NET WORK LISER

Discussion Paper Series

IZA DP No. 19014

October 2026

Competitive Search for Bundles

Xiaoming Cai
Peking University HSBC Business School

Pieter Gautier
Vrije Universiteit Amsterdam,
Tinbergen Institute and IZA@LISER

The IZA Discussion Paper Series (ISSN: 2365-9793) ("Series") is the primary platform for disseminating research produced within the framework of the IZA@LISER Network, an unincorporated international network of labour economists coordinated by the Luxembourg Institute of Socio-Economic Research (LISER). The Series is operated by LISER, a Luxembourg public establishment (établissement public) registered with the Luxembourg Business Registers under number J57, with its registered office at 11, Porte des Sciences, 4366 Esch-sur-Alzette, Grand Duchy of Luxembourg.

Any opinions expressed in this Series are solely those of the author(s). LISER accepts no responsibility or liability for the content of the contributions published herein. LISER adheres to the European Code of Conduct for Research Integrity. Contributions published in this Series present preliminary work intended to foster academic debate. They may be revised, are not definitive, and should be cited accordingly. Copyright remains with the author(s) unless otherwise indicated.

ISSN: 2365-9793

IZA DP No. 19014 October 2026

Competitive Search for Bundles*

Abstract
How does buyers' search for bundles generate sorting among heterogeneous sellers? We study competitive search in which buyers assemble bundles through separate transactions, where prices direct search and determine congestion. With two seller types, buyers choose only same-type bundles below a single threshold in cross-type value and cross-type bundles above it. The threshold lies strictly between geometric and arithmetic means of same-type values, so search produces cross-type matching even when types are complements. With a continuum of types, log-supermodularity and submodularity yield unique stable equilibria with positive and negative assortative matching. Positive sorting extends to any fixed number of visits.

JEL classification
C78, D83

Keywords
competitive search, search for bundles, assortative matching, seller heterogeneity

Corresponding author
Xiaoming Cai
xmingcai@gmail.com

* We thank Jim Albrecht, Susan Vroman, and Philipp Kircher for valuable comments.

1 Introduction

In many markets, buyers search for a combination of goods that have to be obtained in separate transactions. A firm that staffs a project needs workers whose skills fit together, a manufacturer needs components that work together as a system, and a household that buys a house typically also needs a mortgage. In each of these cases, one transaction is only valuable if the other one succeeds as well. When sellers differ in quality or specialization, it also matters from whom the goods are bought, because the value of the final bundle depends on the identity of each seller. A buyer therefore has to decide not only how many sellers to visit, but also which types of sellers to combine.

The O-ring theory of Kremer (1993) is the canonical framework for studying how such complementarities shape the formation of teams, and it has become one of the most influential explanations for why quality clusters in markets. Kremer's starting point is that production is fragile. The value of each component depends on the quality of all other components, in the same way as the failure of a single O-ring destroyed the space shuttle *Challenger*. As a result, high-quality inputs are combined with other high-quality inputs (positive assortative matching), and small differences in quality translate into large differences in output, wages, and productivity across firms and countries. Recent empirical work suggests that these complementarities are important in practice. Herkenhoff et al. (2024) show that hiring two capable workers is not enough, because what matters is whether they work well together, and that getting the match right can raise both current production and learning. Demir et al. (2024) find that the same force affects which firms trade with each other, as skill-intensive Turkish firms trade disproportionately with other skill-intensive firms. In this paper, we add competitive search frictions to the O-ring benchmark and study how prices and congestion shape the formation of teams and bundles in a decentralized market.[1]

Competitive search is a natural framework to study how frictions affect the assembly of bundles. Sellers post terms of trade, buyers decide which sellers to visit, and the resulting queues determine the probability that a transaction takes place. Prices therefore play a double role, as they divide the surplus and direct buyers across sellers. When a buyer needs two goods, the markets for these goods become linked, because the price and the queue at one seller affect the value of visiting the other. Most competitive search models, however, allow each buyer only one visit and one possible trade, so they cannot address how buyers search for bundles.

[1]Buyers may also purchase from several sellers because of a taste for variety, as in Dixit and Stiglitz (1977). In their framework, each variety has value on its own, and the question is how many varieties the market provides. In our model, a single good has no value on its own, and the question is which seller types buyers combine when each component has to be found through search.

We allow buyers to visit several sellers. To simplify the exposition, we start our analysis with two visits per buyer and two types of sellers (low and high). Each seller supplies one indivisible good, and identical buyers enter at a cost. A single good has no value; value arises only from completing a two-good bundle. Bundle value depends on seller types: x_{11} for two low types, x_{22} for two high types, and x_{12} for one of each. We assume $x_{11} < x_{22}$. Seller types are complements in bundle value if $x_{12} < (x_{11} + x_{22})/2$ and substitutes if $x_{12} > (x_{11} + x_{22})/2$, in the sense of Becker (1973).[2] A law firm hiring two associates may value two litigators differently from one transactional lawyer and one litigator, and a club signing two players may value two forwards, two defenders, or one of each differently.

In most real-world markets, firms post a price independent of the buyer's other visit. Our main analysis nevertheless allows a richer contract space in which a seller may condition its price on the type of the buyer's other intended seller: for example, a low-type seller can quote one price to buyers who also visit a low type and another to those who visit a high type. We do so for tractability: when each seller must charge a single price, sellers' payoff functions have kinks and multiple equilibria can arise. The Online Appendix shows that the same broad sorting forces remain under uniform pricing, though the analysis is substantially more involved. Menu pricing lets prices respond directly to the buyer's intended bundle and isolates the main sorting forces from additional informational or contracting frictions.

Our menu-pricing benchmark is in the spirit of Rhodes and Zhou (2024), who study personalized pricing based on consumers' valuations for both a firm's own product and other firms' products. They motivate their benchmark by firms' growing access to rich consumer data, including browsing and purchase histories, which can reveal preferences across products. They treat perfect knowledge of preferences as a limiting case that richer data and better algorithms may allow firms to approach, rather than a literal description of firms' information. We view our menus in the same way. Related examples of pricing contingent on other transactions include: (i) house buyers may pay different prices depending on their mortgage type, see Reher and Valkanov (2024); (ii) drug manufacturers may receive different net payments for the same product depending on the other drugs with which it is prescribed, as illustrated by combination-specific rebates in Switzerland;[3] and (iii) firms may offer discounts when a complementary product from another firm is purchased, as in supermarket–gasoline or airport–car-rental cross-promotions, see Armstrong (2013).

[2] By complementarity, we mean the interaction between seller types across the two visits, not increasing marginal value in the quantity of one good. A valuation with $v(0) = v(1) = 0$ and $v(2) > 0$ is non-concave in quantity but does not imply complementarity across distinct goods.

[3] For example, the January 2020 Swiss reimbursement conditions for Revlimid specified different rebates for its use with ixazomib and carfilzomib. The ixazomib-specific rebate required documented use of both drugs and did not apply to subsequent Revlimid maintenance treatment after ixazomib was discontinued (Swiss Federal Office of Public Health, 2020, p. 24).

Menus, however, also contain prices for unused bundles, and because buyers in our model need two successful visits, such prices can matter artificially. If, for example, no buyer chooses a mixed bundle, mixed-bundle prices are never paid in equilibrium, and sellers could set them very high to make mixed search unattractive, even when the mixed bundle is valuable. Conversely, if all buyers choose mixed bundles, same-type bundle prices are never paid and could be set very high to block same-type search, even when these bundles should attract buyers. Such equilibria are not convincing, because they are supported by prices at which no trade ever takes place. Following the no-blocking idea in the assignment game with transferable utility (Shapley and Shubik, 1972), we hold the equilibrium queues fixed and ask whether every inactive bundle can be priced so that the buyer just breaks even (that is, receives an expected payoff equal to the entry cost, as does a buyer following an equilibrium strategy) and neither seller wants to offer that bundle.

To see what this refinement implies, note that in a standard single-visit model sellers operate in separate submarkets. In our model, the value of attracting a buyer depends on the queue and the price at the buyer's other visit, so the two sellers in a bundle must be considered jointly. Consider, for example, the separating equilibrium, in which buyers always visit two sellers of the same type. A buyer who would instead combine a low-type and a high-type seller is relatively attractive to the high-type seller, because her other visit goes to the less congested low-type market. For the same reason, she is relatively unattractive to the low-type seller, because her other visit goes to the more congested high-type market. Prices can shift the surplus of such a mixed bundle between its two sellers, but not the chance that each visit succeeds, which is fixed by the queues.

Whether a bundle can be priced so that its buyers just break even and neither seller type can improve by offering it therefore depends only on its total value. The highest such value is the queue-adjusted stability bound B, which depends only on the two sellers' queues. An active pair must lie exactly on this bound, because its buyers break even and each seller earns exactly its equilibrium payoff, whereas an inactive pair only needs to lie weakly below it. In the separating equilibrium, the low-low and high-high bundles are active, so the queue lengths λ_1 and λ_2 of low and high types solve $x_{11} = B(\lambda_1, \lambda_1)$ and $x_{22} = B(\lambda_2, \lambda_2)$. The mixed bundle is inactive, so stability requires only that $x_{12} \leq B(\lambda_1, \lambda_2)$, a threshold we denote by $\bar{x}_{12}$.

This threshold governs the unique stable equilibrium: if $x_{12} \leq \bar{x}_{12}$, buyers either choose bundles consisting of high-quality components or bundles in which each component is of low quality (positive assortative pairing), and the two same-type markets operate separately. It lies strictly between $\sqrt{x_{11}x_{22}}$ and $(x_{11} + x_{22})/2$. Search frictions therefore work against same-type matching. A high-type seller's good is only worth something if the buyer also

obtains the other component, and this is more likely when that component comes from a less congested low-type seller. As a result, the threshold below which buyers sort by seller type lies strictly below the arithmetic mean $(x_{11} + x_{22})/2$ at which cross-type matching becomes attractive in a frictionless market.

If $x_{12} > \bar{x}_{12}$, mixed bundles are used in equilibrium. The sorting pattern then depends on the composition of sellers. When high-type sellers are scarce, high types are visited only as part of a mixed bundle, while low types also serve buyers who visit two low-type sellers. For intermediate compositions, all buyers choose mixed bundles. When high-type sellers are abundant, the pattern is reversed. Thus seller composition determines which same-type market, if any, remains active together with mixed search.

We then extend the analysis to a continuum of seller types. Equilibrium sorting is now described by a pairing function T, which specifies that a buyer who visits a seller of type x sends her other visit to a seller of type $T(x)$. The function B again determines which pairs can be supported by prices. With queue lengths $\lambda(x)$, every active pair satisfies $f(x, T(x)) = B(\lambda(x), \lambda(T(x)))$, while every inactive pair satisfies $f(x, y) \leq B(\lambda(x), \lambda(y))$. We derive two sufficient conditions for sorting. If the value of a cross-type bundle never exceeds the geometric mean of the two corresponding same-type values, $f(x, y)^2 \leq f(x, x)f(y, y)$, the stable equilibrium is unique and exhibits positive assortative matching (PAM), so that each type is paired with its own type. If instead f is submodular (the continuum analogue of seller types being substitutes), the stable equilibrium is unique and exhibits negative assortative matching (NAM), in which low types are paired with high types. In the two-type model, these conditions become $x_{12}^2 \leq x_{11}x_{22}$ and $x_{11} + x_{22} \leq 2x_{12}$, respectively, and so lie on either side of the exact threshold $\bar{x}_{12}$. In both models, a low cross-type value therefore leads to separation and a high cross-type value leads to cross-type matching.

The sorting pattern also determines how the market responds when some seller types become scarce, for example after a disruption in the supply of a critical mineral or a specialized input. Under PAM, the queue at a type-x seller is determined by the same-type bundle value $f(x, x)$ alone and is independent of the distribution of seller types. A fall in the supply of some types therefore leaves all prices and queues unchanged. Fewer bundles with the affected types are assembled, but sellers of other types are not affected, because no buyer combines them with the scarce types.[4] Under NAM, in contrast, a change in the relative supply of seller types changes the entire pairing function, so that prices and queues adjust at all types. These results relate to the literature on supply-chain disruptions, which

[4]As in Kremer (1993), this relies on the fact that any seller can fill either position of a bundle. If a bundle instead consists of two distinct types, each supplied by its own population, a fall in the supply of only one good leaves some high-quality sellers of the other good without a partner of the same quality. Queues, prices, and in general the pairing then adjust, unless the supply of both goods falls in the same way.

documents that shocks propagate through production networks (Carvalho et al., 2021; Demir et al., 2024), and to search models of credit markets in which the breakup of lending relationships propagates shocks and can even make financial collapse unavoidable (den Haan et al., 2003). In our model, a shock that hits certain seller types spreads to other types under NAM but not under PAM.

To assess the quantitative importance of re-sorting under NAM, we calibrate the model to sales concentration and a model-implied search-expenditure share in Demir et al. (2024). We consider equally large supply reductions concentrated among high or low types, using the capital-destruction scenario in Carvalho et al. (2021) as a numerical benchmark. In each case, we measure the effect of re-sorting by comparing the new equilibrium with a fixed-partner outcome at the same post-change seller supply, in which prices, queues, and entry adjust, but buyers must keep combining the original seller types.

When high types become scarce, re-sorting completes about 6 percent more bundles and reduces the bottom quintile's per-seller profit loss by more than half, while limiting high types' scarcity gains. Nevertheless, the fall in average bundle value more than offsets the additional completions, leaving gross value about 3 percent lower. Higher entry expenditure also contributes to a reduction in total net surplus of about 4 percent. Under free entry, this surplus equals aggregate seller profits. When low types become scarce instead, re-sorting completes about 2 percent fewer bundles but raises gross value by about 3 percent and net surplus by about 4 percent relative to fixed partners. The market re-sorts in both cases, even when doing so lowers net surplus, because the fixed-partner outcome is not stable. Hence, the number of transactions that survive a disruption can give a misleading picture of its welfare costs.

To sum up, we make four main contributions. First, we introduce a competitive search framework in which buyers assemble a bundle by visiting several heterogeneous sellers, so that prices determine both the division of surplus and the congestion in markets that are linked through buyers' visits. Second, we derive the queue-adjusted stability bound B and use it to characterize the unique stable equilibrium outcome with two seller types. Third, we extend this characterization to a continuum of seller types and provide conditions for PAM and NAM. For PAM, the characterization also holds for any fixed number of visits. Fourth, we show that the sorting pattern determines how the market responds to a change in the supply of seller types, and we quantify the effects of re-sorting under NAM.

Related literature. In a frictionless market, Kremer (1993) shows that complementarity in worker quality leads to PAM. Search frictions affect this result in three ways. First, congestion makes buyers of mixed bundles especially attractive to high-type sellers, so that

complementarity no longer guarantees that sellers are matched with their own type. Second, prices affect not only the division of surplus but also the length of queues. A buyer who has secured one component is less sensitive to the price of the other component, which gives rise to a *sunk-success* pricing wedge. Third, the shares of the two seller types matter, because each mixed bundle uses one seller of each type. When one type is abundant, some of its sellers therefore also trade with each other.

Demir et al. (2024) study O-ring forces in production networks. In their quantitative model, firms choose quality and pay for ads to find suppliers and customers. Firms are more likely to find partners of similar quality, and high-quality firms buy more from high-quality suppliers. These interconnected quality choices amplify an economy-wide export-demand shock. Their focus is a many-to-many network in which firm quality changes, whereas we hold seller types fixed and study how prices and queues determine which bundles arise in equilibrium.

Most of the competitive-search literature studies agents seeking one successful match.[5] Related papers study simultaneous search and multiple applications. Albrecht et al. (2006), Galenianos and Kircher (2009), Kircher (2009), Chade and Smith (2006), and Gautier and Holzner (2017) let workers send multiple applications or buyers pay several visits to improve the chance of one match. In these models, an agent needs only one success, so extra applications mainly insure against rejection. Albrecht et al. (2020) allow firms to use richer mechanisms when workers apply to several jobs. Albrecht et al. (2023) show that the Nash and market-maker versions of competitive search, which coincide when buyers visit only one seller, can differ when buyers make two visits, with buyer preferences ranging from unit demand to demand for both goods. For the latter case, as in our model, Albrecht et al. (2026) show that competitive search resolves nonexistence and price indeterminacy in Walrasian markets and establish finite-to-large-market convergence. In these papers, sellers are homogeneous.

The paper also connects to work on heterogeneity and sorting in competitive search. Shi (2001), Shimer (2005), and Eeckhout and Kircher (2010) study sorting across submarkets with heterogeneity on both sides of the market. Here, only sellers are heterogeneous, while buyers are ex ante identical. Sorting instead arises from buyers' endogenous choice of which seller types to combine, so it pairs seller types with each other rather than with buyer types. This pairing is determined jointly with the prices and queues that support it.

Eeckhout and Kircher (2018) study a Walrasian many-to-one matching model in which heterogeneous firms choose both worker type and workforce size. Endogenous size changes Becker's sorting condition, and differential equations determine the assignment, firm size,

[5] See Wright et al. (2021) for a survey.

and wages. In our model, bundle size is fixed, and congestion rather than firm size changes the sorting condition.

A related literature studies with whom agents work in Walrasian markets when output depends on the composition of a whole team. Kremer and Maskin (1996) show how task asymmetries and skill dispersion determine whether workers of different skill are segregated across firms or work together, providing a frictionless parallel to our threshold between same-type and cross-type matching. Garicano and Rossi-Hansberg (2006) study knowledge-based hierarchies, in which agents sort into workers and managers. Because the knowledge of a manager and that of their workers are complements, matching is positively assortative across hierarchical layers, and team composition follows from the organization of problem solving. Carlier and Ekeland (2010) characterize matching for teams as an optimal transport problem, and Boerma et al. (2025) solve a sorting problem in which a firm hires several workers whose skills are imperfect substitutes under a technology that is submodular in worker skills. In their model, a project of value z fails only if every team member fails, so a team's expected output loss is z times the probability that all of its members fail. They show that optimal sorting combines a mixed region in which intermediate-type firms and workers form teams that all incur the same expected output loss, and countermonotonic regions, including those in which high-skill workers are matched with low-skill coworkers at high-productivity firms. In these papers, the team is assembled in a market without search frictions. Here it is assembled through separate simultaneous directed-search visits, so that prices and congestion determine which seller types are combined. Unlike PAM, negative sorting has no direct counterpart beyond two visits: it would require a joint assignment of all seller types in a bundle rather than a single decreasing partner map. Ahlin (2017), Chade and Eeckhout (2018), and Boerma et al. (2025) illustrate the richer sorting patterns that arise in larger groups even without search frictions. We leave the corresponding competitive-search characterization to future work.

