

Equilibrium Selection in Last-Mile Delivery Markets: Centralized Customer–Provider Assignment under Economies of Scale

Duygu Sili^a, Ali Nalbant^b, Barış Yıldız^b

^a*Department of Economics, Universidad Carlos III de Madrid, Madrid, Spain*

^b*Department of Industrial Engineering, Koç University, Sarıyer, Istanbul, Türkiye*

Abstract

We study a stylized model of a future *logistics market* where an independent market operator assigns last-mile parcel deliveries to competing service providers. Customers value providers differently through attributes such as environmental performance and service quality; each provider’s per-parcel tariff decreases with its number of customers, reflecting routing economies of scale. An assignment is self-enforcing if no customer prefers switching to another provider at the tariff she would actually pay after the switch. The resulting game is a resource-selection game with an exact potential: self-enforcing assignments (pure Nash equilibria) always exist, and every maximal sequence of profitable unilateral switches reaches one. Scale economies, however, make equilibria multiple and very different: a market can be locked into a provider that nobody prefers. With two providers we characterize the equilibrium set completely: the operator reaches the payment-minimizing equilibrium by choosing the order in which providers are activated. With three or more providers no activation order is reliable in general, although sequential activation still implements every concentrated equilibrium and attains the minimum total payment when scale economies dominate preference heterogeneity. We therefore optimize any operator objective over the equilibrium set with a mixed-integer program and a clustering-based reduction scaling beyond 2,500 customers. In a study calibrated to parcel demand in a district of Istanbul, the utility gap between best and worst equilibria ranges from 7% to 18% (median 16%), and the utility-maximizing equilibrium always lies on the environmental–quality efficient frontier: steering the market to its best equilibrium sacrifices neither dimension.

Keywords: last-mile delivery, city logistics, Physical Internet, equilibrium selection, congestion games, scale economies, mixed-integer programming

1. Introduction

Home delivery has become a large and growing share of urban freight: global parcel volumes exceeded 160 billion in 2021 and are projected to approach a quarter of a trillion before the end of the decade (Statista, 2022). This growth has put significant pressure on urban logistics systems, with its effects compounded by the way it has been absorbed: not through coordinated scale-up, but through fragmentation. In dense urban districts, several

Email address: dsili@eco.uc3m.es (Duygu Sili)

carriers typically serve the same streets on the same day, each running partially loaded routes. Because the cost of a delivery route is driven by the density of stops, this fragmentation is costly: a carrier that serves twice as many customers in a neighborhood travels far less than twice the distance (Daganzo, 1984), so that consolidating demand across providers would lower costs, congestion, and emissions.

A structural response to this fragmentation has been under development for over a decade under the banner of the Physical Internet (PI), which extends digital-internet principles (encapsulation, open protocols, interconnected hubs) to the movement of physical goods (Montreuil et al., 2013). What began as a conceptual proposal has grown into a substantial research program, with institutional platforms such as the Alliance for Logistics Innovation through Collaboration in Europe (ALICE)¹ and the Japan Physical Internet Center (JPIC)² coordinating initiatives and pilots spanning facility design, network protocols, and operations (Ballot et al., 2021). Its core proposition is that city logistics should be organized as a service market: specialized providers supply transport and handling capacity, and a coordinating platform matches demand to supply dynamically, rather than fragmenting it along carrier lines. This directly targets the density inefficiency above, since pooling stops across providers, rather than within them, recovers the economies of density that fragmented last-mile networks presently forfeit.

This paper studies one building block of that vision: a centralized *logistics market* for last-mile delivery, run by an *independent market operator* (IMO). We want to be clear at the outset that no such market exists today. Parcel assignment is currently controlled by shippers, marketplaces, and originating carriers, and the institution we analyze is a proposal, in the same spirit as the independent system operators introduced when electricity generation was unbundled from transmission (Cramton, 2017): a central body clears the market on behalf of many independent suppliers, but cannot compel participation, since consumers retain the right to switch to another qualified provider at will. We do not mean to suggest that electricity and parcel delivery are alike as industries. Electricity is a homogeneous, non-storable good balanced in real time, while parcels are heterogeneous, storable, and routed with no such constraint. Only that the two settings share the critical institutional feature of interest in this study: a coordinator seeking an efficient outcome while respecting voluntary participation on the demand side. Our IMO inherits the same feature: the requirement that its assignment be self-enforcing, formalized below, is the logistics-market counterpart of that participation constraint.

Our aim is therefore not to describe an existing practice but to understand, in a deliberately stylized model, two specific questions: when can a centralized assignment be sustained without compelling participation, and if so, what is its potential to reduce the negative externalities of urban freight transport? To study these questions, we have the following arrangement in mind. A neutral entity (e.g., a municipal agency, a consortium of shippers and carriers, or a regulated platform) operates a clearinghouse for a delivery district. Parcels destined for the district arrive at a local consolidation center; in line with PI principles and modern multi-echelon city logistics designs, in which freight moves through successive con-

¹<https://www.etp-logistics.eu/>

²<https://j-pic.or.jp/about-jpic/>

solidation points before reaching customers (Cuda et al., 2015; Crainic and Montreuil, 2016), this center is the last such point before the clearinghouse takes over and assigns customers to providers. The operator periodically (say, daily) assigns the corresponding customers to certified last-mile providers that pick up from the center and deliver to them. Providers are certified on observable service attributes (e.g., emissions class of the fleet, delivery-time guarantees) and report the cost parameters of their operations; the operator has, from historical choices and surveys, an estimate of how much each customer values each provider’s attributes. On this basis, the operator publishes an assignment that matches customers to providers, together with a *tariff* for each provider that decreases with the number of customers assigned to that provider. What the operator cannot do is force customers to accept a provider they would rather not use: a customer can always request another provider, and would pay the tariff that provider charges once she joins it. An assignment is therefore *implementable* only if it is self-enforcing, i.e., if no customer would rather switch to another provider at the tariff she would actually pay after switching. This is the requirement that drives the methodological developments in this paper.

Economies of scale make the self-enforcing outcomes of such a market multiple and very different from one another. If a provider is popular, it becomes cheaper, which makes it more popular; if it is unpopular, it becomes expensive, and nobody wants to be the first to join it. A market can therefore be locked into a provider that few customers prefer, simply because no single customer gains by moving alone, and the same market could equally well settle at a very different assignment with much higher aggregate utility, lower emissions, or better service. A decentralized market has no instrument for choosing among these outcomes; an operator that assigns customers (and that must respect self-enforcement) does. The main questions we study are therefore *how large the set of self-enforcing assignments is, how much the operator can gain by selecting among them, and how it can find the best one.*

To answer these questions, we model the market as a game in which each customer chooses one provider (Section 1.1). A customer’s payoff is her valuation of the provider’s attributes minus the provider’s tariff, and the tariff is a decreasing function of the provider’s load, calibrated to the continuous approximation of routing cost due to Daganzo (1984). This is a resource-selection (singleton congestion) game with player-specific valuations and *decreasing* congestion costs, i.e., with positive externalities. The model deliberately abstracts from the operational detail of last-mile delivery: routing, vehicle capacity, and time windows enter only through the load-dependent cost approximation; the assignment is static; and the operator is assumed to know customers’ valuations. We view these abstractions as the price of isolating the equilibrium-selection question, and we discuss in Section 5 which of them matter for the results.

1.1. Problem Definition

We consider a finite set of customers $\mathcal{C}$, with $|\mathcal{C}| = N$, and a finite set of service providers $\mathcal{L}$, with $|\mathcal{L}| = L$. Each customer $i \in \mathcal{C}$ is to be served by exactly one provider $j \in \mathcal{L}$.

Providers and tariffs. Each provider j is characterized by observable service attributes (in our experiments, an environmental-performance level and a service-quality level) and by a cost parameter $\beta_j > 0$ that summarizes its cost structure, including any cost of delivering its attributes. The operational cost of serving a set of customers scattered over a delivery area

exhibits economies of scale: in the continuous approximation of [Daganzo \(1984\)](#), the routing distance per stop is inversely proportional to the square root of the stop density, so that the average cost per customer of a provider serving m customers in the district is proportional to $1/\sqrt{m}$. Following this approximation, we let the *tariff* that provider j charges each of its customers when it serves m customers be

$$p_j(m) = \min\left\{\bar{\beta}, \frac{\beta_j}{\sqrt{m}}\right\}, \quad m = 0, 1, 2, \dots, \quad (1)$$

where $\bar{\beta}$ is a price cap set by the operator and, throughout the paper, $\beta_j/\sqrt{0} := +\infty$, so that a provider without customers quotes the cap, $p_j(0) = \bar{\beta}$. A provider's load never exceeds N , but we define the schedule for every nonnegative integer because the *entry tariff* of [Section 3.1](#) evaluates it at the load plus one. We interpret $p_j(m)$ as the provider's average cost per customer inclusive of a normal return, i.e., tariffs are cost-based, as in the regulated-tariff logic of electricity system operators; the cap protects customers of a provider with a small customer base. The total payment collected by provider j , $\varphi_j(m) = m p_j(m) = \min\{\bar{\beta} m, \beta_j \sqrt{m}\}$, is increasing and concave in m : the marginal cost of an additional customer is below the average cost, which is precisely why an additional customer exerts a *positive* externality on the provider's other customers (their tariff falls). [Section 4.4](#) validates the functional form in (1)—average cost per customer proportional to $1/\sqrt{m}$ —against a large set of vehicle routing instances on the road network of the study area. Most of our results hold for any tariff schedule that is nonincreasing in the load (existence of equilibrium, [Theorem 1](#); lock-in, [Proposition 1](#); the two-provider structure, [Lemma 2](#)); those that use the specific form (1) are indicated explicitly.

Providers report their cost parameters β_j to the operator, who determines the assignment and the corresponding tariffs according to (1). Providers do not discriminate across customers and, other things equal, prefer to serve more customers. We assume a contestable market in which, for each provider, a potential entrant with identical attributes and cost stands ready to enter; [Proposition 6](#) shows that, when the price cap does not bind and the operator assigns customers in ascending order of reported costs, truthful cost reporting is a Nash equilibrium of a two-stage reporting game that we define formally in [Section 3.3](#): no provider gains by deviating from it unilaterally.

Customers. Each customer $i \in \mathcal{C}$ has a valuation v_{ij} for being served by provider $j \in \mathcal{L}$, determined by the provider's attributes. Her utility from being served by j when j serves m customers (herself included) is quasi-linear in the tariff,

$$u_{ij}(m) = v_{ij} - p_j(m). \quad (2)$$

We normalize the utility of not receiving home delivery to zero and assume throughout that every provider is acceptable to every customer at any tariff the market can charge:

Assumption 1 (Acceptability). $v_{ij} \geq \bar{\beta}$ for all $i \in \mathcal{C}$ and $j \in \mathcal{L}$.

Under [Assumption 1](#), $u_{ij}(m) \geq 0$ for every load m , so every assignment is individually rational and the option of remaining unserved plays no role; this is why we can require that every customer be assigned. (In our experiments the assumption holds by construction: the

valuation of receiving the parcel alone is at least $\bar{\beta}$.) For tractability we assume that the operator knows the valuation matrix $\mathcal{V} = \{v_{ij}\}$, e.g., from predictive analytics or repeated-interaction data; relaxing this assumption leads to a Bayesian mechanism design problem that we leave for future work.

Definition 1 (Matching). *A matching (or assignment) is a function $\mu : \mathcal{C} \rightarrow \mathcal{L}$ that assigns each customer to exactly one provider. Given μ , the load of provider j is*

$$m_j(\mu) = |\{i \in \mathcal{C} \mid \mu(i) = j\}|, \quad (3)$$

its tariff is $p_j(m_j(\mu))$, and customer i 's utility is $u_i(\mu) = v_{i\mu(i)} - p_{\mu(i)}(m_{\mu(i)}(\mu))$.

Let $\mathcal{M}$ denote the set of all matchings. A matching is *self-enforcing* if no customer can increase her utility by unilaterally switching to another provider, taking into account that the provider she joins will then serve one more customer; these are the pure-strategy Nash equilibria of the game defined formally in Section 3.1, and we denote their set by $\bar{\mathcal{M}} \subseteq \mathcal{M}$.

Operator's Problem.. The operator evaluates matchings through an objective function $F : \mathcal{M} \rightarrow \mathbb{R}$ that it seeks to minimize. Depending on the context, F may represent total customer payments, environmental externalities, or the negative of aggregate customer utility (so that utility maximization is written as a minimization). Given an instance $\mathcal{P} = \langle \mathcal{L}, \mathcal{C}, F \rangle$, the operator's *equilibrium selection problem* is

$$\min_{\mu \in \bar{\mathcal{M}}} F(\mu). \quad (4)$$

A solution of (4) is an *operator-optimal equilibrium*. Throughout, *aggregate utility* refers to $U(\mu) = \sum_{i \in \mathcal{C}} u_i(\mu)$, the sum of customers' utilities net of payments, and *total payment* to $R(\mu) = \sum_{j \in \mathcal{L}} m_j(\mu) p_j(m_j(\mu))$; we use these two terms consistently and reserve *public value* for the general objective F .

1.2. Contributions of this study

- We formulate the operator's problem as a resource-selection game: customers choose among providers, trading off valuation against a tariff that *decreases* with load – the reverse of standard congestion games. The game admits an exact potential, so self-enforcing assignments exist and are reachable by unilateral switching (Theorem 1). But scale economies also make equilibria multiple, and not all are good: markets can lock into an inferior provider when heterogeneity is small relative to scale economies (Proposition 1). With two providers we characterize the full equilibrium set and show the payment-minimizing equilibrium is always reachable by one of two simple priority orders (Theorem 2); with three or more, no priority order can guarantee the optimal assignment (Propositions 3, 4), though we delimit exactly what ordering still delivers: every lock-in equilibrium, and the first-best payment whenever scale economies dominate heterogeneity (Proposition 5). Selection in general therefore requires optimizing over the equilibrium set directly.
- We formulate this optimization as a mixed-integer program encoding self-enforcement at post-switch tariffs, valid for any operator objective. A clustering-based reduction

with an individual-level repair step keeps the optimality gap small – 0.4% on average – on instances with over 2,500 customers.

- We illustrate the framework on parcel demand calibrated to a district of Istanbul: the utility gap between the best and worst equilibrium ranges from 7% to 18% across scenarios, and the utility-maximizing equilibrium lies on the environmental–quality efficient frontier in every scenario we test – suggesting that steering the market to its best equilibrium need not trade off against environmental or quality goals.

The remainder of the paper is organized as follows. Section 2 positions the model in the literature on congestion and resource-selection games and on coordination in city logistics. Section 3 defines the game, establishes existence, characterizes the multiplicity of equilibria and the extent to which the activation order can select among them, and presents the optimization approach. Section 4 reports the computational study. Section 5 concludes. All proofs are in [Appendix A](#).

2. Related Literature

Our work lies at the intersection of three literatures. First, it builds on congestion and resource-selection games, where agents choose among alternatives whose payoffs depend on aggregate usage. Second, it relates to coordination and consolidation in city logistics, particularly the Physical Internet view of logistics as a market for services coordinated by a neutral operator. Third, it connects to centralized matching platforms in transportation, which use platform-level control to improve outcomes relative to decentralized choice. In the following, we discuss each of these literature streams and position our contribution relative to them.

2.1. Congestion and resource-selection games

In our model providers make no decisions: the only strategic agents are customers, each of whom selects one provider, and a customer’s payoff depends on her own valuation of that provider and on how many other customers select it. The model is therefore not a two-sided matching market in the sense of [Shapley and Shubik \(1971\)](#) or [Jr. and Crawford \(1982\)](#)—there is no preference or choice on the supply side—but a *singleton congestion game*, also called a resource-selection game ([Rosenthal, 1973](#); [Jeong et al., 2005](#)), with two distinctive features. First, payoffs are player-specific through the valuations v_{ij} ; congestion games with player-specific payoff functions were introduced by [Milchtaich \(1996\)](#), who showed that singleton games of this kind always possess a pure-strategy Nash equilibrium but need not admit a potential function, so that sequences of profitable unilateral deviations may cycle. In our game the player-specific component enters additively and does not depend on the load, and this separability restores an exact potential in the sense of [Monderer and Shapley \(1996\)](#); Lemma 1 makes the construction explicit, and it is the same construction as in [Rosenthal \(1973\)](#). Second, the congestion cost is *decreasing* in the load: joining a provider makes it cheaper for everybody, a positive externality. Games of this type appear as cost-sharing games in network design ([Anshelevich et al., 2008](#)), where each user of a resource pays an equal share of its fixed cost, and as resource-selection games with cost-sharing rules ([Gkatzelis](#)

et al., 2016); in economics, the same structure underlies group-formation games with positive externalities (Konishi et al., 1997) and, more broadly, markets with network effects (Katz and Shapiro, 1985), in which several self-fulfilling outcomes coexist and a market can be locked into an inferior one because no single agent gains by moving alone. The multiplicity of equilibria and the gap between the best and the worst of them are the subject of the price-of-anarchy and price-of-stability literature (Koutsoupas and Papadimitriou, 1999; Roughgarden and Tardos, 2002; Anshelevich et al., 2008; Vöcking, 2007). Our contribution relative to this literature is not the existence result, which is standard once the potential is identified, but the study of *equilibrium selection by a coordinator*: how the set of self-enforcing assignments is structured (Lemma 2, Propositions 1 and 3), to what extent a simple instrument—the order in which providers are activated—selects a good equilibrium (Theorem 2), and how to optimize a general objective over the equilibrium set when it does not (Section 3.4).

2.2. Coordination in city logistics

The inefficiency of fragmented last-mile delivery and the potential of consolidation and horizontal collaboration are well documented (Savelsbergh and Van Woensel, 2016; Gansterer and Hartl, 2018; Cleophas et al., 2019). Most of this literature studies collaboration among carriers—exchange of requests, shared vehicles or depots, and the allocation of the resulting gains—and takes the assignment of customers to carriers as the object of negotiation among the carriers. The Physical Internet and hyperconnected city logistics literature (Montreuil et al., 2013; Crainic and Montreuil, 2016) goes further and envisions city logistics organized as a market for logistics services, in which the matching of demand to specialized providers is performed by a coordinating platform; how such a market should clear is largely open. Our institutional inspiration is the independent system operator of electricity markets (Biggar and Hesamzadeh, 2014; Cramton, 2017), a neutral entity that clears supply and demand and sets prices that internalize system-wide costs. Two features distinguish a logistics market from an electricity market and drive our analysis: customers have heterogeneous valuations over providers, so the service is differentiated rather than homogeneous, and per-unit costs decrease with the volume served by each provider, so the operator’s clearing problem is an equilibrium selection problem rather than a market-clearing problem with unique prices.

