application because they change both ownership incentives and the set of available bundle discounts. We decompose a no-synergy merger into profit internalization, the introduction of new discounts, and the subsequent adjustment of outsiders. Holding outsiders' markups fixed, the merged firm's markup exceeds both merger partners' pre-merger markups, while purchases combining both merging firms' products bear one markup rather than two. Pre-merger demands and lift determine which force dominates: consumer surplus rises if and only if the lift exceeds a unique cutoff, and any such gain is reinforced by outsiders' equilibrium responses. In the full equilibrium, holding outsiders' purchase patterns fixed, the merger raises consumer surplus if and only if the lift exceeds a cutoff that is no higher than the first. Moreover, mergers of sufficiently small firms raise consumer surplus, whereas mergers of sufficiently large firms lower it under the conditions we characterize.

Our framework also provides a unified way to understand the welfare effects of mixed bundling. Under logit demand, deadweight loss is proportional to the Kullback–Leibler divergence between the equilibrium choice distribution and that under marginal-cost pricing. With additive values across markets, the Kullback–Leibler divergence separates within-market misallocation from cross-market dependence. Mixed bundling creates such dependence by favoring same-firm purchases, but it can also correct pre-existing within-market distortions. Accordingly, it lowers aggregate surplus in a symmetric covered market, while sufficiently attractive outside options or sufficiently large firm asymmetry in both markets can reverse this effect.

Our analysis imposes multinomial logit demand. The return is that, within this demand structure, we can accommodate an arbitrary number of firms and arbitrary product values, marginal costs, and product portfolios while retaining a low-dimensional equilibrium characterization. Moving beyond logit would generally require a different analytical approach, as the common markup property and the resulting low-dimensional equilibrium structure need not survive under more general demand. The payoff from this tractability is not limited to the theoretical results developed here: the ability to accommodate rich market structure also makes the framework well suited for structural estimation and counterfactual analysis (see Kawaguchi, Kuroda, Sato, and Ida, 2026 for an application to telecommunications markets).

References

Adams, W. J., and J. L. Yellen (1976): "Commodity Bundling and the Burden of Monopoly," *The Quarterly Journal of Economics*, 90(3), 475–498.

Anderson, S. P., A. de Palma, and J.-F. Thisse (1988): "A Representative Consumer Theory of the Logit Model," *International Economic Review*, 29(3), 461–466.

Anderson, S. P., and L. Leruth (1993): "Why Firms May Prefer Not to Price Discriminate via Mixed Bundling," *International Journal of Industrial Organization*, 11(1), 49–61.

Armstrong, M. (1996): "Multiproduct Nonlinear Pricing," *Econometrica*, 64(1), 51–75.

——— (1999): "Price discrimination by a many-product firm," *The Review of Economic Studies*, 66(1), 151–168.

Armstrong, M., and J. Vickers (2010): "Competitive Non-Linear Pricing and Bundling," *The Review of Economic Studies*, 77(1), 30–60.

Balbuzanov, I., and M. H. Kotowski (2024): "The Property Rights Theory of Production Networks," *Theoretical Economics*, 19(4), 1619–1658.

Boyd, S., and L. Vandenberghe (2004): *Convex Optimization*. Cambridge University Press.

Chen, Z., and P. Rey (2026): "A Theory of Conglomerate Mergers," TSE Working Paper 23-1447, Toulouse School of Economics, First version June 2023; revised April 2026.

Choi, J. P. (2008): "Mergers with Bundling in Complementary Markets," *The Journal of Industrial Economics*, 56(3), 553–577.

Cover, T. M., and J. A. Thomas (2006): *Elements of Information Theory*. John Wiley & Sons, Hoboken, NJ, 2nd edn.

Elliott, M., and B. Golub (2019): "A Network Approach to Public Goods," *Journal of Political Economy*, 127(2), 730–776.

Flores-Fillol, R., and R. Moner-Colonques (2011): "Endogenous Mergers of Complements with Mixed Bundling," *Review of Industrial Organization*, 39(3), 231–251.

Galeotti, A., B. Golub, and S. Goyal (2020): "Targeting Interventions in Networks," *Econometrica*, 88(6), 2445–2471.

Hastie, T., R. Tibshirani, and J. Friedman (2009): *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, Springer Series in Statistics. Springer, New York, 2 edn.

Kawaguchi, K., T. Kuroda, S. Sato, and T. Ida (2026): "Oligopoly Pricing with Bundling and Fringe Entry: The Case of Telecommunications," Mimeo.

Konovalov, A., and Z. Sándor (2010): "On price equilibrium with multi-product firms," *Economic Theory*, 44(2), 271–292.

Manelli, A. M., and D. R. Vincent (2006): "Bundling as an Optimal Selling Mechanism for a Multiple-Good Monopolist," *Journal of Economic Theory*, 127(1), 1–35.

Matutes, C., and P. Regibeau (1992): "Compatibility and Bundling of Complementary Goods in a Duopoly," *The Journal of Industrial Economics*, 40(1), 37–54.

McAfee, R. P., J. McMillan, and M. D. Whinston (1989): "Multiproduct Monopoly, Commodity Bundling, and Correlation of Values," *The Quarterly Journal of Economics*, 104(2), 371–383.

Monderer, D., and L. S. Shapley (1996): "Potential games," *Games and economic behavior*, 14(1), 124–143.

National Association of Insurance Commissioners (2023): "Bundling," Center for

Insurance Policy and Research, last updated October 22, 2023.

New Zealand Commerce Commission (2023): "Product Disclosure – Retail Service Bundling Guidelines (Energy and Telecommunications Bundles)," Discussion paper, New Zealand Commerce Commission, Final version, issued under section 234 of the Telecommunications Act 2001, November 22, 2023.

Nocke, V., and N. Schutz (2018): "Multiproduct-Firm Oligopoly: An Aggregative Games Approach," *Econometrica*, 86(2), 523–557.

——— (2025): "Oligopoly, Complementarities, and Transformed Potentials," CEPR Discussion Paper 20527, Centre for Economic Policy Research, Paris and London.

Ofcom (2026): "Pricing and Consumer Engagement: Trends in the UK Communications Sector," Discussion paper, Office of Communications, London, https://www.ofcom.org.uk/siteassets/resources/documents/research-and-data/multi-sector/pricing/2025/pricing-and-consumer-engagement-report.pdf.

Pemmaraju, S. V., and S. S. Skiena (2003): *Computational Discrete Mathematics: Combinatorics and Graph Theory with Mathematica*. Cambridge University Press.

Rochet, J.-C., and P. Choné (1998): "Ironing, Sweeping, and Multidimensional Screening," *Econometrica*, 66(4), 783–826.

Thanassoulis, J. (2007): "Competitive Mixed Bundling and Consumer Surplus," *Journal of Economics and Management Strategy*, 16(2), 437–467.

——— (2011): "Is Multimedia Convergence to Be Welcomed?," *Journal of Industrial Economics*, 59(2), 225–253.

Zhou, J. (2021): "Mixed Bundling in Oligopoly Markets," *Journal of Economic Theory*, 194, 105257.

A Appendix for "Oligopoly with Mixed Bundling"

Contents

A.1 Equilibrium analysis . . . A1
A.1.1 Proof of Theorem 1 . . . A1
A.1.2 Nonexistence of an ordinal potential in the original pricing game . . . A3
A.2 Comparative statics . . . A4
A.2.1 Preliminaries . . . A4
A.2.2 Proof of Proposition 1 . . . A7
A.2.3 Proof of Proposition 2 . . . A8
A.2.4 Proof of Example 1 . . . A8
A.2.5 Proof of Proposition 3 . . . A9
A.2.6 Proof of Example 2 . . . A10
A.2.7 Proof of Proposition 4 . . . A10
A.2.8 Proof of Example 3 . . . A11
A.3 Mergers . . . A11
A.3.1 Preliminaries . . . A11
A.3.2 Proof of Lemma 2 . . . A12
A.3.3 Merging firms' markup changes in Step 1 . . . A13
A.3.4 Proof of Lemma 3 . . . A15
A.3.5 Proof of Lemma 4 . . . A16
A.3.6 Proof of Proposition 5 . . . A17
A.3.7 Proof of Proposition 6 . . . A20
A.3.8 Proof of Proposition 7 . . . A22
A.3.9 Proof of Example 4 . . . A23
A.3.10 Proof of Proposition 8 . . . A24
A.3.11 Proof of Proposition 9 . . . A25
A.4 Mixed bundling versus separate pricing . . . A27
A.4.1 Preliminaries . . . A27
A.4.2 Proof of Proposition 10 . . . A28
A.4.3 Proof of Proposition 11 . . . A30
A.4.4 Proof of Proposition 12 . . . A32

A.1 Equilibrium analysis

A.1.1 Proof of Theorem 1

We show that the potential Φ in equation (5) is strictly concave and attains a maximum, that its maximizer is an equilibrium of the reduced pricing game, and that every pricing equilibrium has these common markups.

We first establish that Φ is strictly concave. For the product pair chosen by a consumer, let I^i equal one if the pair contains a product supplied by firm i and zero otherwise, and let $\mathbf{I} \equiv (I^i)_{i \in \mathcal{N}}$. Thus $Q^i = \mathbb{E}[I^i]$, where expectations are taken under the choice probabilities generated by $\boldsymbol{\mu}$. After division by $H(\boldsymbol{\mu})$, the terms in $H(\boldsymbol{\mu})$ give the probabilities that the chosen pair contains no firm's product, products supplied only by firm i, or products supplied

by firms i and j. Each term contains the factor $\mathrm{e}^{-\mu^i}$ exactly when the corresponding pair contains a product supplied by firm i. Differentiating H gives, for any $i, j \in \mathcal{N}$,

$$\frac{1}{H(\boldsymbol{\mu})}\frac{\partial H(\boldsymbol{\mu})}{\partial \mu^i} = -\mathbb{E}[I^i],$$
$$\frac{1}{H(\boldsymbol{\mu})}\frac{\partial^2 H(\boldsymbol{\mu})}{\partial \mu^i \partial \mu^j} = \mathbb{E}[I^i I^j].$$

It follows that

$$\frac{\partial^2 \log H(\boldsymbol{\mu})}{\partial \mu^i \partial \mu^j} = \mathbb{E}[I^i I^j] - \mathbb{E}[I^i]\mathbb{E}[I^j].$$

Thus $\nabla^2 \log H(\boldsymbol{\mu})$ is the covariance matrix of $\mathbf{I}$. For every nonzero vector $\mathbf{x} = (x^i)_{i\in\mathcal{N}}$,

$$\begin{aligned}\mathbf{x}^\top \nabla^2 \Phi(\boldsymbol{\mu})\mathbf{x} &= -\sum_{i\in\mathcal{N}} \frac{(x^i)^2}{(\mu^i)^2} - \mathbf{x}^\top \operatorname{Cov}(\mathbf{I})\mathbf{x} \\ &= -\sum_{i\in\mathcal{N}} \frac{(x^i)^2}{(\mu^i)^2} - \operatorname{Var}\left(\sum_{i\in\mathcal{N}} x^i I^i\right) < 0.\end{aligned}$$

Hence Φ is strictly concave.

We next show that Φ has a maximum. Because each firm supplies at least one product that can be purchased, $Q^i(\boldsymbol{\mu}) > 0$ for every finite markup vector. Since $\partial \log H(\boldsymbol{\mu})/\partial \mu^i = -Q^i(\boldsymbol{\mu})$, equation (6) gives

$$\frac{\partial \Phi(\boldsymbol{\mu})}{\partial \mu^i} = \frac{1}{\mu^i} - 1 + Q^i(\boldsymbol{\mu}) > 0 \quad \text{whenever } \mu^i \le 1.$$

Increasing any markup below one to one raises Φ, so it is enough to maximize Φ over $[1,\infty)^{|\mathcal{N}|}$.

Let $M \equiv \max_{i\in\mathcal{N}} \mu^i$, and choose a firm i such that $\mu^i = M$. If $T^0 > 0$, then $H(\boldsymbol{\mu}) \ge T^0$, and therefore

$$\Phi(\boldsymbol{\mu}) \le |\mathcal{N}| \log M - M - \log T^0.$$

If $T^0 = 0$, the maintained restriction on the type network implies that either $T^j > 0$ for some $j \ne i$, or $T^{jk} > 0$ for some distinct $j, k \ne i$. In the first case, $H(\boldsymbol{\mu}) \ge T^j \mathrm{e}^{-\mu^j}$, which gives

$$\Phi(\boldsymbol{\mu}) \le |\mathcal{N}| \log M - M - \log T^j.$$

In the second case, $H(\boldsymbol{\mu}) \ge T^{jk}\mathrm{e}^{-\mu^j-\mu^k}$, which gives

$$\Phi(\boldsymbol{\mu}) \le |\mathcal{N}| \log M - M - \log T^{jk}.$$

In the last two bounds, the markups in the exponent of the selected term in H cancel the corresponding terms in $-\sum_l \mu^l$, while $-\mu^i = -M$ remains. Because only finitely many types are positive, $-\log T$ has a finite upper bound over the types that can appear in these inequalities. Since $|\mathcal{N}| \log M - M \to -\infty$ as $M \to \infty$, there exists $\overline{M} > 1$ such that $\Phi(\boldsymbol{\mu})$ is smaller than its value when every markup equals one whenever $M > \overline{M}$.

Because Φ is continuous, it attains a maximum on the compact box $[1, \overline{M}]^{|\mathcal{N}|}$. This is also its maximum over $[1,\infty)^{|\mathcal{N}|}$ because the box contains the vector with every markup equal

to one, while every vector outside the box gives a lower value. It is also the maximum over $\mathbb{R}^{|\mathcal{N}|}_{++}$ because increasing any markup below one to one raises Φ. Moreover, the derivative of Φ with respect to μ^i is positive at $\mu^i = 1$, so every markup at the maximum is strictly greater than one. Strict concavity makes the maximum unique. At the maximum, equation (1) can equivalently be written as

$$\mu^i = \frac{1}{1 - Q^i(\boldsymbol{\mu})} > 1.$$

We now show that the maximizer is a pricing equilibrium. At this maximum, no firm can raise its profit by changing its common markup. If some other change in a firm's prices or discounts were profitable, its best response would also be profitable and, by Lemma 1, would be summarized by a common markup, which is a contradiction. The prices and discounts satisfying the common-markup equations in Lemma 1 therefore form a pricing equilibrium.

Finally, every pricing equilibrium uses common markups. Equation (6) then implies

$$\frac{\partial \Phi(\boldsymbol{\mu})}{\partial \mu^i} = 0$$

for every firm i. Strict concavity implies that these conditions can hold only at the unique maximizer of Φ. Thus every pricing equilibrium has the same common markups and induces the same price for every available product pair. □

A.1.2 Nonexistence of an ordinal potential in the original pricing game

We construct a two-firm example in which four unilateral changes in price vectors, each strictly profitable for the deviating firm, return to the initial price profile. Such an improvement cycle rules out an ordinal potential.

Consider two firms. Firm 1 supplies a_1 and b_1, and firm 2 supplies a_2 and b_2. We normalize $\lambda = 1$ and all marginal costs to zero, and let $L \equiv \log 2$. Suppose that the only available product pairs are 00, $a_1 0$, $0 b_2$, and $a_2 b_1$, with

$$s_{00} = 0, \quad s_{a_1 0} = s_{0 b_2} = 3L, \quad s_{a_2 b_1} = 2L.$$

All other product-pair values are $-\infty$, and all bundle discounts remain zero. Let firm 1's price vectors, ordered as (p_{a_1}, p_{b_1}), be

$$\mathbf{p}^1 = (3L, 2L), \quad \widehat{\mathbf{p}}^1 = (2L, L),$$

and let firm 2's price vectors, ordered as (p_{a_2}, p_{b_2}), be

$$\mathbf{p}^2 = (3L, 2L), \quad \widehat{\mathbf{p}}^2 = (2L, 2L).$$

Consider the cycle

$$(\mathbf{p}^1, \mathbf{p}^2) \longrightarrow (\widehat{\mathbf{p}}^1, \mathbf{p}^2) \longrightarrow (\widehat{\mathbf{p}}^1, \widehat{\mathbf{p}}^2) \longrightarrow (\mathbf{p}^1, \widehat{\mathbf{p}}^2) \longrightarrow (\mathbf{p}^1, \mathbf{p}^2),$$

where firms 1 and 2 alternate in changing their price vectors. Let $\mathbf{q} \equiv (q_{00}, q_{a_1 0}, q_{0 b_2}, q_{a_2 b_1})$.