The paper also contributes to the theory of commodity bundling. Armstrong (2013) develops a general theory with non-additive valuations, allowing complements and substitutes for both an integrated firm and competing sellers. Lucarelli et al. (2010) study empirically how buyers assemble pharmaceutical cocktails from rival sellers, while Calzolari and Denicolò (2013) analyze competition when prices depend on buyers' purchases from rivals. Matutes and Regibeau (1988, 1992) study compatibility and pricing when buyers can combine components from different sellers, and Economides and Salop (1992) examine integration among producers of complements. Our model adds search frictions to this literature, so that prices also determine congestion. Unlike the sequential add-on purchases in Ellison (2005), our buyers search for both goods simultaneously and trade only when both visits succeed.

Our framework also connects to the literature on search frictions in credit markets, where obtaining financing is often necessary to complete a purchase or undertake production. Allen et al. (2019) show that such frictions give lenders substantial market power in the mortgage market, and Dell'Ariccia and Garibaldi (2005) model bank lending itself as a frictional matching market. Wasmer and Weil (2004) study firms that must find both a lender and a worker, so that frictions in credit and labor markets compound, and Petrosky-Nadeau and Wasmer (2013) show that this double friction amplifies labor-market volatility. In these models the two components are found one after the other through random matching, and partners are identical. Here buyers search for both components at once, sellers differ in type, and posted prices direct search.

The paper is organized as follows. Section 2 presents the model with two seller types and defines equilibrium and stability. Section 3 introduces the function B and characterizes the equilibrium. Section 3.1 reformulates the sellers' problem in terms of queue lengths and previews the main result, Section 3.2 derives the separating equilibrium and the threshold $\bar{x}_{12}$, and Section 3.3 studies cross-type matching. Section 4 extends the analysis to a continuum of seller types and derives conditions for PAM and NAM. Section 4.4 extends the PAM characterization to any fixed number of visits, and Section 4.5 studies how re-sorting under NAM changes the effects of a fall in seller supply. Section 5 concludes. Proofs are in the appendix.

2 The Model

Agents. There is a unit measure of sellers and a large measure of potential homogeneous buyers, each of whom may enter at cost $K > 0$; the measure of active buyers, L, is determined in equilibrium. Each seller owns one unit of an indivisible good and is either a *low type* (indexed by 1) or a *high type* (indexed by 2), with measures $1 - z$ and z, where $z \in (0, 1)$. Each buyer demands two units and directs two visits to sellers. A single unit has zero value; the gross value of two units is x_{11}, x_{12}, or x_{22}, depending on whether the two sellers visited are both type 1, of different types, or both type 2. We assume throughout that $K < x_{11} < x_{22}$. We make no assumptions on the value of x_{12} relative to x_{11} and x_{22}. Seller types are complements in bundle value if $x_{12} < (x_{11} + x_{22})/2$ and substitutes if $x_{12} > (x_{11} + x_{22})/2$.

Menu pricing. Because each buyer visits two sellers, a seller's expected payoff depends on where a visiting buyer's other visit is directed. We therefore allow sellers to post *pair-specific prices* (menus), conditioning on the type of the other seller the buyer visits. We classify buyers by their visiting strategy: an *A-buyer* visits two type-1 sellers (strategy A), a

B-buyer one seller of each type (strategy B), and a *C-buyer* two type-2 sellers (strategy C). Each type-1 seller posts a menu (p_1^A, p_1^B) and each type-2 seller a menu (p_2^B, p_2^C), where p_j^i is the price a type-j seller charges an i-buyer.[6] We look for a symmetric pure-strategy equilibrium in which sellers of the same type post the same menu.

Menu pricing is similar in spirit to many competitive search models in which terms of trade are allowed to depend on publicly observable characteristics. In online markets, sellers often observe digital footprints that reveal part of a buyer's shopping path, for example through cookies, tracking pixels, ad-network data sharing, or browser fingerprinting. We do not, however, interpret menu pricing literally as a seller's ability to perfectly condition on the buyer's search strategy. Rather, it serves as a useful and tractable benchmark for environments in which firms have some information about the buyer's other visits. The Online Appendix analyzes uniform pricing, under which each seller posts a single price to all buyers. The analysis is substantially more complicated because sellers' payoff functions are discontinuous and nonconcave.

Search. Buyers observe all posted menus and seller types before directing their two visits. We assume that each seller can meet at most one buyer. Let λ denote the expected number of buyer visits to a seller, which we refer to as the *queue length*, and let $m(\lambda)$ denote the probability that the seller meets a buyer.[7] We assume that $m(\lambda)$ is twice continuously differentiable, strictly increasing, and strictly concave. We further assume that its elasticity, $\varepsilon_m(\lambda) = \lambda m'(\lambda)/m(\lambda)$, is strictly decreasing. The urn-ball function $m(\lambda) = 1 - e^{-\lambda}$ is a leading example. Since each match involves exactly one buyer and one seller, the probability that an individual buyer visit results in being selected is $q(\lambda) = m(\lambda)/\lambda$. We normalize $m'(0) = 1$, which, by L'Hôpital's rule and $m(0) = 0$, implies that $q(0) = 1$.

Let $\alpha_A, \alpha_B, \alpha_C \geq 0$ with $\alpha_A + \alpha_B + \alpha_C = 1$ denote the shares of buyers choosing strategy A, B, or C, respectively. Since each A-buyer visits two type-1 sellers, each B-buyer visits one seller of each type, and each C-buyer visits two type-2 sellers, the equilibrium queue compositions at each seller type are

$$\lambda_1^A = \frac{2\alpha_A L}{1-z}, \qquad \lambda_1^B = \frac{\alpha_B L}{1-z}, \tag{1}$$

$$\lambda_2^B = \frac{\alpha_B L}{z}, \qquad \lambda_2^C = \frac{2\alpha_C L}{z}, \tag{2}$$

[6] Each menu lists two prices rather than three: a type-1 seller is never visited by a C-buyer, who visits two type-2 sellers, so there is no need for p_1^C; symmetrically, there is no need for p_2^A.

[7] Alternatively, we could allow a seller to meet multiple buyers and to observe each applicant's other visit *before* selecting whom to serve. This would not affect the equilibrium, because whenever a seller attracts two buyer groups in equilibrium, it must be indifferent between them. For example, a type-1 seller attracting both A- and B-buyers satisfies $V_1^A = V_1^B$; see (15).

where L is the measure of buyers, λ_1^A is the queue length of A-buyers at a type-1 seller, and so on. The total queue length at a type-1 seller is $\lambda_1 = \lambda_1^A + \lambda_1^B$, and at a type-2 seller it is $\lambda_2 = \lambda_2^B + \lambda_2^C$.

Buyer payoffs and free entry. A buyer trades only if both of her visits succeed: she then pays both sellers and obtains the bundle, while if either visit fails she pays nothing to either seller. Since the two visits are independent, the probability that both succeed is the product of their individual success probabilities. The expected payoffs from the three strategies are therefore

$$U_A = q(\lambda_1)^2 \, (x_{11} - 2p_1^A), \tag{3}$$

$$U_B = q(\lambda_1) \, q(\lambda_2) \, (x_{12} - p_1^B - p_2^B), \tag{4}$$

$$U_C = q(\lambda_2)^2 \, (x_{22} - 2p_2^C), \tag{5}$$

where an A-buyer pays p_1^A at each of her two type-1 sellers, a B-buyer pays p_1^B and p_2^B at the two sellers she visits, and a C-buyer pays p_2^C at each type-2 seller. From now on we write $q_1 \equiv q(\lambda_1)$ and $q_2 \equiv q(\lambda_2)$.

Buyer free entry requires that each active strategy yields expected payoff equal to the entry cost:

$$U_i = K \;\text{ for each strategy } i \text{ with } \alpha_i > 0, \qquad U_i \leq K \;\text{ otherwise.} \tag{6}$$

Seller payoffs. Consider a type-1 seller. With probability $m(\lambda_1)$, the seller is matched with one buyer drawn from its queue in proportion to the queue composition. Trade occurs only if that buyer's *other* visit also succeeds. An A-buyer's other visit succeeds with probability q_1, a B-buyer's with probability q_2. The expected profit in equilibrium is therefore

$$\pi_1 = m(\lambda_1) \left[\frac{\lambda_1^A}{\lambda_1} \, q_1 \, p_1^A + \frac{\lambda_1^B}{\lambda_1} \, q_2 \, p_1^B \right], \tag{7}$$

where the fractions λ_1^A/λ_1 and λ_1^B/λ_1 are the probabilities that the selected buyer is an A-buyer or a B-buyer, respectively. Similarly, the expected profit of a type-2 seller is

$$\pi_2 = m(\lambda_2) \left[\frac{\lambda_2^B}{\lambda_2} \, q_1 \, p_2^B + \frac{\lambda_2^C}{\lambda_2} \, q_2 \, p_2^C \right].$$

Deviations. Consider a deviating type-1 seller who posts an alternative menu $(\widetilde{p}_1^A, \widetilde{p}_1^B)$ and attracts queue lengths $\widetilde{\lambda}_1^A$ of A-buyers and $\widetilde{\lambda}_1^B$ of B-buyers. Its total queue length is

$\widetilde{\lambda}_1 = \widetilde{\lambda}_1^A + \widetilde{\lambda}_1^B$.

An A-buyer who visits the deviating seller and a non-deviating type-1 seller obtains expected payoff $q(\widetilde{\lambda}_1)\, q_1\,(x_{11} - \widetilde{p}_1^A - p_1^A)$. Similarly, a B-buyer visiting the deviating seller obtains $q(\widetilde{\lambda}_1)\, q_2\,(x_{12} - \widetilde{p}_1^B - p_2^B)$. Since both payoffs share the common factor $q(\widetilde{\lambda}_1)$, the relative attractiveness of strategies A and B at the deviant seller, and hence the composition of the deviant's queue, is independent of $\widetilde{\lambda}_1$ and is pinned down by prices $(\widetilde{p}_1^A, \widetilde{p}_1^B)$ alone:

$$(\widetilde{\lambda}_1^A, \widetilde{\lambda}_1^B) = \begin{cases} (\widetilde{\lambda}_1, 0) & \text{if } q_1\,(x_{11} - \widetilde{p}_1^A - p_1^A) > q_2\,(x_{12} - \widetilde{p}_1^B - p_2^B) \\ (0, \widetilde{\lambda}_1) & \text{if } q_1\,(x_{11} - \widetilde{p}_1^A - p_1^A) < q_2\,(x_{12} - \widetilde{p}_1^B - p_2^B) \\ \text{indeterminate} & \text{if } q_1\,(x_{11} - \widetilde{p}_1^A - p_1^A) = q_2\,(x_{12} - \widetilde{p}_1^B - p_2^B) \end{cases} \tag{8}$$

Given this composition, buyer free entry determines the total queue length $\widetilde{\lambda}_1$:

$$K = q(\widetilde{\lambda}_1) \max\Big\{q_1\,(x_{11} - \widetilde{p}_1^A - p_1^A),\ q_2\,(x_{12} - \widetilde{p}_1^B - p_2^B)\Big\}. \tag{9}$$

If the maximum is below K, the deviant attracts no buyers and $\widetilde{\lambda}_1 = 0$. The deviant seller's expected profit takes the same form as the equilibrium profit (7), with the deviant's prices and queue replacing the equilibrium ones:

$$\pi_1(\widetilde{p}_1^A, \widetilde{p}_1^B) = m(\widetilde{\lambda}_1)\left[\frac{\widetilde{\lambda}_1^A}{\widetilde{\lambda}_1}\, q_1\, \widetilde{p}_1^A + \frac{\widetilde{\lambda}_1^B}{\widetilde{\lambda}_1}\, q_2\, \widetilde{p}_1^B\right]. \tag{10}$$

The problem of a deviant type-2 seller is defined analogously.

Equilibrium definition. We can now formally define our equilibrium concept.[8]

Definition 1. *A* competitive search equilibrium *consists of price menus* (p_1^A, p_1^B) *and* (p_2^B, p_2^C)*, queue lengths* $(\lambda_1^A, \lambda_1^B)$ *and* $(\lambda_2^B, \lambda_2^C)$*, strategy shares* $(\alpha_A, \alpha_B, \alpha_C)$*, and a measure of buyers* L *such that:*

1. ***Buyer optimality.*** *Given prices, the queue lengths are determined by buyer free entry and indifference conditions* (1)–(6)*.*

2. ***Seller optimality.*** *No seller has a profitable deviation. In particular, for any alternative menu* $(\widetilde{p}_1^A, \widetilde{p}_1^B)$ *posted by a type-1 seller, the deviation profit in* (10) *cannot exceed*

[8] We do not characterize the constrained-efficient allocation. Albrecht et al. (2023) show that a homogeneous-seller version of our environment is constrained inefficient because buyer entry is excessive, and that efficiency requires richer instruments, such as fees. Since this inefficiency arises without seller heterogeneity, a welfare analysis here would largely rediscover it. Our quantitative results also show that, when high-type sellers become scarce, the equilibrium has 4.28 percent more buyer entry and 3.95 percent lower net surplus than the fixed-partner outcome (Section 4.5); this reverses when low-type sellers become scarce.

the equilibrium profit in (7)*. The deviating seller's queue composition and total queue length are determined by* (8) *and* (9)*. The condition for type-2 sellers is analogous.*

Multiplicity. Because a bundle is completed only when both visits succeed, a prohibitive off-path price at one seller shuts down an entire search strategy, and a deviating seller cannot undo this by lowering its own price. This generates equilibrium multiplicity. Consider an equilibrium in which no buyer chooses a mixed bundle ($\alpha_B = 0$). Since the cross-type prices p_1^B and p_2^B are never paid, sellers can set them sufficiently high to make strategy B unattractive, regardless of how large x_{12} is. Conversely, if all buyers choose mixed bundles ($\alpha_B = 1$), the unused own-type prices p_1^A and p_2^C can be set sufficiently high to rule out strategies A and C, regardless of how valuable the same-type bundles are. In both cases, the outcome is sustained by off-path prices that are not disciplined by actual trade.

Stability. We resolve this multiplicity by disciplining off-path prices. The motivation draws on the no-blocking logic of classical matching theory, although the analogy is conditional on the equilibrium queues. In the Gale–Shapley model, an unmatched pair blocks an assignment if both agents prefer to match with each other. With transferable utility, as in the assignment game of Shapley and Shubik (1972), transfers determine whether the surplus of a potential pair can be divided so that both agents benefit.[9] In our setting, holding the equilibrium queue lengths fixed, buyer free entry fixes the buyer's expected payoff at K, while the two prices divide the remaining expected surplus between the two seller roles. We therefore ask whether every potential bundle, including an inactive one, can be supported by prices on the buyer's free-entry frontier without violating seller optimality.

Definition 2. *A competitive search equilibrium is* stable *if, at its equilibrium queue lengths, there exists a system of prices for all active and inactive buyer strategies such that $U_i = K$ for every $i \in \{A, B, C\}$, and the resulting price menus continue to satisfy seller optimality in Definition 1.*

The first requirement places every potential bundle on the buyer's free-entry frontier, while the second requires that these prices still support the equilibrium. In particular, stability prevents an inactive strategy from being ruled out by arbitrary prices that are never paid.[10]

[9] See Roth and Sotomayor (1990, Chapter 8) for a textbook treatment.

[10] This requirement is distinct from the market-maker equilibrium studied by Albrecht et al. (2023). A deviating market maker opens a new submarket, coordinates multiple sellers posting deviant prices, and attracts new queue lengths; buyers may then direct one or both visits to that submarket. Our stability requirement neither creates a new submarket nor changes the equilibrium queues. The equilibrium concept

For the own-type strategies, symmetry leaves a unique supporting price at which buyers break even: an A-buyer visits two type-1 sellers who charge the same price, so $U_A = K$ determines p_1^A, and likewise $U_C = K$ determines p_2^C. Solving these conditions gives

$$p_1^A = \frac{x_{11}}{2} - \frac{K}{2q_1^2}, \qquad p_2^C = \frac{x_{22}}{2} - \frac{K}{2q_2^2}. \tag{11}$$

In the separating equilibrium, the A- and C-strategies are active, so these two prices are actually paid and $U_A = U_C = K$ is the usual market-utility condition. Stability then imposes no additional restrictions on the own-type prices and only matters for the prices of the inactive mixed bundle, to which we turn next. Whenever an own-type strategy is inactive, by contrast, (11) pins down its otherwise unused price and rules out the prohibitive prices discussed above.

For the mixed bundle in a separating equilibrium ($\alpha_B = 0$), free entry only determines the sum of the two cross-type prices. Solving $U_B = K$ gives

$$p_1^B + p_2^B = x_{12} - \frac{K}{q_1 q_2}. \tag{12}$$

Different ways of splitting this sum divide the surplus of the mixed bundle differently between the two sellers. The separating equilibrium is stable if and only if the sum can be split such that, at the equilibrium queues, a type-1 seller does not gain by attracting B-buyers instead of A-buyers, and a type-2 seller does not gain by attracting B-buyers instead of C-buyers. To see when this fails, note that the value of a B-buyer to a seller is higher when the other seller in the bundle charges a lower price. For a type-1 seller, this value is $q_2(x_{12} - p_2^B)$, and for a type-2 seller it is $q_1(x_{12} - p_1^B)$ (see Section 3.1). Hence, a type-1 seller is only kept away from B-buyers if p_2^B is sufficiently high, and a type-2 seller only if p_1^B is sufficiently high. Since raising one price lowers the other, these two requirements may be incompatible. If no split satisfies both, then for every split at least one seller type strictly prefers B-buyers to its equilibrium buyers. That seller has a profitable deviation, so the separating outcome cannot be supported as a stable equilibrium.

remains Nash-style: unilateral seller deviations that attract different queues are already ruled out by seller optimality in Definition 1. Stability adds only a restriction on the prices supporting the equilibrium at its equilibrium queue lengths.