2.3. Centralized matching platforms in transportation

Our work also relates to the literature on platform-mediated matching in transportation, which shows that centralized dispatch and pricing can substantially outperform decentralized matching. Bimpikis et al. (2019) show how a ride-sharing platform can use spatially differentiated prices to balance supply and demand across a city, and Castillo et al. (2017) show that centralized matching prevents failures such as the “wild goose chase” that arise when drivers search for rides on their own. In those models the platform’s instrument is the price it sets; in ours prices are cost-based and load-dependent, and the platform’s instrument is the selection of an assignment among the self-enforcing ones. Finally, the survey evidence of Engelhardt (2023) on consumers’ willingness to pay for sustainable last-mile services, which we use to calibrate customer segments, documents the preference heterogeneity that makes equilibrium selection consequential.

3. The assignment game: equilibria, multiplicity, and selection

In this section we define the game played by customers, establish that self-enforcing assignments exist and can be reached by simple adjustment dynamics (Section 3.1), show that they are typically multiple and can differ widely (Section 3.2), and study whether the operator can select a good one through the order in which providers are activated (Section 3.3). With two providers the answer is essentially complete and positive; with three or more it is negative in general, which motivates the optimization approach of Section 3.4.

3.1. The game and the existence of self-enforcing assignments

The game. The strategic agents are the customers. Customer $i \in \mathcal{C}$ chooses a provider $\mu(i) \in \mathcal{L}$; a strategy profile is therefore a matching $\mu \in \mathcal{M}$, and it induces the loads $m_j(\mu)$ of (3). Given μ , customer i obtains

$$u_i(\mu) = v_{i\mu(i)} - p_{\mu(i)}(m_{\mu(i)}(\mu)), \quad (5)$$

where the tariff schedule $p_j(\cdot)$ is given by (1). We denote this game by $\Gamma = \langle \mathcal{C}, \mathcal{L}, (v_{ij}), (p_j) \rangle$. If customer i unilaterally switches from $\mu(i)$ to another provider k , provider k serves $m_k(\mu) + 1$ customers, so the tariff she pays after switching is the *entry tariff*

$$p_k^+(\mu) := p_k(m_k(\mu) + 1) \leq p_k(m_k(\mu)), \quad (6)$$

which is weakly lower than the tariff currently paid by k 's customers because her arrival adds to k 's scale economies (if k is empty, she pays $p_k(1) = \min\{\bar{\beta}, \beta_k\}$, the tariff of a provider with a single customer).

Definition 2 (Self-enforcing assignment / Nash equilibrium). *A matching μ is a **Nash equilibrium** of Γ (equivalently, self-enforcing) if no customer can strictly increase her utility by unilaterally switching providers at the entry tariff:*

$$u_i(\mu) \geq v_{ik} - p_k^+(\mu) \quad \text{for all } i \in \mathcal{C} \text{ and all } k \in \mathcal{L} \setminus \{\mu(i)\}. \quad (7)$$

We write $\bar{\mathcal{M}}$ for the set of Nash equilibria.

Remark 1 (On the deviation tariff). *A weaker notion would compare a customer's current utility with $v_{ik} - p_k(m_k(\mu))$, i.e., evaluate the deviation at provider k 's tariff before she joins it. Since $p_k^+(\mu) \leq p_k(m_k(\mu))$, every Nash equilibrium in the sense of Definition 2 is stable in that weaker sense, but the converse fails: the weaker notion may certify assignments from which some customer would profitably switch once the tariff she would actually pay is taken into account. Because the operator's assignment must be self-enforcing against actual switches, we use Definition 2 throughout, in the equilibrium constraints of the optimization models of Section 3.4 as well as in the adjustment dynamics below.*

By Assumption 1, $u_i(\mu) \geq 0$ for every μ , so individual rationality holds for every matching and does not need to be imposed separately. The stability requirement (7) plays a dual role. In a decentralized market it characterizes the rest points of any process in which customers switch providers when it pays to do so. In the centralized market it is the operator's

implementability constraint: an assignment that violates (7) would unravel as soon as the customer with a profitable switch requests it, whereas an assignment that satisfies (7) is one every customer accepts voluntarily given the published tariff schedules and assignments.

Existence.. The game Γ is a singleton congestion game in which the player-specific component of the payoff, v_{ij} , does not depend on the load. It therefore admits an exact potential function of the type introduced by [Rosenthal \(1973\)](#).

Lemma 1 (Exact potential). *Define, for every matching μ ,*

$$\Phi(\mu) = \sum_{i \in \mathcal{C}} v_{i\mu(i)} - \sum_{j \in \mathcal{L}} \sum_{\ell=1}^{m_j(\mu)} p_j(\ell), \quad (8)$$

with the convention that an empty sum is zero. If μ' is obtained from μ by moving a single customer i from $\mu(i)$ to $k \neq \mu(i)$, then $\Phi(\mu') - \Phi(\mu) = u_i(\mu') - u_i(\mu) = (v_{ik} - p_k^+(\mu)) - u_i(\mu)$. In particular, Φ strictly increases along every profitable unilateral switch, so no matching is visited twice by a sequence of such switches.

A profitable unilateral switch is a *better response*; a sequence of better responses is an *improvement path*. Since $\mathcal{M}$ is finite and Φ strictly increases along every improvement path, every improvement path is finite, and a matching from which no better response exists—the last matching of any improvement path that cannot be extended—is, by definition, a Nash equilibrium. This is the finite improvement property of [Monderer and Shapley \(1996\)](#), and it yields existence immediately.

Theorem 1 (Existence and convergence). *The game Γ has at least one Nash equilibrium. More precisely: (i) every maximizer of Φ over $\mathcal{M}$ is a Nash equilibrium; (ii) every improvement path is finite, and every improvement path that cannot be extended ends at a Nash equilibrium, whatever the initial matching and whatever the rule used to choose which profitable switch is executed next. Both statements hold for arbitrary real-valued tariff schedules $p_j(\cdot)$, not only for (1).*

Theorem 1 says that self-enforcing assignments are not only existent but *reachable*: any decentralized adjustment process in which customers keep switching as long as it pays to do so terminates, after finitely many switches, at one of them. Which one it terminates at, however, depends on where the process starts and on the order in which customers move. This is the source of the multiplicity we study next, and it is also what gives the operator its instrument.

The Sweep Algorithm.. The operator’s instrument is a particular improvement path, which we call the *Sweep Algorithm* (SA). Its input is a *priority ordering* $o = [o(1), \dots, o(L)]$, a permutation of the providers chosen by the operator. All customers are initially assigned to the first provider $o(1)$; providers are then *activated* one at a time in the order o . Whenever a new provider is activated, customers may switch among all currently active providers; a switch is executed only if it strictly increases the switching customer’s utility at the entry tariff (6), and among the profitable switches the one with the largest utility gain is executed first (ties are broken in favor of the provider that appears earlier in o , then in favor of

the customer with the smaller index). Loads, tariffs, and utilities are updated after every switch. When no profitable switch remains, the next provider is activated, and the procedure ends when all providers are active and no profitable switch remains. Algorithm 1 gives the pseudocode; $\mathcal{S}^t$ denotes the matching after t switches and T^* the total number of switches.

Algorithm 1 Sweep Algorithm

Require: Customers $\mathcal{C}$; providers $\mathcal{L}$; priority ordering $o = [o(1), \dots, o(L)]$ of $\mathcal{L}$; parameters $(\bar{\beta}, \beta_j)$, valuations (v_{ij})

- 1: **Init:** $t \leftarrow 0$; $\mu(i) \leftarrow o(1)$ for all $i \in \mathcal{C}$; $m_j \leftarrow |\{i : \mu(i) = j\}|$ for all $j \in \mathcal{L}$
- 2: $\mathcal{A} \leftarrow \{o(1)\}$ ▷ active providers
- 3: **for** $k = 2$ **to** L **do** ▷ providers are activated sequentially according to o
- 4: $\mathcal{A} \leftarrow \mathcal{A} \cup \{o(k)\}$
- 5: **while** $\exists i \in \mathcal{C}, a \in \mathcal{A} \setminus \{\mu(i)\}$ with $v_{ia} - p_a(m_a + 1) > v_{i\mu(i)} - p_{\mu(i)}(m_{\mu(i)})$ **do**
- 6: $(i^*, a^*) \leftarrow \arg \max_{i \in \mathcal{C}, a \in \mathcal{A} \setminus \{\mu(i)\}} \left[(v_{ia} - p_a(m_a + 1)) - (v_{i\mu(i)} - p_{\mu(i)}(m_{\mu(i)})) \right]$ ▷ largest gain
 at entry tariffs; ties: earlier in o , then smaller i
- 7: $j_{\text{old}} \leftarrow \mu(i^*)$; $\mu(i^*) \leftarrow a^*$; $m_{a^*} \leftarrow m_{a^*} + 1$; $m_{j_{\text{old}}} \leftarrow m_{j_{\text{old}}} - 1$
- 8: $t \leftarrow t + 1$; $\mathcal{S}^t \leftarrow \{(i, \mu(i)) : i \in \mathcal{C}\}$
- 9: **end while**
- 10: **end for**
- 11: **return** $\mathcal{S}^{T^*} = \{(i, \mu(i)) : i \in \mathcal{C}\}$

Every switch executed by the SA is a better response in Γ (a switch to an active provider is a unilateral deviation in the full game), so the SA follows an improvement path and Theorem 1 applies to it.

Corollary 1 (Termination of the Sweep Algorithm). *For every priority ordering o , the Sweep Algorithm terminates after finitely many switches, and its terminal matching $\mathcal{S}^{T^*}$ is a Nash equilibrium of Γ .*

We emphasize what Corollary 1 does and does not say. Convergence does not depend on the initialization, on the activation order, or on the largest-gain rule: any improvement path converges. What the ordering and the rule determine is *which* equilibrium is reached, and this is the only reason to consider the SA rather than an arbitrary improvement path: the activation order is a lever that the operator controls, and Section 3.3 asks how much it can achieve with it.

Remark 2 (Complexity bound). *The potential argument bounds the number of switches only by the number of matchings, L^N . We conjecture that the SA terminates within $O(N^2L)$ switches: each provider admits at most $N + 1$ distinct load levels, which suggests that the total number of “effective” tariff adjustments scales polynomially. A proof does not follow from load discreteness alone, because a provider’s load can revisit the same value several times as different customers enter and leave. In all numerical experiments in Section 4, the maximum number of switches across 200 runs (50 instances $\times$ 4 objectives) was below $0.01 \cdot N^2L$, and the average runtime remained around 2 seconds for instances with more than 2,500 customers.*

3.2. Multiplicity of equilibria

Theorem 1 guarantees an equilibrium but not a unique one, and with economies of scale multiplicity is the rule rather than the exception. The simplest manifestation is *lock-in*: if all customers are with the same provider, that provider is as cheap as it can be, while every alternative would be joined by a single customer and is as expensive as it can be. Let $\mu^{(j)}$ denote the matching that assigns every customer to provider j .

Proposition 1 (Lock-in equilibria). *For arbitrary tariff schedules (no monotonicity is required) the matching $\mu^{(j)}$ is a Nash equilibrium if and only if*

$$v_{ik} - v_{ij} \leq p_k(1) - p_j(N) \quad \text{for all } i \in \mathcal{C} \text{ and } k \in \mathcal{L} \setminus \{j\}. \quad (9)$$

In particular, if $\max_{i \in \mathcal{C}} \max_{j, k \in \mathcal{L}} (v_{ik} - v_{ij}) \leq \min_{k \in \mathcal{L}} p_k(1) - \max_{j \in \mathcal{L}} p_j(N)$, then $\mu^{(j)}$ is a Nash equilibrium for every $j \in \mathcal{L}$.

Condition (9) compares preference heterogeneity with the scale economies: whenever no customer's valuation advantage for any alternative exceeds the tariff gap between a provider serving one customer and a provider serving all of them, the market can be locked into *any* provider, including one that no customer ranks first. The aggregate utilities of these lock-in equilibria, $U(\mu^{(j)}) = \sum_i v_{ij} - Np_j(N)$, differ across j by the full difference in valuations and tariffs, so the gap between the best and the worst equilibrium can be as large as the differentiation among providers. This is the formal counterpart of the “bad equilibria” that appear in our experiments when valuations are weak relative to tariffs (Scenario 6 in Section 4.5.2), and it is the coordination failure familiar from markets with network effects (Katz and Shapiro, 1985): an inferior outcome persists because no customer gains by moving first.

With two providers we can describe the equilibrium set completely. Let $\mathcal{L} = \{1, 2\}$ and define, for each customer, the *net preference for provider 2*,

$$\delta_i := v_{i2} - v_{i1}, \quad i \in \mathcal{C},$$

and denote by $\delta_{(1)} \leq \delta_{(2)} \leq \dots \leq \delta_{(N)}$ its order statistics (ties broken arbitrarily but fixed), with the conventions $\delta_{(0)} = -\infty$ and $\delta_{(N+1)} = +\infty$. For $m = 1, \dots, N$ define the *tariff wedge*

$$g(m) := p_2(N - m + 1) - p_1(m), \quad (10)$$

the difference between the tariff a customer would pay as the $(N - m + 1)$ -th customer of provider 2 and the tariff she pays as one of m customers of provider 1; we extend it by $g(0) := g(1)$ and $g(N + 1) := g(N)$ (only the conventions on $\delta_{(0)}$ and $\delta_{(N+1)}$ matter at the boundary). Since tariffs are nonincreasing in the load, g is nondecreasing in m : the more customers provider 1 serves, the more attractive it is relative to provider 2.

Lemma 2 (Structure of equilibria with two providers). *Let $\mathcal{L} = \{1, 2\}$ and let the tariff schedules be nonincreasing in the load.*

- (a) *A matching μ with $m_1(\mu) = m$ is a Nash equilibrium if and only if $\delta_i \leq g(m)$ for every customer i assigned to provider 1 and $\delta_i \geq g(m + 1)$ for every customer assigned to provider 2. In every equilibrium, the customers of provider 1 are m customers with the smallest net preferences for provider 2.*

(b) The set of equilibrium loads of provider 1 is

$$E = \{ m \in \{0, 1, \dots, N\} : \delta_{(m)} \leq g(m) \text{ and } \delta_{(m+1)} \geq g(m+1) \}. \quad (11)$$

(c) The Sweep Algorithm under the ordering $[1, 2]$ terminates with $m_1 = \max E$, and under the ordering $[2, 1]$ with $m_1 = \min E$.

Lemma 2 reduces an equilibrium with two providers to a single number, the load of provider 1, and shows that the two activation orders reach the largest and the smallest equilibrium load: activating a provider first gives it the maximal load consistent with equilibrium. The set E need not consist of consecutive integers (in Example 1 it is $\{0, 3\}$) but every equilibrium load lies between these two extremes. The monotonicity result that motivated our study of priority orderings follows at once.

Proposition 2 (Effect of priority order with two providers). *Let $\mathcal{L} = \{j_1, j_2\}$ and let $m_{j_1}^{[j_1, j_2]}$ and $m_{j_1}^{[j_2, j_1]}$ denote the load of provider j_1 in the terminal matching of the Sweep Algorithm under the orderings $[j_1, j_2]$ and $[j_2, j_1]$, respectively. Then every Nash equilibrium μ satisfies*

$$m_{j_1}^{[j_2, j_1]} \leq m_{j_1}(\mu) \leq m_{j_1}^{[j_1, j_2]},$$

and in particular $m_{j_1}^{[j_2, j_1]} \leq m_{j_1}^{[j_1, j_2]}$.

Giving a provider lower priority can therefore only reduce the number of customers it serves, and the two orderings bracket every equilibrium. This multiplicity is the operator's opportunity: among all self-enforcing assignments, which one best serves its objective, and can it be reached?

A decentralized benchmark.. Theorem 1 also indicates how to model a market *without* an operator. Suppose customers start from an arbitrary initial assignment and switch one at a time, in an arbitrary order, for as long as some customer has a profitable switch. The market then follows an improvement path, and when no profitable switch remains it has reached a Nash equilibrium—which one depends on the initial assignment and on the order of moves. Randomizing over both induces a distribution over equilibria, and a natural measure of the value of coordination is the gap between the operator's objective at the equilibrium it selects and the expected value of that objective at the equilibrium such a process reaches. This gap is bounded above by the spread of the operator's objective over the equilibrium set: every outcome of the process is itself an equilibrium, hence at least as good, for the operator's objective, as the worst equilibrium. Section 4 reports that best–worst range; estimating the distribution within it is a natural next step that we leave for future work.

3.3. Selecting equilibria through the activation order

We now ask how much the operator can achieve with the activation order alone, taking as objective the total payment made by customers.

Definition 3 (Total Customer Payment). *For a matching μ with load vector $\mathbf{m} = (m_1, \dots, m_L)$, the **total customer payment** is*

$$R(\mu) = R(\mathbf{m}) = \sum_{j \in \mathcal{L}} m_j p_j(m_j) = \sum_{j \in \mathcal{L}} \varphi_j(m_j). \quad (12)$$

Total payment is a natural objective for a cost-conscious operator and, for a fixed assignment, reducing it is equivalent to increasing aggregate utility; across assignments the two objectives differ because valuations depend on the assignment (Remark 4). Each $\varphi_j(m) = m p_j(m)$ is concave, so R is a concave function of the load vector and is minimized, absent equilibrium constraints, by concentrating all customers on the provider with the smallest cost parameter (Lemma 3 in the Appendix A.3). The equilibrium constraints are what prevent this, and the question is how close to the unconstrained minimum an activation order gets. With two providers the answer is complete.

Theorem 2 (Equilibrium selection with two providers). *Let $\mathcal{L} = \{1, 2\}$, and let $\mu^{[1,2]}$ and $\mu^{[2,1]}$ be the terminal matchings of the Sweep Algorithm under the two activation orders.*

(a) *A payment-minimizing Nash equilibrium is reached by one of the two orders:*

$$\min_{\mu \in \mathcal{M}} R(\mu) = \min\{R(\mu^{[1,2]}), R(\mu^{[2,1]})\}.$$

(b) *Suppose $\beta_1 \leq \beta_2$. If provider 1 retains at least half of the customers even when it is activated last, i.e., $m_1(\mu^{[2,1]}) \geq N/2$, then the ascending-cost order $[1, 2]$ minimizes total customer payments among all Nash equilibria.*

Part (a) says that the operator never needs more than two runs of the SA to find the payment-minimizing equilibrium with two providers: because total payment is a concave function of the load of provider 1 and every equilibrium load lies between the two SA outcomes (Lemma 2), the minimum over equilibria is attained at one of those two extremes—no equilibrium in between can pay less, whether or not it exists. Part (b) identifies when the endpoint reached by the ascending-cost order is the right one. The intuition is the one behind the unconstrained minimum: total payment decreases as demand concentrates on the cheaper provider, so among the equilibria the one with the largest load of provider 1 pays least—*provided* all equilibria lie in the region where concentrating on provider 1 lowers payments. That region is the half of the load range in which provider 1 already serves at least half of the market, and the condition in (b) says precisely that every equilibrium lies in it (by Proposition 2, the $[2, 1]$ -outcome has the smallest equilibrium load of provider 1). The condition can be checked with a single run of the SA and holds whenever the cheaper provider is sufficiently attractive; when it fails, the ascending order can be strictly worse than the descending order, because concentrating demand on the *other* provider may save more through scale economies than the cost advantage of provider 1 is worth (Example 1 in Appendix A.3).