Direct substitution gives

$$\mathbf{q}(\mathbf{p}^1, \mathbf{p}^2) = \left(\frac{8}{33}, \frac{8}{33}, \frac{16}{33}, \frac{1}{33}\right),$$
$$\mathbf{q}(\widehat{\mathbf{p}}^1, \mathbf{p}^2) = \left(\frac{4}{21}, \frac{8}{21}, \frac{8}{21}, \frac{1}{21}\right),$$
$$\mathbf{q}(\widehat{\mathbf{p}}^1, \widehat{\mathbf{p}}^2) = \left(\frac{2}{11}, \frac{4}{11}, \frac{4}{11}, \frac{1}{11}\right),$$
$$\mathbf{q}(\mathbf{p}^1, \widehat{\mathbf{p}}^2) = \left(\frac{4}{17}, \frac{4}{17}, \frac{8}{17}, \frac{1}{17}\right).$$

Since $Q^1 = q_{a_1 0} + q_{a_2 b_1}$ and $Q^2 = q_{0 b_2} + q_{a_2 b_1}$, the corresponding firm demands, in the same order, are

$$\left(\frac{3}{11}, \frac{17}{33}\right), \quad \left(\frac{3}{7}, \frac{3}{7}\right), \quad \left(\frac{5}{11}, \frac{5}{11}\right), \quad \left(\frac{5}{17}, \frac{9}{17}\right).$$

The corresponding profit pairs are

$$L\left(\frac{26}{33}, \frac{35}{33}\right), \quad L\left(\frac{17}{21}, \frac{19}{21}\right), \quad L\left(\frac{9}{11}, \frac{10}{11}\right), \quad L\left(\frac{14}{17}, \frac{18}{17}\right).$$

The four deviators' profit gains, divided by L, are

$$\frac{5}{231}, \quad \frac{1}{231}, \quad \frac{1}{187}, \quad \frac{1}{561}.$$

Thus every unilateral change is strictly profitable. An ordinal potential would have to increase strictly at every step and yet return to its initial value, which is impossible. Because the four inequalities are strict, the same price changes remain an improvement cycle if the other product pairs are assigned sufficiently low finite gross values. □

A.2 Comparative statics

A.2.1 Preliminaries

This subsection collects three tools: the lift at fixed markups of the other firms, used for Proposition 1 and the merger analysis; a matrix result that signs equilibrium markup responses (Lemma A.1); and the derivatives of equilibrium consumer surplus with respect to individual types, used for Propositions 2–4.

Lift at fixed markups. Fix distinct firms i and j and the markups $(\mu^k)_{k \in \mathcal{N} \setminus \{i,j\}}$ of all other firms. Define

$$\tau^0 \equiv T^0 + \sum_{k \in \mathcal{N} \setminus \{i,j\}} T^k \mathrm{e}^{-\mu^k} + \frac{1}{2} \sum_{k \in \mathcal{N} \setminus \{i,j\}} \sum_{l \in \mathcal{N} \setminus \{i,j,k\}} T^{kl} \mathrm{e}^{-\mu^k - \mu^l}, \tag{A.1}$$

$$\tau^i \equiv T^i + \sum_{k \in \mathcal{N} \setminus \{i,j\}} T^{ik} \mathrm{e}^{-\mu^k}, \quad \tau^j \equiv T^j + \sum_{k \in \mathcal{N} \setminus \{i,j\}} T^{jk} \mathrm{e}^{-\mu^k}. \tag{A.2}$$

These are the contributions to the CS-index, at the fixed markups of the other firms, from product pairs containing neither firm i nor firm j, firm i but not firm j, and firm j but not firm i, respectively.

Substituting into the definition of lift gives

$$\ell^{ij} - 1 = \frac{\tau^0 T^{ij} - \tau^i \tau^j}{(\tau^i + T^{ij} \mathrm{e}^{-\mu^j})(\tau^j + T^{ij} \mathrm{e}^{-\mu^i})}. \tag{A.3}$$

Its sign does not depend on μ^i or μ^j.

Markup responses.

Lemma A.1. *Let* $\mathbf{I}$ *be a random vector whose entries are zero or one, with at most two entries equal to one, and let* $\mathbf{Q} = \mathbb{E}[\mathbf{I}] > \mathbf{0}$*. For any positive diagonal matrix* D*, define*

$$J \equiv D + \mathrm{Cov}(\mathbf{I}), \quad \boldsymbol{\zeta} \equiv J^{-1}\mathbf{Q}.$$

Then, for every i*,*

$$0 < \zeta^i < 1 + \mathbf{Q}^\top \boldsymbol{\zeta}. \tag{A.4}$$

Proof. Define

$$\mathbf{w} \equiv \frac{\boldsymbol{\zeta}}{1 + \mathbf{Q}^\top \boldsymbol{\zeta}}.$$

Because J is positive definite, $\mathbf{Q}^\top \boldsymbol{\zeta} > 0$. Moreover, $\mathbf{w}$ is the unique minimizer of

$$G(\mathbf{w}) \equiv \frac{1}{2}\mathbf{w}^\top D \mathbf{w} + \frac{1}{2}\mathbb{E}\left[(\mathbf{I}^\top \mathbf{w} - 1)^2\right]. \tag{A.5}$$

Indeed, the first-order condition for minimizing G in (A.5) is

$$\left(D + \mathbb{E}[\mathbf{I}\mathbf{I}^\top]\right)\mathbf{w} = \mathbf{Q},$$

which follows from $J\boldsymbol{\zeta} = \mathbf{Q}$ and $\mathrm{Cov}(\mathbf{I}) = \mathbb{E}[\mathbf{I}\mathbf{I}^\top] - \mathbf{Q}\mathbf{Q}^\top$.

For any real number x, let $P(x) \equiv \min\{1, \max\{0, x\}\}$. We have

$$|P(x)| \le |x|, \quad |P(x) - 1| \le |x - 1|, \quad |P(x) + P(y) - 1| \le |x + y - 1|.$$

Because $\mathbf{I}$ has at most two nonzero entries, replacing each w^i by $P(w^i)$ weakly lowers G. The first term falls strictly if any component lies outside $[0, 1]$. Hence $0 \le w^i \le 1$ for every i.

The first-order condition for component i is

$$D_{ii} w^i + \mathbb{E}\left[I^i(\mathbf{I}^\top \mathbf{w} - 1)\right] = 0. \tag{A.6}$$

If $w^i = 1$, both terms on the left-hand side are nonnegative and the first is strictly positive. This argument rules out $w^h = 1$ for every h. If $w^i = 0$, whenever $I^i = 1$, at most one other component of $\mathbf{I}$ is nonzero and that component is strictly below one. Hence $\mathbf{I}^\top \mathbf{w} - 1 < 0$. Since $Q^i > 0$, the second term in (A.6) is then strictly negative. Both cases contradict equation (A.6). Thus $0 < w^i < 1$, which is equivalent to inequalities (A.4). □

Unless otherwise stated, all quantities in the rest of Appendix A.2 are evaluated at the equilibrium associated with the current type network. Let $\mathbf{Q} \equiv (Q^i)_{i\in\mathcal{N}}$. As in the proof of Theorem 1, for the product pair chosen by a consumer, let I^i equal one if the pair contains a product supplied by firm i and zero otherwise, and let $\mathbf{I} \equiv (I^i)_{i\in\mathcal{N}}$. Thus $\mathbf{Q} = \mathbb{E}[\mathbf{I}]$. Let $\Sigma \equiv \mathrm{Cov}(\mathbf{I})$ and define

$$J \equiv \mathrm{diag}\left((1-Q^i)^2\right)_{i\in\mathcal{N}} + \Sigma.$$

The equilibrium conditions are

$$\frac{1}{\mu^i} - 1 + Q^i = 0, \quad i \in \mathcal{N}.$$

Since the derivative of $\mathbf{Q}$ with respect to $\boldsymbol{\mu}$ is $-\Sigma$, J is the negative of the Jacobian of these conditions with respect to $\boldsymbol{\mu}$. The matrix Σ is positive semidefinite, and $Q^i < 1$ for every firm. Hence J is positive definite.

Define $\boldsymbol{\zeta} \equiv J^{-1}\mathbf{Q}$. Lemma A.1, applied with $D = \mathrm{diag}((1-Q^i)^2)_{i\in\mathcal{N}}$, gives, for every firm i,

$$0 < \zeta^i < 1 + \mathbf{Q}^\top \boldsymbol{\zeta}. \tag{A.7}$$

Changes in individual types. Holding markups fixed,

$$\frac{\partial \log H}{\partial T^i} = \frac{\mathrm{e}^{-\mu^i}}{H}, \quad \frac{\partial \mathbf{Q}}{\partial T^i} = \frac{\mathrm{e}^{-\mu^i}}{H}\left(\mathbf{e}_i - \mathbf{Q}\right),$$
$$\frac{\partial \log H}{\partial T^{ij}} = \frac{\mathrm{e}^{-\mu^i-\mu^j}}{H}, \quad \frac{\partial \mathbf{Q}}{\partial T^{ij}} = \frac{\mathrm{e}^{-\mu^i-\mu^j}}{H}\left(\mathbf{e}_i + \mathbf{e}_j - \mathbf{Q}\right), \quad i \neq j,$$

where $\mathbf{e}_i$ is the ith unit vector. Applying the implicit function theorem to the equilibrium conditions and accounting for the induced markup changes gives

$$\frac{\mathrm{d}\log H^*}{\mathrm{d}\,T^i} = \frac{\mathrm{e}^{-\mu^i}}{H}\left(1 + \mathbf{Q}^\top\boldsymbol{\zeta} - \zeta^i\right), \tag{A.8}$$
$$\frac{\mathrm{d}\log H^*}{\mathrm{d}\,T^{ij}} = \frac{\mathrm{e}^{-\mu^i-\mu^j}}{H}\left(1 + \mathbf{Q}^\top\boldsymbol{\zeta} - \zeta^i - \zeta^j\right), \quad i \neq j. \tag{A.9}$$

For example,

$$\frac{\mathrm{d}\,\boldsymbol{\mu}}{\mathrm{d}\,T^i} = \frac{\mathrm{e}^{-\mu^i}}{H} J^{-1}\left(\mathbf{e}_i - \mathbf{Q}\right).$$

Substituting this expression into the total derivative of $\log H$ and using $\partial \log H/\partial\boldsymbol{\mu} = -\mathbf{Q}$ gives equation (A.8). The derivation of equation (A.9) is identical. The same expressions give the relevant right derivatives when the type being increased is initially zero.

An increase in the outside-option type gives $\partial \log H/\partial T^0 = 1/H$ and $\partial\mathbf{Q}/\partial T^0 = -\mathbf{Q}/H$ at fixed markups. Hence $\mathrm{d}\,\boldsymbol{\mu}/\,\mathrm{d}\,T^0 = -\boldsymbol{\zeta}/H$, and

$$\frac{\mathrm{d}\log H^*}{\mathrm{d}\,T^0} = \frac{1 + \mathbf{Q}^\top\boldsymbol{\zeta}}{H} > 0. \tag{A.10}$$

By inequalities (A.7), a more attractive outside option lowers every equilibrium markup and raises consumer surplus.

A.2.2 Proof of Proposition 1

We first show that the slope of firm i's best response has the sign of $1-\ell^{ij}$. We then establish the sign of $\ell^{ij}-1$ in each case of the proposition.

Applying the implicit function theorem to firm i's optimality condition $1/\mu^i - 1 + Q^i(\boldsymbol{\mu}) = 0$ gives

$$\frac{\mathrm{d}\,\mu^i}{\mathrm{d}\,\mu^j} = \frac{\partial Q^i/\partial \mu^j}{1/(\mu^i)^2 - \partial Q^i/\partial \mu^i}.$$

Because $\partial Q^i/\partial \mu^i = -Q^i(1-Q^i) < 0$, the denominator is positive. As shown in the main text, $\partial Q^i/\partial \mu^j = Q^i Q^j (1-\ell^{ij})$. It follows that firm i's best-response markup increases with μ^j when $\ell^{ij} < 1$ and decreases with μ^j when $\ell^{ij} > 1$.

The remark after Proposition 1 follows from equation (A.3), whose sign does not depend on μ^i or μ^j. When $\mathcal{N} = \{i, j\}$, $\tau^0 = T^0$, $\tau^i = T^i$ and $\tau^j = T^j$. If $T^0 > 0$, the sign therefore changes once, at $T^{ij} = T^i T^j / T^0$. If $T^0 = 0$, the maintained restriction on the type network implies $T^i, T^j > 0$, so the lift is always below one.

For the last statement of the proposition, we allow all equilibrium markups to adjust as T^{ij} changes, holding all other types fixed. Denote the resulting equilibrium by $\boldsymbol{\mu}^*(T^{ij})$.

Consider first sufficiently small T^{ij}. Suppose that each of firms i and j supplies a product that can be purchased without any product of the other firm, so that both firms have positive equilibrium demand at $T^{ij} = 0$. At $T^{ij} = 0$, $Q^{ij*} = 0$ and hence $\ell^{ij*} = 0$. The Hessian of the potential is negative definite, so the Jacobian of the equilibrium conditions is nonsingular. The implicit function theorem therefore implies that the equilibrium markups and demands vary continuously with T^{ij} near zero. It follows that $\ell^{ij*} < 1$ for all sufficiently small positive T^{ij}.

For the large-T^{ij} result, suppose that some available product pair involves neither firm i nor firm j. We first show that

$$\mu^{i*}(T^{ij}) \longrightarrow \infty \quad \text{and} \quad \mu^{j*}(T^{ij}) \longrightarrow \infty \quad \text{as } T^{ij} \to \infty.$$

Consider any sequence $T_n^{ij} \to \infty$. If both markups remained bounded along a subsequence, the weight $T_n^{ij} \mathrm{e}^{-\mu^{i*}(T_n^{ij}) - \mu^{j*}(T_n^{ij})}$ of the ij pair would diverge, while every other weight would remain bounded. The choice probability of the ij pair would therefore converge to one. This would imply that both firms' demands converge to one and, by $\mu^{k*} = 1/(1-Q^{k*})$, that both markups diverge, contradicting their boundedness.

Suppose instead that $\mu^{i*}(T_n^{ij}) \to \infty$ while $\mu^{j*}(T_n^{ij})$ remains bounded. Then $Q^{i*}(T_n^{ij}) \to 1$, and

$$\frac{Q^{ij*}(T_n^{ij})}{Q^{i*}(T_n^{ij})} = \frac{T_n^{ij} \mathrm{e}^{-\mu^{j*}(T_n^{ij})}}{T^i + \sum_{k \in \mathcal{N} \setminus \{i,j\}} T^{ik} \mathrm{e}^{-\mu^{k*}(T_n^{ij})} + T_n^{ij} \mathrm{e}^{-\mu^{j*}(T_n^{ij})}} \longrightarrow 1.$$

It follows that $Q^{ij*}(T_n^{ij}) \to 1$ and hence $\mu^{j*}(T_n^{ij}) \to \infty$, again a contradiction. The same argument with i and j reversed rules out a bounded markup for firm i when firm j's markup diverges. Therefore both markups diverge as $T^{ij} \to \infty$.