3 Equilibrium Characterization

3.1 Sellers' problem reformulated

We now reformulate the seller's problem in terms of queue lengths rather than prices. This is a standard approach in the competitive search literature, and it greatly simplifies the analysis of sorting. We then preview the main result.

Reformulation. Consider a type-1 deviant seller who posts an alternative menu $(\widetilde{p}_1^A, \widetilde{p}_1^B)$, attracting queues $(\widetilde{\lambda}_1^A, \widetilde{\lambda}_1^B)$, with total queue length $\widetilde{\lambda}_1 = \widetilde{\lambda}_1^A + \widetilde{\lambda}_1^B$. When both queues are strictly positive, both A-buyers and B-buyers earn exactly the entry cost K from visiting the deviant. Solving equation (9) for the deviant's prices gives

$$\widetilde{p}_1^A = x_{11} - p_1^A - \frac{K}{q(\widetilde{\lambda}_1)\, q_1} \tag{13}$$

$$\widetilde{p}_1^B = x_{12} - p_2^B - \frac{K}{q(\widetilde{\lambda}_1)\, q_2}. \tag{14}$$

Next, define the expected surpluses for a type-1 seller:

$$V_1^A \equiv q_1(x_{11} - p_1^A), \qquad V_1^B \equiv q_2(x_{12} - p_2^B), \qquad V_1 \equiv \max\{V_1^A,\, V_1^B\}. \tag{15}$$

The term V_1^A captures the expected surplus that a type-1 seller can extract from an A-buyer: the gross value x_{11} minus the A-buyer's payment p_1^A to the *other* type-1 seller, weighted by the probability q_1 that the buyer's other visit succeeds. The term V_1^B has an analogous interpretation for B-buyers. Both surpluses are pinned down by the equilibrium prices at *other* sellers and are independent of the deviant's choices.

Substituting the prices in (13) and (14) into (10), and using $m(\lambda) = \lambda q(\lambda)$, the deviant type-1 seller's problem can be written in terms of queue lengths as

$$\max_{\widetilde{\lambda}_1^A, \widetilde{\lambda}_1^B} \pi_1(\widetilde{\lambda}_1^A, \widetilde{\lambda}_1^B) = m(\widetilde{\lambda}_1) \left[\frac{\widetilde{\lambda}_1^A}{\widetilde{\lambda}_1} V_1^A + \frac{\widetilde{\lambda}_1^B}{\widetilde{\lambda}_1} V_1^B \right] - \widetilde{\lambda}_1 K. \tag{16}$$

This reformulation expresses the deviant's problem in terms of the queue lengths $(\widetilde{\lambda}_1^A, \widetilde{\lambda}_1^B)$ rather than the prices $(\widetilde{p}_1^A, \widetilde{p}_1^B)$. The two formulations are equivalent: given the optimal queue lengths, the corresponding prices follow from (13) and (14).[11]

The deviant seller's problem can now be simplified further. For any fixed total queue

[11] The derivation assumes both queues are strictly positive, but (16) continues to hold at corner solutions. When $\widetilde{\lambda}_1^B = 0$, it reduces to $m(\widetilde{\lambda}_1^A)\, V_1^A - \widetilde{\lambda}_1^A\, K$ and symmetrically when $\widetilde{\lambda}_1^A = 0$.

length, the bracketed term is a weighted average of the expected surpluses generated by A- and B-buyers. The seller therefore allocates the entire queue to the group with the higher surplus. If $V_1^A > V_1^B$, it sets $\widetilde{\lambda}_1^B = 0$; if $V_1^A < V_1^B$, it sets $\widetilde{\lambda}_1^A = 0$; and if $V_1^A = V_1^B$, only the total queue $\widetilde{\lambda}_1$ matters. Thus, the seller's problem reduces to a one-dimensional choice of total queue length.

Lemma 1. *A deviant type-1 seller's problem can be written as*

$$\max_{\widetilde{\lambda}_1 \geq 0} \; m(\widetilde{\lambda}_1) V_1 - \widetilde{\lambda}_1 K, \tag{17}$$

where $V_1 = \max\{V_1^A, V_1^B\}$.

Proof. See the above discussion. □

Lemma 1 reduces the deviation problem to the familiar one-dimensional choice of a total queue length. The difference from a single-visit model is that V_1 is determined by the buyer's other visit and therefore by prices, queues, and sorting elsewhere in the market. It is through V_1 that a seller's problem depends on the rest of the market, and this is what distinguishes the two-visit environment from a standard competitive search model.

The problem of a deviant type-2 seller is analogous. Define

$$V_2^B \equiv q_1(x_{12} - p_1^B), \qquad V_2^C \equiv q_2(x_{22} - p_2^C), \qquad V_2 \equiv \max\{V_2^B,\, V_2^C\}. \tag{18}$$

Its deviation problem is

$$\max_{\widetilde{\lambda}_2^B, \widetilde{\lambda}_2^C} \; \pi_2(\widetilde{\lambda}_2^B, \widetilde{\lambda}_2^C) = m(\widetilde{\lambda}_2)\left[\frac{\widetilde{\lambda}_2^B}{\widetilde{\lambda}_2} V_2^B + \frac{\widetilde{\lambda}_2^C}{\widetilde{\lambda}_2} V_2^C\right] - \widetilde{\lambda}_2 K. \tag{19}$$

The same weighted-average argument reduces this problem to $\max_{\widetilde{\lambda}_2 \geq 0} m(\widetilde{\lambda}_2) V_2 - \widetilde{\lambda}_2 K$.

Equilibrium profits. For either seller type $j = 1, 2$, the reformulation and first-order condition $m'(\lambda_j) V_j = K$ imply $\pi_j = m(\lambda_j) V_j - K\lambda_j = \Pi(\lambda_j)$, where

$$\Pi(\lambda) = K\lambda\left(\frac{1}{\varepsilon_m(\lambda)} - 1\right). \tag{20}$$

Since $\varepsilon_m \in (0, 1)$ is strictly decreasing, expected profits increase strictly with the seller's equilibrium queue, regardless of the sorting regime.

Preview of results. A single threshold $\bar{x}_{12} \in (\sqrt{x_{11}x_{22}}, (x_{11}+x_{22})/2)$ governs equilibrium sorting. Below it, buyers choose only same-type bundles ($\alpha_B = 0$): the two markets operate separately, and queues are independent of z and x_{12}. Above it, cross-type buyers are active ($\alpha_B > 0$). They coexist with A-buyers when type-1 sellers are relatively abundant ($z < z^*$), with C-buyers when type-2 sellers are relatively abundant ($z > z^{**}$), and are the only active buyers for intermediate compositions ($z^* \leq z \leq z^{**}$). Away from the knife edge, the stable symmetric pure-strategy equilibrium is unique, apart from unused cross-type prices in the separating regime.[12]

Section 3.2 characterizes separation and derives $\bar{x}_{12}$; Section 3.3 characterizes cross-type matching and the composition thresholds z^* and z^{**}.

3.2 The Separating Equilibrium

We first characterize the separating equilibrium, in which only A- and C-buyers are active ($\alpha_B = 0$) and the two submarkets operate in isolation.

Own-type visitor values. Given the equilibrium queue lengths λ_1 and λ_2, substituting the prices p_1^A and p_2^C from (11) into the definitions of the own-type visitor values (15) and (18), we express V_1^A and V_2^C as functions of queue lengths alone:

$$V_1^A = \frac{q_1\, x_{11}}{2} + \frac{K}{2q_1}, \qquad V_2^C = \frac{q_2\, x_{22}}{2} + \frac{K}{2q_2}. \tag{21}$$

The cross-type visitor values V_1^B and V_2^B are not yet determined, since no B-buyers are active; we return to them when analyzing stability.

Queue lengths. By Lemma 1, the queue length λ_j must satisfy the first-order condition

$$K = m'(\lambda_j)\, V_j. \tag{22}$$

Since $V_1 = V_1^A$ and $V_2 = V_2^C$, substituting (21) into the above first-order condition yields

$$x_{jj} = \frac{K}{q(\lambda_j^{\text{sep}})^2} \left(\frac{2}{\varepsilon_m(\lambda_j^{\text{sep}})} - 1 \right) \tag{23}$$

where we denote the equilibrium queue lengths by λ_j^{sep}, $j = 1, 2$. Since both $q(\lambda)$ and $\varepsilon_m(\lambda)$ are strictly decreasing, the right-hand side of (23) is strictly increasing in λ. Because

[12] At $x_{12} = \bar{x}_{12}$, queues and payoffs are unique, but buyer strategy shares need not be. Any feasible mixture consistent with market clearing can be supported.

$K < x_{jj}$, equation (23) has a unique solution. Finally, $x_{22} > x_{11}$ implies $\lambda_2^{\text{sep}} > \lambda_1^{\text{sep}}$. Since the two submarkets operate in isolation, λ_1^{sep} and λ_2^{sep} depend only on (x_{11}, K) and (x_{22}, K), respectively, and are independent of the seller composition z and the cross-type value x_{12}.

Equilibrium prices in elasticity form. Substituting (23) into the free-entry condition (11) to eliminate K yields the equilibrium on-path prices in a compact form:

$$p_j^{\text{sep}} = \frac{-\varepsilon_q(\lambda_j^{\text{sep}})}{1 - \varepsilon_q(\lambda_j^{\text{sep}})}\, x_{jj}, \tag{24}$$

where $p_1^{\text{sep}} = p_1^A$ and $p_2^{\text{sep}} = p_2^C$, and $\varepsilon_q(\lambda) \equiv \lambda\, q'(\lambda)/q(\lambda) = \varepsilon_m(\lambda) - 1$ is the elasticity of q with respect to λ. Since $\varepsilon_q \in (-1, 0)$, the price lies strictly between 0 and $x_{jj}/2$.

We now compare the pricing rule in (24) with the single-visit benchmark, in which the Hosios rule gives $p = -\varepsilon_q(\lambda)x = (1 - \varepsilon_m(\lambda))x$. With two visits, each seller charges a share $-\varepsilon_q/(1 - \varepsilon_q)$ of the same-type bundle value, so the total price share is $-2\varepsilon_q/(1 - \varepsilon_q)$, which strictly exceeds $-\varepsilon_q$. This difference reflects a *sunk-success* effect. Conditional on one visit succeeding, the other visit determines whether the buyer completes the bundle and obtains any value at all. Buyers are therefore less sensitive to a price increase at either seller than in the single-visit benchmark: when a seller raises its price, its queue thins more slowly because buyers place a high value on completing the bundle. The equilibrium total price is correspondingly higher.

Stability of the separating equilibrium. The analysis so far has not pinned down the off-path prices p_1^B and p_2^B, since no B-buyers are present to discipline them. Under Definition 2, stability asks whether these prices can be chosen so that the inactive strategy B lies on the buyer's free-entry frontier, while the resulting menus continue to satisfy seller optimality. Thus the free entry condition (12) must hold. Using $V_1^B = q_2(x_{12} - p_2^B)$ and $V_2^B = q_1(x_{12} - p_1^B)$, we can equivalently express this condition in terms of V_1^B and V_2^B:

$$x_{12} = \frac{1}{q_1 q_2}\left(q_1 V_1^B + q_2 V_2^B - K\right). \tag{25}$$

Equation (25) traces a downward-sloping line in (V_1^B, V_2^B) space. Stability requires a point on this line satisfying $V_1^B \le V_1^A$ and $V_2^B \le V_2^C$. Because the two cross-type prices divide the available surplus between the two groups of sellers, such a point on (25) exists if and only if $x_{12} \le \frac{1}{q_1 q_2}(q_1 V_1^A + q_2 V_2^C - K)$.

At the separating queues, the seller first-order conditions (22) imply $q(\lambda_1^{\text{sep}})V_1^A =$

$K/\varepsilon_m(\lambda_1^{\text{sep}})$ and $q(\lambda_2^{\text{sep}})V_2^C = K/\varepsilon_m(\lambda_2^{\text{sep}})$. Hence the preceding inequality is equivalent to

$$x_{12} \leq B(\lambda_1^{\text{sep}}, \lambda_2^{\text{sep}}), \tag{26}$$

where the function B is defined by

$$B(\lambda_1, \lambda_2) \equiv \frac{K}{q(\lambda_1)q(\lambda_2)} \left[\frac{1}{\varepsilon_m(\lambda_1)} + \frac{1}{\varepsilon_m(\lambda_2)} - 1 \right]. \tag{27}$$

The B function and the stability requirement. The function $B(\lambda_1, \lambda_2)$ gives the highest bundle value that can be supported at queues (λ_1, λ_2) without creating a blocking pair. If the bundle value is no greater than this bound, prices can be chosen so that the buyer breaks even and neither seller prefers the candidate pair to its current buyers. Thus the mixed pair does not block the equilibrium if and only if $x_{12} \leq B(\lambda_1, \lambda_2)$.

In the separating equilibrium, each same-type pair is active and therefore lies exactly on the stability bound. Hence

$$x_{jj} = B(\lambda_j^{\text{sep}}, \lambda_j^{\text{sep}}), \qquad j = 1, 2. \tag{28}$$

This is the separating first-order condition (23), written in terms of B. By contrast, (26) requires the inactive mixed pair to lie weakly below the bound. The following lemma records the properties of B used below and later in the continuum analysis.

Lemma 2. *The function $B(\lambda_1, \lambda_2)$ is symmetric, strictly increasing in each argument, and strictly supermodular, while $\log B$ is strictly submodular. In particular, if $\lambda_1 \neq \lambda_2$, then*

$$B(\lambda_1, \lambda_2)^2 > B(\lambda_1, \lambda_1)B(\lambda_2, \lambda_2). \tag{29}$$

Proof. See Appendix A.1. □

Define $\bar{x}_{12} \equiv B(\lambda_1^{\text{sep}}, \lambda_2^{\text{sep}})$. The analysis above establishes that separation is stable if and only if $x_{12} \leq \bar{x}_{12}$. The curvature of B places the separation threshold strictly between the geometric and arithmetic means. Since $\lambda_1^{\text{sep}} \neq \lambda_2^{\text{sep}}$, strict supermodularity gives $2\bar{x}_{12} = 2B(\lambda_1^{\text{sep}}, \lambda_2^{\text{sep}}) < B(\lambda_1^{\text{sep}}, \lambda_1^{\text{sep}}) + B(\lambda_2^{\text{sep}}, \lambda_2^{\text{sep}}) = x_{11} + x_{22}$, while strict log-submodularity gives $\bar{x}_{12}^2 = B(\lambda_1^{\text{sep}}, \lambda_2^{\text{sep}})^2 > B(\lambda_1^{\text{sep}}, \lambda_1^{\text{sep}})B(\lambda_2^{\text{sep}}, \lambda_2^{\text{sep}}) = x_{11}x_{22}$.

We can now summarize our results on the separating equilibrium.

Proposition 1. *A stable separating equilibrium exists if and only if $x_{12} \leq \bar{x}_{12} \equiv B(\lambda_1^{\text{sep}}, \lambda_2^{\text{sep}})$, where $x_{11} = B(\lambda_1^{\text{sep}}, \lambda_1^{\text{sep}})$, $x_{22} = B(\lambda_2^{\text{sep}}, \lambda_2^{\text{sep}})$, and the function B is defined by equation (27). Furthermore, $\sqrt{x_{11}\, x_{22}} < \bar{x}_{12} < (x_{11} + x_{22})/2$.*

When it exists, the separating equilibrium is unique: queue lengths are given by (23), on-path prices by (24), and off-path prices (p_1^B, p_2^B) must satisfy (12), $V_1^B \leq V_1^A$, and $V_2^B \leq V_2^C$.

Proof. See the above discussion. □

Thus ordinary supermodularity, $2x_{12} < x_{11} + x_{22}$, does not by itself guarantee separation. The reason is a congestion asymmetry that makes cross-type buyers particularly attractive to type-2 sellers. In the separating equilibrium, type-2 sellers attract longer queues, $\lambda_2^{\text{sep}} > \lambda_1^{\text{sep}}$, and buyers therefore have a lower success probability, $q_2 < q_1$. A B-buyer's other visit is directed to the less-congested type-1 market and succeeds with probability q_1, whereas a C-buyer's other visit is directed to another type-2 seller and succeeds only with probability q_2. B-buyers are therefore relatively more valuable to type-2 sellers, causing separation to break down before x_{12} reaches the arithmetic mean. This force is absent in frictionless models such as Kremer (1993), in which every agent who wants to trade at the equilibrium prices can do so. Without differences in success probabilities, separation holds whenever seller types are complements, $2x_{12} \leq x_{11} + x_{22}$, as in Becker (1973).

To understand why log-supermodularity is sufficient, note that a seller's payoff from attracting a buyer depends not only on the value of the bundle, but also on the probability that the buyer's other visit succeeds. Congestion therefore affects sellers' incentives in a multiplicative way, and log-supermodularity is the corresponding multiplicative notion of complementarity. It ensures that the high-high bundle is sufficiently valuable relative to the mixed bundle to offset the stronger incentive of type-2 sellers to attract B-buyers, so that prices can be chosen such that neither seller type prefers B-buyers to its same-type buyers.

3.3 Equilibrium with Cross-Type Matching

We now turn to equilibria in which B-buyers are active ($\alpha_B > 0$).