With three or more providers neither part of Theorem 2 extends, and the following proposition shows that the failure is structural rather than a matter of choosing the right order.

Proposition 3 (Sequential activation with three or more providers). *(i) For every $L \geq 3$ there are instances in which the ascending-cost order does not minimize total customer payments among Nash equilibria, even though $\beta_1 < \beta_2 < \dots < \beta_L$ and the price cap never binds; the terminal matching of the ascending order pays strictly more than another equilibrium while its load vector is lexicographically larger.*

- (ii) *More strongly, there are instances—with as few as three providers—in which no priority ordering reaches a payment-minimizing Nash equilibrium: the payment-minimizing equilibrium is not the terminal matching of the Sweep Algorithm under any of the $L!$ orderings.*

The examples that establish Proposition 3 are small (three providers, with three and six customers) and are worked out in Appendix A.3, where part (i) is extended to every $L \geq 3$ by appending providers that no customer finds attractive. Both turn on the shape of total payment. When the cap does not bind, $\varphi_j(m) = \beta_j \sqrt{m}$, so R rewards two things at once: serving customers through a provider with a small cost parameter, and serving them *together*, since concavity makes a fragmented market pay more than a concentrated one. With two providers these are the same question—an equilibrium is described by the load of provider 1, and the two activation orders move that single number to its extremes—which is why Theorem 2 settles the case. With three or more, the question becomes *which* provider the market should gather at, and the answer need not be the cheapest one.

The first example makes the point. Its cost parameters are $\beta = (1.3, 1.8, 1.9)$ with $N = 3$. Serving all three customers through provider 1 would pay only $1.3\sqrt{3} = 2.252$, but it is not an equilibrium: the third customer values provider 2 at 4.0 against 2.3 for provider 1, and leaves. The ascending order keeps the other two at provider 1 and lets her go, terminating at loads $(2, 1, 0)$; she is then alone at provider 2 and pays $p_2(1) = 1.8$, the highest tariff that provider’s schedule can produce, so the market pays $1.3\sqrt{2} + 1.8 = 3.638$. Concentrating all three at provider 2 instead—the provider with the highest aggregate valuation—forgoes provider 1’s cost advantage but removes the fragmentation, and pays $1.8\sqrt{3} = 3.118$, some 17% less. Note also that the ascending outcome $(2, 1, 0)$ is lexicographically larger than $(0, 3, 0)$: the natural greedy rule—fill the cheapest provider as far as equilibrium allows, then the next—selects precisely the assignment that pays more.

The second example shows that this is not a defect of one particular order, and its mechanism is a different one: a provider can retain customers only once it already serves enough of them. There $\beta = (1.2, 1.6, 1.9)$ and $N = 6$, and the payment-minimizing equilibrium splits the market between the two cheaper providers, four customers paying $p_2(4) = 0.800$ at provider 2 and two paying $p_1(2) = 0.849$ at provider 1. Those four stay only because the tariff is already low: three of them value provider 2 at 3.9 and provider 3 at 4.8, so each obtains $3.9 - 0.800 = 3.1$ and would obtain $4.8 - p_3(1) = 2.9$ by switching, whereas if provider 2 served a single customer the same comparison would read $3.9 - 1.600 = 2.3$ against 2.9, and she would switch. Sequential activation cannot manufacture that head start: it opens providers one at a time and lets the market settle among the active ones before opening the next, so by the time provider 3 is opened, provider 2 holds one customer, retains nobody, and the market drains into provider 3. All six orders terminate at that equilibrium, which pays 19% more than the one none of them reaches.

The two examples pull in opposite directions—in the first, concentrating on a costlier provider beats splitting across cheaper ones; in the second, splitting across the two cheaper providers beats concentrating on the costliest—so no rule that ranks providers before seeing the valuations can be right in both. Proposition 3 thus draws a sharp line: the activation order is a complete selection instrument with two providers and an incomplete one with more, which is why Section 3.4 optimizes directly over the equilibrium set.

The incompleteness is in fact more basic than the payment objective. Because the SA starts from a lock-in assignment and executes only strictly profitable switches, it cannot leave a lock-in equilibrium, and there may be equilibria—including the only ones that separate customers according to their preferences—that no ordering reaches.

Proposition 4 (Incompleteness of the Sweep Algorithm). *For every $L \geq 3$ there are instances of Γ with Nash equilibria that are not the terminal matching of the Sweep Algorithm under any of the $L!$ priority orderings. Indeed, whenever the lock-in condition of Proposition 1 holds for every provider, the Sweep Algorithm can reach only the L lock-in equilibria, whatever the ordering.*

These negative results delimit the activation order as a selection instrument; they do not make it powerless. The next proposition collects what sequential activation *can* guarantee with any number of providers: it implements every lock-in equilibrium exactly, it attains the unconstrained minimum of total payments whenever scale economies dominate preference heterogeneity, and every run comes with a performance certificate that the operator can evaluate before running it. Recall that $\mu^{(j)}$ denotes the matching that assigns every customer to provider j and that $\varphi_j(m) = m p_j(m)$ is the total payment collected by provider j at load m .

Proposition 5 (What sequential activation guarantees).

- (i) Concentration is implementable. *For arbitrary tariff schedules, the Sweep Algorithm under an ordering with $o(1) = j$ terminates at $\mu^{(j)}$ if and only if $\mu^{(j)}$ is a Nash equilibrium—whatever the rest of the ordering. Consequently, L runs of the algorithm, one with each provider activated first, reach every lock-in equilibrium of the instance, and their best terminal matching has total payment at most $\min\{R(\mu^{(j)}) : \mu^{(j)} \in \bar{\mathcal{M}}\}$ whenever some lock-in equilibrium exists.*
- (ii) First-best payments under dominant scale economies. *For any nonincreasing tariff schedules, every matching μ satisfies $R(\mu) \geq \min_{j \in \mathcal{L}} \varphi_j(N)$. Hence, if the lock-in condition (9) holds for some $j^* \in \arg \min_{j \in \mathcal{L}} \varphi_j(N)$, the single run that activates j^* first terminates at $\mu^{(j^*)}$ and minimizes total customer payment over all matchings, not merely over equilibria: in this regime the self-enforcement constraint costs the operator nothing.*
- (iii) An ex-ante certificate. *Under the schedule (1), for every ordering o the terminal matching μ^T satisfies*

$$U(\mu^T) \geq \Phi(\mu^{(o(1))}) = \sum_{i \in \mathcal{C}} v_{i,o(1)} - \sum_{\ell=1}^N p_{o(1)}(\ell) \geq \sum_{i \in \mathcal{C}} v_{i,o(1)} - \min\{\bar{\beta}N, 2\beta_{o(1)}\sqrt{N}\},$$

a lower bound on realized aggregate utility that the operator can evaluate for every candidate first mover.

Part (i) turns the mechanism behind Proposition 4 into a guarantee: the same inertia that traps the Sweep Algorithm in lock-in equilibria also means that any lock-in equilibrium the operator prefers can be implemented exactly, with a single run and no computation beyond checking condition (9). Part (ii) identifies the regime in which this is the optimum itself. The

condition—the provider that is cheapest at full scale can hold the whole market—is precisely a dominance of scale economies over preference heterogeneity, is checkable in $O(NL)$ time by Proposition 1, and when it holds the sweep attains the minimum of R over all matchings, so no optimization over the equilibrium set can improve on it. This is also the regime our experiments with the minimum total payment objective turn out to inhabit: in Section 4.5.1 the best-of-orderings sweep matches the exact MIP in every payment-minimization instance. Part (iii) is weaker but universal: whatever the ordering, the operator knows in advance a level of aggregate utility that the terminal equilibrium is guaranteed to deliver (nonnegative under Assumption 1). It is a floor, not an optimality statement; closing the gap between this floor and the operator-optimal equilibrium is what the optimization of Section 3.4 is for.

Remark 3 (The stakes of selection). *Proposition 5(ii) and the cap place the payment of every matching in the interval $[\min_j \varphi_j(N), \bar{\beta}N]$, and both ends can be occupied by equilibria of the same instance. With two providers, $\beta_1 \leq \bar{\beta}$, $\beta_2 \geq \bar{\beta}\sqrt{N}$, and every customer valuing them at $v_{i1} = \bar{\beta}$ and $v_{i2} = 2\bar{\beta} - \beta_1$ (so that Assumption 1 holds), both lock-in conditions are satisfied— $v_{i2} - v_{i1} = \bar{\beta} - \beta_1 \leq p_2(1) - p_1(N) = \bar{\beta} - \beta_1/\sqrt{N}$, and $v_{i1} - v_{i2} = \beta_1 - \bar{\beta} = p_1(1) - p_2(N)$ —so $\mu^{(1)}$ and $\mu^{(2)}$ are both equilibria, with*

$$\frac{R(\mu^{(2)})}{R(\mu^{(1)})} = \frac{\bar{\beta}N}{\beta_1\sqrt{N}} = \frac{\bar{\beta}}{\beta_1} \sqrt{N}.$$

By the same two bounds this is the worst case: no two equilibria of any instance differ in total payment by more than the factor $\bar{\beta}N/\min_j \varphi_j(N)$. Equilibrium selection therefore matters by a factor of order $\sqrt{N}$, not by a constant, which is the analytical counterpart of the utility ranges documented in Section 4.5.2.

Propositions 3–5 together clarify the role of the SA in the rest of the paper. As a computational device it is fast, its outcomes are always equilibria, and it is provably optimal in the regime of Proposition 5(ii); in the experiments of Section 4 we use the best of its $L!$ outcomes as a heuristic. But outside that regime, enumerating orderings is neither exhaustive nor guaranteed to contain the operator-optimal equilibrium, and the number of orderings grows factorially. As a mechanism, the SA has a natural implementation—providers are opened sequentially and customers are allowed to switch—and, as we show next, its ascending-cost version has a desirable incentive property on the supply side.

Provider incentives.. The operator relies on the cost parameters that providers report, and we close the model with an explicit *reporting game*. The players are the L incumbents and, for each incumbent j , one *potential entrant* $e(j)$ with the same attributes ($v_{ie(j)} = v_{ij}$ for all i) and the same cost parameter β_j ; the market is thus contestable in the sense of Baumol et al. (1982). The game has two stages. In the first, each incumbent simultaneously reports a parameter $\hat{\beta}_j$; in the second, each entrant observes the reports and either stays out or enters and reports a parameter of its own; entry is costless. The operator then runs the SA in ascending order of reported parameters over the participating providers (ties broken in favor of incumbents, then by index) with the tariff schedules $\hat{p}_j(m) = \min\{\bar{\beta}, \hat{\beta}_j/\sqrt{m}\}$ for $m \geq 1$ (and $\hat{p}_j(0) = \bar{\beta}$). A provider’s average cost per customer, inclusive of a normal return, is $\beta_j/\sqrt{m}$ when it serves $m \geq 1$ customers, so its payoff is the profit $\pi_j = m_j(\hat{p}_j(m_j) - \beta_j/\sqrt{m_j})$,

equal to zero if it serves no customer or stays out; providers rank outcomes lexicographically, first by profit and then by the number of customers served. The profile we analyze is the one in which every incumbent reports truthfully and every entrant follows the *zero-profit entry rule* σ : enter and report truthfully ($\hat{\beta}_{e(j)} = \beta_j$) if the incumbent's report exceeds its cost ($\hat{\beta}_j > \beta_j$), and stay out otherwise. The rule is the natural contestability response: a truthful report never causes the entrant a loss, and it takes over the customers an overstating incumbent would have served.

Proposition 6 (Truthful reporting is an equilibrium under contestable entry). *Suppose the price cap does not bind for any provider serving at least one customer ($\bar{\beta} \geq \beta_j$ for all j), and consider the two-stage reporting game and the entry rule σ just described. Against σ :*

- (i) *If incumbent j understates its cost ($\hat{\beta}_j < \beta_j$), its profit is strictly negative whenever it serves at least one customer.*
- (ii) *If incumbent j overstates its cost ($\hat{\beta}_j > \beta_j$), its potential entrant enters and the Sweep Algorithm assigns no customer to j : its load and profit are zero.*
- (iii) *Truthful reporting yields zero profit and a weakly positive load.*

Moreover, (iv) *when all incumbents report truthfully, no entrant obtains a lexicographically better outcome by deviating from σ . Consequently, truthful reporting is a best response for every incumbent against the entry rule, and the profile in which every incumbent reports truthfully and every entrant follows σ is a Nash equilibrium of the two-stage reporting game.*

The argument is Bertrand-like: under free entry, an incumbent that inflates its cost is undercut by an identical entrant that is activated earlier and is cheaper at every load, and an incumbent that deflates its cost serves at a loss. We emphasize what is and is not claimed. First, Proposition 6 establishes a Nash equilibrium of the reporting game—truthfulness is a best response *against the entry rule*—not strategy-proofness, since truthful reporting need not be a dominant strategy; in a closed market with a fixed set of providers and no entry threat, providers would have incentives to overstate, and the operator would need an auction-like mechanism (Section 5). Second, the equilibrium concept is Nash, with the entry rule part of the profile; we do not claim subgame perfection. Off the equilibrium path, after an overstated report $\hat{\beta}_j > \beta_j$, the entrant could earn a strictly positive margin by reporting just below $\hat{\beta}_j$ rather than at cost—the familiar openness of Bertrand best responses—so the zero-profit rule is best viewed as the standard contestability benchmark rather than as the entrant's unique optimal conduct.

Remark 4 (Payments and aggregate utility). *Aggregate utility decomposes as $U(\mu) = \sum_i v_{i\mu(i)} - R(\mu)$. With two providers, if μ and μ' are equilibria with $m_1(\mu') < m_1(\mu)$ and D is the set of customers served by provider 1 under μ but by provider 2 under μ' , then $U(\mu) - U(\mu') = \sum_{i \in D} (v_{i1} - v_{i2}) - [R(\mu) - R(\mu')]$: the comparison of two equilibria is the sum of a valuation effect, which depends on who the contested customers are, and a payment effect, which depends only on loads. Theorem 2 controls the second effect; the first can go either way, which is why the payment-minimizing and the utility-maximizing equilibria generally differ (Section 4.5.1) and why the utility objective requires the optimization approach that follows. The potential also bounds how good the best equilibrium must be. For every matching μ , $U(\mu) \geq \Phi(\mu) \geq U(\mu) - R(\mu)$, because the payment term $\sum_j \sum_{\ell=1}^{m_j} p_j(\ell)$ of (8) lies between $R(\mu)$ and $2R(\mu)$ (nonincreasing tariffs give the lower*

bound, and $\sum_{\ell=1}^m p_j(\ell) \leq \min\{\bar{\beta}m, 2\beta_j\sqrt{m}\} \leq 2\varphi_j(m)$, by the telescoping bound in the proof of Proposition 5(iii), the upper). A maximizer μ° of Φ is a Nash equilibrium by Theorem 1(i), and chaining the two inequalities gives $U(\mu^\circ) \geq \Phi(\mu^\circ) \geq \Phi(\mu) \geq U(\mu) - R(\mu)$ for every matching μ : the equilibrium set always contains an assignment whose aggregate utility is within total payments of the first-best, and the optimization of Section 3.4 is guaranteed to find one at least this good.

3.4. Optimization over the equilibrium set

For general operator objectives, and for three or more providers, we solve the equilibrium selection problem (4) directly by embedding the self-enforcement conditions (7) as constraints of a mathematical program. The tariff relation makes this program a nonconvex mixed-integer *nonlinear* program; following common usage, and because a standard solver handles it through the spatial branch-and-bound described below, we refer to it as a MIP throughout. Although the framework accommodates any objective that is a function of the assignment (total payment, emissions, realized quality), we present the aggregate-utility version for concreteness.

Binary variables $x_{ij} \in \{0, 1\}$ indicate whether customer $i \in \mathcal{C}$ is assigned to provider $j \in \mathcal{L}$. For each provider, the continuous variable p_j is the tariff at the assigned load and p_j^+ is the entry tariff a switching customer would pay; u_{ij} and u_{ij}^+ are the corresponding utilities. The utility-maximization problem $\mathcal{Z}^0$ reads

$$\max \sum_{i \in \mathcal{C}} \sum_{j \in \mathcal{L}} u_{ij} x_{ij}, \quad (13)$$

$$\text{s.t.} \quad \sum_{j \in \mathcal{L}} x_{ij} = 1 \quad \forall i \in \mathcal{C} \quad (14)$$

$$p_j = \min\left\{\bar{\beta}, \frac{\beta_j}{\sqrt{\sum_{i \in \mathcal{C}} x_{ij}}}\right\} \quad \forall j \in \mathcal{L} \quad (15)$$

$$p_j^+ = \min\left\{\bar{\beta}, \frac{\beta_j}{\sqrt{1 + \sum_{i \in \mathcal{C}} x_{ij}}}\right\} \quad \forall j \in \mathcal{L} \quad (16)$$

$$u_{ij} = v_{ij} - p_j, \quad u_{ij}^+ = v_{ij} - p_j^+ \quad \forall i \in \mathcal{C}, j \in \mathcal{L} \quad (17)$$

$$u_{ij} \geq u_{ij'}^+ - M_{ijj'}(1 - x_{ij}) \quad \forall i \in \mathcal{C}, j, j' \in \mathcal{L}, j \neq j' \quad (18)$$

$$0 \leq p_j^+ \leq p_j \leq \bar{\beta} \quad \forall j \in \mathcal{L} \quad (19)$$

$$x_{ij} \in \{0, 1\} \quad \forall i \in \mathcal{C}, j \in \mathcal{L} \quad (20)$$

The objective maximizes aggregate utility. Constraint (14) assigns each customer to exactly one provider. Constraints (15) and (16) are the schedule (1) evaluated at the provider's assigned load and at that load plus one; note that the load in (16) is the load *excluding* the switching customer, since a customer who considers switching to j' is not assigned to it, so its argument is at least one. In (15) the argument is zero for a provider that serves no customer, and the convention $\beta_j/\sqrt{0} = +\infty$ of (1) then gives $p_j = \bar{\beta}$. This value never affects the optimization: p_j at a zero load enters only the utilities u_{ij} of customers not assigned to j , which appear in the objective with coefficient $x_{ij} = 0$ and in (18) only in the deactivated form,

and it satisfies (19) because $p_j^+ \leq \bar{\beta}$. Section 4 describes how the two cases are separated in the implementation with an indicator of positive load. Constraint (17) defines utilities, and (18) is the self-enforcement condition (7): if i is assigned to j , her utility must be at least what she would obtain by switching to any j' at j' 's entry tariff. Individual rationality holds automatically under Assumption 1 and is not imposed. The constant $M_{ijj'}$ must exceed the largest possible value of $u_{ij'}^+ - u_{ij}$; since $u_{ij} \geq v_{ij} - \bar{\beta}$ and $u_{ij'}^+ \leq v_{ij'} - p_{j'}(N+1)$, it suffices to take

$$M_{ijj'} = \max\left\{0, v_{ij'} - v_{ij} + \bar{\beta} - \min\left\{\bar{\beta}, \frac{\beta_{j'}}{\sqrt{N+1}}\right\}\right\}.$$

Finally, (19) bounds tariffs and (20) defines the binary domain. Every feasible solution of $\mathcal{Z}^0$ is a Nash equilibrium of Γ and conversely, so $\mathcal{Z}^0$ optimizes exactly over $\bar{\mathcal{M}}$; replacing the objective yields the operator-optimal equilibrium for any other objective, and minimizing (13) instead of maximizing it yields the worst equilibrium, which we use in Section 4 to bound the range of equilibrium outcomes.

