Article

"Incentive Design for Heterogeneous Gig Delivery Workers: A Stackelberg Framework with Windfall Decomposition"

Sixuan Li [1,*]

1 McCallum Business School, Bentley University, Waltham, Massachusetts, United State
* Correspondence: Sixuan Li, McCallum Business School, Bentley University, Waltham, Massachusetts, United State

Abstract: The rapid expansion of the gig economy has made the efficient management of independent contractors a critical operational challenge. Gig delivery platforms frequently spend heavily on zone-level financial bonuses to stabilize workforce supply. However, a substantial share of this spending inadvertently accrues to workers who would have participated without additional compensation—an inefficiency we term a 'windfall' transfer. To address this, we model the incentive-allocation problem as a rigorous Stackelberg game. The platform acts as the leader, setting zone-specific bonuses under a budget constraint, while heterogeneous workers respond through a logit congestion game that produces endogenous spatial competition. Our analytical results establish two critical boundaries: a no-spend threshold, representing a workforce composition beyond which all bonus spending becomes wasteful, and a zone-concentration condition, defining the demand asymmetry beyond which optimal spending concentrates in a single zone. Comprehensive computational experiments on synthetic instances, parameterized using published elasticity estimates and empirical labor survey data, yield two primary findings. First, the windfall share is governed almost entirely by workforce composition rather than the allocation policy, demonstrating that spatial targeting cannot effectively screen out infra-marginal workers. Second, the Stackelberg optimizer's advantage stems from selectively withholding spending in zones where windfall costs exceed activation benefits, improving overall profit by approximately 40% over the no-bonus baseline. Ultimately, these results suggest that the binding constraint on incentive effectiveness is the inherent share of already-committed workers in the labor pool, rather than the absolute size or spatial allocation of the incentive budget.

Keywords: gig economy; stackelberg game; bilevel optimization; incentive design; last-mile delivery; labor heterogeneity

Received: 05 March 2026
Revised: 14 April 2026
Accepted: 27 April 2026
Published: 03 May 2026

1. Introduction

Last-mile delivery has become one of the most operationally demanding segments of urban logistics. As consumer expectations for speed have intensified, platforms have shifted from dedicated fleet models toward hybrid systems that combine a core of employed drivers with a flexible pool of independent contractors—gig workers who self-schedule and self-select into delivery tasks through mobile applications. This workforce architecture introduces a fundamental design challenge: the platform cannot command workers to appear at specific times and places, but it can attempt to shape their behavior through financial incentives [1].

In practice, platforms deploy a variety of incentive instruments—zone bonuses, surge multipliers, guaranteed-earnings floors, and referral bonuses—to attract workers into zones and periods where demand exceeds organic supply. These programs are operationally central. Yet accumulating empirical evidence from the ride-hailing literature suggests a persistent inefficiency: a large fraction of incentive spending flows to workers who were already planning to participate. Studies document substantial heterogeneity in driver labor-supply elasticities, with a significant share of drivers exhibiting near-zero responsiveness to short-run pay changes [1]. Other findings indicate reference-dependent earnings targeting among drivers, implying that bonus payments to workers already at their earnings target do not generate additional supply. These patterns suggest that platform incentive programs may transfer rents to infra-marginal workers at least as much as they activate marginal ones.

The starting point of our analysis is that the incentive-design problem is better understood as a game between a platform and a heterogeneous workforce than as a budget allocation exercise with an exogenous supply response. The platform moves first by posting zone-level bonuses. Workers then decide whether to participate and where to work, taking into account the incentive, the base wage, their individual outside options, and the congestion they anticipate from other workers choosing the same zone. Because workers differ in their baseline attachment to the platform—some are effectively full-time and would work regardless, others are marginal and responsive to bonuses, and still others are unreachable at any feasible incentive level—the design problem involves adverse selection layered on top of a spatial congestion game.

The Stackelberg structure captures a strategic asymmetry that single-agent optimization models miss. When the platform raises the bonus in a zone, it simultaneously attracts responsive workers (the intended effect), pays a windfall to always-on workers already in that zone (an unintended transfer), and induces congestion that partially offsets the benefit for responsive entrants (an equilibrium externality). These three effects interact in ways that depend on the composition of the workforce, the spatial distribution of demand, and the severity of congestion. A framework that accounts for all three is needed to identify when incentive spending creates value and when it is wasted [1].

The paper makes four contributions [2].

First, it formulates a Stackelberg incentive-design model that embeds three worker types—always-on, responsive, and dormant—within a logit-based spatial congestion game and links this equilibrium to a platform profit-maximization problem through a bilevel program (Section 3). Unlike models that treat labor supply as a smooth, increasing function of compensation, this formulation explicitly tracks which worker types absorb bonus payments, enabling direct quantification of windfall inefficiency [3].

Second, it derives analytical results for tractable special cases that provide structural insight without computation (Section 4). A no-spend threshold identifies the workforce composition beyond which any bonus spending generates net-negative returns. A companion targeting result shows that when spending is warranted, the optimal allocation generically concentrates in a single zone once demand asymmetry is sufficiently large [4].

Third, it introduces a windfall decomposition metric that partitions total incentive expenditure into the share absorbed by infra-marginal workers and the share that activates marginal participants (Section 3.5). This metric is observable in principle from platform data and provides a governance tool for monitoring the incremental value of incentive programs over time [3].

Fourth, it demonstrates through computational experiments on synthetic instances with empirically grounded parameter ranges that Stackelberg-optimal targeting achieves substantial improvements over benchmark policies across a range of operating conditions (Section 6). The experiments reveal that the return on incentive spending depends primarily on workforce composition, not budget scale, and that uniform bonus programs are especially wasteful when demand is spatially concentrated [5].

The analysis is organized around three research questions: (1) How should a platform allocate a fixed incentive budget across zones when workers are strategic and heterogeneous in responsiveness? (2) Under what workforce and demand conditions does bonus spending expand realized coverage versus create windfall transfers? (3) How do practical benchmark policies compare with Stackelberg-optimal targeting, and what drives the performance gap?

The remainder of the paper proceeds as follows. Section 2 reviews related work. Section 3 presents the model. Section 4 states theoretical results. Section 5 describes calibration and computational design. Section 6 presents experimental findings. Section 7 discusses implications [5]. Section 8 concludes. All proofs appear in the Appendix.

2. Related Literature

The work connects to four research streams: platform incentive design and pricing, gig labor-supply behavior, bilevel optimization in transportation, and crowdsourced delivery operations [6]. Each stream is reviewed to identify specific gaps that motivate the present framework.

2.1. Platform Pricing and Incentive Design

Existing models of platform compensation design capture important aspects of the problem but uniformly miss the interaction between worker heterogeneity and windfall. Surge pricing with self-scheduling capacity has been analyzed, but assumptions about a representative worker with a single reservation wage prevent addressing which workers capture the surplus from price adjustments. Spatially differentiated pricing has been shown to dominate uniform pricing—a result confirmed by experiments—but models often assume homogeneous drivers and fail to analyze windfall. Foundations for pricing in networks with externalities have been established, and studies have examined delay-sensitive customers, self-scheduling workers, dynamic type matching, and the potential misinterpretation of surge signals. For a broad review, see related discussions [7].

In gig delivery specifically, guaranteed minimum earnings have been found to improve welfare, but the gains depend on driver heterogeneity—a dependency documented but not structurally modeled [4]. Pricing with differentiated service levels has been analyzed, but assumptions of symmetric drivers limit the scope of these studies. The gap across this literature is a formal treatment of how worker-type heterogeneity interacts with spatial incentive allocation to create windfall transfers—a gap this paper addresses.