The equilibrium conditions then imply $Q^{i*}(T^{ij}), Q^{j*}(T^{ij}) \to 1$. Since Q^{ij} is the probability that both firms appear in the chosen pair,

$$Q^{ij*}(T^{ij}) \geq Q^{i*}(T^{ij}) + Q^{j*}(T^{ij}) - 1 \longrightarrow 1.$$

Any product pair involving an outsider differs from the ij pair. Therefore, for every $k \in \mathcal{N} \setminus \{i, j\}$,

$$Q^{k*}(T^{ij}) \leq 1 - Q^{ij*}(T^{ij}) \longrightarrow 0, \quad \mu^{k*}(T^{ij}) \longrightarrow 1.$$

By equation (A.3), $\ell^{ij} > 1$ if and only if $T^{ij}\tau^0 > \tau^i\tau^j$. The term τ^0 consequently converges to

$$T^0 + \mathrm{e}^{-1} \sum_{k \in \mathcal{N} \setminus \{i,j\}} T^k + \mathrm{e}^{-2} \sum_{\substack{k<l \\ k,l \in \mathcal{N} \setminus \{i,j\}}} T^{kl},$$

which is strictly positive by assumption. The other two terms in that condition converge to finite numbers. Multiplying the strictly positive term by T^{ij} therefore makes this inequality hold strictly for all sufficiently large T^{ij}. Hence $\ell^{ij*} > 1$ for all sufficiently large T^{ij}.

Finally, if every available product pair involves firm i or firm j, then $\tau^0 = 0$, and equation (A.3) gives $\ell^{ij} \leq 1$ at every markup vector and for every T^{ij}. If firm i sells only together with firm j, then $\tau^i = 0$, and equation (A.3) gives $\ell^{ij} \geq 1$ for every $T^{ij} > 0$; the same holds with the roles of i and j exchanged. □

A.2.3 Proof of Proposition 2

We show that equilibrium consumer surplus strictly increases in T^i.

By inequalities (A.7), $1 + \mathbf{Q}^\top \boldsymbol{\zeta} - \zeta^i > 0$. Equation (A.8) therefore implies that $\log H^*$, and hence consumer surplus $\lambda \log H^*$, is strictly increasing in T^i. Integrating the derivative between any two values of T^i proves the result. □

A.2.4 Proof of Example 1

We derive the equilibrium conditions stated in Example 1 and show that H^* strictly decreases in x.

The symmetry and equilibrium uniqueness imply $\mu^1 = \mu^2 = \mu$. The equilibrium CS-index and firm 1's demand are

$$H^* = 2\mathrm{e}^{-\mu} + x\mathrm{e}^{-2\mu}, \quad Q^1 = \frac{\mathrm{e}^{-\mu} + x\mathrm{e}^{-2\mu}}{H^*}.$$

The equilibrium condition $\mu(1 - Q^1) = 1$ gives

$$\mu = 2 + x\mathrm{e}^{-\mu}, \quad H^* = \mu\mathrm{e}^{-\mu},$$

as stated in (7). Implicit differentiation yields

$$\frac{\mathrm{d}\,\mu}{\mathrm{d}\,x} = \frac{\mathrm{e}^{-\mu}}{1 + x\mathrm{e}^{-\mu}} = \frac{\mathrm{e}^{-\mu}}{\mu - 1} > 0$$

and

$$\frac{\mathrm{d}\log H^*}{\mathrm{d}\,x} = \left(\frac{1}{\mu} - 1\right)\frac{\mathrm{d}\,\mu}{\mathrm{d}\,x} = -\frac{\mathrm{e}^{-\mu}}{\mu} < 0. \tag{A.11}$$

Thus an increase in $T^{12} = x$ strictly lowers consumer surplus. Since the equilibrium and the derivative in (A.11) vary continuously with T^0, the same local effect holds, for every fixed $x \geq 0$, when T^0 is positive and sufficiently small. □

A.2.5 Proof of Proposition 3

We show that a uniform reduction in the marginal costs of firm i's products strictly raises consumer surplus. We first derive the effect of an auxiliary change that scales up every type involving firm i. The proof then splits the cost reduction into this change and an increase in T^i.

Consider an auxiliary change that multiplies by e^t the social value of every product pair involving firm i, counting a pair containing two of its products only once. Equivalently, it multiplies T^i and every T^{ij}, $j \neq i$, by e^t and leaves all other types fixed. Holding markups fixed,

$$\frac{\partial \mathbf{Q}}{\partial t} = \Sigma \mathbf{e}_i, \quad \frac{\partial \log H}{\partial t} = Q^i.$$

The induced markup response is

$$\frac{\mathrm{d}\,\boldsymbol{\mu}}{\mathrm{d}\,t} = J^{-1}\Sigma \mathbf{e}_i = \mathbf{e}_i - J^{-1}\,\mathrm{diag}\left((1-Q^k)^2\right)_{k\in\mathcal{N}}\mathbf{e}_i.$$

It follows that

$$\frac{\mathrm{d}\log H^*}{\mathrm{d}\,t} = Q^i - \mathbf{Q}^\top \frac{\mathrm{d}\,\boldsymbol{\mu}}{\mathrm{d}\,t} = (1-Q^i)^2\zeta^i. \tag{A.12}$$

This derivative is strictly positive by inequalities (A.7).

Parameterize the common cost reduction by

$$c_a(\theta) = c_a - \lambda\theta, \quad c_b(\theta) = c_b - \lambda\theta \quad \text{for every } a \in \mathcal{A}^i \text{ and } b \in \mathcal{B}^i.$$

At each point on this path, define

$$q^{ii} \equiv \sum_{a\in\mathcal{A}^i}\sum_{b\in\mathcal{B}^i} q_{ab}.$$

Thus q^{ii} is the probability that the chosen pair contains two products supplied by firm i.

A change in θ makes every product pair involving firm i more attractive once and makes a pair counted by q^{ii} more attractive a second time. The first component is the auxiliary change in (A.12). The second component changes only T^i, and its direct contribution satisfies

$$\frac{\mathrm{e}^{-\mu^i}}{H}\frac{\mathrm{d}\,T^i}{\mathrm{d}\,\theta} = q^{ii},$$

where the derivative on the left refers only to this second component. Equations (A.12) and (A.8) therefore give

$$\frac{\mathrm{d}\log H^*(\theta)}{\mathrm{d}\,\theta} = (1-Q^i)^2\zeta^i + q^{ii}\left(1 + \mathbf{Q}^\top\boldsymbol{\zeta} - \zeta^i\right) > 0. \tag{A.13}$$

The first term is strictly positive by inequalities (A.7), and the second is nonnegative. Integrating equation (A.13) over the cost-reduction path proves the proposition. □

A.2.6 Proof of Example 2

We specify primitives that generate the type network of Example 1. We then show that reducing c_a lowers consumer surplus, whereas reducing $c_{a'}$ raises it.

We normalize $\lambda = 1$. Market A contains products a and a', both supplied by firm 1, while market B contains product b, supplied by firm 2. Let $c_a = c_{a'} = c_b = 1$. Suppose that the only available product pairs are ab, $a'0$, and $0b$, with gross surpluses

$$s_{ab} = 2 + \log x, \quad s_{a'0} = 1, \quad s_{0b} = 1, \quad x > 0.$$

All other product-pair values are $-\infty$. Thus, product a can be consumed only with product b, whereas product a' can be consumed only with the outside option in market B. The corresponding type network is

$$(T^0, T^1, T^2, T^{12}) = (0, 1, 1, x).$$

This is precisely the type network in Example 1.

First, reduce the cost of product a according to $c_a(\theta) = 1 - \theta$, holding all other primitives fixed. Because ab is the only available pair containing a, the resulting type network is

$$(T^0, T^1, T^2, T^{12}(\theta)) = (0, 1, 1, x\mathrm{e}^{\theta}).$$

Example 1 shows that equilibrium H^* is strictly decreasing in T^{12}. Hence

$$\left.\frac{\mathrm{d}\log H^*}{\mathrm{d}\,\theta}\right|_{\theta=0} = x\frac{\mathrm{d}\log H^*}{\mathrm{d}\,x} < 0.$$

A marginal reduction in c_a therefore lowers consumer surplus.

Next, reduce the cost of product a' according to $c_{a'}(\theta) = 1 - \theta$. Because $a'0$ is the only available pair containing a', the resulting type network is

$$(T^0, T^1(\theta), T^2, T^{12}) = (0, \mathrm{e}^{\theta}, 1, x).$$

This change increases only firm 1's standalone type. Proposition 2 therefore implies

$$\left.\frac{\mathrm{d}\log H^*}{\mathrm{d}\,\theta}\right|_{\theta=0} > 0.$$

A marginal reduction in $c_{a'}$ therefore raises consumer surplus. □

A.2.7 Proof of Proposition 4

We show that entry strictly raises consumer surplus. To do so, we connect the pre-entry and post-entry economies by raising the entrant's marginal costs, and we sign the derivative of $\log H^*$ along this path.

Starting from the post-entry economy, add $\lambda\theta$ to the marginal cost of every product supplied by the entrant, where $\theta \geq 0$. A product pair containing one entrant product has its social value multiplied by $\mathrm{e}^{-\theta}$, while a pair containing two entrant products has its social value multiplied by $\mathrm{e}^{-2\theta}$. All types involving only incumbents remain fixed, and every type involving

the entrant converges to zero as $\theta \to \infty$. The post-entry economy is obtained at $\theta = 0$. At every finite θ, at least one type involving the entrant remains positive, so Proposition 3 implies that $H(\theta)$ is strictly decreasing in θ.

To identify the limit, let $\alpha \equiv \mathrm{e}^{-\theta}$ and allow $\alpha = 0$. Each type involving the entrant is a finite sum of terms proportional to either α or α^2, so the equilibrium conditions extend continuously to $\alpha = 0$. At $\alpha = 0$, the pre-entry markups together with $\mu^E = 1$ solve the equilibrium conditions for the firm set $\mathcal{N} \cup \{E\}$. The negative Jacobian of these conditions is block diagonal, with the pre-entry matrix J_{pre} and the scalar 1 on its diagonal, and is therefore positive definite. The implicit function theorem gives a unique continuous solution for sufficiently small $\alpha \geq 0$, and equilibrium uniqueness identifies it with the path above. Hence the incumbent markups converge to their pre-entry values, $\mu^E(\theta) \to 1$, and $H(\theta) \to H^*$ as $\theta \to \infty$. For any fixed $\theta_0 > 0$,

$$\overline{H}^* = H(0) > H(\theta_0) \geq \lim_{\theta \to \infty} H(\theta) = H^*,$$

which proves the result. □

A.2.8 Proof of Example 3

We show that entry raises firm 1's markup if and only if $T^0 T^{1E} > T^1 T^E$, which is equivalent to $\ell^{1E} > 1$.

Before entry, firm 1's first-order condition $\mu^1(1-Q^1) = 1$, with $Q^1 = T^1 \mathrm{e}^{-\mu^1}/(T^0 + T^1 \mathrm{e}^{-\mu^1})$, gives $(\mu^{1*} - 1)\mathrm{e}^{\mu^{1*}} = T^1/T^0$. After entry, $1 - \overline{Q}^{1*} = (T^0 + T^E \mathrm{e}^{-\overline{\mu}^{E*}})/\overline{H}^*$, so the same condition gives

$$(\overline{\mu}^{1*} - 1)\mathrm{e}^{\overline{\mu}^{1*}} = \frac{T^1 + T^{1E} y}{T^0 + T^E y}, \quad y \equiv \mathrm{e}^{-\overline{\mu}^{E*}} > 0.$$

Since $(\mu - 1)\mathrm{e}^{\mu}$ is strictly increasing for $\mu > 1$, $\overline{\mu}^{1*} > \mu^{1*}$ if and only if $(T^1 + T^{1E} y)/(T^0 + T^E y) > T^1/T^0$, that is, if and only if $T^0 T^{1E} > T^1 T^E$. By the remark after Proposition 1, this holds if and only if $\ell^{1E} > 1$. In that case, the total markup on every product pair supplied only by firm 1 rises, so consumers who buy only firm 1's products are worse off. □

A.3 Mergers

A.3.1 Preliminaries

This subsection collects definitions and identities used throughout Appendix A.3: the Step-1 and Step-2 problems, the cross-demand derivatives behind equation (8), the merging firms' first-order conditions in Steps 1 and 2, and the lift formula of Appendix A.2.1 applied to the merging firms.

Hold the prices and discounts of the outsiders fixed at their pre-merger levels. In Step 1, the merging firms maximize their combined profit subject to the restriction that the products previously supplied by each firm carry one common markup and that no discount is available on a pair containing products previously supplied by both merging firms. Let $(\breve{\mu}^f, \breve{\mu}^g)$ denote a Step-1 optimum. In Step 2, the merged firm chooses all prices and discounts available under post-merger ownership.

For distinct firms i and j, logit demand gives

$$\frac{\partial Q^j}{\partial \mu^i} = Q^i Q^j - Q^{ij} = Q^i Q^j (1 - \ell^{ij}), \quad \frac{\partial \Pi^j}{\partial \mu^i} = \lambda \mu^j Q^i Q^j (1 - \ell^{ij}). \tag{A.14}$$

At the pre-merger equilibrium, $\partial \Pi^i / \partial \mu^i = 0$. With $\{i, j\} = \{f, g\}$, adding this condition to the second equality gives equation (8) in the main text.

Let $\breve{Q}^i$ be the Step-1 probability that the chosen pair contains a product previously supplied by firm $i \in \{f, g\}$, and let $\breve{Q}^{fg}$ be the probability that it contains products previously supplied by both merging firms. Let $\breve{\Pi}^M$ and $\widetilde{\Pi}^M$ denote the merging firms' combined profits in Steps 1 and 2. Because a pair containing products previously supplied by both firms is counted in both $\breve{Q}^f$ and $\breve{Q}^g$, the Step-1 first-order condition and the Step-2 common-markup condition give

$$\breve{\mu}^i + \breve{\mu}^j \frac{\breve{Q}^{fg}}{\breve{Q}^i} = 1 + \frac{\breve{\Pi}^M}{\lambda}, \quad \{i, j\} = \{f, g\}, \tag{A.15}$$

$$\widetilde{\mu}^M = 1 + \frac{\widetilde{\Pi}^M}{\lambda}. \tag{A.16}$$

We use equations (A.1)–(A.3) with $\{i, j\} = \{f, g\}$, so that the fixed markups are those of the outsiders, $(\mu^k)_{k \in \mathcal{O}}$. The maintained restriction on the type network, applied after the merger, implies $\tau^0 > 0$.

A.3.2 Proof of Lemma 2

We show that $\breve{H} - H^*$ has the sign of $\ell^{fg*} - 1$. After treating the case in which a merging firm sells only together with the other, we show that the Step-1 markups are positive. We then sign each merging firm's unilateral first-order condition at the Step-1 optimum and integrate the change in $\log H$ along a path from the pre-merger equilibrium to that optimum.

Suppose first that $\tau^f = 0$ in equation (A.2), so that every available product pair containing products of firm f also contains products of firm g. Then $Q^f = Q^{fg} \leq Q^g$ at every markup vector, so $\ell^{fg*} = 1/Q^{g*} = 1/\max\{Q^{f*}, Q^{g*}\} > 1$. In Step 1, firm f's markup affects only the prices of pairs involving both merging firms, so the merging firms can choose independently the markup on pairs involving firm g but not firm f and the total markup on pairs involving both firms. Every Step-1 optimum therefore induces the Step-2 prices, so $\breve{H} = \widetilde{H}$. Since $\widehat{\ell}(Q^{f*}, Q^{g*}) < 1/\max\{Q^{f*}, Q^{g*}\}$, Lemma 4 gives $\widetilde{H} > H^*$. The case $\tau^g = 0$ is symmetric. We assume henceforth that $\tau^f, \tau^g > 0$, so that $Q^f > Q^{fg}$ and $Q^g > Q^{fg}$ at every markup vector.