Although formulation $\mathcal{Z}^0$ is exact, it becomes computationally prohibitive for large instances. To improve tractability, we introduce a clustering-based approximation, denoted by $\mathcal{Z}^c$. Customers with similar valuation patterns (such as comparable preferences for environmental performance or service quality) are grouped into behavioral clusters, on the premise that customers with similar preference profiles (Ziari et al., 2025) are likely to choose the same provider in an equilibrium. Let $\mathcal{G}$ denote the set of clusters, where each $g \in \mathcal{G}$ represents a subset of customers $\mathcal{C}_g \subseteq \mathcal{C}$ of size $n_g = |\mathcal{C}_g|$, and let $V_g^j = \sum_{i \in \mathcal{C}_g} v_{ij}$ be the aggregate valuation of cluster g for provider j . The decision variables are $x_g^j \in \{0, 1\}$, equal to 1 if cluster g is served by provider j ; p_j is the tariff at the assigned load; p_j^{+g} is the tariff provider j would charge if the whole cluster g joined it; and $U_g^j = V_g^j - n_g p_j$ and $U_g^{+j} = V_g^j - n_g p_j^{+g}$ are the corresponding cluster-level utilities. The cluster-level problem $\mathcal{Z}^c$ is

$$\max \quad \sum_{g \in \mathcal{G}} \sum_{j \in \mathcal{L}} U_g^j x_g^j \quad (21)$$

$$\text{s.t.} \quad \sum_{j \in \mathcal{L}} x_g^j = 1 \quad \forall g \in \mathcal{G} \quad (22)$$

$$p_j = \min\left\{\bar{\beta}, \frac{\beta_j}{\sqrt{\sum_{g \in \mathcal{G}} n_g x_g^j}}\right\} \quad \forall j \in \mathcal{L} \quad (23)$$

$$p_j^{+g} = \min\left\{\bar{\beta}, \frac{\beta_j}{\sqrt{n_g + \sum_{g' \in \mathcal{G}} n_{g'} x_{g'}^j}}\right\} \quad \forall g \in \mathcal{G}, j \in \mathcal{L} \quad (24)$$

$$U_g^j = V_g^j - n_g p_j, \quad U_g^{+j} = V_g^j - n_g p_j^{+g} \quad \forall g \in \mathcal{G}, j \in \mathcal{L} \quad (25)$$

$$U_g^j \geq U_g^{+j'} - M_g(1 - x_g^j) \quad \forall g \in \mathcal{G}, j, j' \in \mathcal{L}, j \neq j' \quad (26)$$

$$0 \leq p_j \leq \bar{\beta}, \quad 0 \leq p_j^{+g} \leq \bar{\beta} \quad \forall g \in \mathcal{G}, j \in \mathcal{L} \quad (27)$$

$$x_g^j \in \{0, 1\} \quad \forall g \in \mathcal{G}, j \in \mathcal{L} \quad (28)$$

The objective (21) maximizes aggregate utility over cluster-provider pairs, (22) assigns each cluster to one provider, (23)–(24) define assigned and entry tariffs through (1) (with the

same convention at a zero load), (25) defines cluster utilities, and (26) is the self-enforcement condition at the cluster level, where the constant must dominate $U_g^{+j'} - U_g^j$. Since $p_j \leq \bar{\beta}$ and the argument of $p_{j'}^{+g}$ never exceeds $N + n_g$, it suffices to replace M_g in (26) by

$$M_g^{jj'} = \max \left\{ 0, V_g^{j'} - V_g^j + n_g \left(\bar{\beta} - \min \left\{ \bar{\beta}, \frac{\beta_{j'}}{\sqrt{N+n_g}} \right\} \right) \right\}.$$

Solving $\mathcal{Z}^c$ yields a tractable approximation that drastically reduces the decision space, at the price of treating each cluster as if it switched as a block.

Since formulation $\mathcal{Z}^c$ aggregates customers, the resulting allocation may violate the individual-level conditions (7). To restore self-enforcement at the individual level, we introduce a refinement problem $\mathcal{Z}^f$ that re-optimizes individual assignments. The parameter $y_{ij} \in \{0, 1\}$ indicates whether customer i was assigned to provider j at the cluster level, and binary variables $d_i \in \{0, 1\}$ capture whether this assignment changes during refinement. The model minimizes the number of such deviations:

$$\min \sum_{i \in \mathcal{C}} d_i \tag{29}$$

$$(14)-(20) \tag{30}$$

$$d_i \geq y_{ij} - x_{ij} \quad \forall i \in \mathcal{C}, \forall j \in \{ \ell \in \mathcal{L} : y_{i\ell} = 1 \} \tag{31}$$

$$d_i \in \{0, 1\} \quad \forall i \in \mathcal{C} \tag{32}$$

Constraints (14)–(20) are the constraint set of $\mathcal{Z}^0$, so every feasible solution of $\mathcal{Z}^f$ is a Nash equilibrium; constraint (31) identifies the customers whose assignment deviates from the cluster-level outcome, and (32) imposes the binary domain on d_i .

Although the formulations $\mathcal{Z}^0$, $\mathcal{Z}^c$, and $\mathcal{Z}^f$ differ in structure, they share the nonlinear tariff relations, which induce a nonconvex dependence between assignment and tariff variables. We solve the formulations in this compact nonlinear form by spatial branch-and-bound, the standard framework for the global optimization of nonconvex MINLPs: convex underestimators provide valid lower-bounding relaxations of the nonlinear terms, and a systematic partition of the variable domain tightens these relaxations as the partitions shrink (Smith and Pantelides, 1999; Grossmann, 2002). Modern solvers such as GUROBI implement this framework directly, delivering rigorous primal and dual bounds throughout the search (Gurobi Optimization, LLC, 2025); their difference defines a certified optimality gap, which we report in Section 4. Spatial branch-and-bound is a choice of solution technology rather than a necessity, however: because loads are integers between 0 and N , the tariff relations can also be reformulated exactly with linear constraints.

Remark 5 (An exact linear reformulation). *For $j \in \mathcal{L}$ and $r \in \{0, 1, \dots, N\}$, introduce binaries $z_{jr} \in \{0, 1\}$ with*

$$\sum_{r=0}^N z_{jr} = 1, \quad \sum_{r=0}^N r z_{jr} = \sum_{i \in \mathcal{C}} x_{ij},$$

so that $z_{jr} = 1$ exactly when provider j serves r customers. Two substitutions then remove

every nonlinearity from $\mathcal{Z}^0$. First, the tariff relations (15)–(16) are replaced by the linear equations

$$p_j = \sum_{r=0}^N p_j(r) z_{jr}, \quad p_j^+ = \sum_{r=0}^N p_j(r+1) z_{jr},$$

in which the coefficients $p_j(r)$ are constants precomputed from the schedule (1). Second, the objective (13) is bilinear as written, since both u_{ij} and x_{ij} are variables; using the identity $\sum_{i,j} u_{ij} x_{ij} = \sum_{i,j} v_{ij} x_{ij} - \sum_j m_j p_j(m_j)$, it is replaced by the linear objective

$$\max \sum_{i \in \mathcal{C}} \sum_{j \in \mathcal{L}} v_{ij} x_{ij} - \sum_{j \in \mathcal{L}} \sum_{r=0}^N r p_j(r) z_{jr},$$

whose second term is total customer payment. The self-enforcement constraints (18) are linear in $(u_{ij}, u_{ij}^+, x_{ij})$ and are kept as they are. This turns $\mathcal{Z}^0$ into an equivalent mixed-integer linear program with $L(N+1)$ additional binaries, solvable by linear branch-and-bound alone. For $\mathcal{Z}^f$ the tariff substitution alone suffices, since its objective $\sum_i d_i$ is already linear. For $\mathcal{Z}^c$, the load equations become $\sum_r r z_{jr} = \sum_{g \in \mathcal{G}} n_g x_g^j$, the objective is replaced by $\sum_{g,j} V_g^j x_g^j - \sum_j \sum_r r p_j(r) z_{jr}$ in the same way, and the cluster entry tariffs (24) become $p_j^{+g} = \sum_{r=0}^N p_j(r+n_g) z_{jr}$ for each $g \in \mathcal{G}$ and $j \in \mathcal{L}$, again with precomputed coefficients. We keep the compact nonlinear forms in Section 4, where the solver’s spatial branch-and-bound handles them directly and certifies the reported gaps; the linear reformulation is the alternative of choice when only a MILP solver is available, at the cost of a larger model.

Implementing a selected equilibrium.. Once the operator has solved (4)—by the SA with two providers, or by the optimization models above—it must implement the selected equilibrium μ^* . In the institution described in Section 1 this requires no instrument beyond the ones the operator already has. The operator publishes each provider’s tariff schedule $p_j(\cdot)$ and the assignment μ^* , and each customer is billed the tariff $p_{\mu^*(i)}(m_{\mu^*(i)}(\mu^*))$. A customer who requests a different provider k is served by k at the entry tariff $p_k^+(\mu^*)$, which she can compute from the published schedule and loads. Because μ^* satisfies (7), no customer has an incentive to make such a request, so the assignment is self-enforcing without any restriction on customers’ freedom to choose. The operator’s role is thus that of a coordinator recommending one of several self-fulfilling outcomes—the coordination role that Proposition 1 shows a decentralized market cannot perform on its own—and its only informational requirements are the ones stated in Section 1.1: the providers’ cost parameters and the customers’ valuations. On the supply side, Proposition 6 establishes that truthful reporting is an equilibrium under the ascending-order Sweep Algorithm; whether an operator that selects equilibria through the optimization models above can likewise be made robust to misreporting is open, and we return to it in Section 5. Relaxing the latter is the Bayesian extension discussed in Section 5.

4. Numerical Experiments

In this section, we present the results of our numerical experiments that are designed to evaluate both the computational performance of the proposed solution approaches and

the managerial implications of centralized decision making with respect to customer utility, environmental friendliness, and service quality.

In particular, our numerical experiments aim to address the following research questions:

- **R1:** To what extent can the cluster-based MIP formulation and the sweep algorithm approximate the solutions of the exact MIP model across different IMO objectives, and what trade-offs arise between solution quality and computational time?
- **R2:** What level of total utility improvement can centralized decision making provide under different customer preference distributions with respect to environmental friendliness, service price, and service quality?
- **R3:** What trade-offs arise among alternative objectives available to the IMO, namely maximizing total service utility, environmental friendliness, and service quality?

The remainder of this section is organized as follows. We first describe the data used to generate the problem instances, present the experimental design, provide implementation details, and validate the approximation (1) used to represent delivery costs in our experiments. Then, we discuss the main computational and managerial results structured around Research Questions R1–R3. Details on the problem setup and its treatment for R3 are provided in [Appendix B.1](#).

4.1. Data

Focusing on a realistic urban delivery environment, we study the Sariyer district of Istanbul—specifically the neighborhoods of Merkez and Yeni Mahalle, which together form one of the most densely populated residential zones and exhibit correspondingly high parcel-generation activity. Using publicly available demographic statistics and national e-commerce transaction records, we generate the neighborhood-level parcel delivery demand (see [Table Appendix B1](#) in [Appendix B.2](#)). The resulting range, approximately 2,595–2,866 parcels, centers around a mean estimate of 2,731.

To generate the spatial input set for the model, we sample 500 points from Merkez and 140 points from Yeni Mahalle, yielding 640 candidate delivery locations, all of which correspond to residential areas, whose geographic distribution is shown in [Figure 1](#). These locations lie within the administrative boundaries of the two neighborhoods and serve as discrete demand units for the matching and routing experiments. Each sampled point may correspond to multiple parcels, capturing the practical reality that a single stop often serves several households or apartments clustered around a short walking radius.




Figure 1: Spatial distribution of customers in Yeni Mahalle and Merkez, Sariyer

Each location is associated with latitude–longitude coordinates, but these coordinates are not used directly when computing marginal routing costs in the matching problem. Instead, route lengths are generated using the Daganzo approximation introduced in Equation (1). This surrogate is calibrated and then validated against actual road-network distances and a set of sample routes within the target region. In Section 4.4, we report the ex-post comparison between the approximation results and the observed route lengths, confirming that the approximation is quite accurate for the purposes of our study.

4.2. Experiment design

To examine how heterogeneity in customer composition and differences in valuations of service quality and environmental friendliness influence equilibrium outcomes and the computational performance of the proposed methods, the experimental design is structured along two controlled dimensions: (i) customer segment composition, (ii) valuation intensity for environmental and quality attributes.

Service providers: We consider a last-mile delivery market with four service providers, indexed by $p \in \mathcal{P} = \{1, 2, 3, 4\}$. Each provider is characterized by a cost structure and a service capability profile defined along two latent service dimensions: environmental friendliness (e) and delivery quality (q). The environmental dimension captures sustainability-related attributes such as emissions performance, fleet composition, and packaging practices, while the quality dimension reflects operational service features including punctuality, parcel handling, and customer interface quality. A third component, successful receipt of the package (r), is assumed to be fulfilled by all providers.

Each provider $j \in \mathcal{P}$ is associated with a capability pair (ϵ_j, ρ_j) , where ϵ_j denotes its environmental performance level and ρ_j its service-quality performance level. The provider set is constructed to represent distinct strategic profiles with an ordinal dominance structure along the two dimensions:

$$(\text{Environment}) \quad \epsilon_2 > \epsilon_1 > \epsilon_3 > \epsilon_4, \quad (\text{Quality}) \quad \rho_3 > \rho_1 > \rho_2 > \rho_4.$$

Operationally, this corresponds to four stylized provider types: P1 performs strongly on both dimensions, representing a high-standard and environmentally responsible operator;

P2 specializes in sustainability without quality differentiation; P3 emphasizes delivery quality with minimal environmental commitment; and P4 provides a baseline service without enhancements along either dimension. Provider capability levels used in the experiments are summarized in Table 1. These values determine the fraction of each customer’s attribute valuation that is realized when matched with a given provider.

Table 1: Provider Capability for Value Realization

Provider	Environmental (ϵ_j)	Quality (ρ_j)
P1	0.50	0.50
P2	0.75	0.25
P3	0.25	0.75
P4	0	0

Customers characteristics: We consider heterogeneous customers, characterized by a unique valuation vector (e_i, q_i, r_i) representing the intrinsic willingness to pay for environmental performance, delivery quality, and successful receipt, for a customer i . Using the customer segments identified by Engelhardt (2023) based on a nationally representative survey conducted in Germany, customers are partitioned into four behavioral groups: S1–S4 (see Appendix B.3 for details).

For each customer segment, individual customer valuations are sampled from triangular distributions calibrated to the segment’s preference orientation. Environment-oriented segments assign higher weights to the environmental attribute e , while quality-oriented segments assign higher weights to service quality q . All customers have strictly positive receipt valuations r_i , representing the base utility of obtaining home delivery. To depict the value and utility procedure, we first sample each customer’s attribute valuations (e_i, q_i) from the segment-specific triangular distributions reported in Appendix B.3. Given a draw (e_i, q_i, r_i) and a candidate provider j , the realized value is

$$v_{ij} = r_i + \epsilon_j e_i + \rho_j q_i,$$

so that $v_{i1} = r_i + 0.5(e_i + q_i)$ while $v_{i2} = r_i + (0.75e_i) + (0.25q_i)$ (and analogously $v_{i3} = r_i + (0.75q_i) + (0.25e_i)$, $v_{i4} = r_i$). Tariffs p_j then translate realized value into utility via $u_{ij} = v_{ij} - p_j$. Since the receipt valuation satisfies $r_i \geq \bar{\beta}$ for every customer (see Appendix B.3 and the choice of $\bar{\beta}$ below), Assumption 1 holds in all instances.

Segment definitions capture the *relative preference orientation* of customers — that is, whether they value environmental performance or service quality more strongly compared to other attributes and other segments. However, beyond this relative structure, the *absolute strength* of these valuations with respect to the base delivery utility also plays a critical role in equilibrium outcomes and optimal allocations. Customers with stronger attribute valuations react more sensitively to environmental and quality differences, which can substantially change equilibrium structures. To systematically control this effect across experiments, we introduce three levels of preference strength for the environmental and quality valuation distributions: low (L), medium (M), and high (H). The medium level (M) corresponds to the base-case parameterization of the triangular distributions. The medium level

Table 2: Experiment groups and population compositions (segment shares in %)

Experiment group	ID	Configuration	S1	S2	S3	S4	Valuation shift (e,q)
BC (base case)	0	Reference mix	34.08	16.76	16.48	32.68	(M,M)
E1 (segment composition)	1	Balanced-type skew	40	20	20	20	(M,M)
	2	Environment-focused	20	40	20	20	(M,M)
	3	Quality-focused	20	20	40	20	(M,M)
	4	Low-valuation skew	20	20	20	40	(M,M)
	5	Even mix	25	25	25	25	(M,M)
E2 (preference strength)	6	Low	34.08	16.76	16.48	32.68	(L,L)
	7	Environment-high	34.08	16.76	16.48	32.68	(H,M)
	8	Quality-high	34.08	16.76	16.48	32.68	(M,H)
	9	High	34.08	16.76	16.48	32.68	(H,H)

(M) corresponds to the base-case triangular parameterization reported in [Appendix B.3](#) Table [Appendix B.4](#). Relative to this base case, the high level (H) doubles the parameters of the corresponding distribution for all segments (e.g., Env=H multiplies each segment’s environmental triangular parameters by two), whereas the low level (L) halves them.

Using these segment definitions and preference strength levels, we organize the experiments into two groups with a total of ten different scenarios as presented in [Table 2](#).

- **Base case:** The base case (BC) defines the reference configuration for the scenario design and serves as the baseline for one-way analysis. It represents a plausible scenario based on an empirically observed segment distribution, with all other parameters set to realistic mid-range values. In the experiment groups we define below, either the segment composition or the preference strength is varied while the other is kept fixed to the BC values.

E1 Experiments: The first experiment group focuses on variation in customer segment composition while holding service quality and environmental friendliness valuations fixed at their base-case levels (M,M). In these experiments, the shapes and parameterizations of the triangular distributions governing environmental and quality valuations are held constant across segments, while only the population shares of the four customer types (S1–S4) vary. Our goal with this design is to isolate the effect of market composition from valuation intensity, allowing differences in equilibrium outcomes and computational performance to be attributed solely to shifts in segment prevalence.