2.2. Gig Worker Labor Supply and Behavioral Heterogeneity

The empirical case for worker heterogeneity is strong, but the theoretical models that use it for incentive optimization are sparse. Estimates suggest labor supply elasticities of 0.4–0.6 on the intensive margin, with considerable cross-driver variation. Importantly, a substantial fraction exhibit near-zero elasticity, consistent with the always-on type [8]. Roughly half of drivers remain active for over a year, while many churn within months, supporting the persistent/transient distinction. Reference-dependent earnings targeting has been documented, where drivers stop working after reaching an implicit income goal. This implies that bonuses paid to drivers at their target generate no additional supply—a mechanism closely related to windfall. Similar patterns have been observed among taxi and ride-hailing drivers. The behavioral game theory literature provides additional foundations for modeling bounded rationality in strategic settings.

In delivery specifically, driver arrival patterns are shown to be stochastic and heterogeneous. Scheduling methods have been developed to account for driver availability uncertainty, and distinctions have been made between dedicated and opportunistic gig drivers—a categorization paralleling the always-on and responsive types. The empirical evidence collectively motivates a model with discrete worker types, heterogeneous responsiveness, and congestion externalities. What the literature lacks is a formal framework connecting these features to incentive allocation and windfall measurement [5].

2.3. Bilevel Optimization and Stackelberg Games in Transportation

Transportation science has a mature tradition of bilevel and Stackelberg models for problems where a system authority sets policy parameters and users respond at equilibrium [7]. The foundational application is toll design: a bilevel model is formulated in which a government sets highway tolls and drivers choose routes, solving the resulting mathematical program with equilibrium constraints (MPEC) to find welfare-maximizing tolls. This approach has been extended to multicommodity networks with toll revenue maximization and to second-best congestion pricing under user equilibrium. For comprehensive treatments of bilevel programming and MPEC theory, several foundational works are available.

More recent work has broadened the application scope. Comprehensive treatments of bilevel optimization in transportation network design have been developed. Network design problems with equilibrium constraints have been applied to transit frequency setting, parking pricing, and electric vehicle charging infrastructure [9]. In each case, the bilevel structure captures the same fundamental tension present in our problem: the leader optimizes a system-level objective, but realized outcomes depend on decentralized follower decisions that the leader can influence but not control.

The application of bilevel models to workforce incentive design in delivery networks is, to our knowledge, new. The closest precedent is a spatial pricing model, which shares the geographic equilibrium structure but models a ride-hailing market rather than a delivery network with fixed demand.

2.4. Crowdsourced Delivery Operations

A growing operations research literature studies planning and operational problems in crowdsourced delivery systems. Vehicle routing problems with occasional drivers have been explored, with the development of exact and heuristic solution methods [1]. Same-day delivery systems incorporating crowd-shippers with heterogeneous willingness to participate have demonstrated that accounting for worker heterogeneity significantly improves delivery performance. Additionally, crowdsource-enabled systems with relay points have been designed, modeling the matching between crowd-shippers and delivery tasks.

On the workforce management side, robust optimization approaches have been developed to address workforce management in crowdsourced delivery, using distributionally robust formulations to handle uncertainty in driver supply. Comprehensive updates on challenges and opportunities in crowdsourced delivery planning have identified workforce heterogeneity and incentive design as key research directions. Task assignment, pricing, and capacity planning for hybrid fleets of centralized and decentralized couriers have also been studied, with a focus on operational routing rather than strategic incentive allocation [3].

The present paper complements this operational literature by addressing a question that sits upstream of routing and assignment: given that the platform cannot directly control worker supply, how should it allocate a fixed incentive budget to shape supply toward demand? The answer depends on workforce composition in a way that purely operational models do not capture [10].

Table 1 summarizes the positioning. The defining combination of this work is the joint treatment of (i) a Stackelberg leader-follower structure, (ii) discrete worker-type heterogeneity with incomplete information, (iii) zone-level congestion effects in the follower game, and (iv) a windfall decomposition that measures the structural inefficiency of incentive spending [10]. No prior paper in our survey combines all four elements, though each has been studied individually.

Table 1. Literature positioning.

	Stackelberg	**Discrete**	**Congestion**	**Windfall**	**Public**
	structure	types	game	metric	calibration
[8]	✓				
[6]	✓		✓		
[10]		✓			

[14]		✓			
[23]		✓			
[21]	✓		✓		
[1]		✓			
This paper	✓	✓	✓	✓	✓

3. Model

At a high level, the model captures a three-way tradeoff that platform managers face every time they design a bonus campaign. Bonuses attract additional workers to underserved zones (the activation benefit), but they also pay workers who were coming anyway (the windfall cost), and the workers who do respond create crowding that reduces everyone's earnings (the congestion externality). The platform moves first by choosing how much bonus to offer in each zone. Workers then observe the bonuses and decide whether to participate and where to go, accounting for the bonus, their personal costs, and the congestion they expect from other workers making similar calculations. The model formalizes this sequence and makes the tradeoffs precise enough to optimize effectively [2].

3.1. Setting and Timing

Consider a gig delivery platform operating in a metropolitan area divided into Z geographic zones indexed by $z \in \mathcal{Z} = \{1, \dots, Z\}$ over T time periods indexed by $t \in \mathcal{T} = \{1, \dots, T\}$. Demand is uncertain, with scenario $\omega \in \Omega$ realized with probability p_ω. In scenario ω, zone z in period t has demand d_{zt}^{ω} measured in delivery tasks [11].

The workforce consists of N delivery partners. Each worker belongs to one of three types $k \in \mathcal{K} = \{A, R, D\}$. The type structure reflects empirical evidence that gig workers are not a homogeneous population but rather a mixture of persistent high-attachment workers, elastic occasional workers, and inactive registrants [12].

Remark 1 (Discrete versus continuous heterogeneity). The three-type structure is a modeling choice that warrants justification. A continuous heterogeneity model—where reservation utilities are drawn from a smooth distribution—would be more general but would obscure the economic mechanism the analysis aims to isolate: the interaction between infra-marginal workers who capture windfall and marginal workers who generate coverage. The discrete structure makes this interaction transparent by assigning each worker to a named category with interpretable behavior. Importantly, the computational model introduces within-type heterogeneity (individual costs, zone preferences, and reservation utilities drawn from type-specific distributions), so the discrete types serve as anchors for a richer individual-level specification rather than as a rigid partition. The key results—windfall monotonicity, zone-concentration thresholds, spending restraint—hold under any specification that distinguishes workers by their responsiveness to bonuses, whether that distinction is discrete or continuous [13].

Always-on workers (type A) have low reservation wages and participate in most periods regardless of bonus availability. They represent the stable core of the delivery workforce—workers for whom platform delivery is a primary income source. Responsive workers (type R) have moderate reservation wages and adjust their participation based on expected net earnings [14]. They represent the margin that incentive programs target: workers with alternative uses of their time who can be drawn into the platform when conditions are sufficiently attractive. Dormant workers (type D) are registered on the platform but do not participate at any feasible incentive level; they represent the long tail of inactive accounts observed on most gig platforms.

The platform observes aggregate type shares (N_A, N_R, N_D) through historical data analysis but cannot identify individual worker types ex ante [15]. Let $\alpha = N_A/N$ denote the always-on share.