Subtracting the two equations in (A.15) gives

$$\Big(1 - \frac{\breve{Q}^{fg}}{\breve{Q}^g}\Big)\breve{\mu}^f = \Big(1 - \frac{\breve{Q}^{fg}}{\breve{Q}^f}\Big)\breve{\mu}^g.$$

Both coefficients are positive, so the two markups have the same sign. Since the merging firms can reproduce the pre-merger markups in Step 1, $\breve{\Pi}^M \geq \Pi^{f*} + \Pi^{g*} > 0$, and equation (A.15) then rules out nonpositive markups.

At the Step-1 optimum, the left-hand side of each merging firm's unilateral first-order condition $1/\mu^i - 1 + Q^i = 0$ has the sign of $\ell^{fg*} - 1$. To see this, let $\gamma^i \equiv 1/\breve{\mu}^i - 1 + \breve{Q}^i$. Since

$\partial\Pi^i/\partial\mu^i = \lambda\mu^i Q^i(1/\mu^i - 1 + Q^i)$, the first-order condition for μ^i and equation (A.14) give

$$\gamma^i = \frac{\breve{\mu}^j \breve{Q}^j}{\breve{\mu}^i}\big(\breve{\ell}^{fg} - 1\big), \quad \{i,j\} = \{f,g\}.$$

By equation (A.3), γ^f and γ^g have the sign of $\ell^{fg*} - 1$.

The change in the CS-index between the pre-merger equilibrium and the Step-1 optimum can be obtained by integrating its derivative along a path connecting the two. View the potential Φ in (5) as a function of (μ^f, μ^g), holding the outsiders' markups fixed, so that $\partial\Phi/\partial\mu^i = 1/\mu^i - 1 + Q^i$. For $t \in [0,1]$, let $\boldsymbol{\mu}(t)$ maximize $\Phi(\mu^f, \mu^g) - t(\gamma^f\mu^f + \gamma^g\mu^g)$. This function is strictly concave and tends to $-\infty$ as either markup tends to zero or to infinity, because $-\log H \le -\log\tau^0$ and $1 + t\gamma^i > 0$ since $\gamma^i > -1$. Hence $\boldsymbol{\mu}(t)$ is unique, positive, and continuously differentiable, and it solves $\partial\Phi/\partial\mu^i = t\gamma^i$ for $i \in \{f,g\}$. Thus $\boldsymbol{\mu}(0) = (\mu^{f*}, \mu^{g*})$, and $\boldsymbol{\mu}(1) = (\breve{\mu}^f, \breve{\mu}^g)$ by the definition of γ^i.

Along this path, the CS-index increases in t if $\ell^{fg*} > 1$ and decreases if $\ell^{fg*} < 1$. Let $\nabla^2\Phi$ denote the Hessian of Φ with respect to (μ^f, μ^g):

$$-\nabla^2\Phi = \begin{pmatrix} (\mu^f)^{-2} + Q^f(1-Q^f) & Q^{fg} - Q^fQ^g \\ Q^{fg} - Q^fQ^g & (\mu^g)^{-2} + Q^g(1-Q^g) \end{pmatrix}.$$

Differentiating $\partial\Phi/\partial\mu^i = t\gamma^i$ with respect to t gives $\nabla^2\Phi\,\boldsymbol{\mu}'(t) = (\gamma^f, \gamma^g)^\top$, and $\mathrm{d}\log H/\mathrm{d}\,t = -Q^f\mu^{f\prime}(t) - Q^g\mu^{g\prime}(t)$. Solving for $\boldsymbol{\mu}'(t)$,

$$\det\big(-\nabla^2\Phi\big)\frac{\mathrm{d}\log H}{\mathrm{d}\,t} = \gamma^f\left[\frac{Q^f}{(\mu^g)^2} + Q^g\big(Q^f - Q^{fg}\big)\right] + \gamma^g\left[\frac{Q^g}{(\mu^f)^2} + Q^f\big(Q^g - Q^{fg}\big)\right],$$

with all quantities other than γ^f, γ^g evaluated at $\boldsymbol{\mu}(t)$. The determinant is positive by strict concavity, and both brackets are positive because $Q^i > Q^{fg}$. Hence $\mathrm{d}\log H/\mathrm{d}\,t$ has the sign of $\ell^{fg*} - 1$ at every t, and integrating from $t = 0$ to $t = 1$ proves equation (9). □

A.3.3 Merging firms' markup changes in Step 1

The following lemma makes precise the description of Step-1 markup changes in the main text after equation (8).

Lemma A.2 (Markup changes under joint profit maximization)**.** *If $\ell^{fg*} < 1$, then $\breve{\mu}^f > \mu^{f*}$ and $\breve{\mu}^g > \mu^{g*}$. If instead $\ell^{fg*} > 1$, then $\breve{\mu}^i < \mu^{i*}$ for at least one $i \in \{f,g\}$, and both markups fall if the merging firms and the Step-1 optimum are symmetric. In the latter case, if $Q^{f*} > Q^{g*}$, then $\breve{\mu}^g < \mu^{g*}$; hence only firm f's markup can rise. There exist type networks for which $\breve{\mu}^f > \mu^{f*}$.*

Proof. Suppose first that $\tau^f = 0$ or $\tau^g = 0$. By the proof of Lemma 2, $\ell^{fg*} > 1$ and every Step-1 optimum induces the Step-2 prices. If exactly one of τ^f and τ^g is zero, say $\tau^i = 0$, then $\breve{\mu}^i = 0 < \mu^{i*}$ and $Q^{i*} = Q^{fg*} < Q^{j*}$, which proves the claims for $\ell^{fg*} > 1$. If $\tau^f = \tau^g = 0$, the merging firms' products are purchased only together, so $Q^{f*} = Q^{g*} = Q^{fg*}$, and Step 2 sets the single markup $\widetilde{\mu}^M$ on these purchases. Holding the outsiders' markups fixed, $s\big(1 - Q^{fg}(s)\big)$ is strictly increasing in the total markup s on these purchases. Since $\widetilde{\mu}^M(1 - \widetilde{Q}^{fg}) = 1$ and $(\mu^{f*} + \mu^{g*})(1 - Q^{fg*}) = 2$, we obtain $\breve{\mu}^f + \breve{\mu}^g = \widetilde{\mu}^M < \mu^{f*} + \mu^{g*}$. Hence at least one markup falls, and both fall at a symmetric optimum. We assume henceforth that $\tau^f, \tau^g > 0$.

By the proof of Lemma 2, both Step-1 markups are positive, and by equation (A.3), $\ell^{fg} - 1$ has the sign of $\ell^{fg*} - 1$ at every markup vector. The Step-1 first-order condition is

$$\frac{\partial \Pi^i}{\partial \mu^i} = -\frac{\partial \Pi^j}{\partial \mu^i}, \quad \{i, j\} = \{f, g\}. \tag{A.17}$$

Firm i's pre-merger best-response markup solves $1/\mu^i - 1 + Q^i(\boldsymbol{\mu}) = 0$. The implicit function theorem gives

$$\frac{\mathrm{d}\,\mu^i}{\mathrm{d}\,\mu^j} = \frac{Q^i Q^j (1 - \ell^{ij})}{1/(\mu^i)^2 + Q^i(1 - Q^i)}. \tag{A.18}$$

If $\ell^{fg*} < 1$, equation (A.17) implies that each Step-1 markup exceeds the corresponding pre-merger best response to the other firm's Step-1 markup, while equation (A.18) shows that these best responses are increasing. Starting from the Step-1 markups, best-response iteration therefore gives a decreasing sequence. Every best-response markup exceeds one, so the sequence converges. By continuity, its limit is a fixed point; uniqueness identifies it with the pre-merger equilibrium. Hence $\breve{\mu}^i > \mu^{i*}$ for both merging firms.

If $\ell^{fg*} > 1$, each Step-1 markup is below the corresponding pre-merger best response, and the best responses are decreasing in the other merging firm's markup. If $\breve{\mu}^i \geq \mu^{i*}$ for both merging firms, firm i's best response to the other firm's Step-1 markup is no larger than μ^{i*}. This contradicts $\breve{\mu}^i$ being strictly below that best response. At least one markup therefore falls. Symmetry makes the same conclusion apply to both merging firms.

Now suppose that $\ell^{fg*} > 1$ and $Q^{f*} > Q^{g*}$, and let $\Delta^i \equiv \breve{\mu}^i - \mu^{i*}$. Solving the two Step-1 first-order conditions gives

$$\frac{\breve{\mu}^f}{\breve{\mu}^g} = \frac{1 + \breve{Q}^{fg}/(\breve{Q}^g - \breve{Q}^{fg})}{1 + \breve{Q}^{fg}/(\breve{Q}^f - \breve{Q}^{fg})}.$$

Logit demand implies

$$\frac{\breve{Q}^{fg}}{\breve{Q}^g - \breve{Q}^{fg}} = \frac{Q^{fg*}}{Q^{g*} - Q^{fg*}} \mathrm{e}^{-\Delta^f}, \quad \frac{\breve{Q}^{fg}}{\breve{Q}^f - \breve{Q}^{fg}} = \frac{Q^{fg*}}{Q^{f*} - Q^{fg*}} \mathrm{e}^{-\Delta^g}.$$

Suppose, to the contrary, that $\Delta^g \geq 0$. The preceding argument then gives $\Delta^f < 0$, and hence

$$\frac{\breve{\mu}^f}{\breve{\mu}^g} > \frac{Q^{g*}(Q^{f*} - Q^{fg*})}{Q^{f*}(Q^{g*} - Q^{fg*})}.$$

Since $\mu^{i*} = 1/(1 - Q^{i*})$,

$$\frac{Q^{g*}(Q^{f*} - Q^{fg*})}{Q^{f*}(Q^{g*} - Q^{fg*})} - \frac{\mu^{f*}}{\mu^{g*}} = \frac{(Q^{f*} - Q^{g*})(Q^{fg*} - Q^{f*}Q^{g*})}{Q^{f*}(Q^{g*} - Q^{fg*})(1 - Q^{f*})} > 0,$$

because $Q^{f*} > Q^{g*}$ and $\ell^{fg*} > 1$. This contradicts $\Delta^f < 0 \leq \Delta^g$, which implies $\breve{\mu}^f/\breve{\mu}^g < \mu^{f*}/\mu^{g*}$. Therefore $\breve{\mu}^g < \mu^{g*}$.

Finally, we construct a type network in which firm f's markup rises. Let $\{f, g\} = \{1, 2\}$ and consider a market with no outsiders and

$$(T^0, T^1, T^2, T^{12}) = (1, 1, \varepsilon, 1), \quad \varepsilon > 0.$$

The lift exceeds one for every vector of markups whenever $\varepsilon < 1$. Consider first $\varepsilon = 0$. The only purchases involving firm 2 then also involve firm 1, so Step 1 can choose independently the markups on purchases involving only firm 1 and on purchases involving both firms. The common markup property gives

$$\breve{\mu}^1 = \mu^\dagger, \quad \breve{\mu}^2 = 0, \quad (\mu^\dagger - 1)\mathrm{e}^{\mu^\dagger} = 2.$$

At $\varepsilon = 0$, firm 1's pre-merger first-order condition gives

$$(\mu^{1*} - 1)\mathrm{e}^{\mu^{1*}} = 1 + \mathrm{e}^{-\mu^{2*}} < 2.$$

Since $(u-1)\mathrm{e}^u$ is strictly increasing for $u > 1$, $\mu^\dagger > \mu^{1*}$. Moreover,

$$Q^{1*} - Q^{2*} = \frac{\mathrm{e}^{-\mu^{1*}}}{H^*} > 0$$

at $\varepsilon = 0$. In terms of the independently chosen markups $x = \mu^1$ and $y = \mu^1 + \mu^2$, the limiting Step-1 problem is the standard two-alternative logit monopoly problem. Its unique global maximizer is $(x, y) = (\mu^\dagger, \mu^\dagger)$, and its value strictly exceeds the value obtained when either alternative is removed. For $\varepsilon > 0$, the term for purchases involving only firm 2 has coefficient ε, and its positive contribution to the profit numerator is bounded by ε/e. Because the limiting maximum strictly exceeds the value obtained when either alternative is removed, this bound implies that every sequence of global Step-1 maximizers as $\varepsilon \downarrow 0$ remains in a common compact set. Hence global maximizers converge to $(\mu^\dagger, \mu^\dagger)$ in (x, y), or $(\mu^\dagger, 0)$ in (μ^1, μ^2), as $\varepsilon \downarrow 0$. The pre-merger equilibrium also varies continuously at zero.

For $\varepsilon > 0$, both firms have positive demand from purchases that do not involve the other firm. Let $R \equiv 1 + (\Pi^1 + \Pi^2)/\lambda > 0$. The Step-1 first-order conditions are

$$\breve{\mu}^1 + \frac{\breve{Q}^{12}}{\breve{Q}^1}\breve{\mu}^2 = R, \quad \breve{\mu}^2 + \frac{\breve{Q}^{12}}{\breve{Q}^2}\breve{\mu}^1 = R.$$

The optimum earns positive profit, and both demand ratios lie strictly between zero and one. Solving these equations therefore gives positive markups. It follows that $Q^{1*} > Q^{2*}$ and $\breve{\mu}^1 > \mu^{1*}$ for all sufficiently small $\varepsilon > 0$. □

A.3.4 Proof of Lemma 3

At the pre-merger equilibrium, firm i's first-order condition gives

$$\Pi^{i*} = \lambda \mu^{i*} Q^{i*} = \lambda(\mu^{i*} - 1).$$

The merged firm can reproduce all pre-merger prices and discounts in Step 2 by setting every newly available discount between products previously supplied by distinct merging firms equal to zero. Hence

$$\widetilde{\Pi}^M \geq \Pi^{f*} + \Pi^{g*}.$$

Combining these two equations with equation (A.16) gives

$$\widetilde{\mu}^M \geq \mu^{f*} + \mu^{g*} - 1.$$

Positive pre-merger demand implies $\mu^{f*}, \mu^{g*} > 1$, so the right-hand side, $\mu^{f*} + \mu^{g*} - 1$, strictly exceeds both pre-merger markups. This proves the lemma.

□

A.3.5 Proof of Lemma 4

Let $x \equiv Q^{fg*}$, so that $\ell^{fg*} = x/(Q^{f*}Q^{g*})$. We show that the sign of $\widetilde{H} - H^*$ equals the sign of a strictly increasing function $F(x)$ that is negative at $x = Q^{f*}Q^{g*}$, where $\ell^{fg*} = 1$, and positive at the largest feasible x.