As summarized in [Table 2](#), Scenario 0 serves as the *base case* and defines the reference population mix with the population shares reported in [Engelhardt \(2023\)](#). The remaining experiments in this group introduce controlled skews toward specific customer types. In particular, the *balanced-type skew* increases the share of customers with intermediate valuations across both environmental and quality dimensions. The *environment-focused* and *quality-focused* configurations concentrate the population in segments that place dominant weight on environmental friendliness and service quality,

respectively. The *low-valuation skew* increases the share of customers with weak environmental and quality valuations. Finally, the *even-mix* configuration assigns equal shares to all segments.

- **E2 Experiments:** The second experiment group shifts attention from population composition to preference strength. In these experiments, segment shares are held fixed at the base-case levels from Engelhardt (2023), while the triangular distributions governing environmental and quality valuations are selectively rescaled. This preserves the relative ordering of segments by preference orientation but changes the overall strength of customer valuations. As a result, observed differences in equilibrium outcomes can be attributed to intensity effects rather than compositional effects.

As reported in Table 2, the *Low* configuration uniformly contracts both environmental and quality valuation distributions across all segments, producing a market in which customers are systematically less sensitive to both attributes. The *Environment-high* and *Quality-high* configurations selectively amplify one valuation dimension while leaving the other at its base level, thereby creating asymmetric preference pressure on provider offerings. The *High* configuration scales both dimensions upward simultaneously, representing a market with strongly expressed preferences for both service quality and environmental performance relative to the pure utility of home delivery.

Across two experiment groups, both of which include the base case, each configuration is evaluated over five independently generated instances, yielding a total of fifty experimental runs.

4.3. Implementation details

Having specified the customer population, their value realizations, and the ten scenario compositions described in Section 4.1 and 4.2, the remaining structural component is the set of provider parameters β_j . Recall that these parameters encode how prices adjust with workload and therefore determine the competitive positioning of each provider in terms of cost. A direct empirical estimation of these dynamics is beyond the scope of this study. However, to enable a meaningful equilibrium analysis, the parameters β_j are calibrated using the base-case customer segment distribution so that this distribution constitutes an equilibrium. Details of this calibration procedure are provided in Appendix B.4.

The resulting parameter values are:

$$\beta_{P1} = 3.59, \quad \beta_{P2} = 3.10, \quad \beta_{P3} = 3.02, \quad \beta_{P4} = 0.75,$$

We set $\bar{\beta} = 0.5$, the largest value for which Assumption 1 holds in every instance, since the receipt valuation r_i is at least 0.5 for every customer (Appendix B.3).

With the β_j parameters given, the solution procedure proceeds according to the three-stage framework summarized in Figure 2. In implementation, a k -means procedure is first applied to customers based on their realized valuation attributes (e, q) . The resulting clusters serve as inputs to Stage $\mathcal{Z}^c$, which assigns each cluster to a single common provider. Specifically, Stage $\mathcal{Z}^c$ determines the provider choice for an entire cluster based on the aggregate valuation of that cluster. The subsequent stage then refines these assignments to ensure




Figure 2: MIP Framework

individual rationality. To restore individual-level equilibrium behavior, Stage $\mathcal{Z}^f$ refines the allocation by minimizing the number of customer reassignments required to achieve a Nash-consistent outcome. In this stage, utility dominance constraints are imposed to guarantee that each customer's assigned provider yields maximal utility among all feasible options.

We conduct a variance analysis to determine the number of clusters to consider in $\mathcal{Z}^c$, as shown in Figure 3. Based on these results, we set the number of clusters to $k = 40$. Increasing the number of clusters to this level eliminates approximately 90% of the total within-group variance observed at $k = 5$, providing a favorable balance between computational tractability and model fidelity. The performance of the approximation framework with 40 clusters is evaluated in Section 4.5.1, where we report its optimality gap across all scenarios.




Figure 3: Total in-group variance with clustering

During optimization, the tariff rule (15) requires care at a zero load, where the argument of the square root vanishes and (1) returns the cap. We introduce the load $\ell_j = \sum_{i \in \mathcal{C}} x_{ij}$ and a binary indicator a_j of a positive load,

$$a_j \leq \ell_j \leq N a_j, \quad a_j \in \{0, 1\}, \quad \forall j \in \mathcal{L}, \quad (33)$$

so that $a_j = 1$ if and only if provider j serves at least one customer. Writing

$$q_j = \frac{\beta_j}{\sqrt{\ell_j + 1 - a_j}} + \bar{\beta}(1 - a_j), \quad \forall j \in \mathcal{L}, \quad (34)$$

gives $q_j = \beta_j / \sqrt{\ell_j}$ when j serves at least one customer and $q_j = \beta_j + \bar{\beta} > \bar{\beta}$ when j is empty, so that

$$p_j = \min\{\bar{\beta}, q_j\}, \quad \forall j \in \mathcal{L} \quad (35)$$

reproduces (15) exactly, including $p_j = \bar{\beta}$ at a zero load. The min in (35) is then linearized through a big- M with a binary variable t_j :

$$p_j \leq \bar{\beta} \quad \forall j \in \mathcal{L}, \quad (36)$$

$$p_j \leq q_j \quad \forall j \in \mathcal{L}, \quad (37)$$

$$p_j \geq \bar{\beta} - M(1 - t_j) \quad \forall j \in \mathcal{L}, \quad (38)$$

$$p_j \geq q_j - Mt_j \quad \forall j \in \mathcal{L}, \quad (39)$$

$$q_j - \bar{\beta} \leq Mt_j \quad \forall j \in \mathcal{L}, \quad (40)$$

$$\bar{\beta} - q_j \leq M(1 - t_j) \quad \forall j \in \mathcal{L}, \quad (41)$$

$$t_j \in \{0, 1\} \quad \forall j \in \mathcal{L}. \quad (42)$$

Here $M = \max_{j \in \mathcal{L}} \beta_j + \bar{\beta}$ is used as the big- M constant; note that the constant $\bar{\beta}$ that inflates q_j in (34) must be kept separate from M , since $q_j \leq \max_j \beta_j + \bar{\beta} = M$ is what makes (40) and (39) slack in their inactive branch. The entry tariffs p_j^+ of constraint (16) are linearized in the same way with $\beta_j / \sqrt{\ell_j + 1}$ in place of q_j and a second binary variable per provider; no indicator is needed there, since the post-switch load is at least one.

All algorithms are implemented in Python, and the computational experiments are conducted on a Linux workstation equipped with two Intel Xeon Gold 5418Y processors (2.00 GHz) , 96 CPU(s) and 512 GB of RAM. All MIP models are solved using GUROBI 13.0.1.

4.4. Validation of the cost approximation

To test the empirical validity of the routing cost approximation by Equation (1), we solve a large collection of CVRP instances using PyVRP (Wouda et al., 2024), as a state-of-the-art vehicle routing solver. The resulting routing costs are then compared with the Daganzo approximation to assess its empirical accuracy in predicting average routing costs across different problem scales. Each scenario involves a complete solution to a vehicle routing problem defined over the studied urban road network, where depot and customer locations are embedded spatially within the Sariyer district of Istanbul. Pairwise distances between all nodes—both depots and demand points—are computed as the shortest feasible travel paths on the real Sariyer road network. For each instance, the fleet operates under a cardinality constraint, whereby each vehicle is limited to serving at most 250 packages. The resulting total fleet travel distance, denoted $\mathcal{R}$, is recorded and used to compute the observed average per-customer delivery cost. This way we construct a sample of 500 capacitated VRP instances, stratified across ten customer-count regimes of increasing scale. The smallest regime covers instances with 100–250 customers, corresponding to an interval width of 150, while the remaining nine regimes span customer counts from 250 up to 2,500 in uniform increments of 250. Each customer-count regime contains exactly 50 independently generated instances. The fleet size is increased deterministically with problem scale, ranging from one vehicle in the smallest regime to ten vehicles in the largest. Customer coordinates

are sampled within two representative neighborhoods: Yeni Mahalle ($A = 0.588 \text{ km}^2$) and Merkez ($A = 1.377 \text{ km}^2$).

We regressed the observed average cost per customer, $\bar{C} = \mathcal{R}/N$, against $1/\sqrt{N}$:

$$\bar{C} = \frac{\beta}{\sqrt{N}} + \varepsilon,$$

obtaining an estimated scaling coefficient $\hat{\beta} = 1.810$ ($p < 0.001$) and adjusted $R^2 = 0.949$. The regression has exactly the functional form of the tariff schedule (1), which confirms the validity of the cost approximation for our numerical experiments.

4.5. Results

This section presents the results of the numerical experiments and evaluates them with respect to research questions **R1–R3**. Each research question is examined in a dedicated subsection.

4.5.1. Computational Performance Analysis

Table 3 summarizes the computational performance of the exact MIP $\mathcal{Z}^0$, the cluster heuristic ($\mathcal{Z}^c$ followed by $\mathcal{Z}^f$), and the sweep heuristic for the problem instances considered in this study. The sweep heuristic runs the Sweep Algorithm under all $L! = 24$ priority orderings and reports the best terminal equilibrium for the objective at hand; by Corollary 1 every one of its outputs is a Nash equilibrium, but by Propositions 3 and 4 the best of them need not be the operator-optimal equilibrium, so its gap measures how much the activation order alone can achieve. To provide a comprehensive evaluation, we solve all 50 instances (five per scenario) under four alternative IMO objectives: maximizing aggregate utility, minimizing total customer payments, maximizing total environmental value, and maximizing total quality value. For each approach, the table reports the runtime and the resulting optimality gaps. The optimality gaps are computed using the upper bounds (for the maximization problems) and lower bounds (for the minimization problem) obtained during the solution of the exact MIP model ($\mathcal{Z}^0$). With four providers, Theorem 2 does not apply, so the sweep solutions for the Min. Payment objective are not proven optimal by theory; nevertheless, the sweep heuristic achieves the same objective values and optimality gaps as the MIP. That the best of 24 orderings matches the MIP in all payment-minimization instances, even though Proposition 3 shows this cannot be guaranteed, indicates that in these calibrated instances demand is concentrated enough that sequential activation reaches the payment-minimizing equilibrium, in line with Proposition 5(ii); a complete characterization of the instances for which sequential activation is payment-optimal with more than two providers remains an interesting direction for future work.

The results show that problem difficulty varies substantially across objectives. In particular, the Max. Utility and Min. Payment objectives are considerably more challenging for the MIP solver. For these cases, the average runtime exceeds 5,000 seconds and optimality gaps remain nonzero on average (0.8% and 16.4%, respectively). In contrast, instances with Max. Environment and Max. Quality objectives are solved very quickly, with average runtimes of only 31.5 and 93.6 seconds and zero optimality gaps.

Table 3: Scenarios and Optimality Gaps

Objective	Scenario	MIP ($\mathcal{Z}^0$)		Cluster Heuristic ($\mathcal{Z}^c + \mathcal{Z}^f$)		Sweep	
		Runtime (sec)	Gap (%)	Runtime (sec)	Gap (%)	Runtime (sec)	Gap (%)
Max. Utility	0	9000.0	1.5	39.1	1.5	55.5	2.9
	1	9000.0	1.6	34.6	1.6	38.2	3.5
	2	3431.5	0.0	17.4	0.0	38.1	1.8
	3	2773.2	0.0	20.3	0.0	36.1	1.7
	4	9000.0	1.6	32.4	1.6	36.4	3.5
	5	8780.2	1.2	33.4	1.2	42.3	3.8
	6	9000.0	1.8	36.2	1.7	0.2	1.7
	7	724.7	0.0	4.0	0.0	79.1	0.0
	8	479.1	0.0	3.4	0.0	100.8	0.0
	9	7261.1	0.2	463.0	0.2	53.6	0.2
Overall		5945.0	0.8	68.4	0.8	48.0	1.9
Min. Payment	0	9001.0	42.2	6.7	42.2	55.5	42.2
	1	9001.2	32.0	3.3	32.0	38.2	32.0
	2	8773.7	23.1	2.6	23.1	38.1	23.1
	3	6140.2	10.8	3.9	10.8	36.1	10.8
	4	9000.8	35.0	5.4	35.0	36.4	35.0
	5	8152.3	20.8	4.7	20.8	42.3	20.8
	6	4.7	0.0	0.1	0.0	0.2	0.0
	7	1124.3	0.0	0.2	0.0	79.1	0.0
	8	5479.8	0.0	0.2	0.0	100.8	0.0
	9	6458.5	0.0	1.5	0.0	53.6	0.0
Overall		6313.6	16.4	2.8	16.4	48.0	16.4
Max. Environment	0	9.4	0.0	2.9	0.0	55.5	0.0
	1	10.6	0.0	2.6	0.0	38.2	0.0
	2	3.9	0.0	2.8	0.0	38.1	0.0
	3	7.8	0.0	2.8	0.0	36.1	0.0
	4	7.4	0.0	2.6	0.0	36.4	0.0
	5	8.7	0.0	2.8	0.0	42.3	0.0
	6	7.6	0.0	2.4	0.0	0.2	0.0
	7	5.9	0.0	3.7	0.0	79.1	0.0
	8	52.5	0.0	4.1	0.0	100.8	0.0
	9	201.5	0.0	10.7	0.0	53.6	0.0
Overall		31.5	0.0	3.7	0.0	48.0	0.0
Max. Quality	0	7.1	0.0	2.6	0.0	55.5	0.0
	1	4.8	0.0	2.7	0.0	38.2	0.0
	2	6.2	0.0	2.8	0.0	38.1	0.0
	3	7.3	0.0	3.0	0.0	36.1	0.0
	4	6.7	0.0	3.6	0.0	36.4	0.0
	5	6.6	0.0	3.0	0.0	42.3	0.0
	6	6.6	0.0	2.3	0.0	0.2	0.0
	7	53.7	0.0	4.0	0.0	79.1	42.6
	8	8.0	0.0	3.4	0.0	100.8	0.00
	9	829.3	0.0	13.3	0.0	53.6	5.8
Overall		93.6	0.0	4.1	0.0	48.0	4.8

Despite these differences in difficulty, both heuristics consistently produce high-quality solutions while requiring orders-of-magnitude less computation time than the MIP formulation. For example, under the Max. Utility objective, the cluster heuristic reduces the average runtime from 5945 seconds to 68.4 seconds while maintaining the same average op-

tinality gap of 0.8%. Similarly, under the Min. Payment objective, the cluster heuristic solves instances in only 2.8 seconds on average while achieving the same gap reported by the MIP. Across all objective types, the heuristics therefore provide solutions of comparable quality with drastically reduced computational effort.

Comparing the two heuristics, the reported sweep runtime is the same across objective types because all values refer to the same run, in which every provider ordering is evaluated. In contrast, although the cluster heuristic is faster on average, its runtime varies substantially across objectives and scenarios, primarily because the computational difficulty of obtaining $\mathcal{Z}^f$ differs across instances. For example, under the Max. Utility objective in Scenario 9, the cluster heuristic requires 463.0 seconds, whereas the sweep heuristic requires only 48.0 seconds. These results suggest that while both heuristics are effective, the sweep heuristic can offer a substantial computational advantage for the most challenging instances.

4.5.2. Utility Gains Under Centralized Decision Making

Our goal in this section is to assess the potential benefit of centralized decision-making by the IMO. To this end, we characterize the range of aggregate utility outcomes that may arise when the system settles at an equilibrium under alternative customer preference environments. For each experimental configuration, we compute both the maximum and the minimum aggregate utility attainable among equilibrium outcomes, by maximizing and minimizing the objective of $\mathcal{Z}^0$. Because both extrema are self-enforcing assignments, the difference between them is the widest range over which the outcome of a decentralized market can vary; it is an upper bound on the loss from leaving equilibrium selection to the market, not a prediction of the equilibrium a decentralized market would reach, which depends on where the market starts and on the order in which customers move (Theorem 1). We complement the range with the decentralized benchmark of Section 3.2, which simulates random improvement paths. The range quantifies the potential value of centralized coordination: the IMO can guarantee the top of the range, whereas the market alone can end anywhere in it.

Table 4 reports the equilibrium utility ranges obtained across the considered scenarios. The aggregate utility reported here excludes the baseline utility associated with package receipt, as this component is constant and does not depend on the matching decisions. For each configuration, the table reports the realized aggregate value obtained when $\mathcal{Z}^0$ is solved under the utility-maximizing and utility-minimizing objectives, together with the relative gap between the realized values of these two outcomes.

The results show that the dispersion across equilibria—the potential benefit of IMO-coordinated assignment—is substantial in all scenarios: the gap ranges from 7.6% (Scenario 9) to 18.0% (Scenario 3), with a median of 16.4% over the nine scenarios in which it is defined. Changes in the distribution of customer segments amplify the gap. Comparing Scenarios 1–5 shows that the gap increases as the proportion of customers who place a higher value on service quality or environmental performance rises; in Scenarios 2 and 3, where these segments constitute a larger share of the population, it exceeds 17%. From a practical perspective, this suggests that the value of centralized coordination is most pronounced in markets where customers care strongly about sustainability or high-quality service, because in such markets the assignments that a decentralized market can settle in differ most in how well they serve those preferences. A different pattern emerges for preference intensity.

Table 4: Scenario configurations and realized values

Scenario	Population (%)	ϵ -val.	q -val.	Value Realized (ϵ, q)		Gap
				Max Utility	Min Utility	%
0	34.08–16.76–16.48–32.68	M	M	668.6	583.0	12.8
1	40–20–20–20	M	M	759.5	700.3	7.8
2	20–40–20–20	M	M	913.7	749.7	17.9
3	20–20–40–20	M	M	913.5	749.3	18.0
4	20–20–20–40	M	M	698.7	616.5	11.8
5	25–25–25–25	M	M	820.9	693.1	15.6
6	34.08–16.76–16.48–32.68	L	L	291.4	0.0	-
7	34.08–16.76–16.48–32.68	H	M	1023.6	844.3	17.5
8	34.08–16.76–16.48–32.68	M	H	1018.6	843.5	17.2
9	34.08–16.76–16.48–32.68	H	H	1338.8	1236.7	7.6

When both valuation dimensions are low (Scenario 6), the market can settle in the equilibrium that realizes no environmental or quality value at all (realized value ranges from 0 to 291.4): this is the lock-in of Proposition 1, in which the price-based provider retains all customers because no customer’s valuation advantage exceeds the tariff gap between a provider serving one customer and one serving the whole market. At the opposite extreme, when both environmental and quality preferences are strong (Scenario 9), the relative gap is smaller despite the high level of aggregate utility: when valuations are uniformly strong, every equilibrium generates substantial surplus and lock-in is less severe. Large relative gaps arise in the intermediate cases in which only one preference dimension is strong (Scenarios 7 and 8), where the gap is over 17%. These results indicate that centralized coordination is most valuable when preferences are weak or asymmetric relative to tariffs, which is when economies of scale dominate preference heterogeneity and the equilibrium set is widest.