The $\mathcal{B}$-curve. Recall that $B(\lambda_1, \lambda_2)$ gives the highest bundle value that can be supported at the two queues without creating a blocking pair. Because the mixed pair is now active, this bound holds with equality:

$$x_{12} = B(\lambda_1, \lambda_2), \tag{30}$$

the off-diagonal counterpart of (28). We call this locus the $\mathcal{B}$-curve.[13] Since B is strictly increasing in each argument by Lemma 2, the $\mathcal{B}$-curve is strictly downward-sloping in (λ_1, λ_2) space; see Figure 1. Moreover, because B depends only on (λ_1, λ_2) and K, and not on x_{12}, an increase in x_{12} shifts the curve outward, toward longer queues.

[13] Alternatively, equation (30) can be derived by the same procedure used for (26): solve for p_1^B and p_2^B from (31)–(32) and substitute them into buyer free entry (12).

Equilibrium prices in elasticity form. The cross-type prices can be expressed in a compact form that parallels the separating-equilibrium pricing rule (24). Since $V_1 = V_1^B$ and $V_2 = V_2^B$, the sellers' first-order condition (22) implies

$$\frac{K}{m'(\lambda_1)} = V_1^B = q_2(x_{12} - p_2^B), \tag{31}$$

$$\frac{K}{m'(\lambda_2)} = V_2^B = q_1(x_{12} - p_1^B), \tag{32}$$

where the second equalities use the definitions of V_1^B and V_2^B in (15) and (18), respectively.

Substituting (30) into (31) and (32) eliminates $K/(q_1 q_2)$ and yields two equations in the two unknown prices p_1^B and p_2^B. Solving them gives

$$p_j^B = \frac{\dfrac{1}{\varepsilon_m(\lambda_j)} - 1}{\dfrac{1}{\varepsilon_m(\lambda_1)} + \dfrac{1}{\varepsilon_m(\lambda_2)} - 1} x_{12}, \qquad j = 1, 2. \tag{33}$$

Hence, each seller receives a share of the cross-type value x_{12} that only depends on the two matching elasticities. Holding the other seller's queue fixed, a longer queue at a type-j seller lowers $\varepsilon_m(\lambda_j)$, which raises this seller's share of the bundle value and lowers the share of the other seller. Intuitively, the seller with the longer queue is the more congested side of the bundle. Its probability of meeting a buyer responds less to a small reduction in visits, so it can raise its price without losing much in terms of trading probability. As a result, the division of the bundle value shifts in its favor, and the shares left to the other seller and to the buyer fall. In the limit, $p_j^B / x_{12} \to 1$ as $\lambda_j \to \infty$.

The formula also nests the standard single-visit pricing rule. As $\lambda_2 \to 0$, the buyer's visit to the type-2 seller succeeds with probability approaching one and $\varepsilon_m(\lambda_2) \to 1$. Equation (33) then gives $p_1^B = (1 - \varepsilon_m(\lambda_1))x_{12}$, which is the Hosios pricing rule. When success at the other seller is effectively guaranteed, the type-1 seller's pricing problem reduces to the single-visit benchmark.

Lower bounds on queue lengths. When B-buyers are active, seller optimality requires $V_1^B \geq V_1^A$ and $V_2^B \geq V_2^C$. Equivalently, the same-type bundles must be nonblocking:[14]

$$x_{11} \leq B(\lambda_1, \lambda_1), \qquad x_{22} \leq B(\lambda_2, \lambda_2). \tag{34}$$

[14]The conditions in (34) can also be derived formally as in (26). Because stability imposes $U_A = K$ whether strategy A is active or inactive, it fixes p_1^A at (11) and hence V_1^A at (21), while V_1^B is given by (31). Hence $V_1^B \geq V_1^A$ is equivalent to $x_{11} \leq \frac{K}{q(\lambda_1)^2}(\frac{2}{\varepsilon_m(\lambda_1)} - 1) = B(\lambda_1, \lambda_1)$. The type-2 inequality follows analogously from $V_2^B \geq V_2^C$.

Since $x_{jj} = B(\lambda_j^{\text{sep}}, \lambda_j^{\text{sep}})$ (see (28)) and B is strictly increasing in each argument, these inequalities are equivalent, respectively, to $\lambda_1 \geq \lambda_1^{\text{sep}}$ and $\lambda_2 \geq \lambda_2^{\text{sep}}$. Equality holds exactly when the corresponding seller type is indifferent between its same-type buyers and B-buyers. This yields the following lemma.

Lemma 3. *In any stable equilibrium with $\alpha_B > 0$, $\lambda_1 \geq \lambda_1^{\text{sep}}$ and $\lambda_2 \geq \lambda_2^{\text{sep}}$. The first inequality binds if and only if $V_1^A = V_1^B$, and the second binds if and only if $V_2^B = V_2^C$.*

Proof. See the discussion above. □

Stability ensures that each seller can turn to its own-type market, even when that market is inactive in equilibrium. Each seller must therefore earn at least its separating profit. Since equilibrium profits rise with queues, the separating queue length λ_j^{sep} acts as a floor. The floor binds exactly when a seller type is indifferent between its same-type buyers and B-buyers, that is, when the two groups can coexist.

Three regimes. By Lemma 3, any cross-type equilibrium must lie on the portion of the $\mathcal{B}$-curve inside the feasible region $[\lambda_1^{\text{sep}}, \infty) \times [\lambda_2^{\text{sep}}, \infty)$. Since $\bar{x}_{12} \equiv B(\lambda_1^{\text{sep}}, \lambda_2^{\text{sep}})$, if $x_{12} < \bar{x}_{12}$, no point on the curve is feasible; if $x_{12} = \bar{x}_{12}$, only the corner $(\lambda_1^{\text{sep}}, \lambda_2^{\text{sep}})$ is feasible; and if $x_{12} > \bar{x}_{12}$, the feasible portion of the curve is a nondegenerate segment.

For $x_{12} > \bar{x}_{12}$, the two endpoints are defined by

$$B(\lambda_1^{\text{sep}}, \lambda_2^{AB}) = x_{12}, \qquad B(\lambda_1^{BC}, \lambda_2^{\text{sep}}) = x_{12}.$$

These are the blue and red points in Figure 1, and the segment of the $\mathcal{B}$-curve between them is exactly the set of feasible cross-type queue pairs.

At the blue endpoint, $\lambda_1 = \lambda_1^{\text{sep}}$, so type-1 sellers are indifferent between A- and B-buyers. Since $\lambda_2^{AB} > \lambda_2^{\text{sep}}$, type-2 sellers strictly prefer B-buyers to C-buyers. This endpoint therefore supports the AB regime. Symmetrically, the red endpoint supports the BC regime. At every interior point of the feasible segment, both queues lie strictly above their floors, so both seller types strictly prefer B-buyers to their same-type buyers. Only B-buyers are then active. Thus, when $x_{12} > \bar{x}_{12}$, the only possible configurations are AB, pure-B, and BC; in particular, A- and C-buyers cannot coexist.

Seller composition selects among these three possibilities. Since every B-buyer visits one seller of each type, a pure-B equilibrium requires

$$(1 - z)\lambda_1 = z\lambda_2. \tag{35}$$

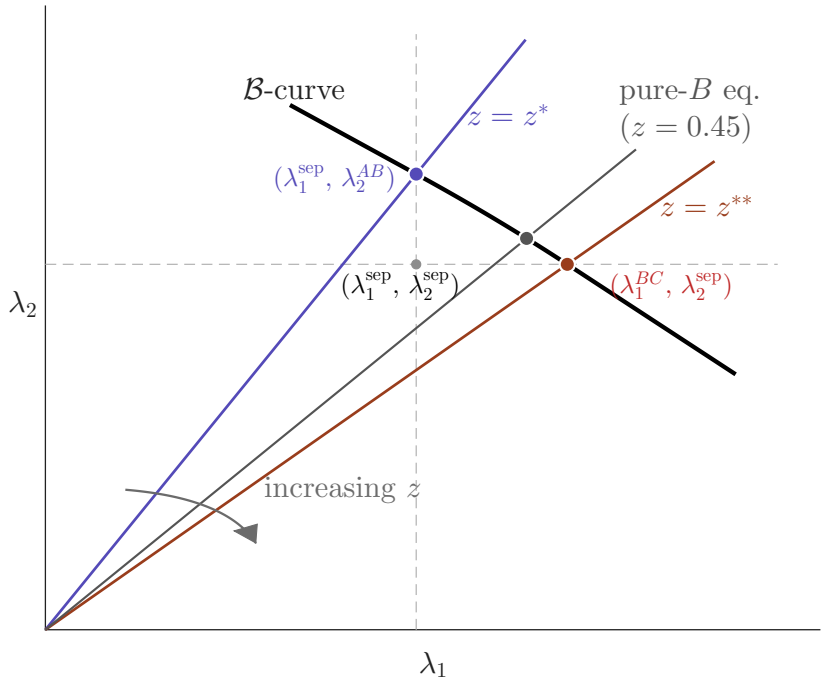


Figure 1: The $\mathcal{B}$-curve is the locus of queue-length pairs (λ_1, λ_2) consistent with seller optimality and buyer free entry when $\alpha_B > 0$. The dashed lines mark the queue floors λ_j^{sep}, and the feasible portion of the curve is the segment between the blue AB point $(\lambda_1^{\text{sep}}, \lambda_2^{AB})$ and the red BC point $(\lambda_1^{BC}, \lambda_2^{\text{sep}})$. Each ray from the origin corresponds to a seller composition z, with slope $(1-z)/z$. Rays intersecting the interior support pure-B equilibria, while the endpoint rays define z^* and z^{**}. As z increases, the ray rotates clockwise. Parameters: $x_{11} = 1$, $x_{22} = 2$, $x_{12} = 1.9$, $K = 0.2$, and $m(\lambda) = 1 - e^{-\lambda}$, yielding $\bar{x}_{12} = 1.435$, $\lambda_1^{\text{sep}} = 0.88$, $\lambda_2^{\text{sep}} = 1.31$, $\lambda_2^{AB} = 1.63$, $\lambda_1^{BC} = 1.24$, $z^* = 0.35$, and $z^{**} = 0.49$.

This condition defines the composition ray in Figure 1. Because the $\mathcal{B}$-curve is strictly downward-sloping, the ray intersects the feasible segment at most once. The rays through its two endpoints define

$$z^* = \frac{\lambda_1^{\text{sep}}}{\lambda_1^{\text{sep}} + \lambda_2^{AB}}, \qquad z^{**} = \frac{\lambda_1^{BC}}{\lambda_1^{BC} + \lambda_2^{\text{sep}}}. \tag{36}$$

For $z \in [z^*, z^{**}]$, the ray intersects the feasible segment, yielding a pure-B equilibrium. As z increases, the intersection moves from the blue endpoint toward the red endpoint: λ_1 rises and λ_2 falls. For $z < z^*$, a pure-B allocation would require $\lambda_1 < \lambda_1^{\text{sep}}$. A-buyers therefore enter, keeping the queue pair at $(\lambda_1^{\text{sep}}, \lambda_2^{AB})$ and absorbing the excess type-1 capacity. Symmetrically, for $z > z^{**}$, C-buyers enter and the queue pair remains at $(\lambda_1^{BC}, \lambda_2^{\text{sep}})$.

The following proposition collects the results.

Proposition 2. *A stable equilibrium with $\alpha_B > 0$ exists if and only if $x_{12} > \bar{x}_{12} \equiv B(\lambda_1^{\text{sep}}, \lambda_2^{\text{sep}})$. When it exists, the equilibrium is unique and takes one of three forms de-*

pending on the seller composition z:

(i) ***AB regime*** *($z < z^*$). Type-1 sellers attract both A- and B-buyers; type-2 sellers attract only B-buyers. The total queue length at type-1 sellers is $\lambda_1 = \lambda_1^{\text{sep}}$, determined by (23) with $j = 1$; the total queue length at type-2 sellers is $\lambda_2 = \lambda_2^{AB}$, determined by the requirement that $(\lambda_1^{\text{sep}}, \lambda_2^{AB})$ lies on the $\mathcal{B}$-curve (30). Both queue lengths are independent of z. The strategy share α_A decreases with z, falling to zero as $z \uparrow z^*$.*

(ii) ***Pure-B regime*** *($z^* \leq z \leq z^{**}$). All buyers adopt strategy B ($\alpha_B = 1$). The queue-length pair (λ_1, λ_2) is the unique intersection of the $\mathcal{B}$-curve and the composition ray $(1-z)\,\lambda_1 = z\,\lambda_2$. Queue lengths vary with z along the $\mathcal{B}$-curve: λ_1 is strictly increasing and λ_2 is strictly decreasing in z.*

(iii) ***BC regime*** *($z > z^{**}$). Type-1 sellers attract only B-buyers; type-2 sellers attract both B- and C-buyers. The total queue length at type-2 sellers is $\lambda_2 = \lambda_2^{\text{sep}}$, determined by (23) with $j = 2$; the total queue length at type-1 sellers is $\lambda_1 = \lambda_1^{BC}$, determined by the requirement that $(\lambda_1^{BC}, \lambda_2^{\text{sep}})$ lies on the $\mathcal{B}$-curve. Both queue lengths are independent of z. The strategy share α_C increases with z, starting from zero at $z = z^{**}$.*

In all three regimes, the own-type prices p_1^A and p_2^C are given by (11), and the cross-type prices p_1^B and p_2^B are given by (33).

The regime boundaries satisfy $0 < z^* < z^{**} < 1$. The threshold z^* is strictly decreasing in x_{12} and z^{**} is strictly increasing in x_{12}. As $x_{12} \downarrow \bar{x}_{12}$, both thresholds converge to $\lambda_1^{\text{sep}}/(\lambda_1^{\text{sep}} + \lambda_2^{\text{sep}})$; as $x_{12} \to \infty$, $z^* \to 0$ and $z^{**} \to 1$. At $x_{12} = x_{22}$, $z^{**} = 1/2$.

Proof. See Appendix A.2. □

The cross-type equilibrium can be understood as follows. In any cross-type equilibrium, queue lengths cannot fall below their values in the separating equilibrium, so that λ_j^{sep} acts as a floor. When a market with only B-buyers generates queues above this floor for both seller types, no other buyers enter and we obtain the pure-B regime. When the composition z is such that a market with only B-buyers would require a queue below the floor for one seller type, same-type buyers enter at that seller type until its queue reaches the floor (the AB or BC regime).

4 Sorting with a Continuum of Seller Types

We now allow seller types to vary continuously. Apart from replacing the two seller types by an interval, the setup is unchanged: a buyer chooses two sellers to visit, each seller posts a

price conditional on the buyer's other visit, and both visits must succeed for trade to occur. The new issue is that buyers now choose from a continuum of possible type pairs rather than the three strategies A, B, and C.

Seller types are distributed on $\mathcal{X} = [\underline{x}, \overline{x}] \subset \mathbb{R}_+$, with cumulative distribution function G and a strictly positive, continuously differentiable density g. A buyer who trades with sellers of types x and y obtains gross bundle value $f(x, y)$. We assume that f is symmetric and twice continuously differentiable, with partial derivatives $f_1(x, y) > 0$ and $f_2(x, y) > 0$ for all $(x, y) \in \mathcal{X}^2$. We also assume that $f(\underline{x}, \underline{x}) > K$.

A type-x seller posts a continuous menu $p_x(y)$, where $p_x(y)$ is the price charged to a buyer whose other visit is directed to a seller of type y.[15] We focus on deterministic matching: buyers who visit a type-x seller always direct their other visit to sellers of a single type. Let $T : \mathcal{X} \to \mathcal{X}$ denote this partner map. Because the order of the two sellers does not matter, if type x is paired with type $T(x)$, then type $T(x)$ must be paired with type x. Hence T is an involution, $T(T(x)) = x$, and therefore a bijection.[16]

Definitions 1 (equilibrium) and 2 (stability) extend directly to the continuum type space. An active bundle must give buyers exactly their entry cost K, while stability also disciplines the prices of bundles that are not used. As in the two-type model, the function B converts these pricing requirements into a queue-adjusted pairwise stability condition.

4.1 Equilibrium characterization

Visit balance. Let $\lambda(x) > 0$ denote the queue length at a type-x seller. Consider a small interval $[x, x + dx]$. The sellers in this interval receive $\lambda(x)g(x)\,dx$ visits. Their partners lie near $T(x)$ in an interval of length $|T'(x)|\,dx$ and receive $\lambda(T(x))g(T(x))|T'(x)|\,dx$ visits. These two flows come from the same buyers, so market clearing requires

$$g(x)\lambda(x) = g(T(x))|T'(x)|\lambda(T(x)). \tag{37}$$

The absolute value $|T'(x)|$ is needed because T may be decreasing, as under negative assortative matching. This condition makes buyer flows consistent along the proposed assignment. We next ask whether the assignment can be supported by prices.

Stable equilibrium. The two-type analysis provides a convenient way to state the price-support condition. For any candidate pair (x, y), $B(\lambda(x), \lambda(y))$ is the highest bundle value

[15] The function $p_x(y)$ plays the role of the discrete menu entries $p_1^A, p_1^B, p_2^B, p_2^C$.

[16] Suppose $T(x_1) = T(x_2) = y$ for some x_1, x_2. Applying T to both sides and using $T(T(x)) = x$ gives $x_1 = x_2$, so T is one-to-one. Moreover, for any y, $y = T(T(y))$, so y is in the range of T.

that can be supported at these queues by prices on the buyer's free-entry frontier without violating seller optimality. An active pair, with $y = T(x)$, must lie exactly on this bound: $f(x, T(x)) = B(\lambda(x), \lambda(T(x)))$. An inactive pair need only lie weakly below it: $f(x, y) \leq B(\lambda(x), \lambda(y))$. To express both conditions uniformly, define the slack

$$\Phi(x, y) \equiv B\Big(\lambda(x), \lambda(y)\Big) - f(x, y). \tag{38}$$

Active pairs therefore have zero slack, while inactive pairs must have nonnegative slack. The following lemma combines this price-support requirement with visit balance.