3.2. Game Structure

The interaction between the platform and its workforce follows the structure of a two-stage game with incomplete information [16].

Definition 1 (Incentive design game). The incentive design game $\mathcal{G} = (\mathcal{P}, \mathcal{S}, \mathcal{U})$ consists of:

1. Players $\mathcal{P}$: one platform (leader) and N workers (followers), indexed by $i = 1, \dots, N$, each with private type $k_i \in \mathcal{K}$.
2. Strategy sets $\mathcal{S}$: the platform chooses $b = (b_{zt}) \in \mathbb{R}_+^{ZT}$ subject to $\sum_{z,t} b_{zt} s_{zt} \leq B$. Each worker i chooses an action $a_i \in \mathcal{Z} \times \mathcal{T} \cup \{0\}$, where (z, t) represents entering zone z in period t and 0 denotes the outside option (non-participation).
3. Payoffs $\mathcal{U}$: the platform receives the objective (4). Worker i of type k choosing (z, t) receives utility $U_{zt}^k(b, s)$ defined in (1); choosing the outside option yields utility zero.

The game has three features that distinguish it from standard Stackelberg formulations in the transportation literature [14].

First, it is a Stackelberg game: the platform commits to b before workers make decisions, and workers observe b before choosing actions. This timing reflects the operational reality that platforms post zone bonuses in advance and workers see them before deciding whether to log in and where to deliver [17].

Second, the follower game among workers is a congestion game in the sense that each worker's payoff from choosing zone z depends on the total number of workers choosing that zone through the congestion function $\Gamma_{zt}(s_{zt})$. This creates strategic interdependence among workers—each worker's optimal zone depends on what others choose—and the relevant equilibrium concept for the follower game is a Nash equilibrium (or, under the logit specification, a quantal response equilibrium).

Third, the platform faces incomplete information about worker types. While the platform knows the type distribution (N_A, N_R, N_D), it does not observe individual types and therefore cannot condition bonuses on type. The bonus b_{zt} is paid to every worker in zone z, regardless of type [18]. This creates an adverse selection structure: the platform would prefer to pay only responsive workers (whose participation is contingent on the bonus) but cannot screen out always-on workers (who collect the bonus without changing behavior).

Definition 2 (Stackelberg equilibrium). A Stackelberg equilibrium of the incentive design game is a pair $(b^\star, (x^\star, s^\star))$ such that:

1. $(x^\star, s^\star) \in \mathcal{E}(b^\star)$: the worker flow-supply pair constitutes a Nash equilibrium of the follower congestion game given $b^\star$.
2. $b^\star$ solves the platform's problem (4) subject to the equilibrium constraint.

3.3. Worker Decisions and Equilibrium

The worker's decision involves two components: whether to participate at all and, if so, which zone-period to enter [8]. These choices are modeled jointly using a multinomial logit specification that includes an outside option, representing non-participation.

Given incentive vector b and aggregate supply vector $s = (s_{zt})$, a type-k worker's utility from selecting zone-period (z, t) is determined by several factors [19].

$U_{zt}^k(b, s) = \bar{w}_{zt} + b_{zt} - c_{zt}^k - \Gamma_{zt}(s_{zt}) - r_k,$

Here, $\bar{w}_{zt}$ represents the base payment per delivery, b_{zt} denotes the bonus, c_{zt}^k captures the type-specific effort and commute cost, $\Gamma_{zt}(s_{zt})$ reflects the congestion disutility, and r_k is the reservation utility [20]. The outside option is normalized to yield zero utility.

The congestion term $\Gamma_{zt}(s_{zt})$ encapsulates a key economic characteristic of gig delivery: as more workers enter a zone, each worker receives fewer delivery tasks, thereby reducing their effective hourly earnings [21]. A linear specification $\Gamma_{zt}(s_{zt}) = \gamma_{zt} s_{zt}$ is adopted, where $\gamma_{zt} > 0$ represents the congestion slope. This creates a classic congestion game structure, where each worker's utility decreases as more workers choose the same option.

The choice probabilities are determined by the multinomial logit model:

$$P_{zt}^k(\boldsymbol{U}^k) = \frac{\exp(\mu\, U_{zt}^k)}{1+\sum_{z',t'}\exp(\mu\, U_{z't'}^k)},$$

In equilibrium, supply must be self-consistent:

$$x_{zt}^k = N_k P_{zt}^k(U^k(b,s)), \qquad s_{zt} = \sum_{k\in\mathcal{K}} x_{zt}^k\,, \quad \forall z,t.$$

Let $\mathcal{E}(b)$ represent the set of supply vectors that satisfy this condition.

The economic implications of the equilibrium condition are significant. When the platform increases b_{zt}, three effects occur simultaneously [22]. First is the direct attraction effect, where responsive workers outside zone z find it more appealing, leading to an increase in x_{zt}^R. Second is the windfall transfer effect, where always-on workers already in zone z receive the bonus without altering their behavior, raising platform costs without increasing supply. Third is the congestion feedback effect, where the influx of workers into zone z raises Γ_{zt}, partially offsetting the initial attraction. The net impact on coverage and cost depends on the interplay of these forces, which vary based on workforce composition, demand levels, and congestion intensity.

3.4. Platform Optimization

The platform maximizes expected profit net of base wages, incentive costs, and unmet-demand penalties [16].

$$\max_{b\geq 0} \quad \sum_\omega p_\omega \sum_{z,t}\left[R_{zt}\lambda_{zt} - \bar{w}_{zt}\lambda_{zt} - b_{zt}s_{zt} - C^u(d_{zt}^\omega - \lambda_{zt})^+\right]$$

$$\text{s.t.} \quad \sum_{z,t} b_{zt}\, s_{zt} \leq B, \quad (x,s)\in\mathcal{E}(b),$$

where $\lambda_{zt} = \min(s_{zt}, d_{zt}^\omega)$ is the number of completed deliveries and C^u is the unmet-demand penalty [23]. The budget constraint $\sum_{z,t} b_{zt}\, s_{zt} \leq B$ bounds total incentive expenditure. The constraint $(x,s)\in\mathcal{E}(b)$ makes this a bilevel program: the supply that determines both revenue and cost is itself an equilibrium outcome of the follower game.

A critical modeling choice is that the bonus cost $b_{zt}s_{zt}$ is paid to all workers in zone z, not only responsive entrants. This is realistic—platforms cannot condition bonuses on counterfactual participation—and it is precisely the feature that generates windfall transfers [24]. The always-on component of s_{zt} collects the bonus without any change in behavior, creating a structural tension between the coverage benefit and the transfer cost of incentive spending.

3.5. Windfall Decomposition

To measure this tension, the windfall share is defined as follows:

$$\text{WindfallShare}(b) = \frac{\sum_{z,t} b_{zt}x_{zt}^A}{\sum_{z,t} b_{zt}s_{zt}}.$$

This metric represents the fraction of total bonus expenditure absorbed by always-on workers [25]. A value close to one indicates that nearly all spending is allocated to workers who would have participated regardless. Conversely, a value near zero suggests that spending is concentrated in areas where responsive workers dominate the supply. The windfall share serves as a summary statistic for evaluating the structural efficiency of an incentive program, which can, in principle, be estimated using platform data by identifying persistently active workers and calculating their share of bonus payouts.

3.6. Game-Theoretic Structure

The model incorporates three layers of strategic structure, each contributing uniquely to different aspects of the results [5].