4.5.3. Tradeoffs Among IMO Objectives

The goal of this subsection is to examine the tradeoffs between alternative objectives that the IMO may pursue when determining matches. In particular, we study how prioritizing environmental performance versus service quality affects outcomes along the other dimension. To analyze these tradeoffs, we adopt an efficient frontier perspective that characterizes the set of Pareto efficient allocations with respect to these two objectives and visualizes the range of outcomes attainable across different equilibrium cases. Within this framework, we are also interested in understanding how the third objective, maximizing aggregate customer utility, performs relative to these two dimensions. To this end, we identify the equilibria that yield the maximum and minimum aggregate utility and mark their locations relative to the efficient frontier. Our aim with this is to see whether assignments that maximize aggregate utility also perform well in terms of environmental friendliness and service quality, and how much society would lose on these dimensions if it were to end up in a “bad” (for instance, utility-minimizing) equilibrium.

Figure 4 reports the efficient frontiers corresponding to Instance 1 of each scenario, the instance on which the cost parameters β_j were calibrated (Appendix B.4). We do not average frontier points across replications because averaging would blur the shape of the tradeoff curve and obscure the frontier structure; the remaining four instances of each scenario reveal no qualitatively different patterns. Each frontier represents the set of Pareto efficient outcomes in terms of realized environmental performance and service quality among self-




Figure 4: Efficient Frontiers

enforcing assignments. Circle markers denote the maximum utility Nash equilibria, cross markers denote the minimum utility equilibria, and star markers indicate the upper bounds for both quality and environmental performance. These bounds are the highest levels of each dimension attainable by a self-enforcing assignment.

Two critical insights emerge from the figure. First, in all scenarios the utility-maximizing Nash equilibrium lies on the efficient frontier. This indicates that maximizing aggregate utility does not inherently conflict with environmental performance or service quality. Rather, inefficiencies arise from equilibrium selection: a decentralized market may settle in equilibria that lie strictly inside the attainable frontier, leaving realizable environmental or quality gains unexploited.

Second, the minimum-utility equilibria often lie substantially below the efficient frontier (4 out of 10 scenarios), indicating that a decentralized market may end in assignments with large losses in realized environmental and quality performance. These losses reflect equilibria in which the market overemphasizes one dimension while severely underrealizing the other, highlighting the role that coordinated equilibrium selection by the IMO can play in steering the market toward balanced and efficient outcomes. Finally, we observe that, across all scenarios, the upper bounds (points marked by stars, i.e., the theoretical maxima obtained by summing 75% of all values for the respective dimension) are closely approached by the frontier equilibria, showing that self-enforcement does not prevent the market from reaching the highest attainable levels of environmental and quality performance.

The figure further illustrates how these patterns vary across the scenario groups. In the segment distribution group (Figure 4a), the quality and environment focused scenarios (Scenarios 2 and 3) behave differently from the rest. As expected, their frontiers extend much further toward their respective dominant dimensions. Interestingly, however, in these scenarios the utility-minimizing equilibrium also lies on the frontier, indicating that the pull toward the dominant dimension is sufficiently strong to prevent dominated equilibria from arising. The remaining scenarios exhibit generally similar frontier shapes despite differences in preference composition, yet their utility-minimizing equilibria may occur at markedly different locations relative to the frontier.

In the preference strength group (Figure 4b), valuation intensity becomes the key driver. When both environmental and quality preferences are weak (Scenario 6), the system collapses to a corner outcome where neither dimension is meaningfully realized. At the opposite extreme, when both preferences are strong (Scenario 9), even the worst case equilibrium remains above the efficiency envelope of several other scenarios, indicating that strong valuations can partially safeguard performance. Interestingly, when only one dimension receives strong emphasis (Scenario 7, Environment-high, and Scenario 8, Quality-high), utility maximizing equilibria move toward the corresponding frontier corner, while worst case equilibria shift toward the opposite corner. This pattern illustrates how asymmetric preferences may lead a decentralized market to emphasize the wrong dimension, producing inefficient outcomes despite the presence of strong valuations.

5. Conclusion

This paper studies a stylized model of a centralized last-mile delivery market in which an independent operator assigns customers to competing providers whose tariffs decrease with the number of customers they serve. The model is a resource-selection game with player-specific valuations and positive load externalities, and its central feature is the multiplicity of self-enforcing assignments that economies of scale create. We showed that such assignments always exist and are reached by any maximal sequence of profitable switches (Theorem 1); that a market can be locked into any provider when preference heterogeneity is small relative to the scale economies (Proposition 1); that with two providers every equilibrium is described by the load of one provider and the two activation orders reach the largest and the smallest equilibrium load, so that the payment-minimizing equilibrium is always one of two candidates and, under a simple checkable condition, the one obtained by activating the cheaper provider first (Theorem 2); and that with three or more providers the activation order is no longer a reliable instrument: the ascending-cost order can be strictly suboptimal, and the payment-minimizing equilibrium—indeed, some equilibria altogether—may be unreachable under every order (Propositions 3 and 4). What the order retains, with any number of providers, is delimited by Proposition 5: every lock-in equilibrium is implementable with a single run, and when scale economies dominate preference heterogeneity that run attains the minimum total payment over all assignments. Equilibrium selection in general therefore requires optimizing over the equilibrium set, which we do with a mixed-integer program whose constraints encode self-enforcement at post-switch tariffs, made scalable by a clustering-based reduction with an individual-level repair step. On the supply side, we showed that under contestable entry, when the operator uses the ascending-order Sweep Algorithm and the price cap does not bind, truthful cost reporting is a Nash equilibrium of the two-stage reporting game in which each incumbent faces a costless identical potential entrant (Proposition 6).

The computational study, calibrated to parcel demand in a district of Istanbul, quantifies what is at stake. The aggregate-utility gap between the best and the worst equilibrium ranges from 7% to 18% across scenarios and is largest when customers care strongly about a single service dimension or when valuations are weak relative to tariffs, which is when lock-in is most severe. The utility-maximizing equilibrium lies on the environmental-quality efficient frontier in every scenario: steering the market toward its best equilibrium does not force a

trade-off between aggregate utility and either environmental performance or service quality; the trade-offs arise only at the dominated equilibria a decentralized market may settle in.

We have deliberately kept the model abstract, and its abstractions delimit the conclusions. Routing, vehicle capacity, and time windows enter only through the load-dependent tariff, whose square-root form we validated on routing instances but which ignores the geography of individual customers; a spatially explicit version of the tariff would make a provider’s cost depend on *which* customers it serves, not only on how many, and would turn the game into one with non-anonymous externalities for which the potential structure of Theorem 1 no longer holds automatically. The assignment is static, whereas an actual market would clear repeatedly, with providers’ scale carrying over across periods; in such a dynamic setting the lock-in phenomenon of Proposition 1 suggests that early coordination matters most. The operator is assumed to know customers’ valuations; with private valuations the operator faces a Bayesian mechanism design problem in which the assignment must be incentive compatible for customers as well as self-enforcing, and it is an open question how much of the coordination gain survives the resulting information rents. Finally, the truthful-reporting result relies on contestability; in a market with a fixed set of providers, provider incentives would have to be addressed by an auction-like mechanism in which providers bid tariff schedules rather than report costs, drawing on the theory of multi-attribute procurement auctions. More broadly, this work is a first step toward the logistics markets envisioned by the Physical Internet, and extending it to multi-stage supply chains that span mid-mile and last-mile operations would bring the analysis closer to the operational reality that vision demands.

References

- Anshelevich, E., Dasgupta, A., Kleinberg, J., Tardos, É., Wexler, T., and Roughgarden, T. (2008). The price of stability for network design with fair cost allocation. *SIAM Journal on Computing*, 38(4):1602–1623.
- Ballot, E., Montreuil, B., and Zacharia, Z. G. (2021). Physical internet: First results and next challenges. *Journal of business logistics*, 42(1):101.
- Baumol, W. J., Panzar, J. C., and Willig, R. D. (1982). *Contestable Markets and the Theory of Industry Structure*. Harcourt Brace Jovanovich, New York.
- Biggar, D. R. and Hesamzadeh, M. R. (2014). *The economics of electricity markets*. John Wiley & Sons.
- Bimpikis, K., Candogan, O., and Saban, D. (2019). Spatial pricing in ride-sharing networks. *Operations Research*, 67(3):744–769.
- Castillo, J. C., Knoepfle, D., and Weyl, E. G. (2017). Surge pricing solves the wild goose chase. Technical report, University of Chicago. Working Paper.
- Cleophas, C., Cottrill, C., Ehmke, J. F., and Tierney, K. (2019). Collaborative urban transportation: Recent advances in theory and practice. *European Journal of Operational Research*, 273(3):801–816.

- Crainic, T. G. and Montreuil, B. (2016). Physical internet enabled hyperconnected city logistics. *Transportation Research Procedia*, 12:383–398.
- Cramton, P. (2017). Electricity market design. *Oxford Review of Economic Policy*, 33(4):589–612.
- Cuda, R., Guastaroba, G., and Speranza, M. G. (2015). A survey on two-echelon routing problems. *Computers & Operations Research*, 55:185–199.
- Daganzo, C. F. (1984). The distance traveled to visit n points with a maximum of c stops per vehicle: An analytic model and an application. *Transportation science*, 18(4):331–350.
- Engelhardt, M. (2023). Who is willing-to-pay for sustainable last mile innovations? *Transportation Research Procedia*, 69:910–917.
- Gansterer, M. and Hartl, R. F. (2018). Collaborative vehicle routing: A survey. *European Journal of Operational Research*, 268(1):1–12.
- Gkatzelis, V., Kollias, K., and Roughgarden, T. (2016). Optimal cost-sharing in general resource selection games. *Operations Research*, 64(6):1230–1238.
- Grossmann, I. E. (2002). Review of nonlinear mixed-integer and disjunctive programming techniques. *Optimization and engineering*, 3(3):227–252.
- Gurobi Optimization, LLC (2025). *Nonlinear Constraints — Gurobi Optimizer Reference Manual*. Gurobi Optimization, LLC. Version 12.0, last updated Oct 31, 2025.
- Ieong, S., McGrew, R., Nudelman, E., Shoham, Y., and Sun, Q. (2005). Fast and compact: A simple class of congestion games. In *Proceedings of the 20th National Conference on Artificial Intelligence (AAAI-05)*, pages 489–494. AAAI Press.
- Jr., A. S. K. and Crawford, V. P. (1982). Job matching, coalition formation, and gross substitutes. *Econometrica*, 50(6):1483–1504.
- Katz, M. L. and Shapiro, C. (1985). Network externalities, competition, and compatibility. *American Economic Review*, 75(3):424–440.
- Konishi, H., Le Breton, M., and Weber, S. (1997). Pure strategy Nash equilibrium in a group formation game with positive externalities. *Games and Economic Behavior*, 21(1–2):161–182.
- Koutsoupias, E. and Papadimitriou, C. (1999). Worst-case equilibria. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS 99)*, volume 1563 of *Lecture Notes in Computer Science*, pages 404–413. Springer.
- Masin, M. and Bukchin, Y. (2008). Diversity maximization approach for multiobjective optimization. *Operations Research*, 56(2):411–424.
- Milchtaich, I. (1996). Congestion games with player-specific payoff functions. *Games and Economic Behavior*, 13(1):111–124.

- Monderer, D. and Shapley, L. S. (1996). Potential games. *Games and Economic Behavior*, 14(1):124–143.
- Montreuil, B., Meller, R. D., and Ballot, E. (2013). Physical internet foundations. In *Service orientation in holonic and multi agent manufacturing and robotics*, pages 151–166. Springer.
- Rosenthal, R. W. (1973). A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2(1):65–67.
- Roughgarden, T. and Tardos, É. (2002). How bad is selfish routing? *Journal of the ACM*, 49(2):236–259.
- Savelsbergh, M. and Van Woensel, T. (2016). 50th anniversary invited article—City logistics: Challenges and opportunities. *Transportation Science*, 50(2):579–590.
- Shapley, L. S. and Shubik, M. (1971). The assignment game i: The core. *International Journal of Game Theory*, 1(1):111–130.
- Smith, E. M. and Pantelides, C. C. (1999). A symbolic reformulation/spatial branch-and-bound algorithm for the global optimisation of nonconvex minlps. *Computers & chemical engineering*, 23(4-5):457–478.
- Statista (2022). Global parcel shipping volume between 2013 and 2027 (in billion parcels)*. <https://www.statista.com/statistics/1139910/parcel-shipping-volume-worldwide/>. Accessed on May 01, 2024.
- Vöcking, B. (2007). Selfish load balancing. In Nisan, N., Roughgarden, T., Tardos, É., and Vazirani, V. V., editors, *Algorithmic Game Theory*, pages 517–542. Cambridge University Press.
- Wouda, N. A., Lan, L., and Kool, W. (2024). Pyvrp: A high-performance vrp solver package. *INFORMS Journal on Computing*, 36(4):943–955.
- Ziari, M., Ghomi-Avili, M., and Taleizadeh, A. A. (2025). Metropolitan city supply chain network design under consumer behavior and partial disruptions: A machine learning approach. *Computers & Operations Research*, page 107367.

Appendix A. Proofs of Results

Appendix A.1. Proofs from Section 3.1

Proof of Lemma 1. Let μ' be obtained from μ by moving customer i from $j = \mu(i)$ to $k \neq j$, and write $m_j = m_j(\mu)$, $m_k = m_k(\mu)$. Under μ' the loads are $m_j - 1$ and $m_k + 1$, all other loads being unchanged. The first sum in (8) changes by $v_{ik} - v_{ij}$. The second sum changes only in the terms of providers j and k :

$$\left(\sum_{\ell=1}^{m_j-1} p_j(\ell) + \sum_{\ell=1}^{m_k+1} p_k(\ell) \right) - \left(\sum_{\ell=1}^{m_j} p_j(\ell) + \sum_{\ell=1}^{m_k} p_k(\ell) \right) = p_k(m_k + 1) - p_j(m_j).$$

Hence $\Phi(\mu') - \Phi(\mu) = (v_{ik} - v_{ij}) - (p_k(m_k + 1) - p_j(m_j)) = (v_{ik} - p_k^+(\mu)) - (v_{ij} - p_j(m_j)) = u_i(\mu') - u_i(\mu)$, where we used $u_i(\mu') = v_{ik} - p_k(m_k + 1)$. If the switch is profitable, $u_i(\mu') > u_i(\mu)$ and Φ strictly increases. Along a sequence of profitable switches Φ is therefore strictly increasing, so no matching can appear twice. The argument uses only that $p_j(\cdot)$ is a function of the load. ■

Proof of Theorem 1. (i) Let μ maximize Φ over the finite set $\mathcal{M}$, and let μ' be obtained from μ by a unilateral switch of customer i . By Lemma 1, $u_i(\mu') - u_i(\mu) = \Phi(\mu') - \Phi(\mu) \leq 0$, so no unilateral switch is profitable and μ satisfies (7). (ii) Let $\mu^0, \mu^1, \dots$ be an improvement path. By Lemma 1, $\Phi(\mu^0) < \Phi(\mu^1) < \dots$, so the matchings along the path are pairwise distinct; since $\mathcal{M}$ is finite, the path is finite. If the path cannot be extended, its last matching admits no profitable unilateral switch, hence it is a Nash equilibrium; and since every improvement path is finite, every improvement path can be extended to one that cannot. Neither statement depends on the initial matching, on the rule selecting the next switch, or on the functional form of the tariff schedules. ■

Proof of Corollary 1. Every switch executed in the **while** loop of Algorithm 1 moves one customer to a provider that strictly increases her utility at the entry tariff; since a switch to an active provider is a unilateral deviation in Γ , the sequence of matchings generated by the SA is an improvement path of Γ and is finite by Theorem 1(ii). At termination all providers are active and the stopping condition of the loop states that no customer has a profitable switch to any provider; the terminal matching therefore satisfies (7). ■

Appendix A.2. Proofs from Section 3.2

Proof of Proposition 1. Under $\mu^{(j)}$ every customer is served by j at the tariff $p_j(N)$, and every other provider k has load zero, so the entry tariff of k is $p_k^+(\mu^{(j)}) = p_k(1)$. Condition (7) for customer i and provider $k \neq j$ reads $v_{ij} - p_j(N) \geq v_{ik} - p_k(1)$, which is (9). If $v_{ik} - v_{ij} \leq \min_{k'} p_{k'}(1) - \max_{j'} p_{j'}(N)$ for all i, j, k , then (9) holds for every j , since $\min_{k'} p_{k'}(1) - \max_{j'} p_{j'}(N) \leq p_k(1) - p_j(N)$. ■

Proof of Lemma 2. (a) Let $m = m_1(\mu)$. A customer i with $\mu(i) = 1$ pays $p_1(m)$ and would pay the entry tariff $p_2(N - m + 1)$ at provider 2, so (7) for her reads $v_{i1} - p_1(m) \geq v_{i2} - p_2(N - m + 1)$, i.e., $\delta_i \leq g(m)$. A customer with $\mu(i) = 2$ pays $p_2(N - m)$ and would pay $p_1(m + 1)$ at provider 1, so (7) reads $v_{i2} - p_2(N - m) \geq v_{i1} - p_1(m + 1)$, i.e., $\delta_i \geq p_2(N - m) - p_1(m + 1) = g(m + 1)$. Since $g(m) \leq g(m + 1)$, in an equilibrium every

customer of provider 1 has a net preference $\delta_i \leq g(m) \leq g(m+1) \leq \delta_{i'}$ for every customer i' of provider 2, so the customers of provider 1 hold the m smallest values of δ .

(b) If μ is an equilibrium with $m_1(\mu) = m$, then by (a) the customers of provider 1 hold the m smallest net preferences, so $\delta_{(m)} \leq \max_{i:\mu(i)=1} \delta_i \leq g(m)$ (if $m \geq 1$) and $\delta_{(m+1)} \geq \min_{i:\mu(i)=2} \delta_i \geq g(m+1)$ (if $m \leq N-1$); hence $m \in E$. Conversely, if $m \in E$, the matching that assigns the customers with the m smallest net preferences to provider 1 and the others to provider 2 satisfies the conditions in (a): $\delta_i \leq \delta_{(m)} \leq g(m)$ for the former and $\delta_i \geq \delta_{(m+1)} \geq g(m+1)$ for the latter.