Recall that $\mu^{i*} = 1/(1 - Q^{i*})$. The pre-merger probabilities that the chosen pair contains neither firm, only firm f, only firm g, and both firms are, respectively,

$$q^{-M}(x) \equiv 1 - Q^{f*} - Q^{g*} + x, \quad Q^{f*} - x, \quad Q^{g*} - x, \quad x.$$

Define

$$W(x) \equiv (Q^{f*} - x)\mathrm{e}^{\mu^{f*}} + (Q^{g*} - x)\mathrm{e}^{\mu^{g*}} + x\mathrm{e}^{\mu^{f*}+\mu^{g*}}.$$

Holding the outsiders' markups fixed and letting μ be a candidate markup for the merged firm gives

$$\frac{H(\mu)}{H^*} = q^{-M}(x) + W(x)\mathrm{e}^{-\mu}.$$

The probability that the chosen pair contains a product supplied by the merged firm is $1 - q^{-M}(x)H^*/H(\mu)$. The merged firm's first-order condition therefore gives

$$\frac{\widetilde{H}}{H^*} = \widetilde{\mu}^M q^{-M}(x), \quad W(x) = q^{-M}(x)(\widetilde{\mu}^M - 1)\mathrm{e}^{\widetilde{\mu}^M}.$$

Since $(\mu - 1)\mathrm{e}^{\mu}$ is strictly increasing for $\mu > 1$, these identities imply

$$\operatorname{sign}(\widetilde{H} - H^*) = \operatorname{sign} F(x), \quad F(x) \equiv W(x) - (1 - q^{-M}(x))\mathrm{e}^{1/q^{-M}(x)}.$$

The function F is strictly increasing. Indeed,

$$W'(x) = \mathrm{e}^{\mu^{f*}+\mu^{g*}} - \mathrm{e}^{\mu^{f*}} - \mathrm{e}^{\mu^{g*}} > 0,$$

because $\mu^{f*}, \mu^{g*} > 1$, while, for $R(q) \equiv (1 - q)\mathrm{e}^{1/q}$,

$$R'(q) = -\mathrm{e}^{1/q}\left(1 + \frac{1 - q}{q^2}\right) < 0.$$

At $x = Q^{f*}Q^{g*}$, which corresponds to $\ell^{fg*} = 1$,

$$q^{-M}(x) = (1 - Q^{f*})(1 - Q^{g*}) = \frac{1}{\mu^{f*}\mu^{g*}}.$$

Let $\psi(u) \equiv 1 + (u-1)\mathrm{e}^u$. Direct substitution gives

$$W(Q^{f*}Q^{g*}) = \frac{\psi(\mu^{f*})\psi(\mu^{g*}) - 1}{\mu^{f*}\mu^{g*}},$$
$$(1 - q^{-M}(Q^{f*}Q^{g*}))\mathrm{e}^{1/q^{-M}(Q^{f*}Q^{g*})} = \frac{\psi(\mu^{f*}\mu^{g*}) - 1}{\mu^{f*}\mu^{g*}}.$$

For any $u, v > 1$, $\psi(uv) > \psi(u)\psi(v)$. To prove this inequality, fix $u > 1$ and let $D(v) = \psi(uv) - \psi(u)\psi(v)$. We have $D(1) = 0$ and

$$D'(v) = v\mathrm{e}^v \left[u^2\mathrm{e}^{(u-1)v} - 1 - (u-1)\mathrm{e}^u\right].$$

The expression in brackets is larger than

$$u^2\mathrm{e}^{u-1} - 1 - (u-1)\mathrm{e}^u > 0.$$

The last expression equals zero at $u = 1$ and has derivative

$$u\mathrm{e}^{u-1}(u + 2 - \mathrm{e}) > 0,$$

for $u > 1$, because $u + 2 - \mathrm{e} > 3 - \mathrm{e} > 0$. Hence $D'(v) > 0$ and the inequality follows. Therefore, $F(Q^{f*}Q^{g*}) < 0$.

At the largest feasible value $x = \min\{Q^{f*}, Q^{g*}\}$, suppose without loss of generality that $Q^{f*} \leq Q^{g*}$. Then $q^{-M}(x) = 1 - Q^{g*} = 1/\mu^{g*}$ and

$$W(x) = \mathrm{e}^{\mu^{g*}}\left[Q^{g*} + Q^{f*}(\mathrm{e}^{\mu^{f*}} - 1)\right] > Q^{g*}\mathrm{e}^{\mu^{g*}} = (1 - q^{-M}(x))\mathrm{e}^{1/q^{-M}(x)}.$$

Thus $F(x) > 0$ at the upper bound. Continuity and strict monotonicity imply a unique cutoff $\widehat{x} \in (Q^{f*}Q^{g*}, \min\{Q^{f*}, Q^{g*}\})$. Since $\ell^{fg*} = x/(Q^{f*}Q^{g*})$, the cutoff $\widehat{\ell}(Q^{f*}, Q^{g*}) = \widehat{x}/(Q^{f*}Q^{g*})$ has the properties stated in the lemma. □

A.3.6 Proof of Proposition 5

We first show by contradiction that $\overline{\mu}^{M*} > \min\{\mu^{f*}, \mu^{g*}\}$, and then that some consumers are strictly worse off. After introducing notation, we suppose the opposite inequality and derive three claims about demand changes, which are mutually inconsistent.

In this proof, we omit equilibrium asterisks. Let $\mathcal{O} = \mathcal{N} \setminus \{f, g\}$. Since the maintained restriction on the type network holds both before and after the merger, every firm is absent from at least one available product pair in each case, so all equilibrium markups are finite. We continue to index each product pair by its pre-merger suppliers. Before the merger, define

$$q^0 \equiv \frac{T^0}{H}, \quad q^i \equiv \frac{T^i\mathrm{e}^{-\mu^i}}{H}, \quad Q^{ij} \equiv \frac{T^{ij}\mathrm{e}^{-\mu^i-\mu^j}}{H}.$$

Here q^i is the probability of choosing a product pair supplied only by firm i, and $Q^{ij} = Q^{ji}$ is the probability of choosing a pair containing products from firms i and j. After the merger,

define the probabilities associated with the same pre-merger types by

$$\overline{q}^0 \equiv \frac{T^0}{\overline{H}}, \qquad \overline{Q}^{fg} \equiv \frac{T^{fg}\mathrm{e}^{-\overline{\mu}^M}}{\overline{H}},$$
$$\overline{q}^i \equiv \frac{T^i\mathrm{e}^{-\overline{\mu}^M}}{\overline{H}} \quad (i \in \{f,g\}), \qquad \overline{q}^k \equiv \frac{T^k\mathrm{e}^{-\overline{\mu}^k}}{\overline{H}} \quad (k \in \mathcal{O}),$$
$$\overline{Q}^{ik} \equiv \frac{T^{ik}\mathrm{e}^{-\overline{\mu}^M-\overline{\mu}^k}}{\overline{H}} \quad (i \in \{f,g\},\ k \in \mathcal{O}), \qquad \overline{Q}^{kl} \equiv \frac{T^{kl}\mathrm{e}^{-\overline{\mu}^k-\overline{\mu}^l}}{\overline{H}} \quad (k,l \in \mathcal{O},\ k \neq l).$$

Let $\overline{Q}^{ki} \equiv \overline{Q}^{ik}$. For every pre-merger firm i, let

$$Q^i \equiv q^i + \sum_{j \in \mathcal{N}\setminus\{i\}} Q^{ij}, \quad \overline{Q}^i \equiv \overline{q}^i + \sum_{j \in \mathcal{N}\setminus\{i\}} \overline{Q}^{ij}. \tag{A.19}$$

For an outsider, $\overline{Q}^k$ is its post-merger demand. For $i \in \{f,g\}$, $\overline{Q}^i$ is instead the post-merger probability that the chosen pair contains a product that firm i supplied before the merger.

Define markup changes and the change in consumer surplus by

$$\Delta^i \equiv \mu^i - \overline{\mu}^M \quad (i \in \{f,g\}), \quad \Delta^k \equiv \mu^k - \overline{\mu}^k \quad (k \in \mathcal{O}), \quad \Delta\mathrm{CS} \equiv \lambda \log \frac{\overline{H}}{H}.$$

Whenever the corresponding type is positive, the ratios of post-merger to pre-merger probabilities are

$$\frac{\overline{q}^i}{q^i} = \mathrm{e}^{\Delta^i - \Delta\mathrm{CS}/\lambda}, \quad i \in \mathcal{N}, \tag{A.20}$$

$$\frac{\overline{Q}^{ij}}{Q^{ij}} = \begin{cases} \mathrm{e}^{\mu^f+\mu^g-\overline{\mu}^M-\Delta\mathrm{CS}/\lambda} & \text{if } \{i,j\} = \{f,g\}, \\ \mathrm{e}^{\Delta^i+\Delta^j-\Delta\mathrm{CS}/\lambda} & \text{otherwise.} \end{cases} \tag{A.21}$$

Suppose, toward a contradiction, that

$$\overline{\mu}^M \leq \min\{\mu^f, \mu^g\}, \tag{A.22}$$

so $\Delta^f, \Delta^g \geq 0$.

Claim 1. The first-order conditions imply

$$\overline{Q}^i < Q^i, \quad i \in \{f,g\}, \quad \overline{Q}^k - Q^k = -\frac{\Delta^k}{\mu^k\overline{\mu}^k}, \quad k \in \mathcal{O}.$$

Proof of Claim 1. Every product pair counted in $\overline{Q}^i$ is also counted in $\overline{Q}^M$. For the other merging firm j, the merged firm's demand additionally counts the product pairs supplied only by firm j, which are available because $T^j > 0$. Hence $\overline{Q}^i < \overline{Q}^M$. Using the first-order conditions and inequality (A.22),

$$\overline{Q}^i < \overline{Q}^M = 1 - \frac{1}{\overline{\mu}^M} \leq 1 - \frac{1}{\mu^i} = Q^i.$$

For an outsider k, subtracting its two first-order conditions gives

$$\overline{Q}^k - Q^k = \left(1 - \frac{1}{\overline{\mu}^k}\right) - \left(1 - \frac{1}{\mu^k}\right) = -\frac{\Delta^k}{\mu^k \overline{\mu}^k}.$$

This proves Claim 1.

Claim 2. Under inequality (A.22), it is impossible that every outsider weakly lowers its markup.

Proof of Claim 2. Suppose $\Delta^k \geq 0$ for every outsider. The total markup on every product pair other than 00 then weakly falls, so $\overline{H} \geq H$ and $\Delta\text{CS} \geq 0$.

We first show that $\overline{q}^l \leq q^l$ for every pre-merger firm l. The conclusion is immediate if $T^l = 0$. Otherwise, if $\overline{q}^l > q^l$, equation (A.20) gives $\Delta^l > \Delta\text{CS}/\lambda$. For a pair containing products from l and another firm i, the reduction in total markup is $\Delta^l + \Delta^i > \Delta\text{CS}/\lambda$ unless $\{i, l\} = \{f, g\}$. In that case, the reduction is

$$\mu^i + \mu^l - \overline{\mu}^M = \mu^i + \Delta^l > \frac{\Delta\text{CS}}{\lambda}.$$

Thus every product pair counted in $\overline{Q}^l$ has weakly higher probability, and the product pairs represented by T^l have strictly higher probability. Hence $\overline{Q}^l > Q^l$. This contradicts Claim 1 if $l \in \{f, g\}$; if $l \in \mathcal{O}$, it contradicts Claim 1 because $\Delta^l \geq 0$ implies $\overline{Q}^l \leq Q^l$. Therefore $\overline{q}^l \leq q^l$ for every l.

Similarly, $\overline{q}^0 \leq q^0$; when $T^0 = 0$, both probabilities are zero. The sum $\sum_{l \in \mathcal{N}} Q^l$ is the expected number of pre-merger firms supplying the chosen product pair. Since all choice probabilities sum to one,

$$\sum_{l \in \mathcal{N}} (\overline{Q}^l - Q^l) = 2(q^0 - \overline{q}^0) + \sum_{l \in \mathcal{N}} (q^l - \overline{q}^l) \geq 0.$$

Claim 1 instead makes every term for a merging firm strictly negative and every outsider's term weakly negative. Their sum must be negative, a contradiction. This proves Claim 2.

If $\mathcal{O} = \varnothing$, Claim 2 already contradicts inequality (A.22) and proves $\overline{\mu}^M > \min\{\mu^f, \mu^g\}$. Hence suppose $\mathcal{O} \neq \varnothing$. Claim 2 implies that some outsider has $\Delta^k < 0$.

Claim 3. Choose an outsider k with the smallest Δ^k. Then $\overline{Q}^k \leq Q^k$.

Proof of Claim 3. Suppose instead that $\overline{Q}^k > Q^k$. By equation (A.19), either $\overline{q}^k > q^k$ or $\overline{Q}^{kl} > Q^{kl}$ for some firm l.

First suppose that $\overline{q}^k > q^k$, or that $\overline{Q}^{kl} > Q^{kl}$ for a firm l with $\Delta^l < 0$. The ratios above and the minimality of Δ^k imply

$$\frac{\Delta\text{CS}}{\lambda} < \Delta^k < 0.$$

For $i \in \{f, g\}$, we have $\overline{q}^i > q^i$ because $T^i > 0$ and $\Delta^i \geq 0 > \Delta\text{CS}/\lambda$. For every outsider h, the minimality of Δ^k gives

$$\Delta^i + \Delta^h \geq \Delta^k > \frac{\Delta\text{CS}}{\lambda}.$$

Thus every available pair containing products from firm i and an outsider has higher probability.

The fg pair also has higher probability because

$$\mu^f + \mu^g - \overline{\mu}^M = \Delta^f + \mu^g > 0 > \frac{\Delta \text{CS}}{\lambda}.$$

Consequently, $\overline{Q}^f > Q^f$ and $\overline{Q}^g > Q^g$, contradicting Claim 1.

It remains to consider $\overline{Q}^{kl} > Q^{kl}$ for a firm l with $\Delta^l \geq 0$. Equation (A.21) gives

$$\Delta^l > \frac{\Delta \text{CS}}{\lambda} - \Delta^k > \frac{\Delta \text{CS}}{\lambda}.$$

Every other pre-merger firm i satisfies $\Delta^i \geq \Delta^k$. If i and l are not both merging firms, the reduction in their pair's total markup satisfies

$$\Delta^l + \Delta^i \geq \Delta^l + \Delta^k > \frac{\Delta \text{CS}}{\lambda}.$$

If $\{i, l\} = \{f, g\}$, the reduction is

$$\mu^i + \mu^l - \overline{\mu}^M = \mu^i + \Delta^l > \frac{\Delta \text{CS}}{\lambda}.$$

Hence every available pair containing products from firm l has higher probability; the pair containing products from k and l has strictly higher probability. We also have $\overline{q}^l > q^l$ whenever $T^l > 0$. It follows that $\overline{Q}^l > Q^l$. This contradicts Claim 1 whether l is a merging firm or an outsider. This proves Claim 3.

Claim 2 supplies an outsider k with $\Delta^k < 0$. Claim 1 then gives $\overline{Q}^k > Q^k$, whereas Claim 3 gives $\overline{Q}^k \leq Q^k$. This contradiction proves

$$\overline{\mu}^M > \min\{\mu^f, \mu^g\}.$$

To show that some consumers are strictly worse off, let $i \in \{f, g\}$ satisfy $\mu^i = \min\{\mu^f, \mu^g\}$. Since $T^i > 0$, at least one available product pair is supplied only by firm i before the merger. Its total markup rises from μ^i to $\overline{\mu}^M$, so its deterministic utility falls. Because the menu is finite and the utility shocks have full support, there is an open set of shock realizations for which this pair is uniquely optimal both before and after the merger. This set has positive probability, and every consumer in it is strictly worse off. □

A.3.7 Proof of Proposition 6

We show that $\widetilde{H} > H^*$ implies $\overline{H}^* \geq \widetilde{H}$, with strict inequality if $\mathcal{O} \neq \emptyset$. The proof uses the following lemma, which relates the outsiders' demand changes to the change in the CS-index from the end of Step 2 to the post-merger equilibrium. Let $\widetilde{Q}^k$ denote outsider k's demand at the end of Step 2.

Lemma A.3 (Outsider demand changes and consumer surplus)**.** *Suppose that $\mathcal{O} \neq \emptyset$. If*

$$\widetilde{Q}^k \leq Q^{k*} \quad \textit{for every } k \in \mathcal{O},$$

then $\overline{H}^* \geq \widetilde{H}$. If $\widetilde{Q}^k \geq Q^{k*}$ for every $k \in \mathcal{O}$, then $\overline{H}^* \leq \widetilde{H}$. In either case, the inequality is

strict if at least one outsider's demand changes strictly.

Proof. We connect the Step-2 markups to the post-merger equilibrium by a path of markup vectors, each maximizing the post-merger potential adjusted by a linear term, and we sign the derivative of $\log \overline{H}$ along this path.

Let $\widetilde{\boldsymbol{\mu}}$ denote the markup vector at the end of Step 2:

$$\widetilde{\boldsymbol{\mu}} \equiv \left(\widetilde{\mu}^M, (\mu^{k*})_{k\in\mathcal{O}}\right).$$

The merged firm's Step-2 first-order condition and the outsiders' pre-merger first-order conditions imply

$$\nabla\overline{\Phi}(\widetilde{\boldsymbol{\mu}}) = \mathbf{g} \equiv \left(0, (\widetilde{Q}^k - Q^{k*})_{k\in\mathcal{O}}\right)^\top, \tag{A.23}$$

where $\overline{\Phi}$ is the post-merger potential.