3.6.1. Why This Is a Congestion Game.

For any fixed bonus vector b, workers compete for delivery tasks in the same zones [26]. Each worker's payoff from entering zone z depends on how many others choose the same zone—more entrants mean fewer tasks per worker and lower effective earnings. This is the defining feature of a congestion game. Under the logit choice model (2), the game admits a potential function.

$$\Phi(\boldsymbol{s}) = \sum_{k\in\mathcal{K}} \frac{N_k}{\mu} \log\left(1 + \sum_{z,t} \exp\left(\mu\, U_{zt}^k(\boldsymbol{b},\boldsymbol{s})\right)\right) - \sum_{z,t} \int_0^{s_{zt}} \Gamma_{zt}(\sigma)\, d\sigma.$$

The equilibrium (3) corresponds to a stationary point of Φ. Convexity of Γ_{zt} and concavity of the logit log-sum term make Φ strictly concave, guaranteeing both existence and uniqueness of the follower equilibrium [27]. Uniqueness ensures that the Stackelberg leader faces a well-defined response: for each b, there is a single equilibrium $s(b)$.

3.6.2. Why Crowded Zones Are Bad Targets.

Workers in the same zone are strategic substitutes: entry by one worker increases congestion, reducing utility for all others in that zone. The aggregate welfare loss from this externality is $\gamma_{zt} \cdot s_{zt}$ per additional entrant, growing linearly with zone supply. Zones that are already well-served impose significant externalities on entrants, making further entry socially costly even if privately appealing. This asymmetry is central to the optimizer's strategy: bonuses in crowded zones attract workers who impose large externalities, while bonuses in under-served zones attract workers who impose smaller ones.

3.6.3. How Bonuses Propagate through the System.

The equilibrium response to bonus changes can be decomposed as follows [28]. Increasing b_{zt} raises supply in zone z ($\partial s_{zt}/\partial b_{zt} > 0$, the direct attraction), weakly reduces supply in other zones ($\partial s_{z't}/\partial b_{zt} \leq 0$, substitution), and increases total participation ($\sum_{z'} \partial s_{z't}/\partial b_{zt} > 0$, the net effect). The supply response $\partial s_{zt}/\partial b_{zt}$ decreases with γ_{zt} and s_{zt}: bonuses are less effective in zones that are already congested. This diminishing-returns property forms the formal basis for the budget-saturation patterns observed in the experiments.

3.6.4. Why the Platform Cannot Avoid Paying Windfall.

The platform's type-ignorance creates an adverse-selection problem that it cannot resolve through contract design. In standard mechanism design, a principal facing heterogeneous agents offers a menu of contracts to induce self-selection. The institutional constraint of gig platforms—all workers in a zone receive the same bonus—eliminates this possibility. Any bonus designed to attract responsive workers simultaneously pays always-on workers in the same zone. The windfall share (5) measures the cost of pooling: the unavoidable transfer to infra-marginal agents under a pooling contract. This cost is structural, not incidental; it cannot be eliminated by better optimization, only mitigated by concentrating bonuses in zones where the responsive-to-always-on ratio is highest [29].

3.6.5. Why Anticipation Matters.

The appropriate solution concept is Stackelberg equilibrium. The platform selects b to maximize its objective, anticipating the unique follower response $s(b)$. The timing reflects institutional reality: platforms post incentive schedules before workers decide, and posted bonuses are commitments [30]. Because $s(b)$ is continuous in b (as shown by the implicit function theorem applied to the potential's first-order conditions) and the feasible set is compact, the existence of a Stackelberg equilibrium follows from the Weierstrass theorem. The computational challenge lies in the objective being non-concave in $\boldsymbol{b}$, necessitating multi-start methods for the numerical solution.

4. Theoretical Results

This section establishes structural properties of the Stackelberg game. The results provide qualitative insight that holds independently of computational experiments and grounds the interpretation of numerical findings. Proofs appear in the Appendix.

4.1. Regularity Conditions

Assumption 1 (Congestion regularity). For each (z,t), $\Gamma_{zt}: [0, \bar{S}_{zt}] \to \mathbb{R}_+$ is continuous, convex, and nondecreasing, with $\Gamma_{zt}(0) = 0$.

Assumption 2 (Type ordering). Reservation utilities satisfy $r_A < r_R < r_D$, with r_D sufficiently large that type-D participation is zero at any feasible bonus [31].

Assumption 3 (Bounded productivity). Completed deliveries satisfy $\lambda_{zt}(s,d) = \min(s,d)$, which is continuous and nondecreasing in s.

4.2. Existence and Structure of Equilibrium

Proposition 1 (Existence). Under the assumption of congestion, for any feasible $b \geq 0$, the follower game admits at least one equilibrium in $\mathcal{E}(b)$.

The proof (Appendix 9.1) applies Brouwer's theorem to the continuous self-map defined by (3) on the compact set $[0, N]^{ZT}$. The logit structure ensures continuity of the best-response mapping, and convexity of congestion prevents pathological cycling [32]. The practical import of uniqueness is that the Stackelberg leader faces a well-defined response: for each bonus vector b, there is exactly one equilibrium supply $s(b)$, so the leader's optimization problem is well-posed.

Proposition 2 (Monotonicity of responsive supply). Under the assumptions of congestion and types, total type-R participation is weakly increasing in the bonus vector b (componentwise).

This result is intuitive—higher bonuses make participation more attractive—but it is not trivial, because the congestion feedback from additional workers partially offsets the direct incentive effect. The proof (Appendix 9.2) shows that the direct effect dominates the congestion feedback under the logit structure, ensuring that bonuses never reduce responsive supply.

4.3. Windfall Characterization

The following results demonstrate how workforce composition influences the structural efficiency of incentive spending [33].

Proposition 3 (Windfall monotonicity). Under specific assumptions, the windfall share (5) increases with the always-on share α. Additionally, a higher congestion slope γ_{zt} amplifies the windfall share by reducing the marginal coverage gain per bonus dollar.

The economic reasoning is straightforward. When α is large, always-on workers represent a significant portion of zone-level supply, causing a substantial fraction of any bonus payment to flow to them. At the same time, the responsive pool is smaller, making the marginal worker activated by a bonus dollar more costly to engage. Congestion intensifies both effects: in a congested zone, additional responsive workers contribute less to coverage (as many deliveries are already managed by always-on workers), while the bonus paid to all workers in the zone remains unchanged.

4.4. Two-Zone Specialization

To obtain the sharpest possible results, consider $Z = 2$ zones, $T = 1$ period, two types (A and R), linear congestion with common slope γ, and deterministic demand [34].

Proposition 4 (No-spend threshold). There exists $\alpha^\star \in (0,1)$ such that for $\alpha > \alpha^\star$, the Stackelberg-optimal bonus is $\boldsymbol{b}^\star = \boldsymbol{0}$. The threshold satisfies $\alpha^\star = 1 - \frac{\gamma \bar{d}}{\mu^{-1} + \Delta w}$, where $\bar{d}$ is mean demand and Δw is the inter-zone wage differential.

Proposition 4 formalizes the central insight of the paper: there is a workforce-composition boundary beyond which no incentive program is worthwhile. Below $\alpha^\star$, the responsive pool is large enough that targeted bonuses can activate enough workers to justify the windfall cost. Above $\alpha^\star$, the always-on workforce alone provides sufficient coverage, and any spending generates net-negative returns because windfall costs dominate activation benefits.

Proposition 5 (Zone concentration). Under the same specialization, when positive spending is optimal, there exists a demand-asymmetry threshold $\delta^\star > 0$ such that if

$|d_1 - d_2| > \delta^\star$, the optimal allocation is a corner solution concentrating the entire budget in one zone [35].