(c) *Ordering* [1, 2]. All customers start at provider 1, so $m = N$. We claim that, throughout the run, when the load of provider 1 is m , (α) the customers of provider 1 hold the m smallest values of δ and (β) every customer of provider 2 satisfies $\delta_i > g(m+1)$. Both hold initially. Suppose they hold at load m . A customer i of provider 2 would gain $(v_{i1} - p_1(m+1)) - (v_{i2} - p_2(N-m)) = g(m+1) - \delta_i < 0$ from returning to provider 1, so no such switch is executed. A customer i of provider 1 gains $(v_{i2} - p_2(N-m+1)) - (v_{i1} - p_1(m)) = \delta_i - g(m)$ from switching to provider 2, which is largest for the largest δ among provider 1's customers, namely $\delta_{(m)}$ by (α). Hence a switch occurs if and only if $\delta_{(m)} > g(m)$, in which case a customer with $\delta_i = \delta_{(m)}$ moves (ties among equal values of δ are immaterial for loads), (α) is preserved at load $m-1$, and (β) is preserved as well: the mover satisfies $\delta_i = \delta_{(m)} > g(m) = g((m-1)+1)$, and the customers already at provider 2 satisfy $\delta_{i'} > g(m+1) \geq g(m)$ because g is nondecreasing. The run therefore stops at the largest $m \in \{0, \dots, N\}$ with $\delta_{(m)} \leq g(m)$, which we denote m^A ($m = 0$ always qualifies by the convention $\delta_{(0)} = -\infty$). Since every $m' > m^A$ was passed, $\delta_{(m^A+1)} > g(m^A+1)$ whenever $m^A < N$, so $m^A \in E$; and every $m \in E$ satisfies $\delta_{(m)} \leq g(m)$, hence $m \leq m^A$. Thus $m^A = \max E$.

Ordering [2, 1]. All customers start at provider 2, so $m = 0$. The invariant is now (α) as above and (β') every customer of provider 1 satisfies $\delta_i < g(m)$. A customer i of provider 1 would gain $\delta_i - g(m) < 0$ from returning to provider 2. A customer of provider 2 gains $g(m+1) - \delta_i$ from switching to provider 1, which is largest for the smallest δ among provider 2's customers, namely $\delta_{(m+1)}$; a switch occurs if and only if $\delta_{(m+1)} < g(m+1)$; the mover satisfies $\delta_i < g(m+1)$ and the customers already at provider 1 satisfy $\delta_{i'} < g(m) \leq g(m+1)$, so both invariants are preserved at load $m+1$. The run stops at the smallest m with $\delta_{(m+1)} \geq g(m+1)$, denoted m^B ($m = N$ qualifies by the convention $\delta_{(N+1)} = +\infty$). Since every $m' < m^B$ was passed, $\delta_{(m^B)} < g(m^B)$ whenever $m^B \geq 1$, so $m^B \in E$; every $m \in E$ satisfies $\delta_{(m+1)} \geq g(m+1)$, hence $m \geq m^B$. Thus $m^B = \min E$. ■

Proof of Proposition 2. Relabel j_1 as provider 1 and j_2 as provider 2. By Lemma 2(b) the load of provider 1 in any Nash equilibrium belongs to E , and by Lemma 2(c) the two orderings yield $\max E$ and $\min E$. ■

Appendix A.3. Auxiliary Results and Proofs from Section 3.3

For $\beta > 0$ let $f(x) = \sqrt{x}$ and $\psi_\beta(x) = \min\{\beta x, \beta f(x)\}$ for $x \geq 0$, so that the total payment collected by a provider with cost parameter β serving m customers is $\varphi(m) = m p(m) = \psi_\beta(m)$ for every $m \geq 0$ (both sides vanish at $m = 0$, whatever the convention for $p(0)$). For a function h on the nonnegative integers write $\Delta h(x) = h(x+1) - h(x)$.

- Lemma 3** (Properties of f and ψ_β). (i) $f(0) = 0$, f is strictly increasing and strictly concave on $[0, \infty)$: $f'(x) = \frac{1}{2\sqrt{x}} > 0$ and $f''(x) = -\frac{1}{4x^{3/2}} < 0$ for $x > 0$.
- (ii) For every $\beta > 0$, ψ_β is nondecreasing and concave on $[0, \infty)$ with $\psi_\beta(0) = 0$; in particular its increments $\Delta\psi_\beta(x)$ are nonnegative and nonincreasing in x , and $\Delta\psi_\beta(x) \leq \psi_\beta(x+1)/(x+1) \leq \bar{\beta}$.
- (iii) If $0 < \beta \leq \beta'$, then $\Delta\psi_\beta(x) \leq \Delta\psi_{\beta'}(x)$ for every integer $x \geq 0$.
- (iv) For any load vector, $R(\mathbf{m}) = \sum_j \psi_{\beta_j}(m_j)$ is concave on the simplex $\{\mathbf{m} \in \mathbb{R}_{\geq 0}^L : \sum_j m_j = N\}$ and attains its minimum at the vertex $N\mathbf{e}_{j^*}$, where $j^* \in \arg \min_j \beta_j$: absent equilibrium constraints, total payment is minimized by assigning every customer to the provider with the smallest cost parameter.

Proof. (i) Direct computation. (ii) ψ_β is the pointwise minimum of two nondecreasing concave functions vanishing at zero, hence itself nondecreasing, concave, and zero at zero (the hypograph of a pointwise minimum of two functions is the intersection of their hypographs, and the intersection of two convex sets is convex). Concavity gives nonincreasing increments, so $\psi_\beta(x+1) = \sum_{t=0}^x \Delta\psi_\beta(t) \geq (x+1)\Delta\psi_\beta(x)$; and $\psi_\beta(x+1) \leq \bar{\beta}(x+1)$. (iii) Let $\hat{x}(\beta) := (\beta/\bar{\beta})^2$, so that $\psi_\beta(x) = \bar{\beta}x$ for $0 \leq x \leq \hat{x}(\beta)$ and $\psi_\beta(x) = \beta\sqrt{x}$ for $x \geq \hat{x}(\beta)$, and note $\hat{x}(\beta) \leq \hat{x}(\beta')$. On $[0, \infty)$, ψ_β is $\bar{\beta}$ -Lipschitz—its derivative, defined for all $t > 0$ except the kink $\hat{x}(\beta)$, is

$$\psi'_\beta(t) = \begin{cases} \bar{\beta}, & 0 < t < \hat{x}(\beta), \\ \frac{\beta}{2\sqrt{t}}, & t > \hat{x}(\beta), \end{cases}$$

and is bounded by $\bar{\beta}$ —hence absolutely continuous, so $\Delta\psi_\beta(x) = \int_x^{x+1} \psi'_\beta(t) dt$ for every $x \geq 0$. It therefore suffices to show that $\psi'_\beta(t) \leq \psi'_{\beta'}(t)$ at every t where both derivatives are defined. For $t < \hat{x}(\beta)$ both equal $\bar{\beta}$; for $\hat{x}(\beta) < t < \hat{x}(\beta')$ we have $\sqrt{t} > \beta/\bar{\beta}$, so $\psi'_\beta(t) = \beta/(2\sqrt{t}) < \bar{\beta}/2 < \bar{\beta} = \psi'_{\beta'}(t)$; and for $t > \hat{x}(\beta')$, $\psi'_\beta(t) = \beta/(2\sqrt{t}) \leq \beta'/(2\sqrt{t}) = \psi'_{\beta'}(t)$. Integrating over $[x, x+1]$ gives $\Delta\psi_\beta(x) \leq \Delta\psi_{\beta'}(x)$. (iv) R is a sum of concave functions of separate coordinates, hence concave; a concave function on a polytope attains its minimum at some vertex; the vertices are $N\mathbf{e}_k$ with $R(N\mathbf{e}_k) = \psi_{\beta_k}(N)$, which is nondecreasing in β_k . $\square$

Proof of Theorem 2. By Lemma 2, total payment in an equilibrium μ depends on μ only through $m = m_1(\mu)$:

$$R(\mu) = \rho(m) := \psi_{\beta_1}(m) + \psi_{\beta_2}(N - m), \quad m \in \{0, 1, \dots, N\},$$

and the equilibrium loads satisfy $E \subseteq \{m^B, \dots, m^A\}$ with $m^B = \min E = m_1(\mu^{[2,1]})$ and $m^A = \max E = m_1(\mu^{[1,2]})$.

(a) The increments of ρ are $\Delta\rho(m) = \Delta\psi_{\beta_1}(m) - \Delta\psi_{\beta_2}(N - m - 1)$. As m increases, the first term is nonincreasing and $\Delta\psi_{\beta_2}(N - m - 1)$ is nondecreasing (Lemma 3(ii)), so $\Delta\rho$ is nonincreasing: ρ is concave on $\{0, \dots, N\}$. A function with nonincreasing increments attains its minimum over an integer interval at an endpoint: if $\rho(m) < \rho(m^B)$ for some $m^B < m < m^A$, some increment $\Delta\rho(t)$ with $t < m$ is negative, hence all later increments are negative and $\rho(m^A) < \rho(m)$. Therefore $\rho(m) \geq \min\{\rho(m^B), \rho(m^A)\}$ for every $m \in E$, and since $m^A, m^B \in E$, the minimum of R over $\bar{\mathcal{M}}$ equals $\min\{R(\mu^{[1,2]}), R(\mu^{[2,1]})\}$.

(b) Let $\beta_1 \leq \beta_2$ and let m be an integer with $N/2 \leq m \leq N-1$. Then $2m \geq N > N-1$, so $m \geq N-m-1$, and

$$\Delta\rho(m) = \Delta\psi_{\beta_1}(m) - \Delta\psi_{\beta_2}(N-m-1) \leq \Delta\psi_{\beta_1}(N-m-1) - \Delta\psi_{\beta_2}(N-m-1) \leq 0,$$

where the first inequality uses that the increments of ψ_{β_1} are nonincreasing (Lemma 3(ii)) and the second uses Lemma 3(iii). Hence ρ is nonincreasing on $\{\lceil N/2 \rceil, \dots, N\}$. If $m^B \geq N/2$, then $E \subseteq \{m^B, \dots, m^A\} \subseteq \{\lceil N/2 \rceil, \dots, N\}$, so $\rho(m^A) \leq \rho(m)$ for every $m \in E$: the ascending order $[1, 2]$ minimizes total payment among all equilibria. ■

Example 1 (Two providers: the ascending order can fail without the condition of Theorem 2(b)). Let $\beta_1 = 1$, $\beta_2 = 1.2$, $\bar{\beta} = 2$ (the cap never binds for a provider with at least one customer), and $N = 6$: three customers (“type a”) have valuations $(v_{i1}, v_{i2}) = (2.2, 2.0)$ and three customers (“type b”) have $(2.0, 3.0)$. The tariff schedules are $p_1(m) = 1/\sqrt{m}$ and $p_2(m) = 1.2/\sqrt{m}$ for $m \geq 1$. Under the ascending order $[1, 2]$, the three type-b customers leave provider 1 one by one (the first gains $(3.0 - 1.2) - (2.0 - 1/\sqrt{6}) = 0.208$, and each subsequent switch is more profitable), after which no type-a customer wants to follow: with loads $(3, 3)$ a type-a customer obtains $2.2 - 1/\sqrt{3} = 1.623$ at provider 1 and would obtain $2.0 - 1.2/\sqrt{4} = 1.400$ at provider 2, while a type-b customer obtains $3.0 - 1.2/\sqrt{3} = 2.307$ and would obtain $2.0 - 1/\sqrt{4} = 1.500$. The terminal matching has loads $(3, 3)$ and total payment $R = 3 \cdot 0.577 + 3 \cdot 0.693 = 3.811$. Under the order $[2, 1]$ nobody moves: with loads $(0, 6)$ a type-a customer obtains $2.0 - 1.2/\sqrt{6} = 1.510$ and would obtain $2.2 - 1.0 = 1.200$ by moving alone to provider 1 (a type-b customer would lose even more), so the terminal matching has loads $(0, 6)$ and $R = 6 \cdot 0.490 = 2.939$. Both matchings are equilibria and they are the only ones ($E = \{0, 3\}$). Here $m^B = 0 < N/2$, the condition of Theorem 2(b) fails, and the ascending order pays 30% more than the descending one: concentrating all demand on the costlier provider saves more through scale economies than provider 1’s cost advantage is worth. (This equilibrium also has the higher aggregate utility, 12.06 versus 11.79.)

Proof of Proposition 3. We give one instance for each statement; in both, $\bar{\beta} = 2$, Assumption 1 holds, and the cap never binds for a provider with at least one customer because $\beta_j \leq 2$ for all j , so $p_j(m) = \beta_j/\sqrt{m}$ for $m \geq 1$. All utilities are reported to three decimals and all comparisons hold with margins larger than the rounding error. The set of all Nash equilibria of each instance was also verified by exhaustive enumeration; the enumeration and verification code is provided as supplementary material.

(i) *The ascending order need not minimize payments.* Let $L = 3$, $\beta = (1.3, 1.8, 1.9)$ and $N = 3$, with valuations

$$(v_{ij}) = \begin{pmatrix} 2.7 & 2.7 & 2.0 \\ 3.9 & 4.0 & 3.2 \\ 2.3 & 4.0 & 3.5 \end{pmatrix}.$$

The relevant tariffs are $p_1(2) = 0.919$, $p_1(3) = 0.751$, $p_2(1) = 1.8$, $p_2(3) = 1.039$, and $p_3(1) = 1.9$. The ascending order $[1, 2, 3]$ proceeds as follows. From loads $(3, 0, 0)$, provider 2 is activated; customer 3 (utility $2.3 - 0.751 = 1.549$ at provider 1) is the only one with a profitable switch and moves to provider 2 (utility $4.0 - 1.8 = 2.200$); at loads $(2, 1, 0)$ no further switch is profitable, and none becomes profitable when provider 3 is activated. The

terminal matching $\mu^{[1,2,3]}$ has loads $(2, 1, 0)$ and total payment $R = 2 \cdot 0.919 + 1.8 = 3.638$; its equilibrium property is verified in Table [Appendix A1](#). The matching $\mu^{(2)}$ that assigns all three customers to provider 2 is also an equilibrium (Table [Appendix A1](#)) with total payment $R = 3 \cdot 1.039 = 3.118$, so $\mu^{[1,2,3]}$ pays 17% more, and these two are the only equilibria of the instance. The load vector $(2, 1, 0)$ of the ascending order is lexicographically larger than $(0, 3, 0)$. The orders $[2, 1, 3]$, $[2, 3, 1]$ and $[3, 2, 1]$ reach $\mu^{(2)}$; the orders $[1, 3, 2]$ and $[3, 1, 2]$ reach $\mu^{[1,2,3]}$.

For $L > 3$, append $L - 3$ providers with pairwise distinct cost parameters $\beta_4 < \dots < \beta_L$ in the interval $(1.9, 2)$ and with $v_{ij} = \bar{\beta} = 2$ for every customer i and every appended provider j . Assumption 1 continues to hold, the cost parameters remain strictly ascending, and the cap still never binds for a provider with at least one customer ($\beta_j < 2$ for all j). Switching to an appended provider j would yield utility $\bar{\beta} - p_j(1) = 2 - \beta_j < 0.1$, while every customer's utility at every point of the ascending sweep and at the two equilibria above is at least 1.549; hence no customer ever joins an appended provider, the ascending order activates the appended providers last and executes no further switches, and both matchings above remain equilibria with the appended providers empty. Loads and payments are unchanged, so the displayed comparison establishes part (i) for every $L \geq 3$.

Table Appendix A1: Example (i): utilities at the two equilibria (“dev. k ” is the utility from switching to provider k at its entry tariff).

Customer	Loads $(2, 1, 0)$: tariffs $(0.919, 1.800, -)$			Loads $(0, 3, 0)$: tariffs $(-, 1.039, -)$		
	provider, utility	dev. to	utility	provider, utility	dev. to	utility
1	1, 1.781	2; 3	1.427; 0.100	2, 1.661	1; 3	1.400; 0.100
2	1, 2.981	2; 3	2.727; 1.300	2, 2.961	1; 3	2.600; 1.300
3	2, 2.200	1; 3	1.549; 1.600	2, 2.961	1; 3	1.000; 1.600
R	3.638			3.118		

(ii) *No ordering need reach a payment-minimizing equilibrium.* Let $L = 3$, $\beta = (1.2, 1.6, 1.9)$ and $N = 6$, with three customers of type a (valuations $(3.4, 3.9, 4.8)$), two of type b (valuations $(2.4, 2.0, 3.4)$) and one of type c (valuations $(2.8, 4.8, 2.4)$). Exhaustive enumeration shows that the instance has exactly two equilibrium load vectors: $(2, 4, 0)$, with type b at provider 1 and types a, c at provider 2 (Table [Appendix A2](#)), with $R = 2 \cdot 0.849 + 4 \cdot 0.800 = 4.897$; and $(0, 1, 5)$, with type c at provider 2 and everybody else at provider 3, with $R = 1.6 + 5 \cdot 0.850 = 5.849$. The payment-minimizing equilibrium is $(2, 4, 0)$, and the Sweep Algorithm never reaches it: all six orderings terminate at $(0, 1, 5)$. Under $[1, 2, 3]$, for instance, only the type- c customer moves when provider 2 is activated ($2.310 \rightarrow 3.200$); when provider 3 is activated, the three type- a customers move there one after the other ($2.863 \rightarrow 2.900$, $2.800 \rightarrow 3.456$, $2.707 \rightarrow 3.703$), which makes provider 3 cheap enough for the two type- b customers to follow ($1.551 \rightarrow 2.450$, $1.200 \rightarrow 2.550$). Under the orders starting with provider 3, only the type- c customer leaves provider 3 (for provider 2), and under those starting with provider 2 all customers except type c migrate to provider 3. The equilibrium in which types a and c share provider 2 while type b uses provider 1 requires provider 2 to be populated *before* provider 3 becomes attractive, which no sequential activation achieves.

Table Appendix A2: Example (ii): utilities at the payment-minimizing equilibrium, loads (2, 4, 0), tariffs (0.849, 0.800, −), entry tariffs (0.693, 0.716, 1.900).

Customer (type)	provider	utility	dev. to	utility
1, 2, 3 (a)	2	3.100	1; 3	2.707; 2.900
4, 5 (b)	1	1.551	2; 3	1.284; 1.500
6 (c)	2	4.000	1; 3	2.107; 0.500

■

Proof of Proposition 4. We prove the second statement first; combined with an instance in which the lock-in condition holds for every provider and a further equilibrium exists—which we then exhibit for every $L \geq 3$ —it yields the first. Suppose (9) holds for every provider j . Under any ordering o , the Sweep Algorithm starts from $\mu^{(o(1))}$, which is a Nash equilibrium of Γ by Proposition 1; hence no customer has a profitable switch to *any* provider, in particular to any active provider, and no switch is ever executed: the terminal matching is $\mu^{(o(1))}$. Thus only the L lock-in equilibria are reachable.