Lemma 4. *Let $T : \mathcal{X} \to \mathcal{X}$ be a differentiable map satisfying $T(T(x)) = x$, and let $\lambda : \mathcal{X} \to (0, \infty)$ be continuous. There exists a system of price menus supporting a stable equilibrium with assignment T and queue schedule λ if and only if*

(i) *the visit-balance condition (37) holds;*

(ii) *$\Phi(x, T(x)) = 0$ for every active pair;*

(iii) *$\Phi(x, y) \geq 0$ for every inactive pair $y \neq T(x)$.*

Proof. See Appendix A.3. □

The lemma distinguishes two types of conditions. Visit balance ensures that both sides of each active pair receive visits from the same buyers, while the slack conditions ensure that prices support the active pairs without violating seller optimality for any inactive pair. Before turning to PAM and NAM, we record an implication of stability used in both cases.

Local dynamics and queue monotonicity. Stability requires $\Phi(x, y) \geq 0$ for every y, with equality at the active partner $y = T(x)$. Because x is the active partner of $T(x)$, it minimizes $z \mapsto \Phi(z, T(x))$. At an interior type, this yields the following first-order condition.

$$\Phi_1(x, T(x)) = 0. \tag{39}$$

Differentiating the definition of Φ in (38) gives $\Phi_1(x, y) = B_1(\lambda(x), \lambda(y))\lambda'(x) - f_1(x, y)$. Evaluating at $y = T(x)$ and solving (39) for λ',

$$\lambda'(x) = \frac{f_1(x, T(x))}{B_1(\lambda(x), \lambda(T(x)))} > 0, \tag{40}$$

where the sign follows from $f_1 > 0$ and, by Lemma 2, $B_1 > 0$. Thus every differentiable deterministic stable equilibrium has a strictly increasing queue schedule. Holding the partner

fixed, a higher type raises the bundle value and must attract a longer queue for the stability bound to rise by the same amount.

Lemma 5. *In any differentiable deterministic stable equilibrium, $\lambda(x)$ is strictly increasing.*

Proof. See the above discussion. □

4.2 Positive assortative matching

We first consider PAM, where monotonicity pins down the assignment. Under PAM, the partner map is increasing. An increasing involution must be the identity, so $T(x) = x$ for every x.[17] Thus every seller type is paired with itself, making PAM the continuum counterpart of the separating equilibrium in the two-type model. Visit balance is then automatic, and the active-pair condition determines the queue of each type independently.

As in the separating two-type equilibrium, the active same-type pairs lie exactly on the stability bound. For each type x, let $\lambda^{\text{sep}}(x)$ denote the equilibrium queue determined by $\Phi(x, x) = 0$:

$$f(x, x) = B\Big(\lambda^{\text{sep}}(x), \lambda^{\text{sep}}(x)\Big). \tag{41}$$

Because $B(\lambda, \lambda)$ is strictly increasing from K to infinity and $f(x, x) > K$, this solution exists and is unique. Moreover, $f(x, x)$ is strictly increasing in x, so $\lambda^{\text{sep}}(x)$ is also strictly increasing. Higher types produce more in a same-type bundle and can therefore support longer queues.

It remains to check whether every off-diagonal pair (x, y) can be supported as inactive. Such a pair is the continuum analogue of the mixed bundle in the two-type model, and stability requires $\Phi(x, y) \geq 0$, or equivalently $f(x, y) \leq B(\lambda^{\text{sep}}(x), \lambda^{\text{sep}}(y))$. By Lemma 4, this gives the following characterization.

Proposition 3. *A stable PAM equilibrium exists if and only if, for all $x, y \in \mathcal{X}$, $f(x, y) \leq B(\lambda^{\text{sep}}(x), \lambda^{\text{sep}}(y))$. A queue-independent sufficient condition is*

$$f(x, y)^2 \leq f(x, x) f(y, y) \qquad \text{for all } x, y \in \mathcal{X}. \tag{42}$$

Under (42), every deterministic stable equilibrium exhibits PAM. Consequently, the deterministic stable equilibrium outcome is unique, with $T(x) = x$ and $\lambda(x) = \lambda^{\text{sep}}(x)$.

Proof. See Appendix A.4. □

[17] If $T(x_0) < x_0$, applying the increasing map T gives $T(T(x_0)) < T(x_0)$, or $x_0 < T(x_0)$, a contradiction. The case $T(x_0) > x_0$ is symmetric.

The proposition above generalizes the separating-equilibrium result for two seller types. As in that case, if every cross-pair value does not exceed the geometric mean of the corresponding diagonal values, PAM is stable. Moreover, this condition rules out any active off-diagonal pair, so PAM is the unique stable equilibrium outcome. Finally, condition (42) is implied by log-supermodularity but is weaker, because it compares only a cross pair with its two associated diagonal pairs rather than imposing all four-point comparisons.[18]

Under PAM, the active-pair condition $B(\lambda^{\text{sep}}(x), \lambda^{\text{sep}}(x)) = f(x, x)$ determines the queue of each type independently of the seller distribution g. If some types lose capacity, prices and queues therefore remain unchanged throughout the market, while entry and completed trade at each affected type fall in proportion to its supply.

4.3 Negative assortative matching

Negative assortative matching means that T is strictly decreasing, so lower types are paired with higher types. Unlike the two-type model, where there is only one mixed pair, the continuum model must determine a partner for every type. Moreover, unlike an increasing involution, a decreasing involution is not uniquely determined.[19]

Below we assume that f is submodular. The argument proceeds in three steps. First, we show that stability itself forces negative assortative matching. Second, we characterize the equilibrium by two differential equations and two boundary conditions. Importantly, the assignment equation alone makes T an involution, so this property need not be imposed separately. Finally, we establish existence and uniqueness.

Supermodularity of Φ. The first step is to show that Φ is strictly supermodular. This property will determine the ordering of partners and later help establish sufficiency. Differentiating the definition of Φ in (38) twice gives

$$\Phi_{12}(x, y) = \lambda'(x)\lambda'(y)B_{12}(\lambda(x), \lambda(y)) - f_{12}(x, y)$$

By Lemma 2, $B_{12} > 0$. Hence, if λ is strictly increasing and f is submodular, both terms work in the same direction and $\Phi_{12} > 0$. This yields the following lemma.

[18] Condition (42) compares each cross-type bundle only with the two associated same-type bundles. For example, let $\mathcal{X} = [0, 2]$ and $f(x, y) = Ce^{2(x+y)}/(1 + (x - y)^2)$, with $C > K$. This function is positive, symmetric, and increasing, and satisfies $f(x, y)^2 \leq f(x, x)f(y, y)$ for all x, y. However, $\partial_{xy} \log f(x, y) = 2(1 - (x - y)^2)/(1 + (x - y)^2)^2 < 0$ whenever $|x - y| > 1$, so f is not log-supermodular.

[19] For any continuously differentiable, strictly increasing bijection $H : \mathcal{X} \to [0, 1]$ with $H' > 0$, the map $T(x) = H^{-1}(1 - H(x))$ is a strictly decreasing involution.

Lemma 6. *Suppose that f is submodular. If $\lambda(x)$ is strictly increasing, then $\Phi(x, y)$ defined by (38) is strictly supermodular.*

Proof. See the above discussion. □

Necessity of NAM. Strict supermodularity of Φ then implies that the equilibrium must exhibit NAM. Take $x_1 < x_2$ and write $y_i = T(x_i)$. Note that $y_1 \neq y_2$ since T is a bijection. If $y_1 < y_2$, strict supermodularity gives

$$\Phi(x_1, y_1) + \Phi(x_2, y_2) > \Phi(x_1, y_2) + \Phi(x_2, y_1).$$

The left-hand side is zero because both pairs are active. Stability, however, requires the right-hand side to be nonnegative because the crossed pairs are inactive. This is impossible. Hence $y_1 > y_2$ whenever $x_1 < x_2$, so T is strictly decreasing.

Lemma 7. *Suppose that f is submodular. Then every differentiable deterministic stable equilibrium exhibits NAM.*

Proof. See the above discussion. □

Having determined the direction of matching, we can now remove the absolute value from visit balance and combine it with the queue equation.

The differential system. Since $T'(x) < 0$, the visit-balance condition (37) becomes

$$T'(x) = -\frac{g(x)\lambda(x)}{g(T(x))\lambda(T(x))}. \tag{43}$$

Equations (43) and (40) form a closed differential system for the assignment and queue-length functions (T, λ), which we solve subject to the boundary conditions

$$T(\underline{x}) = \overline{x}, \qquad \Phi(\underline{x}, \overline{x}) = 0. \tag{44}$$

The first condition pairs the lowest type with the highest, while the second places this active pair on the stability bound. The latter provides one relation between $\lambda(\underline{x})$ and $\lambda(\overline{x})$, but does not determine either queue separately.

Sufficiency. The differential equations and boundary conditions above are necessary. We now show that strict supermodularity of Φ also makes them sufficient.

We first record that the assignment equation alone implies the involution property.

Lemma 8. *Let $T : \mathcal{X} \to \mathcal{X}$ be differentiable. Suppose that $T'(x)$ is given by equation (43). Then T is an involution, i.e., $T(T(x)) = x$ for any $x \in \mathcal{X}$.*

Proof. See Appendix A.5. □

The following proposition characterizes equilibrium using the differential system and the two boundary conditions.

Proposition 4. *Suppose that f is submodular. Let $T : \mathcal{X} \to \mathcal{X}$ and $\lambda : \mathcal{X} \to (0, \infty)$ be differentiable. Then (T, λ) is a stable deterministic equilibrium if and only if it satisfies (43), (40), and the boundary conditions (44).*

Proof. See Appendix A.6. □

The proof contains two steps. First, the differential equations and the boundary conditions imply that $\Phi(x, T(x)) = 0$ for all $x \in \mathcal{X}$. Next, strict supermodularity of Φ implies that $\Phi(x, y) > 0$ for any $y \neq T(x)$. Hence, Lemma 4 yields an equilibrium. The argument closely parallels standard arguments in classical matching theory. To see the second step, note that, since $\Phi_1(x, \cdot)$ is strictly increasing, if $T(x) > y$, then $\Phi_1(x, y) < \Phi_1(x, T(x)) = 0$, whereas if $T(x) < y$, the inequality is reversed. Because T is a strictly decreasing involution, these two cases are equivalent to $x < T(y)$ and $x > T(y)$, respectively. Hence, $x \mapsto \Phi(x, y)$ decreases until $x = T(y)$ and increases thereafter. The active type $T(y)$ is therefore its unique global minimizer. Since $\Phi(T(y), y) = 0$, every inactive pair has strictly positive slack.

Equilibrium existence and uniqueness. Proposition 4 reduces the equilibrium to a boundary-value problem. We now show that this problem has exactly one solution.

Proposition 5. *Suppose f is submodular. Then there exists a unique differentiable deterministic stable equilibrium (T, λ), and it exhibits NAM.*

Proof. See Appendix A.7. □

This result extends the logic of cross-type matching in the two-type model. With two types, the active mixed pair lies on a single $\mathcal{B}$-curve, and the seller composition determines where on this curve the equilibrium queues are. With a continuum of types, each active pair lies on its own stability bound. The first-order condition (40) determines how queues change with seller type, while visit balance determines which types are paired, in the sense that the visit-weighted mass of types below x must equal the visit-weighted mass of types above $T(x)$. Together, these conditions determine a unique partner map and queue schedule.

The proof idea. Because the endpoint condition $\Phi(\underline{x}, \overline{x}) = 0$ does not pin down either endpoint queue, the appendix starts instead from the unique fixed point c of the decreasing involution, defined by $T(c) = c$. At this point the active pair is diagonal, so its queue is uniquely determined by $\lambda(c) = \lambda^{\text{sep}}(c)$; see (41).

Each conjectured value of c therefore gives a fully specified initial-value problem. The proof integrates toward lower types, recovers the upper branch from the involution, and varies c until $T(\underline{x}) = \overline{x}$. If c is close to $\underline{x}$, the solution reaches the lower endpoint first; if c is close to $\overline{x}$, its partner reaches the upper endpoint first. Continuity therefore yields an intermediate value of c for which both endpoints are reached simultaneously. Finally, solutions starting from different values of c remain ordered and cannot cross, so at most one value of c can satisfy the endpoint condition.

Comparative statics. Holding f and K fixed, consider two seller densities g_a and g_b, with unique stable NAM equilibria (T_i, λ_i), self-match types c_i defined by $T_i(c_i) = c_i$, and profits π_i, for $i = a, b$.

Proposition 6. *Suppose f is submodular. If g_b/g_a is nondecreasing and nonconstant, then $c_b > c_a$ and there is a unique $x^* \in (c_a, c_b)$ such that, for $x < x^*$, $\lambda_b(x) > \lambda_a(x)$, and $\pi_b(x) > \pi_a(x)$, with equality at x^* and the reversed inequalities for $x > x^*$.*

If g_b/g_a is nonincreasing and nonconstant, then $c_b < c_a$ and a unique $x^* \in (c_b, c_a)$ satisfies the opposite queue and profit inequalities.

Proof. See Appendix A.8. □

When higher types become relatively more abundant, the self-match type rises, and sellers between the old and new self-match types switch from lower-type to higher-type partners. Queues and profits increase below x^* and decrease above it, reflecting the redistribution of trading opportunities across seller types.

4.4 PAM under N buyer visits

The preceding analysis assumes that each buyer visits two sellers. We next show that the PAM argument extends directly to any fixed number of visits.

Suppose that each buyer visits exactly $N \geq 2$ sellers and obtains value only if all N visits succeed. Let $(x_1, \ldots, x_N) \in \mathcal{X}^N$ denote the seller types in a bundle, and let $f(x_1, \ldots, x_N) > 0$ denote its value. We assume that f is symmetric, strictly increasing in every argument, and satisfies $f(\underline{x}, \ldots, \underline{x}) > K$. As before, each seller's price may depend on the types of the other $N - 1$ sellers in the bundle.

The same argument as in the two-visit model yields the following N-variate stability bound; see the proof of Proposition 7 for details:

$$B(\lambda_1, \ldots, \lambda_N) \equiv \frac{K}{\prod_{i=1}^{N} q(\lambda_i)} \left[\sum_{i=1}^{N} \frac{1}{\varepsilon_m(\lambda_i)} - (N-1) \right]. \tag{45}$$

Its interpretation is unchanged: a bundle can be supported at the given queues if and only if its value does not exceed $B(\lambda_1, \ldots, \lambda_N)$, with equality for an active bundle. The function B is symmetric and strictly increasing in every argument. It is also strictly supermodular and strictly log-submodular in every pair of arguments.

We focus on PAM, whose meaning extends directly to this setting.[20] Under PAM, all N sellers in a bundle have the same type. Let $\lambda^{\text{sep}}(x)$ denote the unique queue satisfying

$$f(x, \ldots, x) = B\Big(\lambda^{\text{sep}}(x), \ldots, \lambda^{\text{sep}}(x)\Big). \tag{46}$$

The following proposition extends the two-visit PAM result.

Proposition 7. *A stable N-visit PAM equilibrium exists if and only if for every $(x_1, \ldots, x_N)$,*

$$f(x_1, \ldots, x_N) \leq B\Big(\lambda^{\text{sep}}(x_1), \ldots, \lambda^{\text{sep}}(x_N)\Big). \tag{47}$$

A sufficient condition is

$$f(x_1, \ldots, x_N)^N \leq \prod_{i=1}^{N} f(x_i, \ldots, x_i) \qquad \text{for every } (x_1, \ldots, x_N) \in \mathcal{X}^N. \tag{48}$$

This condition is implied by log-supermodularity of f.

Moreover, under (48) the stable equilibrium is unique and exhibits PAM.

Proof. See Appendix A.9. □

The intuition is the same as with two visits. Under PAM, each diagonal bundle value determines the queue for that type, and (47) asks whether these queues can support every non-diagonal bundle as inactive. Hence, also with N visits, the complementarity condition

[20] With two visits, deterministic NAM is summarized by a decreasing involution T. With $N \geq 3$, each seller must instead be assigned $N-1$ partners, so no single order-reversing map describes the resulting team. For example, Boerma et al. (2025) characterize negative sorting in a frictionless team-production model with a submodular technology. Their assignment combines a mixed region, where teams have equal expected output losses, with countermonotonic regions, including those in which high-skill workers join low-skill coworkers at high-productivity firms. In the mixed region, identical firms may hire different worker combinations, so sorting need not pin down a unique team composition for each firm type.

that guarantees PAM under competitive search is multiplicative, and ordinary supermodularity is not sufficient.

4.5 Re-sorting after a change in seller supply

In this subsection, we compare equally large supply reductions concentrated among high or low types. We ask whether re-sorting reduces the resulting losses and who benefits. Korovkin et al. (2026) show that the formation of new production links reduces the losses from disruptions. We focus on changes in the two supplier types a buyer combines. We consider the NAM case and compare the new stable equilibrium with an outcome in which buyers have to keep combining the original type pairs, while prices, queues, and entry adjust in both cases. We do not consider PAM, because under PAM each type remains paired with the same seller type and its queue does not depend on the relative supply of seller types (see Section 4.2).

Specification. In the baseline, sellers have unit mass, and their types are normalized to be uniformly distributed on $[0, 1]$. As before, each seller supplies one unit, buyers pay K to enter and make two visits, and trade only takes place if both visits succeed. Sellers post price menus and have no costs. Meetings follow the urn-ball function $m(\lambda) = 1 - e^{-\lambda}$, so that the success probability of a visit is $q(\lambda) = m(\lambda)/\lambda$. Bundle values are given by

$$v(x) = v_0 + (1 - v_0)\frac{(x + b)^\gamma - b^\gamma}{(1 + b)^\gamma - b^\gamma}, \qquad f(x, y) = \frac{v(x) + v(y)}{2}, \qquad b = .05. \tag{49}$$

Hence, $v(0) = v_0$ and $v(1) = 1$. Since $f_{12} = 0$, the economy is in the NAM case. The parameter γ determines how steeply values increase with seller type. The shift $b > 0$ ensures that the slope of v is positive at the lowest type and affects its curvature, but not its values at the endpoints. The value $b = .05$ is our own choice, not taken from the literature. The qualitative results for both supply shocks also hold with $b = .01$ and $b = .10$ after recalibrating (v_0, γ, K) to the same baseline targets.