This result contradicts a natural managerial intuition that incentive budgets should be spread across zones to diversify coverage. Instead, when demand is sufficiently uneven, the platform should focus all spending on the high-demand zone where the coverage gap is largest and the marginal return per bonus dollar is highest [36]. The threshold $\delta^\star$ decreases in γ, meaning that higher congestion favors concentration even at moderate demand asymmetry, because congestion dissipates the value of spreading workers thinly across zones.

5. Calibration and Computational Design

5.1. Parameterization

Instances are synthetic, with parameter ranges informed by publicly available empirical evidence [37]. This approach follows the convention in the bilevel optimization literature, where synthetic instances with empirically grounded parameters are standard. Table 2 summarizes the evidence used to set parameter ranges.

Table 2. Parameterization sources.

Parameter	Evidence base	Range used
Type shares	BLS contingent worker surveys; [17] driver attachment patterns	$\alpha \in [0.10, 0.70]$; dormant share 20%
Participation elasticity	[9] intensive-margin estimates for ride-hailing drivers	Logit μ set to produce elasticities of 0.4–0.6
Demand heterogeneity	Documented variation in delivery volumes across urban zones [28, 31]	Three asymmetry regimes
Compensation	Publicly reported gig delivery earnings	Revenue \$9–\$13 per delivery; base wage 50–65%

Workforce type shares are the most consequential parameters [38]. Reports indicate that approximately 10% of workers in certain regions participate in gig platform work as a primary income source, with a larger share engaging intermittently. Studies document that roughly half of platform drivers remain active for over a year, while a substantial fraction churns within months, supporting the distinction between persistent and transient attachment. Baseline type shares are set at $\alpha = 0.30$ (always-on), $\alpha_R = 0.50$ (responsive), $\alpha_D = 0.20$ (dormant), and systematically vary α from 0.10 to 0.70 in the experimental design. The logit scale parameter μ and reservation utilities are jointly set to produce participation elasticities in the 0.4–0.6 range.

Zone-level demand is generated synthetically across three asymmetry regimes. In the low-asymmetry regime, demand is approximately uniform across zones. In the moderate regime, demand follows a descending gradient, reflecting the pattern of a metropolitan core with declining density. In the high regime, one zone concentrates 40% of total demand, representing a commercial district or logistics hub. These stylized patterns are not derived from a specific dataset but reflect the qualitative variation in delivery volumes documented in the last-mile delivery operations literature [39].

5.2. Instance Generation and Policy Benchmarks

Each instance specifies Z zones, $T = 1$ period, and three demand scenarios with multipliers $(0.7, 1.0, 1.3)$ at equal probability. Revenue per delivery is drawn uniformly from [\$8, \$12]; base wages are set at 55–70% of revenue. Effort costs vary by type: always-on workers face 30% lower costs (reflecting experience), while dormant workers face 50% higher costs [40].

Three policies are compared: (i) Stackelberg-optimal, solving (4) via multi-start L-BFGS-B with the equilibrium computed by damped fixed-point iteration; (ii) Uniform, setting equal bonuses across all zones scaled to exhaust the budget; (iii) Demand-proportional, allocating bonuses proportional to expected demand scaled to exhaust the budget. A no-bonus baseline serves as a reference [8].

Five experiments vary different parameters: instance size ($Z \in \{3,5\}$), always-on share ($\alpha \in \{0.10, \ldots, 0.70\}$), congestion intensity, demand asymmetry, and budget level. Each combination is replicated over three to seven random seeds; results are reported as means with standard deviations to characterize instance-to-instance variation.

5.3. Illustrative Scenario

To ground the model in operational terms, consider a delivery platform operating across five zones in a mid-size metropolitan area during a weekday evening shift. The platform has approximately 60 active delivery partners registered for the shift. Of these, roughly 18 (30%) are regulars who deliver most evenings regardless of bonus availability. These individuals have established routines, prefer zones near their homes, and rely on the platform as their primary source of income. Another 30 (50%) are occasional workers who have day jobs, attend school, or work on other platforms. Their participation depends on whether the bonus makes the shift worthwhile after accounting for commute time and opportunity costs. The remaining 12 (20%) are registered but effectively inactive.

The platform faces approximately 50 delivery tasks unevenly distributed: a commercial district generates 16 tasks, nearby residential zones generate 12 and 9, while outer zones generate 8 and 5. The platform's incentive budget is $45 for the shift, roughly 15% of the total wage bill. Under a uniform $0.98-per-delivery bonus spread across all zones, the platform spends the full budget, but most of it reaches the 18 regulars who would have delivered anyway. Under Stackelberg-optimal allocation, the optimizer concentrates a $2.13 bonus in the commercial district (where the coverage gap is largest) and a $0.98 bonus in one residential zone, while setting the remaining three zones to zero. Total spending drops to $36, as the optimizer chooses to leave $9 unspent rather than allocate it to zones where the bonus would be a pure windfall.

5.4. Scaling to Realistic Operations

To test whether these patterns hold at realistic scale, we simulate a large metropolitan delivery market with 10,000 registered workers across 30 zones, with $\alpha = 0.60$ (a mature platform where the majority of active workers are committed participants). Demand is spatially skewed: downtown zones generate 3--4 $\times$ more tasks than suburban zones [33]. Table 3 reports the results.

Table 3. Large-scale simulation: 10,000 workers, 30 zones, $\alpha = 0.60$.

Policy	Objective	Coverage	Windfall	Zones bonused
No bonus	12,023	0.754	-	0/30
Uniform	6,600	0.761	0.714	30/30
Targeted (top 15)	16,879	0.860	0.710	15/30

Three features of this result merit emphasis. First, the uniform policy destroys 45% of baseline profit: spending in all 30 zones transfers most of the budget to committed workers who were already delivering. Second, the targeted policy improves profit by 40% over no-bonus by concentrating spending in the 15 zones with the largest coverage gaps and leaving the other 15 zones unsubsidized. Third, the windfall share is nearly identical under uniform (0.714) and targeted (0.710) allocation, confirming that windfall invariance is not an artifact of small instances but persists at operational scale. The optimizer does not win by reducing windfall; it wins by spending only where the activation benefit exceeds the windfall cost.

For a mature platform where $\alpha \approx 0.60$, the practical implication is that roughly 71 cents of every bonus dollar will reach workers who were coming anyway, regardless of how the budget is allocated. The question is not whether windfall can be avoided (it cannot), but whether the remaining 29 cents of activation value per dollar justifies the expenditure in each zone. In 15 of 30 zones, the answer is yes; in the other 15, it is no.

6. Results

The results are organized around three findings that correspond to the research questions. Instead of presenting the experiments individually, the evidence is structured thematically, integrating data from multiple experiments to support each argument [30].

6.1. Selective Restraint and the Stackelberg Advantage

The most consistent pattern across all experimental conditions is that the Stackelberg-optimal policy achieves higher objective values than benchmark policies despite often producing lower coverage. Table 4 illustrates this for the five-zone baseline.

Table 4. Policy comparison, five-zone baseline (mean over replications).