For $t \in [0,1]$, let $\boldsymbol{\mu}(t)$ maximize

$$\overline{\Phi}(\boldsymbol{\mu}) - t\mathbf{g}^\top\boldsymbol{\mu}. \tag{A.24}$$

At $t = 1$, $\widetilde{\boldsymbol{\mu}}$ is the unique maximizer because equation (A.23) is the first-order condition for this problem. The objective for any $t \in (0,1)$ is the weighted average of $\overline{\Phi}$ and $\overline{\Phi} - \mathbf{g}^\top\boldsymbol{\mu}$. The first function tends to minus infinity when a markup tends to zero, because of the terms $\log\mu^i$, or when the markup vector becomes unbounded, as shown in the proof of Theorem 1, and the second is bounded above by its value at $\widetilde{\boldsymbol{\mu}}$. The objective in (A.24) therefore has a finite interior maximizer. Strict concavity makes it unique. Equation (A.23) implies

$$\boldsymbol{\mu}(0) = \overline{\boldsymbol{\mu}}^*, \quad \boldsymbol{\mu}(1) = \widetilde{\boldsymbol{\mu}}, \quad \nabla\overline{\Phi}(\boldsymbol{\mu}(t)) = t\mathbf{g}. \tag{A.25}$$

The Hessian is negative definite, so the implicit function theorem implies that $\boldsymbol{\mu}(t)$ is continuously differentiable.

For the product pair chosen by a consumer, let I^h equal one if the pair contains a product supplied by post-merger firm h and zero otherwise. Let $\mathbf{I} \equiv (I^h)_{h\in\overline{\mathcal{N}}}$ and $\mathbf{Q}(t) \equiv \mathbb{E}_t[\mathbf{I}]$. The negative Hessian of the potential along the path is

$$J(t) \equiv -\nabla^2\overline{\Phi}(\boldsymbol{\mu}(t)) = \operatorname{diag}\left(\frac{1}{(\mu^h(t))^2}\right)_{h\in\overline{\mathcal{N}}} + \operatorname{Cov}_t(\mathbf{I}).$$

Differentiating equation (A.25) gives

$$\boldsymbol{\mu}'(t) = -J(t)^{-1}\mathbf{g}.$$

Since $\nabla \log\overline{H} = -\mathbf{Q}$,

$$\frac{\mathrm{d}}{\mathrm{d}\,t}\log\overline{H}(\boldsymbol{\mu}(t)) = \mathbf{Q}(t)^\top J(t)^{-1}\mathbf{g}. \tag{A.26}$$

Every post-merger firm supplies at least one product that can be purchased, so $\mathbf{Q}(t) > \mathbf{0}$. Lemma A.1 therefore implies that $J(t)^{-1}\mathbf{Q}(t)$ is componentwise strictly positive. If $\widetilde{Q}^k \le Q^{k*}$ for every $k \in \mathcal{O}$, then $\mathbf{g} \le \mathbf{0}$, so the right-hand side of (A.26) is weakly negative and is strictly negative if the inequality is strict for at least one outsider. Integrating from $t = 0$ to $t = 1$ gives $\overline{H}^* \ge \widetilde{H}$, with strict inequality in the latter case. Reversing the inequalities gives the

other conclusion.

The proof also gives the exact representation

$$\lambda \log \frac{\overline{H}^*}{\widetilde{H}} = -\sum_{k \in \mathcal{O}} a_k (\widetilde{\Pi}^k - \Pi^{k*}), \quad a_k \equiv \frac{1}{\mu^{k*}} \int_0^1 [J(t)^{-1} \mathbf{Q}(t)]_k \, \mathrm{d}\, t > 0,$$

where we have used $\widetilde{\Pi}^k - \Pi^{k*} = \lambda \mu^{k*} (\widetilde{Q}^k - Q^{k*})$, which follows because outsider k's markup is unchanged through Step 2. □

If $\mathcal{O} = \emptyset$, Step 2 is already the post-merger equilibrium. Suppose therefore that $\mathcal{O} \neq \emptyset$, and fix an outsider k. The Step-2 logit weight of a product pair containing a product supplied by k but no product supplied by a merging firm is unchanged. If the pair also contains a product previously supplied by merging firm i, its price rises by $\lambda(\widetilde{\mu}^M - \mu^{i*}) > 0$, by Lemma 3, so its logit weight falls. The total logit weight of pairs containing a product supplied by k therefore weakly falls, while the total weight of all pairs rises strictly from H^* to $\widetilde{H}$. Because Q^k is the ratio of these two weights, $\widetilde{Q}^k < Q^{k*}$. Lemma A.3 then gives $\overline{H}^* > \widetilde{H}$. □

A.3.8 Proof of Proposition 7

We show that, along a family of pre-merger equilibria with the same $(Q^{f*}, Q^{g*}, \mathbf{q}^{\mathcal{O}*})$, the ratio $\overline{H}^*/H^*$ is strictly increasing in ℓ^{fg*}. The cutoff then follows.

Parameterize the family by $x \equiv Q^{fg*}$, so that $\ell^{fg*} = x/(Q^{f*}Q^{g*})$, and let $q^{0*} \equiv T^0/H^*$. Since $Q^{i*} = q^{i*} + \sum_{j \neq i} Q^{ij*}$ for every firm i and the choice probabilities sum to one, holding Q^{f*}, Q^{g*}, and every probability involving an outsider fixed gives

$$\frac{\mathrm{d}\, q^{0*}}{\mathrm{d}\, x} = 1, \quad \frac{\mathrm{d}\, q^{f*}}{\mathrm{d}\, x} = \frac{\mathrm{d}\, q^{g*}}{\mathrm{d}\, x} = -1.$$

Every firm's demand is fixed, so by equation (1) every pre-merger markup $\mu^{i*} = 1/(1 - Q^{i*})$ is fixed along the family.

The pre-merger equilibrium determines the type network up to scale:

$$\frac{T^0}{H^*} = q^{0*}, \quad \frac{T^i}{H^*} = q^{i*} \mathrm{e}^{\mu^{i*}}, \quad \frac{T^{ij}}{H^*} = Q^{ij*} \mathrm{e}^{\mu^{i*} + \mu^{j*}}.$$

Multiplying every type by the same positive constant multiplies H by that constant at every markup vector and leaves all choice probabilities, and hence all equilibrium markups, unchanged. The post-merger network $\overline{\mathbf{T}}/H^*$ therefore has equilibrium CS-index $\overline{H}^*/H^*$. In this network, $\overline{T}^0/H^* = q^{0*}$ and

$$\frac{\overline{T}^M}{H^*} = q^{f*} \mathrm{e}^{\mu^{f*}} + q^{g*} \mathrm{e}^{\mu^{g*}} + x \, \mathrm{e}^{\mu^{f*} + \mu^{g*}},$$

whereas every other type depends only on probabilities and markups that are fixed along the family. Hence

$$\frac{\mathrm{d}}{\mathrm{d}\, x} \frac{\overline{T}^0}{H^*} = 1, \quad \frac{\mathrm{d}}{\mathrm{d}\, x} \frac{\overline{T}^M}{H^*} = \mathrm{e}^{\mu^{f*} + \mu^{g*}} \left(1 - \mathrm{e}^{-\mu^{f*}} - \mathrm{e}^{-\mu^{g*}}\right) > 0,$$

where the inequality holds because $\mu^{f*}, \mu^{g*} > 1$ implies $\mathrm{e}^{-\mu^{f*}} + \mathrm{e}^{-\mu^{g*}} < 2/\mathrm{e} < 1$. By equations (A.10) and (A.8), applied to the post-merger network, the equilibrium CS-index is strictly increasing in $\overline{T}^0$ and in $\overline{T}^M$. The chain rule then shows that $\log(\overline{H}^*/H^*)$ is strictly increasing in x, and hence in ℓ^{fg*}.

Given $(Q^{f*}, Q^{g*}, \mathbf{q}^{\mathcal{O}*})$, let $\ell^\dagger$ be the infimum of the lifts in the family at which the merger raises consumer surplus, or the supremum of the family's lifts if there are none. Strict monotonicity and continuity imply that the merger raises consumer surplus at every lift in the family above $\ell^\dagger$ and lowers it at every lift below $\ell^\dagger$. If $\ell^{fg*} > \widehat{\ell}(Q^{f*}, Q^{g*})$, Lemma 4 gives $\widetilde{H} > H^*$, and Proposition 6 gives $\overline{H}^* > H^*$. Hence the merger raises consumer surplus at every lift in the family above $\widehat{\ell}(Q^{f*}, Q^{g*})$, and $\ell^F \equiv \min\{\ell^\dagger, \widehat{\ell}(Q^{f*}, Q^{g*})\}$ has the properties stated in the proposition. □

A.3.9 Proof of Example 4

We show that $\widetilde{H} < H^*$ for every $x > 0$ and that $\overline{H}^* > H^*$ for all sufficiently large x.

For the types in (10), the pre-merger CS-index factorizes as

$$H = (1 + \mathrm{e}^{-\mu^2})(1 + \mathrm{e}^{-\mu^1} + x\mathrm{e}^{-\mu^3}). \tag{A.27}$$

Firm 2's equilibrium markup is independent of x and solves

$$\mu^{2*} = 1 + \mathrm{e}^{-\mu^{2*}}.$$

The first-order conditions for firms 1 and 3 imply

$$x\mathrm{e}^{-\mu^{3*}} = \frac{\mathrm{e}^{-\mu^{1*}}}{\mu^{1*} - 1} - 1,$$
$$\mu^{3*} = \frac{\mu^{1*}}{(\mu^{1*} - 1)(1 + \mathrm{e}^{\mu^{1*}})}.$$

It follows that $\mu^{1*} \downarrow 1$ and $\mu^{3*} \to \infty$ as $x \to \infty$. Equation (A.27) and the first-order conditions also give

$$H^* = \frac{\mu^{1*}\mu^{2*}\mathrm{e}^{-\mu^{1*}}}{\mu^{1*} - 1} = \mu^{2*}\mu^{3*}(1 + \mathrm{e}^{-\mu^{1*}}). \tag{A.28}$$

The comparison of H^* and $\overline{H}^*$ below uses the inequality

$$\mu^{2*} < 2\mathrm{e}(\mu^{2*} - 1). \tag{A.29}$$

To verify it analytically, note first that $5/4 < \mu^{2*} < 4/3$. Indeed, the increasing function $\mu - 1 - \mathrm{e}^{-\mu}$ is negative at $5/4$ and positive at $4/3$: the first sign follows from the elementary series bounds $\mathrm{e} < 3$ and $\mathrm{e}^{1/4} < 4/3$, and the second from $\mathrm{e}^{4/3} > 1 + 4/3 + (4/3)^2/2 > 3$. Since $\mathrm{e} > 8/3$, these bounds imply

$$2\mathrm{e}(\mu^{2*} - 1) > \frac{4}{3} > \mu^{2*}.$$

This proves inequality (A.29).

We now show that $\widetilde{H} < H^*$. By equation (A.27), demands for firms 1 and 2 are independent, so $\ell^{12*} = 1$. Since the cutoff in Lemma 4 exceeds one, $\widetilde{H} < H^*$ for every $x > 0$.

It remains to show that $\overline{H}^* > H^*$ when x is large. In the full post-merger equilibrium,

$$\overline{\mu}^{M*} = 1 + \frac{3 + x\mathrm{e}^{-\overline{\mu}^{3*}}}{1 + x\mathrm{e}^{-\overline{\mu}^{3*}}}\mathrm{e}^{-\overline{\mu}^{M*}},$$
$$x\mathrm{e}^{-\overline{\mu}^{3*}} = (\overline{\mu}^{3*} - 1)\frac{1 + 3\mathrm{e}^{-\overline{\mu}^{M*}}}{1 + \mathrm{e}^{-\overline{\mu}^{M*}}}.$$

These conditions imply that $\overline{\mu}^{3*} \to \infty$ and $\overline{\mu}^{M*} \to \mu^{2*}$ as $x \to \infty$. The pre- and post-merger conditions for firm 3 can also be written as

$$x = (\mu^{3*} - 1)\mathrm{e}^{\mu^{3*}}(1 + \mathrm{e}^{-\mu^{1*}}),$$
$$x = (\overline{\mu}^{3*} - 1)\mathrm{e}^{\overline{\mu}^{3*}}\frac{1 + 3\mathrm{e}^{-\overline{\mu}^{M*}}}{1 + \mathrm{e}^{-\overline{\mu}^{M*}}}.$$

Taking logarithms shows that both $\mu^{3*}/\log x$ and $\overline{\mu}^{3*}/\log x$ converge to one, and hence $\overline{\mu}^{3*}/\mu^{3*} \to 1$. Finally, firm 3's pre- and post-merger first-order conditions give

$$\overline{H}^* = \overline{\mu}^{3*}(1 + 3\mathrm{e}^{-\overline{\mu}^{M*}}).$$

Combining this identity with equation (A.28) yields

$$\frac{\overline{H}^*}{H^*} \longrightarrow \frac{1 + 3\mathrm{e}^{-\mu^{2*}}}{\mu^{2*}(1 + \mathrm{e}^{-1})} = \frac{3\mu^{2*} - 2}{\mu^{2*}(1 + \mathrm{e}^{-1})} > 1,$$

where the last inequality is inequality (A.29). Thus $H^* < \overline{H}^*$ for all sufficiently large x. Together with $\widetilde{H} < H^*$, shown above, this completes the proof. □

A.3.10 Proof of Proposition 8

We show that $\widetilde{H} > H^*$ when Q^{M*} is small. Proposition 6 then gives $\overline{H}^* > H^*$.

Let $Q \equiv Q^{M*}$. Since $Q^{i*} \leq Q$ for every $i \in \mathcal{M}$, the pre-merger first-order condition gives

$$\mu^{i*} = \frac{1}{1 - Q^{i*}} = 1 + O(Q). \tag{A.30}$$

The Step-2 markup also satisfies $\widetilde{\mu}^M = 1 + O(Q)$. To see this, observe that product pairs containing no product supplied by a merging firm retain their pre-merger logit weights in Step 2, and their total weight is $(1 - Q)H^*$. Moreover, $\widetilde{\mu}^M \geq 1$, and equation (A.30) implies that the ratio of the Step-2 logit weight to the pre-merger logit weight is bounded for every pair containing a product supplied by a merging firm. The total Step-2 weight of these pairs is therefore $O(QH^*)$. Hence $\widetilde{Q}^M = O(Q)$, and the Step-2 first-order condition gives

$$\widetilde{\mu}^M = \frac{1}{1 - \widetilde{Q}^M} = 1 + O(Q). \tag{A.31}$$

Holding outsiders' markups at their pre-merger levels, the Step-2 CS-index can be written in terms of the merging firms' pre-merger demands and the probability of buying from both

of them as

$$\frac{\widetilde{H}}{H^*} = 1 - Q + (Q^{f*} - Q^{fg*})\mathrm{e}^{\mu^{f*} - \widetilde{\mu}^M} + (Q^{g*} - Q^{fg*})\mathrm{e}^{\mu^{g*} - \widetilde{\mu}^M} + Q^{fg*}\mathrm{e}^{\mu^{f*} + \mu^{g*} - \widetilde{\mu}^M}.$$

The coefficients of the second and third terms add up to $Q - Q^{fg*}$. Equations (A.30) and (A.31) therefore imply

$$\begin{aligned}\frac{\widetilde{H}}{H^*} &= 1 - Q + (Q - Q^{fg*})\big(1 + O(Q)\big) + Q^{fg*}\big(\mathrm{e} + O(Q)\big) \\ &= 1 + (\mathrm{e} - 1)Q^{fg*} + O(Q^2). \end{aligned} \tag{A.32}$$

Because Q^{fg*}/Q is bounded away from zero, $Q^{fg*} \geq \rho Q$ for some constant $\rho > 0$. The first-order term in (A.32) therefore dominates the $O(Q^2)$ term when Q is sufficiently small. Hence $\widetilde{H} > H^*$. Proposition 6 then gives $\overline{H}^* \geq \widetilde{H} > H^*$, so the merger raises consumer surplus. □

A.3.11 Proof of Proposition 9

We begin with the post-merger CS-index and demands under the post-merger type network $\overline{\mathbf{T}}$. The CS-index is

$$\overline{H}(\overline{\boldsymbol{\mu}}) = T^0 + \overline{T}^M \mathrm{e}^{-\overline{\mu}^M} + \sum_{k \in \mathcal{O}} T^k \mathrm{e}^{-\overline{\mu}^k} + \sum_{k \in \mathcal{O}} \overline{T}^{Mk} \mathrm{e}^{-\overline{\mu}^M - \overline{\mu}^k} + \frac{1}{2} \sum_{k \in \mathcal{O}} \sum_{l \in \mathcal{O} \setminus \{k\}} T^{kl} \mathrm{e}^{-\overline{\mu}^k - \overline{\mu}^l}.$$

The merged firm's demand and outsider k's demand are

$$\overline{Q}^M(\overline{\boldsymbol{\mu}}) = \frac{\mathrm{e}^{-\overline{\mu}^M}}{\overline{H}(\overline{\boldsymbol{\mu}})} \left[\overline{T}^M + \sum_{k \in \mathcal{O}} \overline{T}^{Mk} \mathrm{e}^{-\overline{\mu}^k} \right],$$

$$\overline{Q}^k(\overline{\boldsymbol{\mu}}) = \frac{\mathrm{e}^{-\overline{\mu}^k}}{\overline{H}(\overline{\boldsymbol{\mu}})} \left[T^k + \overline{T}^{Mk} \mathrm{e}^{-\overline{\mu}^M} + \sum_{l \in \mathcal{O} \setminus \{k\}} T^{kl} \mathrm{e}^{-\overline{\mu}^l} \right].$$

Thus $(1 - \overline{Q}^l)\overline{\mu}^l = 1$ for every $l \in \overline{\mathcal{N}}$.