Seller supply: scale versus composition. We write the density of seller supply as $h(x) = Mg(x)$, where M is the total mass of sellers and g integrates to one. Holding values, the entry cost, and the meeting technology fixed, a proportional change in h leaves partner types, queues, on-path prices, and expected profits per seller unchanged, because the meeting technology has constant returns to scale and buyers enter freely. Only the mass of buyers and aggregate trade change in proportion to M. Re-sorting and gains from scarcity are

therefore driven by changes in the *relative* supply of seller types. We consider

$$h_\delta^H(x) = 1 - \delta x^4, \qquad h_\delta^L(x) = 1 - \delta(1-x)^4, \qquad M_\delta = 1 - \delta/5. \tag{50}$$

The first profile concentrates the reduction in supply among high types and the second among low types. Both reduce the mass of sellers by the same amount, so they only differ in which types become scarce. Below, h_δ stands for either profile and $g_\delta = h_\delta/M_\delta$ is the corresponding density. We set $\delta = .325$, which reduces the mass of sellers by 6.5%. This value of δ is based on the capital-destruction scenario in Carvalho et al. (2021), which we only use as a numerical benchmark. Applying this scenario to seller supply and choosing the exposure profile are our own modeling choices, so they should not be interpreted as estimates of seller exits. Each remaining seller still supplies one unit.

Fixed partner types. Let T_0 denote the NAM partner map in the baseline. We refer to the equilibrium under h_δ as the flexible outcome, denoted by E. In the fixed-partner outcome, denoted by F, only the pairs $(x, T_0(x))$ can be formed, but prices, queues, and entry are still determined endogenously. The queues in this outcome satisfy the active-pair condition $B\big(\lambda^F(x), \lambda^F(T_0(x))\big) = f(x, T_0(x))$, where B is the queue-adjusted stability bound. They also satisfy visit balance, $g_\delta(x)\lambda^F(x) = g_\delta(T_0(x))|T_0'(x)|\lambda^F(T_0(x))$, since each buyer sends one visit to each side of the pair. These two conditions uniquely determine the two queues, because the B-curve is downward sloping while visit balance defines a ray through the origin. The supporting prices then follow from buyer free entry and seller optimality. Note that this is an equilibrium of the restricted economy, so we do not require stability against the excluded pairs. Moreover, only the *types* of partners are fixed, while individual sellers and trading terms can change. Since both outcomes are computed for the same h_δ, comparing them isolates the effect of being able to re-sort.

Outcomes. For both outcomes, let $\lambda(x)$, $\mu(x) = m(\lambda(x))q(\lambda(T(x)))$, and $\pi(x)$ denote the expected number of visits, the probability of trade, and the expected profit of a type-x seller. The aggregate outcomes are then

$$\begin{aligned} Q &= \tfrac{1}{2}\int_0^1 h_\delta(x)\mu(x)\,dx, \quad Y = \tfrac{1}{2}\int_0^1 h_\delta(x)\mu(x)f(x,T(x))\,dx, \\ L &= \tfrac{1}{2}\int_0^1 h_\delta(x)\lambda(x)\,dx, \quad R = \int_0^1 h_\delta(x)\pi(x)\,dx = Y - KL. \end{aligned} \tag{51}$$

Here Q is the mass of completed bundles, Y their total gross value, L the mass of buyers who enter, and R aggregate seller profits. The factors $1/2$ appear because each bundle uses

two sellers and each buyer makes two visits. Finally, $R = Y - KL$ follows from buyer free entry and the absence of seller costs. Dividing these totals by seller mass M_δ gives the corresponding amounts per seller. In particular, Q/M_δ is completed bundles per seller, whereas Y/Q is average value per completed bundle.

Calibration. We choose (v_0, γ, K) jointly to match three targets in the baseline equilibrium: an entry-expenditure share of $KL/Y \simeq 12.9\%$, and profit shares of 3.4% and 78% for the bottom and top quintiles of sellers. The first target determines the size of entry expenditure relative to bundle value, while the other two determine how concentrated profits are across seller types.

These targets are based on Demir et al. (2024). Their Table 7 shows that the 20% of firms that pay the lowest average wages account for 3% of network sales, compared with 78% for the 20% of firms that pay the highest wages. We interpret higher-wage firms as higher-type sellers and use these numbers to capture how strongly sales are concentrated among high types. Since sellers in our model have no costs, their sales revenue equals their profits. For the expenditure target, we use the share of supplier search in the combined cost of finding suppliers and buying inputs that is implied by their model. Applying this share to entry expenditure relative to bundle value is an assumption on our part, since search costs are not directly observed.[21] The fitted values, rounded for presentation, are $v_0 = .015$, $\gamma = 7.29$, and $K = .006 < v_0$.

Aggregate effects. Table 1 distinguishes the effect of a proportional reduction in supply from the effect of the change in composition. A proportional reduction to $h(x) = .935$ leaves all outcomes per seller unchanged and lowers Q, Y, L, and R by 6.5%. Panel A reports the effects when the reduction is concentrated among high types. The number of completed bundles falls by 12.8% with fixed partners and by 7.2% with flexible sorting. The column Q/M removes the common scale effect and shows that the number of completed bundles per seller falls by 6.7% and 0.8%, respectively. Comparing the two outcomes for the same supply after the change, flexible sorting completes 6.4% more bundles than fixed partner types. Panel B reports a reduction of the same size concentrated among low types. Completed

[21] Quintiles weight firms equally. We target 3.4%, which lies within the rounding interval of the reported 3%, because we could not find parameter values that jointly match exactly 3% and the other targets with $K < v_0$. The sales shares are only used to discipline the concentration of profits and are not informative about which suppliers are combined. For the expenditure target, equations (8)–(11) of Demir et al. (2024), with $\alpha_m = .33$, $\sigma = 5$, and $1/\beta_m = .59$, imply that expenditures on supplier search and on manufacturing inputs equal .04 and .26 per unit of revenue. Supplier search therefore accounts for $.039/(.039 + .26) \simeq 12.9\%$ of the combined budget, excluding customer search. We map this share to KL/Y because buyer free entry implies $Y = KL + R$.

bundles now fall by 3.5% with fixed partners and by 5.8% with flexible sorting, so flexible sorting completes 2.4% fewer bundles than fixed partner types.

Table 1: NAM: changes in seller supply, with and without re-sorting

	Completed bundles Q	Bundles per seller Q/M	Gross value Y	Mean value Y/Q	Buyer entry L	Seller profits R
Proportional supply reduction	-6.50	0.00	-6.50	0.00	-6.50	-6.50
Panel A: reduction concentrated among high types						
Fixed partners	-12.77	-6.70	-14.99	-2.54	-17.14	-14.67
Flexible sorting	-7.20	-0.75	-17.47	-11.07	-13.60	-18.04
Flexible relative to fixed	6.38	6.38	-2.92	-8.74	4.28	-3.95
Panel B: reduction concentrated among low types						
Fixed partners	-3.48	3.23	-5.55	-2.14	-1.64	-6.13
Flexible sorting	-5.77	0.78	-2.74	3.21	-3.22	-2.67
Flexible relative to fixed	-2.36	-2.36	2.97	5.46	-1.61	3.68

Percentage changes from the baseline ($M_0 = 1$), except in the rows labelled "flexible relative to fixed". The proportional reduction uses $h = .935$, Panel A uses $h_\delta^H = 1 - .325x^4$, and Panel B uses $h_\delta^L = 1 - .325(1-x)^4$. All have $M = .935$. For each variable Z, the rows labelled "flexible relative to fixed" report $100(Z^E/Z^F - 1)$, not a percentage-point difference. Entries are calculated before rounding.

In Panel A, these additional transactions do not raise total value. The average value of a completed bundle falls by 11.07% with flexible sorting, compared with 2.54% with fixed partners. Flexible sorting therefore lowers gross value by 2.92% relative to fixed partners, even though more bundles are completed. It also raises buyer entry, and hence entry expenditure for a given K, by 4.28%. Together, lower gross value and higher entry expenditure reduce aggregate profits by 3.95% relative to fixed partners. The reason is that re-sorting allows lower types to complete more sales, while the most valuable types complete fewer. Hence, the number of completed bundles rises, but their total value falls.

In Panel B the same forces work in the opposite direction. When low types become scarce, re-sorting moves sellers toward more valuable partners, so that the average value of a completed bundle rises by 3.21%, whereas it falls by 2.14% with fixed partner types. Flexible sorting now raises gross value by 2.97% and aggregate profits by 3.68% relative to fixed partner types, while it completes 2.4% fewer bundles and lowers buyer entry by 1.61%.

In both panels, re-sorting therefore moves the market along the same trade-off between the number of completed bundles and their average value. Which types become scarce determines the direction. When high types are scarce, sellers move to less valuable partners, so that more bundles are completed at a lower average value. When low types are scarce, sellers move to more valuable partners, so that fewer bundles are completed at a higher average value.

Partner adjustment and the distribution of profits. The left panel of Figure 2 shows the partner maps T_0 and T_δ for the high-type supply reduction in Table 1, Panel A. In the calibrated economy, the new map lies below the old one at every interior type, and the self-match type c, defined by $T(c) = c$, falls from .70 to .67. In the fixed-partner outcome, the partner map remains T_0, although the queues change. For the low-type supply reduction in Panel B, not shown in the figure, the new map lies above the old one, and the self-match type rises from .70 to .71.

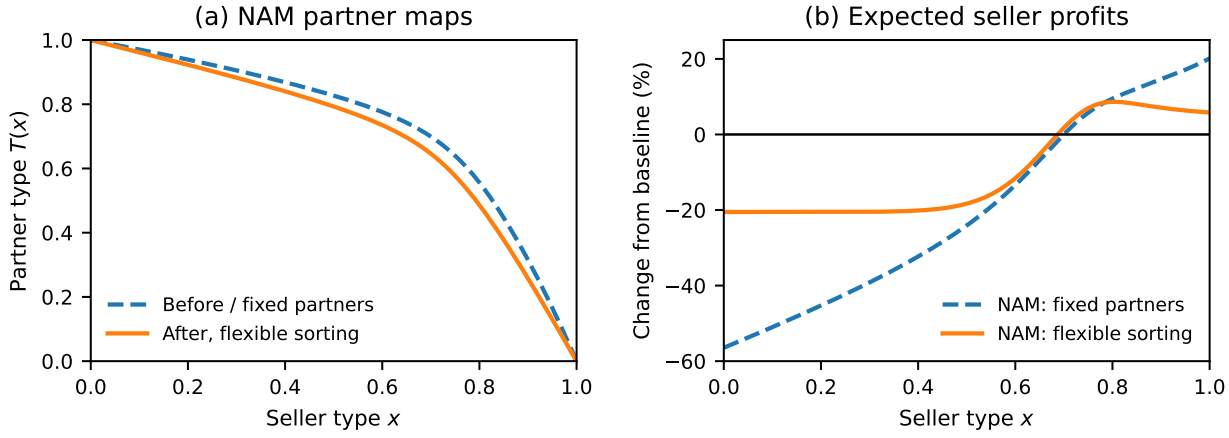


Figure 2: High-type supply reduction (Table 1, Panel A). Left: the baseline partner map, retained with fixed partners, and the new flexible assignment. Right: changes in expected profit per seller relative to baseline; the zero line marks unchanged profits. Both panels hold seller type x fixed. The right panel excludes changes in seller numbers.

The right panel shows who gains and who loses. With fixed partners, the scarcity of high types reduces the trading opportunities of the lower types that are paired with them. Holding the weights of the types at their baseline values, average expected profits in the bottom quintile fall by 50.96% with fixed partners and by 20.52% with flexible sorting, even though the supply of this group falls by only 0.01%. Re-sorting thus substantially limits these indirect losses. For the top quintile, however, re-sorting limits the gains from scarcity. Average expected profits in this group rise by 6.74% with flexible sorting, compared with 16.12% with fixed partners. Note that these gains per seller are not large enough to prevent total profits of the top quintile from falling, as its supply declines.[22]

In Panel B, the incidence is reversed as well. Now the low types become scarce and earn

[22]The same trade-off arises for $\delta = .22$ and $\delta = .43$. In both cases, flexible sorting increases the number of completed bundles, but lowers gross value and aggregate profits relative to fixed partners. The signs of these effects are also unchanged when we vary the calibration restrictions, the entry cost, or the concentration of the supply reduction among high types, or when we make f strictly submodular. Panel B shows that they reverse when the reduction is concentrated among low types instead.

scarcity rents, while the high types that were paired with them lose. With fixed partner types, average expected profits in the bottom quintile rise by 71.81% and those in the top quintile fall by 9.71%. Re-sorting again dampens both effects, to 11.15% and −3.14%, respectively.

Discussion. The two scenarios show why transaction counts alone can misrepresent the losses from a supply disruption. Since buyers earn zero expected profits after paying entry costs and sellers have no costs, aggregate seller profits $R = Y - KL$ equal total net surplus. Relative to fixed partners, re-sorting lowers net surplus by 3.95% when high types become scarce and raises it by 3.68% when low types become scarce. In the former case, more bundles are completed; in the latter, fewer are completed. Thus preserving transactions and preserving net surplus need not go together.

The market nevertheless re-sorts, because the fixed-partner outcome is not stable. The stable equilibrium under h_δ is unique (Proposition 5) and has a different partner map, so under fixed partner types there must be an excluded pair that can be priced such that its buyers break even and both of its sellers gain. Individual sellers therefore have an incentive to form new pairs, whether or not this raises total net surplus. This is in line with Albrecht et al. (2023), who show that competitive search with two visits is constrained inefficient, even when sellers are homogeneous. An important implication is that the number of transactions that survive a disruption can give a misleading picture of its welfare costs. In these two counterfactuals, the outcome with more completed bundles has lower net surplus.

5 Conclusion

This paper studies a competitive search market in which heterogeneous sellers post a menu of prices and buyers direct their search to two sellers in order to purchase a bundle of two goods (or workers). Buyers care not only about obtaining two goods, but also about the types of sellers from which those goods are obtained. This adds a portfolio dimension to the search problem. A buyer's willingness to pay at one seller depends on the type and price of the other seller. This links prices across submarkets in a way that is absent from standard single-visit search models.

We consider two seller types and then generalize to a continuum of types. In the two-type case, the stable equilibrium outcome is unique away from the knife edge and is governed by a single threshold $\bar{x}_{12}$ in the cross-type bundle value. Below this threshold, buyers purchase both goods from the same seller type, and the two same-type markets operate separately. Above the threshold, cross-type matching becomes active with three possible equilibrium configurations. With a continuum of seller types, we derive an exact stability condition

and sufficient conditions for PAM and NAM that also guarantee uniqueness of equilibrium. Although we use two visits to simplify the exposition, we show that the PAM characterization extends to any fixed number of visits. A corresponding multiplicative complementarity condition, implied by log-supermodularity, guarantees a unique stable PAM equilibrium.

Finally, the sorting pattern determines how the market responds to a fall in the supply of some seller types. Under PAM, prices, queues, and expected profits per remaining seller are unchanged. Completed trade and aggregate profits at each affected type fall in proportion to its supply, without affecting other types. Under NAM, buyers re-sort, and prices and queues change at all types. In a calibrated example in which high types become scarce, re-sorting keeps more transactions going and protects the low types that were paired with the scarce high types, but it lowers total net surplus relative to an outcome with fixed partner types. When low types become scarce instead, re-sorting completes fewer bundles of higher value and raises total net surplus. The number of transactions that survive a disruption can therefore give a misleading picture of its welfare costs.

Appendix A Proofs

A.1 Proof of Lemma 2

The definition in (27) implies that $B(\lambda_1, \lambda_2) = B(\lambda_2, \lambda_1)$, i.e., B is symmetric. Since both q and ε_m are strictly decreasing, B is strictly increasing in each argument. Furthermore,

$$B_1(\lambda_1, \lambda_2) = K \frac{1}{q(\lambda_2)} \Big[-\frac{q'(\lambda_1)}{q(\lambda_1)^2} \Big(\frac{1}{\varepsilon_m(\lambda_1)} + \frac{1}{\varepsilon_m(\lambda_2)} - 1 \Big) - \frac{1}{q(\lambda_1)} \frac{\varepsilon_m'(\lambda_1)}{\varepsilon_m(\lambda_1)^2} \Big],$$

which is strictly increasing in λ_2. Hence, B is strictly supermodular. Next,

$$\frac{\partial \log B(\lambda_1, \lambda_2)}{\partial \lambda_1} = -\frac{q'(\lambda_1)}{q(\lambda_1)} + \left(\frac{1}{\varepsilon_m(\lambda_1)} + \frac{1}{\varepsilon_m(\lambda_2)} - 1 \right)^{-1} \frac{d}{d\lambda_1} \left(\frac{1}{\varepsilon_m(\lambda_1)} \right)$$

which is strictly decreasing in λ_2. Hence, $\log B$ is strictly submodular, which then implies that $B(\lambda_1, \lambda_1) B(\lambda_2, \lambda_2) < B(\lambda_1, \lambda_2) B(\lambda_2, \lambda_1) = B(\lambda_1, \lambda_2)^2$. Hence we have (29). □

A.2 Proof of Proposition 2

The main-text argument establishes that, away from the knife edge, $\alpha_B > 0$ requires $x_{12} > \bar{x}_{12}$, identifies the feasible segment of the $\mathcal{B}$-curve, and shows that its two endpoints support the AB and BC regimes while its interior supports pure B. It also gives the unique

composition thresholds z^* and z^{**}. We complete the proof by verifying market clearing, seller optimality, and the comparative statics.