Policy	Coverage	Objective	Windfall	Bonus spent
No bonus	0.732	74.6	-	0.0
Uniform	0.756	48.7	0.381	45.0
Demand-proportional	0.792	76.9	0.373	45.0
Stackelberg-optimal	0.822	107.2	0.378	36.0

The myopic policy is excluded from the main comparison because, lacking anticipation of equilibrium effects, it overspends catastrophically on larger instances. The remaining three bonus policies reveal the central tradeoff clearly. The uniform policy improves coverage modestly (0.756 versus 0.732) but reduces the platform objective from 74.6 to 48.7—a 35% decline relative to the no-bonus baseline. The reason is that the uniform policy pays bonuses in every zone, including zones where supply already exceeds demand and the bonus is pure windfall. The demand-proportional policy performs better (objective 76.9, coverage 0.792), concentrating more spending in high-demand zones, but still wastes a substantial share on zones that are adequately served.

The Stackelberg policy achieves an objective of 107.2—44% above the no-bonus baseline and more than double the uniform policy's objective [33]. It does this while spending less of the budget (36.0 of 45.0) and achieving the highest coverage (0.822). Inspection of the optimal bonus vector reveals the mechanism: the Stackelberg optimizer sets bonuses to zero in zones where supply exceeds demand and concentrates spending in the two zones with the largest demand-supply gaps. This is the zone-targeting behavior predicted by Proposition 5, confirmed numerically.

The key insight is that the Stackelberg advantage does not come from spending more or spending more cleverly everywhere. It comes from not spending in zones where spending is wasteful. The optimizer's restraint—its willingness to leave zones unsubsidized—is what distinguishes it from the benchmark policies. In a sense, the Stackelberg policy solves two problems simultaneously: where to spend, and where not to.

6.2. Workforce Composition as the Binding Constraint

The windfall share—the fraction of bonus spending absorbed by always-on workers—depends primarily on workforce composition α and is, in our experiments, approximately invariant to the allocation policy [21]. Across all three bonus policies, the windfall share rises from approximately 0.19 at $\alpha = 0.10$ to 0.89 at $\alpha = 0.70$, with inter-policy differences of only 1–2 percentage points (within one standard deviation across replications).

This near-invariance has a structural explanation [22, 36]. Always-on workers choose zones based on their individual costs and home-zone preferences, not based on bonuses—by definition, their participation decision is insensitive to bonus levels. Their spatial distribution across zones is therefore approximately the same under any bonus allocation.

The windfall share in a given zone is determined by the ratio of always-on supply to total supply in that zone, which depends on worker population composition rather than on the bonus vector. Spatial targeting can shift responsive workers across zones but cannot move always-on workers out of bonused zones.

Two conditions would break this invariance. First, if always-on workers were strongly concentrated in a small number of zones, a platform could bonus the remaining zones with near-zero always-on presence, achieving substantially lower windfall. Our model includes moderate always-on clustering, which produces the small 1–2 point gap between policies; stronger clustering would widen it [25, 30]. Second, if the platform could condition bonuses on observable worker-level characteristics correlated with type (e.g., historical participation frequency), it could screen more effectively, reducing windfall below the zone-level floor.

We emphasize that the invariance is not a modeling artifact but a consequence of a real institutional constraint: large-scale gig platforms typically offer zone-level bonuses visible to all workers in a zone, without conditioning on individual participation history or counterfactual behavior. Richer contracts—individualized bonuses, history-dependent incentives, or type-contingent offers—could in principle reduce windfall, but such mechanisms raise implementation costs, fairness concerns, and regulatory scrutiny that make them impractical at scale. The near-invariance of windfall under zone-level targeting is therefore a feature of the contracting environment, not a limitation of the model [40]. Recognizing this constraint is itself a managerial insight: it redirects attention from "how to reduce windfall" (which zone-level targeting cannot do) to "where windfall is tolerable" (which Stackelberg optimization answers).

This finding reframes the central question [35]. The Stackelberg advantage does not come from reducing windfall—it comes from optimizing around windfall. Figure 1 presents the core evidence: the net return per bonus dollar, defined as the objective gain relative to the no-bonus baseline divided by total spending.

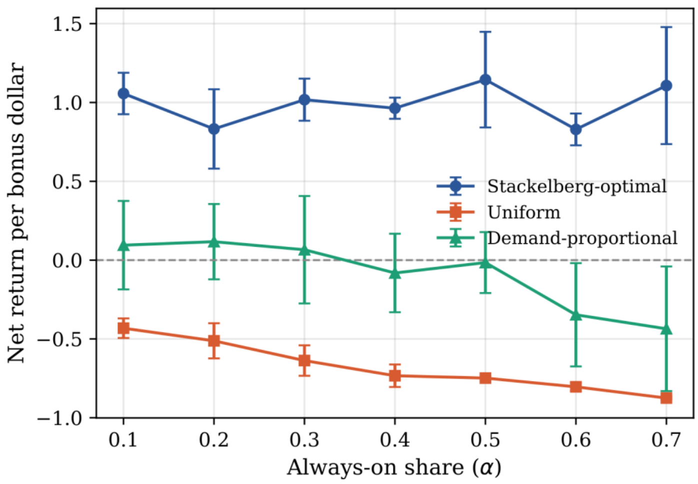

Figure 1. Net return per bonus dollar across workforce compositions. Positive values indicate profitable incentive spending; the Stackelberg policy maintains positive returns across all α levels b...

The Stackelberg policy maintains positive net returns at every α level. The uniform policy produces negative returns at most α levels—it destroys value because it spends in every zone, including zones where supply already exceeds demand. The demand-proportional policy performs moderately. The mechanism is spending restraint: the Stackelberg optimizer spends 36 of its 45-unit budget on the baseline instance, while uniform and demand-proportional exhaust the full budget. The optimizer's willingness to leave budget unspent is the source of its advantage.

Figure 2 shows the complementary pattern in responsive activation. The Stackelberg policy activates fewer responsive workers than uniform or demand-proportional policies—because it deliberately avoids spending in zones where activation would cause congestion—but this restraint generates higher profit per worker activated.

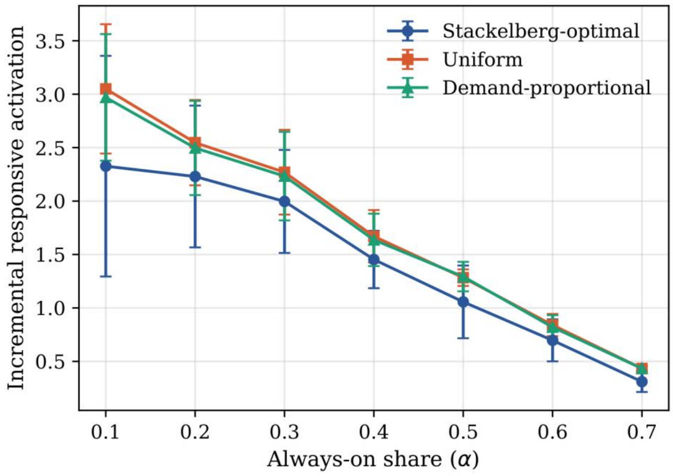

Figure 2. Incremental responsive activation versus always-on share.

The budget-sensitivity experiment (Figure 3) confirms diminishing returns. Increasing the budget from 5% to 10% of the wage bill raises the Stackelberg objective from 75.5 to 84.0. Further increases to 20% and 30% yield only marginal improvement (86.4 and 86.6), because the most cost-effective zone-type combinations are exhausted first. The uniform policy produces objectives below the no-bonus baseline at every budget level.