We show that, as $Q^{M*} \to 1$, the merged firm's markup diverges while purchases involving neither merging firm contribute a vanishing share of the CS-index. Together, these imply $\overline{H}^* < H^*$.

Because q^{f*} and q^{g*} are bounded away from zero, take $\rho > 0$ such that $q^{f*}, q^{g*} \geq \rho$. Every pre-merger firm is absent from at least one of these two sets of purchases. In particular,

$$1 - Q^{h*} \geq \rho \quad \text{for every } h \in \mathcal{N}.$$

The pre-merger first-order conditions therefore imply

$$\mu^{h*} \leq \frac{1}{\rho} \quad \text{for every } h \in \mathcal{N}. \tag{A.33}$$

Let

$$R^* \equiv T^0 + \sum_{k\in\mathcal{O}} T^k \mathrm{e}^{-\mu^{k*}} + \frac{1}{2}\sum_{k\in\mathcal{O}}\sum_{l\in\mathcal{O}\setminus\{k\}} T^{kl}\mathrm{e}^{-\mu^{k*}-\mu^{l*}},$$
$$\overline{R} \equiv T^0 + \sum_{k\in\mathcal{O}} T^k \mathrm{e}^{-\overline{\mu}^{k*}} + \frac{1}{2}\sum_{k\in\mathcal{O}}\sum_{l\in\mathcal{O}\setminus\{k\}} T^{kl}\mathrm{e}^{-\overline{\mu}^{k*}-\overline{\mu}^{l*}}$$

be the pre- and post-merger contributions to the CS-index from product pairs containing no product supplied by a merging firm. By definition,

$$R^* = (1 - Q^{M*})H^*.$$

Every post-merger markup is at least one. Since each product pair contains products from at most two firms, inequality (A.33) gives

$$\frac{\overline{R}}{H^*} \leq \mathrm{e}^{2/\rho}\frac{R^*}{H^*} = \mathrm{e}^{2/\rho}(1 - Q^{M*}). \tag{A.34}$$

Thus $\overline{R}/H^*$ approaches zero as Q^{M*} approaches one.

Define

$$x \equiv \frac{1}{H^*}\left[\overline{T}^M + \sum_{k\in\mathcal{O}} \overline{T}^{Mk}\mathrm{e}^{-\overline{\mu}^{k*}}\right].$$

Because $\overline{T}^M$ includes T^f,

$$x \geq \frac{T^f}{H^*} = q^{f*}\mathrm{e}^{\mu^{f*}} \geq \rho.$$

Moreover, inequality (A.33) and the fact that each pair contains products from at most two firms imply

$$x \leq \frac{\overline{T}^M + \sum_{k\in\mathcal{O}} \overline{T}^{Mk}}{H^*} \leq \mathrm{e}^{2/\rho}. \tag{A.35}$$

Define $r \equiv \overline{R}/H^*$. The post-merger CS-index can be written as

$$\frac{\overline{H}^*}{H^*} = r + x\mathrm{e}^{-\overline{\mu}^{M*}}. \tag{A.36}$$

Because $1 - \overline{Q}^{M*} = \overline{R}/\overline{H}^*$, the merged firm's first-order condition gives

$$\frac{\overline{H}^*}{\overline{H}^*} = \overline{\mu}^{M*}r. \tag{A.37}$$

Combining equations (A.36) and (A.37) yields

$$(\overline{\mu}^{M*} - 1)\mathrm{e}^{\overline{\mu}^{M*}} r = x. \tag{A.38}$$

Inequalities (A.34) and (A.35) and equation (A.38) imply that $\overline{\mu}^{M*}$ diverges as Q^{M*} approaches one. Finally,

$$\frac{\overline{H}^*}{H^*} = x \frac{\overline{\mu}^{M*}}{(\overline{\mu}^{M*} - 1)\mathrm{e}^{\overline{\mu}^{M*}}} \longrightarrow 0.$$

Hence $\overline{H}^* < H^*$ when Q^{M*} is sufficiently close to one, so the merger lowers consumer surplus. □

A.4 Mixed bundling versus separate pricing

A.4.1 Preliminaries

For notational simplicity, throughout the proofs of Section 5 we measure all monetary quantities in units of λ and hence normalize $\lambda = 1$. This is without loss of generality: dividing values, prices, marginal costs, and bundle discounts by λ leaves all choice probabilities unchanged and divides profits, consumer surplus, and aggregate surplus by the same positive factor. Throughout these proofs, superscripts S and B denote outcomes under separate pricing and mixed bundling, respectively.

We impose equation (11). Additive product values make the CS-index and product-pair choice probabilities factor across markets at any separate-pricing profile. Each firm's profit is therefore the sum of its profits in the two single-market logit games. Let $\mathcal{N}_m \equiv \{i \in \mathcal{N} : T_m^i > 0\}$ denote the firms active in market $m \in \{A, B\}$. The standard common-markup result implies that, for each $i \in \mathcal{N}_m$, all of firm i's products in market m carry a common markup μ_m^i.

Using the market-level types defined in equation (12), define

$$H_A(\boldsymbol{\mu}_A) \equiv T_A^0 + \sum_{i \in \mathcal{N}_A} T_A^i \mathrm{e}^{-\mu_A^i},$$
$$H_B(\boldsymbol{\mu}_B) \equiv T_B^0 + \sum_{i \in \mathcal{N}_B} T_B^i \mathrm{e}^{-\mu_B^i}.$$

Under separate pricing,

$$H(\boldsymbol{\mu}_A, \boldsymbol{\mu}_B) = H_A(\boldsymbol{\mu}_A) H_B(\boldsymbol{\mu}_B),$$
$$q_{ab} = q_a q_b.$$

Firm i's profit therefore separates as

$$\Pi^{i,S}(\boldsymbol{\mu}_A, \boldsymbol{\mu}_B) = \mu_A^i Q_A^i(\boldsymbol{\mu}_A) + \mu_B^i Q_B^i(\boldsymbol{\mu}_B), \tag{A.39}$$

where the term for market m is omitted when $i \notin \mathcal{N}_m$ and, for $i \in \mathcal{N}_m$,

$$Q_m^i(\boldsymbol{\mu}_m) = \frac{T_m^i \mathrm{e}^{-\mu_m^i}}{H_m(\boldsymbol{\mu}_m)}.$$

Each term in (A.39) is the profit of a multiproduct firm in a single-market logit oligopoly.

The first-order condition gives

$$\mu_m^i \left[1 - Q_m^i(\boldsymbol{\mu}_m)\right] = 1 \quad \text{for every } m \in \{A, B\} \text{ and } i \in \mathcal{N}_m. \tag{A.40}$$

Under the maintained restriction on the type network, each covered market has at least two active firms. Separate pricing therefore has a unique equilibrium, obtained by combining the unique equilibria of the two single-market logit games (Nocke and Schutz, 2018). In particular, the products that a firm supplies in market B do not enter its pricing problem in market A, and conversely.

At equilibrium under either pricing regime considered below, total markups depend only on the suppliers chosen in the two markets. Conditional on those supplier identities, product choices therefore have their marginal-cost distributions and are independent across markets. Consequently, the relative-entropy and mutual-information terms in equation (14) can be computed from the induced distributions over suppliers and outside options.

A.4.2 Proof of Proposition 10

Suppose there are $N \geq 2$ firms, $T_m^0 = 0$, and $T_m^i = T_m > 0$ for every firm i and $m \in \{A, B\}$. Define

$$w \equiv \log(T_A T_B).$$

We derive both equilibria and then compare profits, aggregate surplus, and consumer surplus.

Under separate pricing, each firm has probability $1/N$ of supplying the product chosen in either market. Equation (A.40) therefore gives

$$\mu^{S*} = \frac{N}{N-1}. \tag{A.41}$$

The CS-index and welfare levels are

$$H^S = N^2 \mathrm{e}^{w - 2\mu^{S*}}, \quad \mathrm{CS}^S = w + 2\log N - 2\mu^{S*}, \quad \Pi^{i,S} = \frac{2\mu^{S*}}{N}.$$

Hence $\mathrm{AS}^S = w + 2\log N$, the aggregate surplus under marginal-cost pricing.

Under mixed bundling, symmetry and Theorem 1 give a unique common markup μ^{B*}. The CS-index and the probability that a product pair contains at least one product supplied by a given firm are

$$\begin{aligned} H^B(\mu) &= N\mathrm{e}^{w-\mu}\left[1 + (N-1)\mathrm{e}^{-\mu}\right], \\ Q(\mu) &= \frac{1 + 2(N-1)\mathrm{e}^{-\mu}}{N\left[1 + (N-1)\mathrm{e}^{-\mu}\right]}. \end{aligned}$$

The first-order condition $[1 - Q(\mu)]\mu = 1$ is equivalent to

$$\mu^{B*} = \frac{N}{N-1} \frac{1 + (N-1)\mathrm{e}^{-\mu^{B*}}}{1 + (N-2)\mathrm{e}^{-\mu^{B*}}}. \tag{A.42}$$

Let

$$r^B \equiv \frac{(N-1)\mathrm{e}^{-\mu^{B*}}}{1+(N-1)\mathrm{e}^{-\mu^{B*}}}$$

be the probability that the two chosen products are supplied by different firms. Under separate pricing, this probability is $r^S = (N-1)/N$. The two first-order conditions imply

$$\frac{1}{\mu^{S*}} = \frac{N-1}{N}, \quad \frac{1}{\mu^{B*}} = \frac{N-1-r^B}{N}.$$

Since $0 < r^B < r^S$, these equations give

$$\mu^{S*} < \mu^{B*} < 2\mu^{S*}.$$

We now compare profits. The first-order condition also gives

$$\Pi^{i,B} = \mu^{B*} - 1 = \frac{1+r^B}{N-1-r^B}.$$

Because $\mu^{B*} > 1$, $\mathrm{e}^{-\mu^{B*}} < 1/2$ and therefore $r^B < (N-1)/(N+1)$. It follows that

$$\Pi^{i,B} < \frac{2}{N-1} = \Pi^{i,S}.$$

Next, we compare aggregate surplus. Symmetry makes the marginal choice probabilities coincide with their marginal-cost counterparts under both regimes. Under separate pricing, choices in the two markets are independent. Under mixed bundling, the probability that both products are supplied by the same firm is $1 - r^B > 1/N$, and hence

$$K\big((q^B_{ab})_{ab} \| (q^B_a q^B_b)_{ab}\big) > 0.$$

Equation (14) therefore gives $\mathrm{AS}^B < \mathrm{AS}^S$.

It remains to compare consumer surplus under the two regimes. When $N = 2$, let $x = (\mu^{B*} - 2)/2$. Equation (A.42) gives $x = \mathrm{e}^{-2-2x}$, so $0 < x < 1$, while $\mu^{S*} = 2$. Hence

$$\mathrm{CS}^B - \mathrm{CS}^S = 2(1-x) + \log\frac{1+x}{2} > 0,$$

where the inequality follows from

$$\log\frac{2}{1+x} < \frac{1-x}{1+x} < 1-x.$$

As $N \to \infty$, equations (A.41) and (A.42) give

$$\mu^{S*} = 1 + \frac{1}{N} + \frac{1}{N^2} + O(N^{-3}),$$
$$\mu^{B*} = 1 + \frac{2}{N} + \frac{4-\mathrm{e}}{N^2} + O(N^{-3}).$$

Substitution into the CS-indices gives

$$\mathrm{CS}^B - \mathrm{CS}^S = \frac{\mathrm{e}-3}{N} + O(N^{-2}) < 0$$

for all sufficiently large N. The same expansions give

$$\sum_{i=1}^{N}(\Pi^{i,S} - \Pi^{i,B}) = \frac{\mathrm{e}-2}{N} + O(N^{-2}).$$

Combining this expression with the preceding expansion of $\mathrm{CS}^B - \mathrm{CS}^S$ yields

$$\mathrm{AS}^B - \mathrm{AS}^S = -\frac{1}{N} + O(N^{-2}).$$

Since the first two terms in (14) are zero under both regimes,

$$K\big((q^B_{ab})_{ab} \| (q^B_a q^B_b)_{ab}\big) = \frac{1}{N} + O(N^{-2}).$$

This completes the proof. □

A.4.3 Proof of Proposition 11

Fix the type shares defined in the main text and suppose that $\sum_i \sigma^i_A \sigma^i_B > 0$.

We expand equilibrium outcomes under both regimes around $t^{\mathcal{N}}_A = t^{\mathcal{N}}_B = 0$ and compare the leading terms.

The equilibrium equations under both pricing regimes extend analytically to $t^{\mathcal{N}}_A = t^{\mathcal{N}}_B = 0$ and are solved there by setting every markup equal to one. Their Jacobian with respect to the markups at this solution is the identity matrix. The implicit function theorem therefore gives analytic solutions for all sufficiently small $t^{\mathcal{N}}_A, t^{\mathcal{N}}_B \geq 0$, and equilibrium uniqueness for positive inside types identifies these solutions with the equilibria characterized in the main text.