In the AB regime, equations (1) and (2) with $\alpha_C = 0$ imply $(1-z)\lambda_1^{\text{sep}} = (1+\alpha_A)L$ and $z\lambda_2^{AB} = (1-\alpha_A)L$. Hence

$$\alpha_A = \frac{(1-z)\lambda_1^{\text{sep}} - z\lambda_2^{AB}}{(1-z)\lambda_1^{\text{sep}} + z\lambda_2^{AB}},$$

which lies in $(0,1)$ exactly when $z < z^*$ and decreases with z. Symmetrically, in the BC regime

$$\alpha_C = \frac{z\lambda_2^{\text{sep}} - (1-z)\lambda_1^{BC}}{z\lambda_2^{\text{sep}} + (1-z)\lambda_1^{BC}},$$

which lies in $(0,1)$ exactly when $z > z^{**}$ and increases with z. At $z = z^*$ and $z = z^{**}$ the corresponding same-type share is zero. In the pure-B regime, market clearing is exactly (35). Thus the stated queues and shares determine a positive buyer mass L in every regime. Both thresholds lie in $(0,1)$, and $z^* < z^{**}$ follows from $\lambda_1^{BC} > \lambda_1^{\text{sep}}$ and $\lambda_2^{AB} > \lambda_2^{\text{sep}}$.

The prices in equations (11) and (33) place all three strategies on their free-entry frontiers. The queue-floor conditions ensure that B-buyers yield each seller type its highest visitor value, with equality for a same-type group precisely when that group is active. The seller first-order condition then holds at the stated queue, and strict concavity of m makes it sufficient. The geometry of the $\mathcal{B}$-curve and the formulas above give uniqueness of the equilibrium outcome.

Finally, strict monotonicity of B implies that λ_2^{AB} and λ_1^{BC} increase with x_{12}, so z^* decreases and z^{**} increases. As $x_{12} \downarrow \bar{x}_{12}$, both endpoints converge to $(\lambda_1^{\text{sep}}, \lambda_2^{\text{sep}})$; as $x_{12} \to \infty$, the free endpoint queues diverge, so $z^* \to 0$ and $z^{**} \to 1$. Moreover, $z^{**} = 1/2$ if and only if $\lambda_1^{BC} = \lambda_2^{\text{sep}}$, which by (28) is equivalent to $x_{12} = x_{22}$. □

A.3 Proof of Lemma 4

Write $q_x = q(\lambda(x))$. For a candidate pair (x,y), define $V_x(y) = q_y[f(x,y) - p_y(x)]$ and $V_y(x) = q_x[f(x,y) - p_x(y)]$. Buyer free entry is equivalent to

$$q_x V_x(y) + q_y V_y(x) = q_x q_y f(x,y) + K. \tag{52}$$

Let $V_x^* = V_x(T(x))$ denote the visitor value delivered by the active partner. Seller optimality requires $V_x(y) \le V_x^*$ for every y and makes $\lambda(x)$ solve $\max_{\widetilde{\lambda}\ge 0}\{m(\widetilde{\lambda})V_x^* - \widetilde{\lambda}K\}$. Since m is strictly concave, this is equivalent to

$$V_x^* = \frac{K}{m'(\lambda(x))}. \tag{53}$$

Visit balance is necessary. For an active pair, with $y = T(x)$, both visitor-value bounds are attained. Hence (52) and (53) give $q_x q_{T(x)} f(x, T(x)) + K = q_x V_x^* + q_{T(x)} V_{T(x)}^* = \frac{K}{\varepsilon_m(\lambda(x))} + \frac{K}{\varepsilon_m(\lambda(T(x)))}$. Rearranging yields $f(x, T(x)) = B(\lambda(x), \lambda(T(x)))$.

Now consider an inactive pair (x, y). A supporting price split must satisfy the free-entry condition (52) and the two visitor-value bounds $V_x(y) \leq V_x^*$, and $V_y(x) \leq V_y^*$. Such a split exists if and only if $q_x q_y f(x, y) + K \leq q_x V_x^* + q_y V_y^*$. Necessity follows by multiplying the two visitor-value bounds by q_x and q_y, respectively, and adding them. Conversely, if the inequality holds, the difference between its right- and left-hand sides can be subtracted from one bound or divided between the two. The resulting visitor values satisfy both bounds and (52), and the corresponding prices follow from the definitions of $V_x(y)$ and $V_y(x)$. Finally, using $q_x V_x^* = \frac{K}{\varepsilon_m(\lambda(x))}$, the inequality is equivalent to $f(x, y) \leq B(\lambda(x), \lambda(y))$. This proves necessity of the three conditions.

Conversely, suppose they hold and define, for every pair (x, y), $q_x V_x(y) = q_x V_x^* - \frac{1}{2} q_x q_y \Phi(x, y)$, and $q_y V_y(x) = q_y V_y^* - \frac{1}{2} q_x q_y \Phi(x, y)$. Because $q_x V_x^* + q_y V_y^* = q_x q_y B(\lambda(x), \lambda(y)) + K$, these values satisfy (52). They equal the on-path bounds when $y = T(x)$ and do not exceed them otherwise. The corresponding prices, obtained from the definitions of $V_x(y)$ and $V_y(x)$, are continuous, and no seller can obtain more than $m(\widetilde{\lambda}) V_x^* - \widetilde{\lambda} K$, which is maximized at $\lambda(x)$. Visit balance supplies the buyer flows, so the resulting menus support a stable equilibrium. □

A.4 Proof of Proposition 3

Under PAM, $T(x) = x$ and the active-pair equations uniquely determine $\lambda(x) = \lambda^{\text{sep}}(x)$. Lemma 4 therefore gives the exact condition in the proposition. Writing $\lambda_x = \lambda^{\text{sep}}(x)$ and $\lambda_y = \lambda^{\text{sep}}(y)$, Lemma 2 and the diagonal active equations imply $B(\lambda_x, \lambda_y)^2 \geq B(\lambda_x, \lambda_x) B(\lambda_y, \lambda_y) = f(x, x) f(y, y)$. Thus (42) is sufficient for PAM.

For the uniqueness result, consider any deterministic stable equilibrium and suppose $y = T(x) \neq x$. The active cross pair and the two inactive diagonal pairs satisfy $f(x, y) = B(\lambda(x), \lambda(y))$, $f(x, x) \leq B(\lambda(x), \lambda(x))$, and $f(y, y) \leq B(\lambda(y), \lambda(y))$. Hence

$$f(x, y)^2 = B(\lambda(x), \lambda(y))^2 \geq B(\lambda(x), \lambda(x)) B(\lambda(y), \lambda(y)) \geq f(x, x) f(y, y).$$

Condition (42) implies equality throughout. If $\lambda(x) \neq \lambda(y)$, the first inequality is strict by Lemma 2, a contradiction. If $\lambda(x) = \lambda(y)$, equality throughout forces $f(x, x) = f(y, y)$, which implies $x = y$ because $f(x, x)$ is strictly increasing. Thus no off-diagonal pair is active. Hence $T(x) = x$. □

A.5 Proof of Lemma 8

Let $\Gamma(x) \equiv \int_{\underline{x}}^{x} g(s)\lambda(s)\,ds$, where $g(s)\lambda(s) > 0$ is the density of visits received by type-s sellers. Thus Γ is strictly increasing and $\Gamma(\underline{x}) = 0$. By (43), $\frac{d}{dx}\Big[\Gamma(T(x)) + \Gamma(x)\Big] = 0$. Hence $\Gamma(T(x)) + \Gamma(x) = C$ for some constant C. Evaluating this identity at the two endpoints gives $\Gamma(T(\underline{x})) = \Gamma(T(\overline{x})) + \Gamma(\overline{x})$. Since $\Gamma(x)$ is strictly increasing, $\Gamma(T(\underline{x})) \leq \Gamma(\overline{x}) \leq \Gamma(T(\overline{x})) + \Gamma(\overline{x})$, where the equalities hold if and only if $T(\underline{x}) = \overline{x}$ and $T(\overline{x}) = \underline{x}$.

Finally, evaluating $\Gamma(T(x)) + \Gamma(x) = C$ at $T(x)$ gives $\Gamma(T(T(x))) + \Gamma(T(x)) = C$. Comparing the two identities yields $\Gamma(T(T(x))) = \Gamma(x)$. Since Γ is strictly increasing, it follows that $T(T(x)) = x$. Thus T is an involution. □

A.6 Proof of Proposition 4

Necessity follows from the main-text arguments: Lemma 7 gives $T' < 0$, visit balance gives (43), tangency gives (40), and the active endpoint pair satisfies (44).

Conversely, suppose the differential equations and (44) hold. Equation (43) gives $T' < 0$, and Lemma 8 gives the involution property. The queue equation gives $\lambda' > 0$ and $\Phi_1(x, T(x)) = 0$. Applying it at $T(x)$ and using symmetry and the involution gives $\Phi_1(T(x), x) = 0$, so $\Phi(x, T(x))$ is constant; the boundary condition sets this constant equal to zero. Lemma 6 then implies that, for fixed y, $x \mapsto \Phi(x, y)$ has its unique minimum at $T(y)$. Thus every inactive pair has strictly positive slack. Finally, (43) is visit balance, so Lemma 4 yields a stable deterministic equilibrium. □

A.7 Proof of Proposition 5

By Proposition 4, it is enough to solve the differential system subject to the endpoint conditions in (44). We use the standard shooting method for boundary-value problems (Hastings and McLeod, 2012). We guess the type c that is matched with itself, so that $T(c) = c$. The diagonal active-pair condition fixes $\lambda(c) = \lambda^{\text{sep}}(c)$, and the equilibrium conditions then determine the assignment and queues toward lower types.

A guess near $\underline{x}$ reaches the lowest type first, whereas a guess near $\overline{x}$ reaches the highest partner type first. Continuous dependence therefore yields an intermediate guess that reaches both endpoints simultaneously. A final no-crossing argument shows that this guess is unique.

Problem setup. Fix $c \in (\underline{x}, \overline{x})$. On the lower branch $x \le c \le T(x)$, write $\lambda_1(x) = \lambda(x)$ and $\lambda_2(x) = \lambda(T(x))$. The state $(T, \lambda_1, \lambda_2)$ follows the vector field

$$T'(x) = -\frac{g(x)\lambda_1(x)}{g(T(x))\lambda_2(x)}, \tag{54}$$

$$\lambda_1'(x) = \frac{f_1(x, T(x))}{B_1(\lambda_1(x), \lambda_2(x))}, \tag{55}$$

$$\lambda_2'(x) = T'(x)\frac{f_2(x, T(x))}{B_2(\lambda_1(x), \lambda_2(x))}. \tag{56}$$

Equation (56) follows from the upper-branch queue equation by the chain rule and symmetry of f and B. At the self-match point, the active-pair condition gives

$$T(c) = c, \qquad \lambda_1(c) = \lambda_2(c) = \lambda^{\text{sep}}(c). \tag{57}$$

Since $T'(x) < 0$, each lower type has a unique upper partner. After solving the lower branch, we match each $T(x)$ back with x and assign it queue $\lambda_2(x)$. This makes T an involution. Equation (56) ensures that the upper-branch queue satisfies the main-text equation, while the inverse of (54) gives visit balance there. The branches join smoothly at c because $T'(c) = -1$ by (54), and the symmetry of f and B makes their queue derivatives agree. Thus the system determines (T, λ) on the entire type space.

A direct differentiation shows that $B(\lambda_1(x), \lambda_2(x)) - f(x, T(x))$ is constant along the system. The initial condition makes this constant zero. Each trajectory therefore stays on the active-pair locus, which determines λ_2 from (x, T, λ_1), so the field is effectively two-dimensional.

Existence. Since $\lambda_1' > 0$, use λ_1 as the independent variable and solve the active-pair equation for λ_2 using $B_2 > 0$. The resulting two-dimensional vector field is continuous and locally Lipschitz in (x, T), so every c generates a unique local shot with a unique maximal extension (Teschl, 2012, Theorems 2.2 and 2.13).

For any $a < c$ before an endpoint is reached, extend the shot to $[a, T(a)]$ by the involution rule above. On this interval, the differential equations and endpoint conditions hold. The sufficiency argument in Proposition 4 therefore gives strictly positive slack for every inactive diagonal pair: $B(\lambda(z), \lambda(z)) > f(z, z) = B\Big(\lambda^{\text{sep}}(z), \lambda^{\text{sep}}(z)\Big)$, for $z \neq c$. Strict monotonicity of B implies the queue barrier, for $x < c$,

$$\lambda_1(x) > \lambda^{\text{sep}}(x), \qquad \lambda_2(x) > \lambda^{\text{sep}}(T(x)).$$

Define Λ_* by $B(\lambda^{\text{sep}}(\underline{x}), \Lambda_*) = f(\overline{x}, \overline{x})$; it exists uniquely because B is strictly increasing and unbounded in its second argument. The queue barrier and active-pair equation give $\lambda_2 \le \Lambda_*$. Since the queues agree at c and $\lambda_1' > 0 > \lambda_2'$, we obtain

$$0 < \lambda^{\text{sep}}(\underline{x}) \le \lambda_1(x) \le \lambda_2(x) \le \Lambda_*.$$

If a maximal shot ended before either endpoint, these bounds would keep the transformed solution in a compact subset of its domain, allowing further extension by Teschl (2012, Corollary 2.15). Thus every shot reaches $x = \underline{x}$ or $T(x) = \overline{x}$, possibly both simultaneously.

Let $\tau(c)$ be the terminal lower type and T_c the partner path. Define the undershoot and overshoot sets by

$$\mathcal{U} = \{c : \tau(c) = \underline{x},\ T_c(\tau(c)) < \overline{x}\},$$
$$\mathcal{O} = \{c : \tau(c) > \underline{x},\ T_c(\tau(c)) = \overline{x}\}.$$

The queue bounds and positivity of g imply that $-T_c'$ is uniformly bounded above and away from zero. Thus guesses near $\underline{x}$ belong to $\mathcal{U}$, and those near $\overline{x}$ belong to $\mathcal{O}$. Continuous dependence on c (Teschl, 2012, Theorem 2.9), together with these slope bounds, makes both sets open. Since $\mathcal{U}$ and $\mathcal{O}$ are disjoint, nonempty, and open, they cannot cover $(\underline{x}, \overline{x})$. A guess c^* outside their union must therefore reach both endpoints simultaneously: $T_{c^*}(\underline{x}) = \overline{x}$. Proposition 4 then yields a stable NAM equilibrium.

Uniqueness. Suppose, to the contrary, that two successful shots have fixed points $c_1 < c_2$. Consider the curve $(T_{c_2}(x) - T_{c_1}(x),\ \lambda_{1,c_2}(x) - \lambda_{1,c_1}(x))$ parametrized by x. At $x = c_1$, monotonicity of T and the queue barrier give

$$T_{c_2}(c_1) > c_2 > c_1 = T_{c_1}(c_1), \qquad \lambda_{1,c_2}(c_1) > \lambda^{\text{sep}}(c_1) = \lambda_{1,c_1}(c_1).$$

Thus $(T_{c_2}(x) - T_{c_1}(x),\ \lambda_{1,c_2}(x) - \lambda_{1,c_1}(x))$ starts strictly inside the positive quadrant. Moving left ($x < c_1$), suppose it first reaches the boundary at $x_0 \in [\underline{x}, c_1)$. Both differences are positive immediately to the right, so any vanishing difference must have a nonnegative right derivative at x_0.

On the horizontal axis, the lower queues agree but $T_{c_2} > T_{c_1}$. The active-pair condition implies $\lambda_{2,c_2} > \lambda_{2,c_1}$. Submodularity of f weakly lowers the positive numerator in (55), while $B_{12} > 0$ strictly raises its denominator. Hence $\lambda_{1,c_2}'(x_0) < \lambda_{1,c_1}'(x_0)$. As x decreases, this points upward into the quadrant, contradicting first contact.

On the vertical axis, the partner types agree but $\lambda_{1,c_2} > \lambda_{1,c_1}$. The active-pair condition now implies $\lambda_{2,c_2} < \lambda_{2,c_1}$. The density terms in (54) agree, whereas the queue ratio is higher

for the second shot. Hence $T'_{c_2}(x_0) < T'_{c_1}(x_0)$. As x decreases, this points rightward into the quadrant, again contradicting first contact.

At the origin, both differences vanish, and the active-pair condition also gives $\lambda_{2,c_2}(x_0) = \lambda_{2,c_1}(x_0)$. The shots therefore pass through the same state. Initial-value uniqueness makes them identical, contradicting their strict ordering at c_1.

No boundary contact is possible, including at $\underline{x}$, where the derivative comparisons are one-sided. But two successful shots require $T_{c_2}(\underline{x}) = T_{c_1}(\underline{x}) = \overline{x}$, placing the difference point on the vertical axis or at the origin. Thus at most one shot is successful. Combined with existence, this proves uniqueness of the differentiable deterministic stable equilibrium. □

A.8 Proof of Proposition 6

Suppose g_b/g_a is nondecreasing and nonconstant. Consider the curve $\Big(T_b(x) - T_a(x), \lambda_b(x) - \lambda_a(x)\Big)$ parametrized by x.

Step 1. We consider equal queues (the horizontal axis) and equal partners (the vertical axis). In both cases, the active-pair condition requires $B(\lambda_i(x), \lambda_i(T_i(x))) = f(x, T_i(x))$, for $i = a, b$. At equal queues, strict monotonicity of B and f orders the partner queues in the same way as the partner types. Since $f_1 > 0$ and $f_{12} \leq 0 < B_{12}$, equation (40) gives

$$\lambda_b(x) = \lambda_a(x) \quad \Longrightarrow \quad \operatorname{sgn}\Big(\lambda'_b(x) - \lambda'_a(x)\Big) = -\operatorname{sgn}\Big(T_b(x) - T_a(x)\Big). \tag{58}$$

When instead $T_a(x) = T_b(x)$, (43) gives

$$\frac{T'_b(x)}{T'_a(x)} = \frac{g_b(x)/g_a(x)}{g_b(T_a(x))/g_a(T_a(x))} \frac{\lambda_b(x)}{\lambda_a(x)} \frac{\lambda_a(T_a(x))}{\lambda_b(T_a(x))}. \tag{59}$$

The active-pair condition implies $\lambda_b(x) \leq \lambda_a(x)$ if and only if $\lambda_b(T_a(x)) \geq \lambda_a(T_a(x))$. If also $x \leq c_a$, then $T_a(x) \geq x$, so every factor is at most one. Thus, $T'_b(x) \geq T'_a(x)$, strictly if the queues differ. For $x \geq c_a$ and $\lambda_b(x) \geq \lambda_a(x)$, the inequalities reverse.