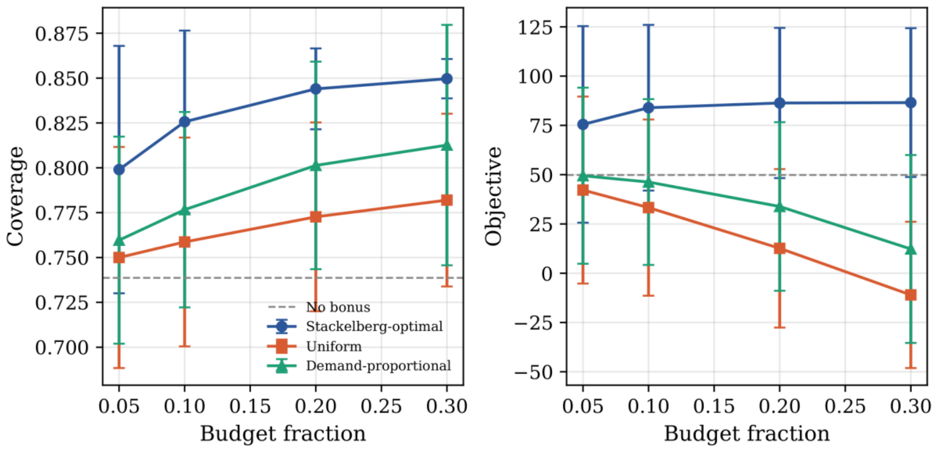

Figure 3. Coverage (left) and objective value (right) across budget levels. Dashed line indicates no-bonus baseline.

6.3. Conditions under Which Targeting Is Most Valuable

Two operating conditions amplify the Stackelberg advantage: congestion intensity and demand asymmetry [40].

Congestion. Under high congestion (Figure 4), the performance gap between Stackelberg and uniform widens because high congestion penalizes the uniform policy's indiscriminate zone allocation more severely. Workers attracted to already-crowded zones by uniform bonuses earn less per task and contribute less to coverage, while the bonus cost remains the same. The Stackelberg optimizer responds by concentrating spending in less-congested zones where the marginal worker faces lower disutility.

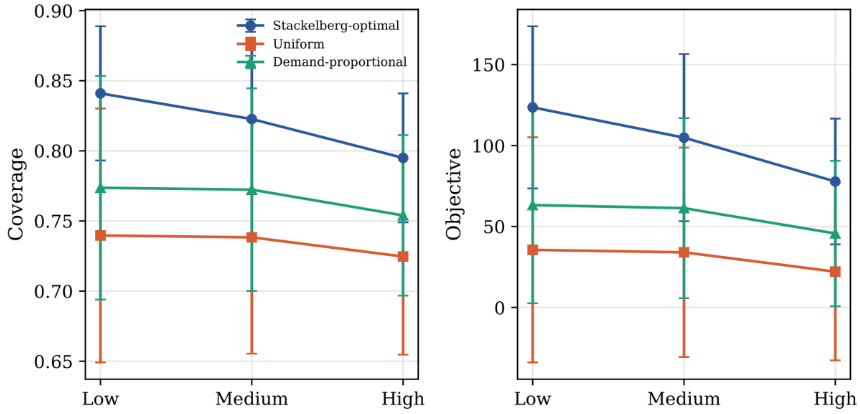

Figure 4. Coverage and windfall across congestion regimes.

Demand asymmetry. Under high asymmetry—where one zone has substantially more demand than others—the Stackelberg optimizer concentrates spending in the high-demand zone, consistent with Proposition 5 (Figure 5). The uniform policy spreads bonuses evenly, over-serving low-demand zones where supply already meets demand and under-serving the high-demand zone where the coverage gap is most costly. In the low-asymmetry regime, the gap narrows, because when demand is balanced across zones, uniform allocation is closer to optimal.

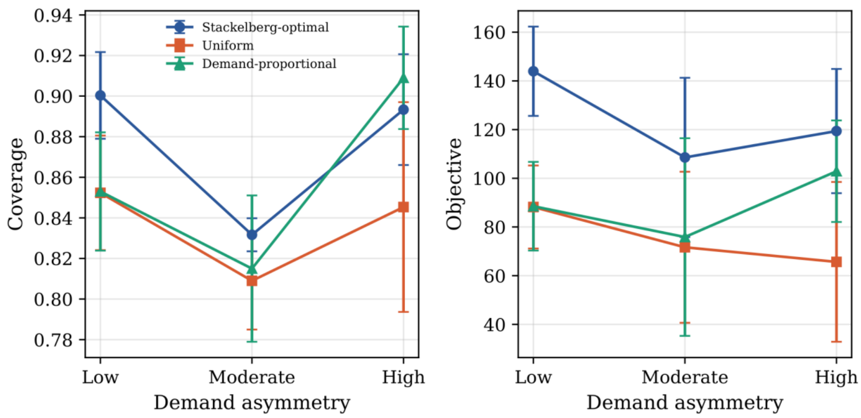

Figure 5. Coverage and windfall across demand-asymmetry regimes.

Two-zone threshold validation. The two-zone validation (Figure 6) confirms the zone-concentration prediction numerically with a striking pattern. At low α (< 0.14), the optimizer allocates bonuses to both zones. At $\alpha \approx 0.20$, the low-demand zone bonus drops to zero—the zone-concentration threshold of Proposition 5. As α increases further, total spending decreases gradually from 35 at $\alpha = 0.05$ to 11 at $\alpha = 0.65$, but does not reach zero. This "soft threshold" behavior—in contrast to the sharp cutoff predicted by the aggregate model—reflects individual-level worker heterogeneity: even at high α, there exist individual responsive workers whose activation is cost-effective.

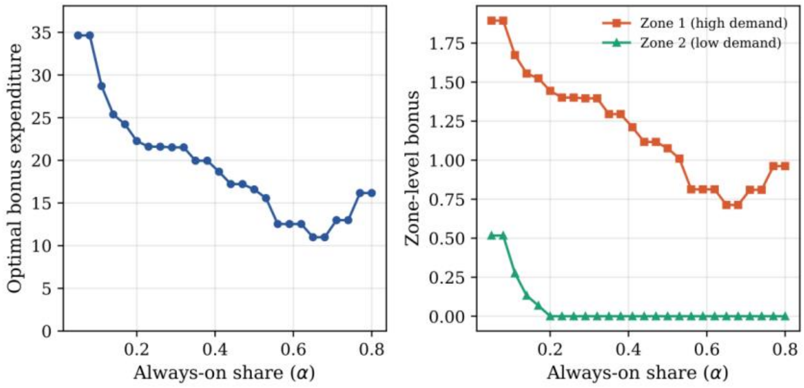

Figure 6. Two-zone validation: optimal bonus expenditure (left) and zone-level allocation (right).

Taken together, these three findings paint a coherent picture. The value of incentive spending in gig delivery is not primarily a question of how much to spend, but of where and to whom. Workforce composition determines the ceiling on what incentive spending can achieve; spatial demand structure determines where spending is most productive; and congestion determines how quickly returns diminish [30]. A framework that jointly accounts for all three factors—as the Stackelberg model does—can identify the narrow conditions under which incentives create value, while avoiding the much broader conditions under which they merely transfer rents.

6.4. Robustness

Table 5 presents the sensitivity of the Stackelberg profit gain to key model parameters. The gain varies significantly across parameter settings, indicating that the baseline result of approximately 44% is specific to the calibrated parameter range and should not be interpreted as a universal constant.

Table 5. Sensitivity of Stackelberg profit gain (% improvement over no-bonus baseline) to model parameters. Mean $\pm$ standard deviation across three replications.