For the first statement, fix $L \geq 3$, let $N = L$, $\beta_j = \beta$ for all j , $\bar{\beta} \geq \beta$ (so that $p_j(m) = \beta/\sqrt{m}$ for every $m \geq 1$), and let customer i value provider i at $v_{ii} = a$ and every other provider at $v_{ij} = a - D$, with $a - D \geq \bar{\beta}$ and

$$\beta\left(1 - \frac{1}{\sqrt{2}}\right) \leq D \leq \beta\left(1 - \frac{1}{\sqrt{L}}\right),$$

an interval of positive length because $L \geq 3$ (e.g., $\beta = 1$, $a = 3$, $D = 0.35$, $\bar{\beta} = 2$ for $L = 3$). Since $v_{ik} - v_{ij} \leq D \leq p(1) - p(N) = \beta(1 - 1/\sqrt{N})$ for all i, j, k , condition (9) holds for every j and all L lock-in matchings are equilibria. The matching $\mu^*(i) = i$ that gives every provider exactly one customer is also an equilibrium: customer i obtains $a - \beta$ and would obtain $a - D - \beta/\sqrt{2}$ by joining any other provider, which is not better because $D \geq \beta(1 - 1/\sqrt{2})$. By the first part, the Sweep Algorithm terminates at $\mu^{(o(1))} \neq \mu^*$ under every ordering o . ■

Proof of Proposition 5. (i) Suppose $\mu^{(j)}$ is a Nash equilibrium and $o(1) = j$. At every point of the run, a switch available to the algorithm moves one customer to an active provider at that provider's entry tariff, and is therefore a unilateral deviation in Γ ; at $\mu^{(j)}$ no unilateral deviation is strictly profitable, so no switch is ever executed—whatever the ordering of the remaining providers—and the terminal matching is $\mu^{(j)}$ (this is the argument in the proof of Proposition 4, which requires no assumption on the schedules). Conversely, suppose $\mu^{(j)}$ is not a Nash equilibrium, so some customer i has a strictly profitable deviation to some provider k : $v_{ik} - p_k(1) > u_i(\mu^{(j)})$. Consider the stage of the run at which k is activated. If the matching at that moment is still $\mu^{(j)}$, the switch of i to k is available and profitable, so the **while** loop executes a switch; if the matching has already changed, a switch was executed earlier. In either case at least one switch occurs, and since Φ strictly increases along the run (Lemma 1), no later matching equals $\mu^{(j)}$: the terminal matching differs from $\mu^{(j)}$. For the last statement, the direct implication shows that for every lock-in equilibrium $\mu^{(j)}$ the run with $o(1) = j$ terminates at $\mu^{(j)}$; hence the best of the L runs pays at most $\min\{R(\mu^{(j)}) : \mu^{(j)} \in \bar{\mathcal{M}}\}$.

(ii) Since each p_j is nonincreasing, $\varphi_j(m) = m p_j(m) \geq m p_j(N) = (m/N) \varphi_j(N)$ for every $0 \leq m \leq N$. Summing over providers,

$$R(\mu) = \sum_{j \in \mathcal{L}} \varphi_j(m_j) \geq \sum_{j \in \mathcal{L}} \frac{m_j}{N} \varphi_j(N) \geq \min_{k \in \mathcal{L}} \varphi_k(N) \cdot \frac{\sum_j m_j}{N} = \min_{k \in \mathcal{L}} \varphi_k(N).$$

If (9) holds for some $j^* \in \arg \min_j \varphi_j(N)$, then $\mu^{(j^*)}$ is a Nash equilibrium by Proposition 1, the run with $o(1) = j^*$ terminates at it by part (i), and its payment $R(\mu^{(j^*)}) = \varphi_{j^*}(N) = \min_j \varphi_j(N)$ attains the lower bound just derived, hence minimizes R over all matchings.

(iii) Every switch executed by the algorithm strictly increases Φ (Lemma 1), so $\Phi(\mu^T) \geq \Phi(\mu^{(o(1))})$. For every matching μ ,

$$U(\mu) - \Phi(\mu) = \sum_{j \in \mathcal{L}} \left[\sum_{\ell=1}^{m_j} p_j(\ell) - m_j p_j(m_j) \right] \geq 0,$$

because each p_j is nonincreasing, so $U(\mu^T) \geq \Phi(\mu^T) \geq \Phi(\mu^{(o(1))})$, and $\Phi(\mu^{(j)}) = \sum_i v_{ij} - \sum_{\ell=1}^N p_j(\ell)$ by (8). Finally, $p_j(\ell) \leq \bar{\beta}$ gives $\sum_{\ell=1}^N p_j(\ell) \leq \bar{\beta}N$, while $p_j(\ell) \leq \beta_j/\sqrt{\ell}$ and $1/\sqrt{\ell} \leq 2(\sqrt{\ell} - \sqrt{\ell-1})$ give, by telescoping, $\sum_{\ell=1}^N p_j(\ell) \leq 2\beta_j\sqrt{N}$. ■

Proof of Proposition 6. Throughout, $\hat{p}_j(m) = \min\{\bar{\beta}, \hat{\beta}_j/\sqrt{m}\}$ for $m \geq 1$ (and $\hat{p}_j(0) = \bar{\beta}$) denotes the tariff schedule induced by the report $\hat{\beta}_j$, and m_j the terminal load of provider j .

(i) If $\hat{\beta}_j < \beta_j$ and $m_j \geq 1$, then $\hat{p}_j(m_j) \leq \hat{\beta}_j/\sqrt{m_j} < \beta_j/\sqrt{m_j}$, so $\pi_j = m_j(\hat{p}_j(m_j) - \beta_j/\sqrt{m_j}) < 0$.

(ii) Let $\hat{\beta}_j > \beta_j$ and let the entrant $e = e(j)$ report $\hat{\beta}_e = \beta_j < \hat{\beta}_j$, so that e precedes j in the ascending order of reports and is activated earlier. We show by induction over the switches of the Sweep Algorithm that j never receives a customer. Initially all customers are at $o(1) \neq j$ (since e precedes j , j is not first). Suppose $m_j = 0$ at some step and consider any customer i and any active provider. If i is at e , her utility is $v_{ie} - \hat{p}_e(m_e)$ with $m_e \geq 1$, and $\hat{p}_e(m_e) = \min\{\bar{\beta}, \beta_j/\sqrt{m_e}\} \leq \min\{\bar{\beta}, \hat{\beta}_j\} = \hat{p}_j(1)$, the entry tariff of j ; since $v_{ie} = v_{ij}$, switching to j is not profitable. If i is at some other provider, the gain from switching to j is $v_{ij} - \hat{p}_j(1) - u_i$, while the gain from switching to e (which is active, being activated before j) is $v_{ie} - \hat{p}_e(m_e + 1) - u_i$; as $\hat{p}_e(m_e + 1) = \min\{\bar{\beta}, \beta_j/\sqrt{m_e + 1}\} \leq \min\{\bar{\beta}, \hat{\beta}_j\} = \hat{p}_j(1)$, the switch to e yields a weakly larger gain, and in case of a tie the algorithm prefers e , which appears earlier in the ordering. Hence the switch executed at this step, if any, is not a switch to j , and $m_j = 0$ is preserved. At termination $m_j = 0$ and $\pi_j = 0$.

(iii) If $\hat{\beta}_j = \beta_j$ and $m_j \geq 1$, then $\hat{p}_j(m_j) = \beta_j/\sqrt{m_j}$ because the cap does not bind for a provider serving at least one customer ($\beta_j/\sqrt{m_j} \leq \beta_j \leq \bar{\beta}$), so $\pi_j = 0$; if $m_j = 0$, $\pi_j = 0$ as well.

(iv) Suppose every incumbent reports truthfully, and consider a unilateral deviation of the entrant $e = e(j)$, whose rule prescribes staying out with payoff $(0, 0)$ (zero profit, zero customers). If e enters with a report $\hat{\beta}_e \geq \beta_j$, then the truthful incumbent j reports the same or a lower parameter and precedes e in the activation order (incumbent-first at equality); since e has j 's attributes and $\hat{p}_j(m) \leq \min\{\bar{\beta}, \hat{\beta}_e\} = \hat{p}_e(1)$ for every $m \geq 1$, the induction of part (ii) applies with the roles of the overstater and the entrant taken by e and j , respectively:

e receives no customer, and its payoff is $(0, 0)$. If e enters with $\hat{\beta}_e < \beta_j$, then by the computation in part (i) applied to e (whose true cost parameter is β_j), its profit is strictly negative if it serves at least one customer and zero otherwise; in the lexicographic order neither outcome improves on $(0, 0)$.

Combining (i)–(iv): given the entry rule, understating gives an incumbent at most zero profit and strictly negative profit whenever it attracts any customer; overstating gives zero profit and no customers; truthful reporting gives zero profit and whatever customers the market assigns to j ; since providers prefer, at equal profit, to serve more customers, truthful reporting is a best response for every incumbent. By (iv), no entrant has a lexicographically improving deviation when incumbents report truthfully. Hence no player has a strictly improving unilateral deviation, and the profile of truthful incumbent reports and the entry rule σ is a Nash equilibrium of the two-stage reporting game. ■

Appendix B. Numerical Experiments Parameter Specifications

Appendix B.1. Efficient Frontier

An additional perspective concerns the interaction between realized performance dimensions—such as environmental sustainability and service quality—which may be conflicting. Improvements in one dimension may come at the expense of the other, motivating an explicit characterization of the resulting trade-offs. The set of desirable solutions can be described in terms of Pareto optimality, namely solutions for which no performance dimension can be improved without degrading another. The image of this set in the objective space is referred to as the efficient (or nondominated) frontier, from which a preferred operating point is selected based on additional considerations (Masin and Bukchin, 2008).

Let q_{ij} denote the realized quality obtained by customer $i \in \mathcal{C}$ when served by provider $j \in \mathcal{L}$, and let e_{ij} denote the corresponding realized environmental contribution. As a benchmark, we define the maximum attainable environmental performance (E_{max}) by solving the baseline model $\mathcal{Z}^0$ with the objective

$$\max \sum_{i \in \mathcal{C}} \sum_{j \in \mathcal{L}} e_{ij} x_{ij}. \quad (\text{Appendix B1})$$

We construct the efficient frontier by solving a parametric family of optimization problems indexed by $\alpha \in [0, 1]$, where each problem maximizes realized service quality subject to a minimum environmental performance constraint. The parameter α is evaluated over the interval $[0, 1]$ in increments of 0.05. The number of distinct equilibrium allocations is smaller than the number of parameter values, as several α levels yield identical assignments. The plotted frontier connects the evaluated points to facilitate comparison across scenarios.

The resulting model is given by

$$\max \sum_{i \in \mathcal{C}} \sum_{j \in \mathcal{L}} q_{ij} x_{ij}, \quad (\text{Appendix B2})$$

$$\text{s.t.} \quad \sum_{i \in \mathcal{C}} \sum_{j \in \mathcal{L}} e_{ij} x_{ij} \geq \alpha E_{max}, \quad (\text{Appendix B3})$$

$$(14)-(20). \quad (\text{Appendix B4})$$

The objective in (Appendix B2) maximizes the aggregate realized quality across all customer–provider matches. Constraint (Appendix B3) imposes the environmental requirement by restricting feasible assignments to those that achieve at least a fraction α of the benchmark environmental level E_{max} . Varying $\alpha \in [0, 1]$ traces the efficient frontier between environmental performance and realized quality. Constraints (Appendix B4) are the assignment, tariff, and self-enforcement constraints of $\mathcal{Z}^0$, so every point of the frontier is a Nash equilibrium.

Table Appendix B1: E-Commerce and Distribution Metrics for Istanbul and Sarıyer

Metric	Value	Source
E-Commerce Volume – Number of Transactions (Turkey, pg. 19)	5,870,000,000	ticaret.gov.tr
E-Commerce Volume – Goods Share (pg. 38)	49.0%	ticaret.gov.tr
Cargo Share – Istanbul	29.7 %	ticaret.gov.tr
Number of Transactions in Istanbul	855,411,620	Computation
Active Delivery Days (2023)	300.5	is-gunu.hesaplama.net
Daily Orders (Istanbul)	2,846,627.7	Computation
Population of Istanbul	15,655,924	data.tuik.gov.tr
Parcels per Capita (Istanbul)	0.18	Computation
Population: Sarıyer – Yeni Mahalle	3,284	nufusune.com
Population: Sarıyer – Merkez	11,734	nufusune.com
Estimated Parcels: Yeni Mahalle	597.1	Computation
Estimated Parcels: Merkez	2,133.5	Computation
Sum of Two Neighborhoods	2,730.6	Computation
Rounded Sum (Two Neighborhoods)	2,731	Rounded
Interval (Two Neighborhoods)	[2595-2866]	Assumption
Demand Points: Merkez	500	Assumption
Demand Points: Yeni Mahalle	139.9	Computation
Rounded Demand Points: Yeni Mahalle	140	Rounded

Appendix B.2. Study Data

Appendix B.3. Customer Segments and Scenarios

Customer segmentation is derived from a nationally representative survey from Germany, conducted in 2020 ($n = 1,071$), analyzed using K-means clustering to identify distinct behavioral archetypes (Engelhardt, 2023). Table Appendix B2 presents the resulting customer groups by income level, monthly order frequency, and willingness to pay (WTP) for both standard and instant delivery options. Each customer segment S_j in Table Appendix B3 is mapped to its same-indexed provider P_j (e.g., $S_1 \leftrightarrow P_1$), reflecting correspondence between segment characteristics and provider positioning.

Table Appendix B2: Summary of Cluster Analysis Results

Segment	Survey Count	Income (EUR)	WTP per Delivery (EUR)	Instant Delivery WTP
Rich Heavy Users	59	3058	1.76	High (3.73)
Conservative Shoppers	117	2562	1.65	Low (1.79)
Willing-to-Pay Professionals	60	2039	4.67	Moderate (2.55)
Young Demanders	122	1929	1.77	High (3.76)

In generating the synthetic instances, we use the segment participation rates in Table Appendix B3 as the population shares. Accordingly, in each scenario we draw customers so that segment S_j represents its stated percentage of the population.

Table Appendix B3: Customer Segments and Survey Distribution

Customer Segment	Analogy	Population %
Rich Heavy Users	S3	16.48%
Conservative Shoppers	S4	32.68%
Willing-to-Pay Professionals	S2	16.76%
Young Demanders	S1	34.08%

These segment archetypes underscore the heterogeneity in delivery preferences, cost sensitivity, and environmental consciousness. *Rich Heavy Users* and *Young Demanders*—despite

divergent income levels—exhibit strong preferences for rapid fulfillment. *Willing-to-Pay Professionals*, while less frequent buyers, exhibit the highest WTP, driven by sustainability motives. *Conservative Shoppers* dominate numerically but offer minimal pricing leverage due to their low service expectations. Tables [Appendix B2](#) and [Appendix B3](#) summarize the empirical segmentation that guides the construction of customer segments in the (e, q) -space. Although a direct numerical mapping between survey-based attributes and model parameters is not possible, the correspondence reflects the qualitative structure observed in the data.

Segment-specific valuation draws follow Table [Appendix B4](#). For each segment S_j , environmental and service-quality valuations are generated from triangular distributions and then scaled by the segment multipliers reported in the E and Q columns. The receipt component R is held fixed at $(0.5, 0.75, 1)$ across all scenarios, since receipt-related heterogeneity is outside the experimental scope. As an illustration, the environmentally oriented segment S_2 uses a higher weight on e (multiplier 0.75) and a lower weight on q (multiplier 0.25).

Table Appendix B4: Base Case (M,M)

S	Group	%	E	Q	R (package)
S1	Young Demanders	34.08	$0.25 \times \mathcal{T}(0.5, 0.75, 1)$	$0.25 \times \mathcal{T}(0.5, 0.75, 1)$	$(0.5, 0.75, 1)$
S2	Willing-to-Pay Professionals	16.76	$0.75 \times \mathcal{T}(0.5, 0.75, 1)$	$0.25 \times \mathcal{T}(0.5, 0.75, 1)$	$(0.5, 0.75, 1)$
S3	Rich Heavy Users	16.48	$0.25 \times \mathcal{T}(0.5, 0.75, 1)$	$0.75 \times \mathcal{T}(0.5, 0.75, 1)$	$(0.5, 0.75, 1)$
S4	Conservative Shoppers	32.68	$0.10 \times \mathcal{T}(0.5, 0.75, 1)$	$0.10 \times \mathcal{T}(0.5, 0.75, 1)$	$(0.5, 0.75, 1)$

The base case is labelled as (M, M) level. When a scenario sets E and/or Q to H or L , we apply a common intensity adjustment to the corresponding triangular distribution parameters: values are doubled under H and halved under L , relative to the base case. Table [2](#) summarizes the (E, Q) level assigned to each scenario.

Appendix B.4. Provider cost parameters

All models— $\mathcal{Z}^0$, $\mathcal{Z}^e$, and $\mathcal{Z}^f$ —require provider-specific parameters β_j . We estimate β once using the base-case setting (Scenario 0) on Instance 1 data, and then hold these values fixed throughout all scenarios and instances. Accordingly, cross-scenario differences in outcomes are driven solely by changes in population composition and valuation distributions.

$$\min \sum_{j \in \mathcal{L}} (n_j - |S_j|)^2 \quad (\text{Appendix B5})$$

$$\text{s.t.} \quad \sum_{j \in \mathcal{L}} x_{ij} = 1 \quad \forall i \in \mathcal{C} \quad (\text{Appendix B6})$$

$$p_j = \min\left\{\bar{\beta}, \frac{\beta_j}{\sqrt{\sum_{i \in \mathcal{C}} x_{ij}}}\right\}, \quad p_j^+ = \min\left\{\bar{\beta}, \frac{\beta_j}{\sqrt{1 + \sum_{i \in \mathcal{C}} x_{ij}}}\right\} \quad \forall j \in \mathcal{L} \quad (\text{Appendix B7})$$

$$u_{ij} = v_{ij} - p_j, \quad u_{ij}^+ = v_{ij} - p_j^+ \quad \forall i \in \mathcal{C}, \forall j \in \mathcal{L} \quad (\text{Appendix B8})$$

$$\begin{aligned}
u_{ij} &\geq u_{ij'}^+ - M(1 - x_{ij}) && \forall i \in \mathcal{C}, \forall j \neq j' \in \mathcal{L} \\
&&& \text{(Appendix B9)} \\
n_j &= \sum_{i \in S_j} x_{ij} && \forall j \in \mathcal{L} \\
&&& \text{(Appendix B10)} \\
x_{ij} &\in \{0, 1\} && \forall i \in \mathcal{C}, \forall j \in \mathcal{L} \\
&&& \text{(Appendix B11)} \\
p_j &\geq 0, \quad p_j^+ \geq 0, \quad \beta_j \geq 0 && \forall j \in \mathcal{L}. \\
&&& \text{(Appendix B12)}
\end{aligned}$$

The objective ([Appendix B5](#)) minimizes squared deviations from target segment sizes. Constraint ([Appendix B6](#)) enforces exclusive assignment. Tariffs and entry tariffs follow ([Appendix B7](#)), i.e., the schedule (1) at the assigned load and at that load plus one; the cap does not bind at the calibrated loads, and the tariff of an empty provider is immaterial here as well. Customer utilities are defined in ([Appendix B8](#)). Constraint ([Appendix B9](#)) is the self-enforcement condition (7), so that the calibrated segment assignment is a Nash equilibrium. The alignment condition ([Appendix B10](#)) reconciles the actual participation with the intended segment size of each provider. Domain restrictions are defined in ([Appendix B11](#))–([Appendix B12](#)).