Step 2. Neither of the following combinations can hold:

$$\begin{aligned} &T_b(x) \leq T_a(x), \quad \lambda_b(x) \leq \lambda_a(x), \quad x \in (\underline{x}, c_a], \\ &T_b(x) \leq T_a(x), \quad \lambda_b(x) \geq \lambda_a(x), \quad x \in [c_a, \overline{x}). \end{aligned} \tag{60}$$

Suppose the first ordering holds at x_0. We first show that it persists toward lower types. Start with both inequalities strict. If an equality occurs at a lower interior type, let $x_1 \in (\underline{x}, x_0)$ be the first such point as x decreases. Any difference that vanishes there is negative immediately to the right, so its derivative at x_1 is nonpositive. If $T_b(x_1) = T_a(x_1)$ and $\lambda_b(x_1) < \lambda_a(x_1)$,

(59) gives $T_b'(x_1) > T_a'(x_1)$. If instead the queues agree and $T_b(x_1) < T_a(x_1)$, (58) gives $\lambda_b'(x_1) > \lambda_a'(x_1)$. Both cases contradict the required derivative sign.

If both variables first agree at x_1, the active-pair condition also equates the partner queues. Equation (59) again gives a contradiction unless g_b/g_a has equal values at x_1 and $T_a(x_1)$. Monotonicity then makes it constant on $[x_1, T_a(x_1)]$. Choose $\eta > 0$ sufficiently small that $x_1 + \eta < x_0$ and $x, T_a(x), T_b(x) \in [x_1, T_a(x_1)]$ for $x \in [x_1, x_1 + \eta]$. On this interval the proportional densities generate the same assignment equation; the queue equation is independent of the density. With partner queues recovered from the active-pair condition, initial-value uniqueness from Proposition 5 makes the solutions agree, contradicting their strict ordering. Thus both strict inequalities persist on $(\underline{x}, x_0]$.

For weak initial inequalities, fix $r \in (\underline{x}, x_0)$. For small $\epsilon > 0$, start the b-system at $T_b^\epsilon(x_0) = T_b(x_0) - \epsilon$ and $\lambda_b^\epsilon(x_0) = \lambda_b(x_0) - \epsilon$. Recover its partner queue from $B(\lambda_b^\epsilon(x), \widehat{\lambda}_b^\epsilon(x)) = f(x, T_b^\epsilon(x))$, and solve

$$(T_b^\epsilon)'(x) = -\frac{g_b(x)\lambda_b^\epsilon(x)}{g_b(T_b^\epsilon(x))\widehat{\lambda}_b^\epsilon(x)},$$
$$(\lambda_b^\epsilon)'(x) = \frac{f_1(x, T_b^\epsilon(x))}{B_1(\lambda_b^\epsilon(x), \widehat{\lambda}_b^\epsilon(x))}.$$

These auxiliary paths need not satisfy the equilibrium endpoint conditions. By continuation and continuous dependence as in Proposition 5, these solutions exist on $[r, x_0]$ for small ϵ and satisfy

$$\sup_{x \in [r, x_0]} (|T_b^\epsilon(x) - T_b(x)| + |\lambda_b^\epsilon(x) - \lambda_b(x)|) \longrightarrow 0 \quad \text{as } \epsilon \downarrow 0.$$

The strict comparison above gives $T_b^\epsilon < T_a$ and $\lambda_b^\epsilon < \lambda_a$ on $[r, x_0]$. Passing to the limit and then using continuity at $\underline{x}$ gives $T_b \leq T_a$ and $\lambda_b \leq \lambda_a$ on $[\underline{x}, x_0]$, since $r > \underline{x}$ is arbitrary.

For both strict and weak initial inequalities, the endpoint condition now gives a contradiction. Both partners equal $\overline{x}$ at $\underline{x}$ by (44), and the active-pair condition implies $\lambda_b(\overline{x}) \geq \lambda_a(\overline{x})$. Equation (59) yields

$$\frac{T_b'(\underline{x})}{T_a'(\underline{x})} \leq \frac{g_b(\underline{x})/g_a(\underline{x})}{g_b(\overline{x})/g_a(\overline{x})} < 1.$$

Strictness follows from the nonconstant monotone density ratio. Hence $T_b'(\underline{x}) > T_a'(\underline{x})$, contradicting $T_b \leq T_a$ immediately to the right of their common endpoint value.

If the second ordering holds at $x_0 \in [c_a, \overline{x})$, the same first-contact and approximation arguments, now increasing x and perturbing the initial pair to $(T_b(x_0) - \epsilon, \lambda_b(x_0) + \epsilon)$, preserve $T_b \leq T_a$ and $\lambda_b \geq \lambda_a$ up to $\overline{x}$. There both partners equal $\underline{x}$, and every factor in (59) is at least one, with the first strictly greater than one. Thus $T_b'(\overline{x}) < T_a'(\overline{x})$, contradicting $T_b \leq T_a$ immediately to the left. Endpoint derivatives are one-sided. In particular, queues and partners cannot both agree at an interior type.

Step 3. At any interior queue equality, Step 2 implies $T_b > T_a$; (58) then gives $\lambda_b' < \lambda_a'$. Every meeting is therefore a crossing from above to below, so there can be at most one.

If $c_b \le c_a$, Proposition 5 gives $\lambda_b(c_b) = \lambda^{\text{sep}}(c_b) \le \lambda_a(c_b)$, while $T_b(c_b) = c_b \le T_a(c_b)$, contradicting Step 2. Hence $c_b > c_a$, and the strict diagonal queue comparison gives $\lambda_b(c_a) > \lambda_a(c_a)$ and $\lambda_b(c_b) < \lambda_a(c_b)$. Continuity yields the unique $x^* \in (c_a, c_b)$, with $\lambda_b > \lambda_a$ below it and $\lambda_b < \lambda_a$ above it in the interior.

Continuity gives weak endpoint inequalities. Equality at either endpoint would imply equality at both by the active-pair condition, and (59) would give $T_b'(\underline{x}) > T_a'(\underline{x})$. Starting the b-system with the same endpoint partner and a slightly smaller lower queue preserves this derivative inequality on a common short right interval by continuous dependence. Its partner thus exceeds T_a, and (58) prevents its queue from crossing above λ_a. Passing to the limit gives $\lambda_b \le \lambda_a$ there, contradicting the interior inequality.

Finally, by (20), profits increase strictly with queues and inherit the same cutoff and inequalities. Interchanging a and b proves the nonincreasing-ratio case. □

A.9 Proof of Proposition 7

We first derive the stability bound. Fix $\mathbf{x} = (x_1, \ldots, x_N)$ and write $q_i = q(\lambda_i)$, $\varepsilon_i = \varepsilon_m(\lambda_i)$, and $\mathcal{Q} = \prod_i q_i$. Buyer free entry requires $\mathcal{Q}\left[f(\mathbf{x}) - \sum_i p_i\right] = K$. For seller i, define $V_i = \left(\prod_{j \ne i} q_j\right)\left[f(\mathbf{x}) - \sum_{j \ne i} p_j\right]$. Then $q_i V_i = K + \mathcal{Q} p_i$, and hence

$$\sum_{i=1}^{N} q_i V_i = \mathcal{Q} f(\mathbf{x}) + (N-1)K. \tag{61}$$

As in the two-visit model, a seller with visitor value V_i chooses its queue to maximize $m(\widetilde{\lambda})V_i - \widetilde{\lambda}K$. Therefore, $V_i^* = \frac{K}{m'(\lambda_i)}$ and $q_i V_i^* = \frac{K}{\varepsilon_i}$.

A bundle can be supported at the fixed queues if and only if the total in (61) can be divided so that $V_i \le V_i^*$ for every i. Because prices can transfer utility across the seller roles, such a division is possible if and only if the total does not exceed the sum of these upper bounds. Thus, $\mathcal{Q} f(\mathbf{x}) + (N-1)K = \sum_{i=1}^N q_i V_i \le K \sum_{i=1}^N \frac{1}{\varepsilon_i}$. By the definition of B, this is equivalent to $f(\mathbf{x}) \le B(\lambda_1, \ldots, \lambda_N)$. Evaluating this support condition at the PAM queues $\lambda_i = \lambda^{\text{sep}}(x_i)$ shows that a PAM equilibrium exists if and only if (47) holds.

We next prove the sufficiency of (48). Consider first the case $N = 3$, and relabel the

queues so that $\lambda_1 \leq \lambda_2 \leq \lambda_3$. By symmetry and strict log-submodularity,

$$\begin{aligned} 3\log B(\lambda_1,\lambda_2,\lambda_3) &= \log B(\lambda_1,\lambda_2,\lambda_3) + \log B(\lambda_3,\lambda_2,\lambda_1) + \log B(\lambda_3,\lambda_1,\lambda_2) \\ &\geq \log B(\lambda_1,\lambda_2,\lambda_1) + \log B(\lambda_2,\lambda_3,\lambda_3) + \log B(\lambda_3,\lambda_1,\lambda_2) \\ &\geq \log B(\lambda_1,\lambda_1,\lambda_1) + \log B(\lambda_3,\lambda_2,\lambda_2) + \log B(\lambda_2,\lambda_3,\lambda_3) \\ &\geq \log B(\lambda_1,\lambda_1,\lambda_1) + \log B(\lambda_2,\lambda_2,\lambda_2) + \log B(\lambda_3,\lambda_3,\lambda_3). \end{aligned}$$

At each step, two vectors are replaced by their coordinatewise minimum and maximum. The first inequality compares $(\lambda_1,\lambda_2,\lambda_3)$ and $(\lambda_3,\lambda_2,\lambda_1)$; the second compares $(\lambda_1,\lambda_2,\lambda_1)$ and $(\lambda_3,\lambda_1,\lambda_2)$; and the third compares $(\lambda_3,\lambda_2,\lambda_2)$ and $(\lambda_2,\lambda_3,\lambda_3)$.

The same argument applies for any N: begin with the N cyclic permutations of $(\lambda_1,\ldots,\lambda_N)$ and repeatedly replace two vectors by their coordinatewise minimum and maximum. This eventually produces the N diagonal vectors. Hence, $B(\lambda_1,\ldots,\lambda_N)^N \geq \prod_{i=1}^N B(\lambda_i,\ldots,\lambda_i)$, with strict inequality unless all the queues are equal. At the PAM queues $\lambda_i = \lambda^{\text{sep}}(x_i)$, the right-hand side equals $\prod_{i=1}^N f(x_i,\ldots,x_i)$. Therefore, (48) implies the exact stability condition (47).

We now establish uniqueness. Let $\mathbf{x} = (x_1,\ldots,x_N)$ be any active bundle in a stable equilibrium, and set $\lambda_i = \lambda(x_i)$. The active-bundle condition, the preceding inequality, stability of the diagonal alternatives, and (48) imply

$$f(\mathbf{x})^N = B(\boldsymbol{\lambda})^N \geq \prod_i B(\lambda_i,\ldots,\lambda_i) \geq \prod_i f(x_i,\ldots,x_i) \geq f(\mathbf{x})^N.$$

All inequalities must therefore bind. It follows that all queues are equal and every diagonal inequality binds. Hence, $f(x_1,\ldots,x_1) = \cdots = f(x_N,\ldots,x_N)$, which implies $x_1 = \cdots = x_N$. Thus every active bundle is diagonal, and the equilibrium is unique and exhibits PAM.

Finally, suppose that $\log f$ is supermodular. Applying the same argument as above, now with the inequalities reversed, gives (48). □

References

Ahlin, C. (2017). Matching patterns when group size exceeds two. *American Economic Journal: Microeconomics*, 9(1):352–384.

Albrecht, J., Cai, X., Gautier, P., and Vroman, S. (2020). Multiple applications, competing mechanisms, and market power. *Journal of Economic Theory*, 190:105121.

Albrecht, J., Cai, X., Gautier, P., and Vroman, S. (2023). On the foundations of compet-

itive search equilibrium with and without market makers. *Journal of Economic Theory*, 208:105605.

Albrecht, J., Cai, X., Gautier, P. A., and Vroman, S. (2026). Competitive search equilibrium with complements. Manuscript.

Albrecht, J. W., Gautier, P. A., and Vroman, S. B. (2006). Equilibrium directed search with multiple applications. *Review of Economic Studies*, 73:869–891.

Allen, J., Clark, R., and Houde, J.-F. (2019). Search frictions and market power in negotiated-price markets. *Journal of Political Economy*, 127(4):1550–1598.

Armstrong, M. (2013). A more general theory of commodity bundling. *Journal of Economic Theory*, 148(2):448–472.

Becker, G. S. (1973). A theory of marriage: Part I. *Journal of Political Economy*, 81(4):813–846.

Boerma, J., Tsyvinski, A., and Zimin, A. P. (2025). Sorting with teams. *Journal of Political Economy*, 133(2):421–454.

Calzolari, G. and Denicolò, V. (2013). Competition with exclusive contracts and market-share discounts. *American Economic Review*, 103(6):2384–2411.

Carlier, G. and Ekeland, I. (2010). Matching for teams. *Economic Theory*, 42(2):397–418.

Carvalho, V. M., Nirei, M., Saito, Y. U., and Tahbaz-Salehi, A. (2021). Supply chain disruptions: Evidence from the Great East Japan earthquake. *Quarterly Journal of Economics*, 136(2):1255–1321.

Chade, H. and Eeckhout, J. (2018). Matching information. *Theoretical Economics*, 13(1):377–414.

Chade, H. and Smith, L. (2006). Simultaneous search. *Econometrica*, 74(5):1293–1307.

Dell'Ariccia, G. and Garibaldi, P. (2005). Gross credit flows. *Review of Economic Studies*, 72(3):665–685.

Demir, B., Fieler, A. C., Xu, D. Y., and Yang, K. K. (2024). O-ring production networks. *Journal of Political Economy*, 132(1):200–247.

den Haan, W. J., Ramey, G., and Watson, J. (2003). Liquidity flows and fragility of business enterprises. *Journal of Monetary Economics*, 50(6):1215–1241.

Dixit, A. K. and Stiglitz, J. E. (1977). Monopolistic competition and optimum product diversity. *American Economic Review*, 67(3):297–308.

Economides, N. and Salop, S. C. (1992). Competition and integration among complements, and network market structure. *Journal of Industrial Economics*, 40(1):105–123.

Eeckhout, J. and Kircher, P. (2010). Sorting and decentralized price competition. *Econometrica*, 78:539–574.

Eeckhout, J. and Kircher, P. (2018). Assortative matching with large firms. *Econometrica*, 86(1):85–132.

Ellison, G. (2005). A model of add-on pricing. *Quarterly Journal of Economics*, 120(2):585–637.

Galenianos, M. and Kircher, P. (2009). Directed search with multiple job applications. *Journal of Economic Theory*, 144:445–471.

Garicano, L. and Rossi-Hansberg, E. (2006). Organization and inequality in a knowledge economy. *Quarterly Journal of Economics*, 121(4):1383–1435.

Gautier, P. A. and Holzner, C. L. (2017). Simultaneous search and efficiency of entry and search intensity. *American Economic Journal: Microeconomics*, 9(3):245–82.

Hastings, S. P. and McLeod, J. B. (2012). *Classical methods in ordinary differential equations: with applications to boundary value problems*, volume 129. American Mathematical Soc.

Herkenhoff, K., Lise, J., Menzio, G., and Phillips, G. M. (2024). Production and learning in teams. *Econometrica*, 92(2):467–504.

Kircher, P. (2009). Efficiency of simultaneous search. *Journal of Political Economy*, 117:861–913.

Korovkin, V., Makarin, A., and Miyauchi, Y. (2026). Supply chain disruption and reorganization: Theory and evidence from Ukraine's war. *Review of Economic Studies*, 93(4):2750–2783.

Kremer, M. (1993). The o-ring theory of economic development. *Quarterly Journal of Economics*, 108(3):551–575.

Kremer, M. and Maskin, E. (1996). Wage inequality and segregation by skill. Working Paper 5718, National Bureau of Economic Research.

Lucarelli, C., Nicholson, S., and Song, M. (2010). Bundling among rivals: A case of pharmaceutical cocktails. Working Paper 16321, National Bureau of Economic Research.

Matutes, C. and Regibeau, P. (1988). Mix and match: Product compatibility without network externalities. *RAND Journal of Economics*, 19(2):221–234.

Matutes, C. and Regibeau, P. (1992). Compatibility and bundling of complementary goods in a duopoly. *Journal of Industrial Economics*, 40(1):37–54.

Petrosky-Nadeau, N. and Wasmer, E. (2013). The cyclical volatility of labor markets under frictional financial markets. *American Economic Journal: Macroeconomics*, 5(1):193–221.

Reher, M. and Valkanov, R. (2024). The mortgage-cash premium puzzle. *Journal of Finance*, 79(5):3149–3201.

Rhodes, A. and Zhou, J. (2024). Personalized pricing and competition. *American Economic Review*, 114(7):2141–2170.

Roth, A. E. and Sotomayor, M. A. O. (1990). *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*. Econometric Society Monographs. Cambridge University Press.

Shapley, L. S. and Shubik, M. (1972). The assignment game I: The core. *International Journal of Game Theory*, 1:111–130.

Shi, S. (2001). Frictional assignment I: Efficiency. *Journal of Economic Theory*, 98:232–260.

Shimer, R. (2005). The assignment of workers to jobs in an economy with coordination frictions. *Journal of Political Economy*, 113(5):996–1025.

Swiss Federal Office of Public Health (2020). BAG-Bulletin 4/2020. 20 January, p. 24. Revlimid reimbursement conditions.

Teschl, G. (2012). *Ordinary differential equations and dynamical systems*, volume 140. American Mathematical Soc.

Wasmer, E. and Weil, P. (2004). The macroeconomics of labor and credit market imperfections. *American Economic Review*, 94(4):944–963.

Wright, R., Kircher, P., Julien, B., and Guerrieri, V. (2021). Directed search and competitive search equilibrium: A guided tour. *Journal of Economic Literature*, 59(1):90–148.