Under separate pricing, the normalized market-m CS-index is

$$\frac{H_m(\boldsymbol{\mu}_m)}{T^0_m} = 1 + t^{\mathcal{N}}_m \sum_{i \in \mathcal{N}} \sigma^i_m \mathrm{e}^{-\mu^i_m}.$$

Expanding the first-order condition (A.40) around $t^{\mathcal{N}}_m = 0$ gives

$$\mu^{i,S}_m - 1 = \frac{t^{\mathcal{N}}_m \sigma^i_m}{\mathrm{e}} - \frac{(t^{\mathcal{N}}_m)^2 \sigma^i_m}{\mathrm{e}^2} + O\left((t^{\mathcal{N}}_m)^3\right). \tag{A.43}$$

Under mixed bundling, equation (13) gives

$$\begin{aligned}\frac{H^B(\boldsymbol{\mu})}{T^0_A T^0_B} = 1 &+ \sum_{i \in \mathcal{N}} \left(t^{\mathcal{N}}_A \sigma^i_A + t^{\mathcal{N}}_B \sigma^i_B + t^{\mathcal{N}}_A t^{\mathcal{N}}_B \sigma^i_A \sigma^i_B\right) \mathrm{e}^{-\mu^i} \\ &+ t^{\mathcal{N}}_A t^{\mathcal{N}}_B \sum_{i<j} \left(\sigma^i_A \sigma^j_B + \sigma^j_A \sigma^i_B\right) \mathrm{e}^{-\mu^i - \mu^j}.\end{aligned}$$

The numerator of firm i's demand satisfies

$$\frac{H^B(\boldsymbol{\mu})}{T_A^0 T_B^0} Q^{i,B}(\boldsymbol{\mu}) = \left(t_A^{\mathcal{N}}\sigma_A^i + t_B^{\mathcal{N}}\sigma_B^i + t_A^{\mathcal{N}}t_B^{\mathcal{N}}\sigma_A^i\sigma_B^i\right)\mathrm{e}^{-\mu^i} + t_A^{\mathcal{N}}t_B^{\mathcal{N}}\sum_{j\neq i}\left(\sigma_A^i\sigma_B^j + \sigma_A^j\sigma_B^i\right)\mathrm{e}^{-\mu^i-\mu^j}.$$

Expanding the first-order condition $\mu^i[1 - Q^{i,B}(\boldsymbol{\mu})] = 1$ gives

$$\mu^{i,B} - 1 = \frac{t_A^{\mathcal{N}}\sigma_A^i + t_B^{\mathcal{N}}\sigma_B^i}{\mathrm{e}} - \frac{(t_A^{\mathcal{N}})^2\sigma_A^i + (t_B^{\mathcal{N}})^2\sigma_B^i}{\mathrm{e}^2} + t_A^{\mathcal{N}}t_B^{\mathcal{N}}\sigma_A^i\sigma_B^i\frac{\mathrm{e}-2}{\mathrm{e}^2} + O\left((t_A^{\mathcal{N}} + t_B^{\mathcal{N}})^3\right). \tag{A.44}$$

At equilibrium, $\Pi^{i,B} = \mu^{i,B} - 1$, while equation (A.40) gives $\Pi^{i,S} = \sum_{m\in\{A,B\}}(\mu_m^{i,S} - 1)$, with the market-$m$ term omitted if $T_m^i = 0$.

The two pricing regimes coincide whenever either $t_A^{\mathcal{N}} = 0$ or $t_B^{\mathcal{N}} = 0$. Analyticity therefore implies that every difference between equilibrium outcomes under the two regimes is divisible by $t_A^{\mathcal{N}}t_B^{\mathcal{N}}$. Subtracting equation (A.43) from equation (A.44) consequently yields

$$\Pi^{i,B} - \Pi^{i,S} = t_A^{\mathcal{N}}t_B^{\mathcal{N}}\sigma_A^i\sigma_B^i\frac{\mathrm{e}-2}{\mathrm{e}^2} + t_A^{\mathcal{N}}t_B^{\mathcal{N}}O\left(t_A^{\mathcal{N}} + t_B^{\mathcal{N}}\right). \tag{A.45}$$

Summing across firms gives

$$\sum_i(\Pi^{i,B} - \Pi^{i,S}) = t_A^{\mathcal{N}}t_B^{\mathcal{N}}\Big(\sum_i \sigma_A^i\sigma_B^i\Big)\frac{\mathrm{e}-2}{\mathrm{e}^2} + t_A^{\mathcal{N}}t_B^{\mathcal{N}}\,O\big(t_A^{\mathcal{N}} + t_B^{\mathcal{N}}\big).$$

Substituting the markup expansions into the CS-indices and using $\mathrm{CS} = \log H$ gives

$$\begin{aligned}\mathrm{CS}^B - \mathrm{CS}^S &= t_A^{\mathcal{N}}t_B^{\mathcal{N}}\Big(\sum_i \sigma_A^i\sigma_B^i\Big)\frac{\mathrm{e}-3}{\mathrm{e}^2} + t_A^{\mathcal{N}}t_B^{\mathcal{N}}\,O\big(t_A^{\mathcal{N}} + t_B^{\mathcal{N}}\big),\\ \mathrm{AS}^B - \mathrm{AS}^S &= t_A^{\mathcal{N}}t_B^{\mathcal{N}}\Big(\sum_i \sigma_A^i\sigma_B^i\Big)\frac{2\mathrm{e}-5}{\mathrm{e}^2} + t_A^{\mathcal{N}}t_B^{\mathcal{N}}\,O\big(t_A^{\mathcal{N}} + t_B^{\mathcal{N}}\big).\end{aligned} \tag{A.46}$$

The same expansion gives, for each $m \in \{A, B\}$,

$$K(\mathbf{q}_m^S\|\widehat{\mathbf{q}}_m) - K(\mathbf{q}_m^B\|\widehat{\mathbf{q}}_m) = t_A^{\mathcal{N}}t_B^{\mathcal{N}}\Big(\sum_i \sigma_A^i\sigma_B^i\Big)\frac{\mathrm{e}-2}{\mathrm{e}^2} + t_A^{\mathcal{N}}t_B^{\mathcal{N}}\,O\big(t_A^{\mathcal{N}} + t_B^{\mathcal{N}}\big), \tag{A.47}$$

$$K\big((q_{ab}^B)_{ab}\|(q_a^B q_b^B)_{ab}\big) = t_A^{\mathcal{N}}t_B^{\mathcal{N}}\Big(\sum_i \sigma_A^i\sigma_B^i\Big)\frac{1}{\mathrm{e}^2} + t_A^{\mathcal{N}}t_B^{\mathcal{N}}\,O\big(t_A^{\mathcal{N}} + t_B^{\mathcal{N}}\big). \tag{A.48}$$

Since $\mathrm{e} - 2 > 0$, $\mathrm{e} - 3 < 0$, and $2\mathrm{e} - 5 > 0$, equations (A.45)–(A.46) prove the proposition. Equations (A.47) and (A.48) also show that the reductions in the two within-market divergence terms exceed the increase in mutual information for all sufficiently small $t_A^{\mathcal{N}}, t_B^{\mathcal{N}} > 0$. □

A.4.4 Proof of Proposition 12

For $i, j \in \{1, 2\}$ and $R \in \{S, B\}$, let

$$Q_{AB}^{ij,R} \equiv \sum_{a \in \mathcal{A}^i} \sum_{b \in \mathcal{B}^j} q_{ab}^R$$

denote the probability that the products chosen in markets A and B are supplied by firms i and j, respectively, under regime R. Order these supplier pairs as $11, 12, 21, 22$. Dividing their contributions to H at marginal-cost pricing by $T_A^1 T_B^1$, and using the normalization $\lambda = 1$ maintained throughout the proofs of Section 5, gives the relative weights

$$\left(1, \mathrm{e}^{-\Delta_B}, \mathrm{e}^{-\Delta_A}, \mathrm{e}^{-(\Delta_A + \Delta_B)}\right).$$

We derive the limits of both equilibria as $\Delta_A, \Delta_B \to \infty$ and then compare profits, aggregate surplus, and consumer surplus.

Under separate pricing, the two markets are independent. For each $m \in \{A, B\}$, define

$$r_m \equiv \mathrm{e}^{-\Delta_m + \mu_m^1 - \mu_m^2}.$$

The probabilities of choosing firms 1 and 2 in market m are $1/(1 + r_m)$ and $r_m/(1 + r_m)$. The first-order conditions in (A.40) give

$$\mu_m^1 = \frac{1 + r_m}{r_m}, \quad \mu_m^2 = 1 + r_m, \quad \log r_m + \Delta_m = \frac{1}{r_m} - r_m. \tag{A.49}$$

For $\Delta_m > 0$, the last equation has a unique solution in $(\mathrm{e}^{-\Delta_m}, 1)$ because its left-hand side strictly increases in r_m and its right-hand side strictly decreases. Separate-pricing profits and supplier-pair probabilities are therefore

$$\Pi^{1,S} = \frac{1}{r_A} + \frac{1}{r_B}, \quad \Pi^{2,S} = r_A + r_B,$$
$$\left(Q_{AB}^{11,S}, Q_{AB}^{12,S}, Q_{AB}^{21,S}, Q_{AB}^{22,S}\right) = \frac{(1, r_B, r_A, r_A r_B)}{(1 + r_A)(1 + r_B)}.$$

As $\min\{\Delta_A, \Delta_B\} \to \infty$, the last equation in (A.49) implies

$$\frac{1}{r_m} = \Delta_m - \log \Delta_m + o(1), \quad m \in \{A, B\}.$$

Consequently,

$$\begin{aligned}
\Pi^{1,S} &= \Delta_A + \Delta_B - \log(\Delta_A \Delta_B) + o(1), \\
\Pi^{2,S} &= \frac{1}{\Delta_A} + \frac{1}{\Delta_B} + o\left(\frac{1}{\Delta_A} + \frac{1}{\Delta_B}\right), \\
Q_{AB}^{11,S} &= 1 - \frac{1}{\Delta_A} - \frac{1}{\Delta_B} + o\left(\frac{1}{\Delta_A} + \frac{1}{\Delta_B}\right), \\
Q_{AB}^{12,S} &= \frac{1}{\Delta_B} + o(\Delta_B^{-1}), \\
Q_{AB}^{21,S} &= \frac{1}{\Delta_A} + o(\Delta_A^{-1}), \\
Q_{AB}^{22,S} &= \frac{1}{\Delta_A \Delta_B} + o\left((\Delta_A \Delta_B)^{-1}\right).
\end{aligned}$$

Under mixed bundling, firm 1 is absent only from supplier pair 22 and firm 2 only from supplier pair 11. Their first-order conditions give

$$\mu^1 Q_{AB}^{22,B} = 1, \quad \mu^2 Q_{AB}^{11,B} = 1, \quad \Pi^{i,B} = \mu^i - 1. \tag{A.50}$$

Taking the ratio of the first two supplier-pair probabilities gives

$$\mu^1 - \mu^2 + \log \frac{\mu^1}{\mu^2} = \Delta_A + \Delta_B. \tag{A.51}$$

Relative to supplier pair 11, the two cross-firm probabilities satisfy

$$Q_{AB}^{12,B} = \frac{\mathrm{e}^{-\Delta_B - \mu^2}}{\mu^2}, \quad Q_{AB}^{21,B} = \frac{\mathrm{e}^{-\Delta_A - \mu^2}}{\mu^2}.$$

Using equations (A.50) and imposing that the four probabilities sum to one therefore gives

$$\frac{1}{\mu^1} + \frac{1 + \mathrm{e}^{-\mu^2}\left(\mathrm{e}^{-\Delta_A} + \mathrm{e}^{-\Delta_B}\right)}{\mu^2} = 1. \tag{A.52}$$

Theorem 1 gives a unique solution. Equations (A.51) and (A.52) imply $\mu^1 \to \infty$ and $\mu^2 \to 1$.

Expanding them at this limit gives

$$\begin{aligned}
\Pi^{1,B} &= \Delta_A + \Delta_B - \log(\Delta_A + \Delta_B) + o(1),\\
\Pi^{2,B} &= \frac{1}{\Delta_A + \Delta_B} + \frac{\mathrm{e}^{-\Delta_A} + \mathrm{e}^{-\Delta_B}}{\mathrm{e}} + o\left(\frac{1}{\Delta_A + \Delta_B} + \mathrm{e}^{-\Delta_A} + \mathrm{e}^{-\Delta_B}\right),\\
Q_{AB}^{11,B} &= 1 - \frac{1}{\Delta_A + \Delta_B} - \frac{\mathrm{e}^{-\Delta_A} + \mathrm{e}^{-\Delta_B}}{\mathrm{e}} + o\left(\frac{1}{\Delta_A + \Delta_B} + \mathrm{e}^{-\Delta_A} + \mathrm{e}^{-\Delta_B}\right),\\
Q_{AB}^{12,B} &= \frac{\mathrm{e}^{-\Delta_B}}{\mathrm{e}} + o(\mathrm{e}^{-\Delta_B}),\\
Q_{AB}^{21,B} &= \frac{\mathrm{e}^{-\Delta_A}}{\mathrm{e}} + o(\mathrm{e}^{-\Delta_A}),\\
Q_{AB}^{22,B} &= \frac{1}{\Delta_A + \Delta_B} + o\left((\Delta_A + \Delta_B)^{-1}\right).
\end{aligned} \tag{A.53}$$

The strong firm's profit gain satisfies

$$\Pi^{1,B} - \Pi^{1,S} = \log\left(\frac{\Delta_A \Delta_B}{\Delta_A + \Delta_B}\right) + o(1) \longrightarrow \infty.$$

Moreover,

$$\frac{1}{\Delta_A} + \frac{1}{\Delta_B} \geq \frac{4}{\Delta_A + \Delta_B},$$

while $\mathrm{e}^{-\Delta_m} = o(\Delta_m^{-1})$ for each m. The expansions above therefore imply $\Pi^{2,B} < \Pi^{2,S}$ for all sufficiently large Δ_A and Δ_B. Firm 2's profit loss converges to zero, so industry profit rises for all sufficiently large Δ_A and Δ_B.

It remains to compare aggregate surplus under the two regimes. Let $Q_m^{i,R}$ denote the probability that firm i supplies the product chosen in market m under regime R, and let $\widehat{Q}_m^i$ denote that probability under marginal-cost pricing. Then

$$\widehat{Q}_m^2 = \frac{\mathrm{e}^{-\Delta_m}}{1 + \mathrm{e}^{-\Delta_m}}.$$

Under separate pricing,

$$Q_m^{2,S} = \frac{r_m}{1 + r_m} = \frac{1}{\Delta_m} + o(\Delta_m^{-1}).$$

The market-m relative-entropy term in equation (14) therefore converges to one for each $m \in \{A, B\}$. Choices are independent under separate pricing, so its mutual-information term is zero.

Under mixed bundling, equations (A.53) give

$$\begin{aligned}
Q_A^{2,B} &= Q_{AB}^{21,B} + Q_{AB}^{22,B} = \frac{1}{\Delta_A + \Delta_B} + \frac{\mathrm{e}^{-\Delta_A}}{\mathrm{e}} + o\left(\frac{1}{\Delta_A + \Delta_B} + \mathrm{e}^{-\Delta_A}\right),\\
Q_B^{2,B} &= Q_{AB}^{12,B} + Q_{AB}^{22,B} = \frac{1}{\Delta_A + \Delta_B} + \frac{\mathrm{e}^{-\Delta_B}}{\mathrm{e}} + o\left(\frac{1}{\Delta_A + \Delta_B} + \mathrm{e}^{-\Delta_B}\right).
\end{aligned}$$

Using the binary relative-entropy formula and $\Delta_m e^{-\Delta_m} \to 0$, these expressions imply

$$K\left((Q_m^{i,B})_{i=1,2} \| (\widehat{Q}_m^i)_{i=1,2}\right) = \frac{\Delta_m}{\Delta_A + \Delta_B} + o(1), \quad m \in \{A, B\}.$$

The two marginal relative-entropy terms therefore sum to one in the limit. Conditional on supplier identities, product choices are independent across markets, so the mutual-information term equals the mutual information between the two supplier identities. It is bounded above by the entropy of the supplier identity in either market. Since $Q_A^{2,B}$ and $Q_B^{2,B}$ both converge to zero, this entropy and hence mutual information converge to zero. Equation (14) therefore gives

$$\mathrm{AS}^B - \mathrm{AS}^S \longrightarrow 1.$$

Finally,

$$\mathrm{CS}^B - \mathrm{CS}^S = (\mathrm{AS}^B - \mathrm{AS}^S) - (\Pi^{1,B} - \Pi^{1,S}) - (\Pi^{2,B} - \Pi^{2,S}).$$

The aggregate-surplus gain converges to one, the gain in firm 1's profit diverges, and firm 2's profit loss converges to zero. Mixed bundling therefore lowers consumer surplus for all sufficiently large Δ_A and Δ_B. This completes the proof. □