Parameter	Setting	Stackelberg gain (%)
Logit scale μ	$\mu = 0.3$ (high noise)	18.0 ± 6.7
	$\mu = 0.6$ (baseline)	63.2 ± 32.8
	$\mu = 1.0$ (moderate)	162.2 ± 97.5
	$\mu = 1.5$ (deterministic)	593.4 ± 486.4
Congestion γ	$0.5 \times$ baseline	71.8 ± 40.8
	$1.0 \times$ baseline	63.2 ± 32.8
	$2.0 \times$ baseline	79.4 ± 50.1
Workforce size N	$N = 30$	377.0 ± 221.0
	$N = 60$ (baseline)	63.2 ± 32.8
	$N = 100$	113.2 ± 105.4

The Stackelberg advantage increases sharply with μ: when workers make more deterministic choices (high μ), targeted bonuses elicit sharper supply responses, enhancing the value of spatial optimization. At low μ (high choice noise), workers distribute more evenly regardless of bonuses, reducing the performance gap between policies. The gain remains relatively stable across congestion levels, suggesting that the spending-restraint mechanism functions independently of congestion intensity. The high variance observed at small workforce sizes ($N = 30$) reflects the noisier type composition in smaller samples; at larger N, the results stabilize.

7. Implications

7.1. Operational Implications

Three operational implications emerge from the analysis.

Segment before spending. The net-return analysis (Figure 1) suggests that the return on incentive spending is governed by the always-on share α, a quantity that can in principle be estimated from historical participation data. Before committing to a bonus program, a platform could estimate α for each zone and time window. If α exceeds the no-spend threshold of Proposition 4, the model predicts that the zone would receive no bonus regardless of demand conditions. This inverts the typical planning sequence, which starts with demand and budgets, not workforce composition, and follows directly from the model's structure.

Concentrate rather than diversify. The analytical targeting result (Proposition 5) and the numerical evidence from the demand-asymmetry experiment both point against the common practice of distributing bonuses broadly. When demand is spatially uneven, the platform should concentrate its budget in the highest-demand zones and leave other zones unsubsidized. The intuition is that bonuses in low-demand zones attract workers to locations where supply already exceeds demand, producing windfall with no coverage benefit. This recommendation strengthens under high congestion, where even moderate worker influx reduces individual utilization.

Track windfall as a governance metric. The windfall share (5) is defined in terms of model types, not directly observable behavior. In practice, an operational proxy can be constructed by comparing participation rates in bonused versus unbonused periods. Specifically, let the empirical windfall rate be the fraction of bonus payouts received by workers who were active in at least k of the last m unbonused periods in the same zone. For sufficiently high k (e.g., $k \geq 3$ of $m = 5$ recent periods), this proxy identifies workers whose participation is insensitive to bonus availability—the behavioral signature of always-on status. The model's windfall share and this empirical proxy will not coincide exactly: the model defines types by structural parameters (reservation utilities, cost distributions), while the empirical proxy defines them by revealed behavior. Nevertheless, the two measures are monotonically related under mild conditions, because workers with low reservation utilities (model type A) are precisely those who participate frequently in unbonused periods (high empirical k). A rising empirical windfall metric over successive incentive cycles signals that the program is losing incremental value—either because the always-on share has grown (as occasional workers become habituated) or because congestion in bonus zones has increased.

Rules of thumb [28]. While the precise thresholds depend on the operating environment, the model suggests several actionable heuristics for incentive governance:

- Windfall above 0.7: reconsider the program. When more than 70 cents of every bonus dollar reach workers who would have shown up anyway, the net return per dollar is low and shrinking. In our experiments, this corresponds to $\alpha \gtrsim 0.5$—a workforce where the majority are committed participants.
- Never run uniform bonuses when demand is spatially concentrated. If the coefficient of variation of demand across zones exceeds 0.3, uniform allocation consistently underperforms the no-bonus baseline in our experiments. Demand-proportional allocation is a strictly better default.
- Estimate α before setting the budget. The conventional planning sequence is: forecast demand, set budget, allocate across zones. Our results suggest inverting the first two steps: estimate the always-on share by zone, determine which zones have α below the no-spend threshold, and budget only for those zones.
- Leave budget unspent if the math says so. The Stackelberg optimizer's most distinctive behavior is spending less than the available budget. Organizational incentives that penalize managers for underspending work against the spending-restraint mechanism that drives the optimizer's advantage.

7.2. Policy and Societal Relevance

The findings also address broader questions that go beyond the operations of individual platforms.

Last-mile capacity and supply chain resilience. The COVID-19 pandemic highlighted the vulnerabilities of last-mile delivery systems, prompting federal policy to prioritize supply chain resilience as a national concern. Understanding how incentive design influences actual versus nominal workforce capacity—and how much of incentive spending is absorbed as windfall rather than activation—is crucial for policymakers evaluating the reliability of platform-based delivery infrastructure.

Labor-market efficiency. The windfall phenomenon suggests that a portion of platform incentive spending is economically inefficient at the system level, as it transfers income to workers who were already planning to work without increasing total labor supply or enhancing service quality. From a labor-policy perspective, distinguishing activation effects from transfer effects helps clarify whether incentive programs improve allocative efficiency or merely redistribute income. This distinction is significant for ongoing regulatory discussions regarding platform compensation practices, minimum-earnings standards, and worker classification.

Open methodology. Since the model relies solely on synthetic instances with parameter ranges supported by publicly available evidence, the analysis can be reproduced and extended by any researcher without requiring access to proprietary platform data. This is important for regulatory and academic purposes, as the platforms holding the most relevant data also have strong incentives to present their incentive programs in a favorable light. An independent analytical framework, grounded in publicly verifiable evidence, provides a necessary counterbalance.

7.3. Economic Conclusions

Beyond operational recommendations, the model and experiments yield five economic conclusions about how gig delivery labor markets function.

Incentive programs face an adverse selection problem that worsens with platform maturity. When a platform posts a zone bonus, it cannot condition payment on whether the worker would have participated without it. Always-on workers collect the bonus as a pure windfall. This is a classic screening failure: the platform wants to pay only responsive workers but cannot distinguish types ex ante. The lifecycle implication is that as a platform matures, its workforce composition shifts—early responsive adopters become habituated and transition into always-on behavior, causing α to drift upward over time. The net-return analysis (Figure 1) therefore describes not only cross-sectional variation but a temporal trajectory: a bonus program that was cost-effective when the platform was young (α low, windfall share near 0.18) becomes progressively wasteful as the workforce matures (α high, windfall share near 0.93). This implies that platforms may benefit from systematically reducing incentive spending as their workforce stabilizes—a pattern that runs counter to the common industry practice of increasing spending to combat churn [31, 37].

The binding constraint on incentive effectiveness is the size of the responsive margin, not the budget. The budget-sensitivity experiment reveals that doubling the incentive budget does not proportionally increase activation; the optimizer converges to essentially the same allocation regardless of budget level. The binding constraint is the responsive pool. Economically, this means that gig delivery labor markets exhibit a form of supply inelasticity at the margin that is invisible to aggregate analysis. The aggregate supply curve may appear elastic—more spending correlates with more workers—but this correlation is driven largely by always-on workers redistributing across zones in response to differential bonuses, not by new workers entering the market. Platforms that interpret aggregate supply response as evidence that bonuses "work" are confounding zone-switching by infra-marginal workers with genuine activation of marginal workers.

Spatial competition among workers imposes a natural ceiling on incentive returns. When a bonus attracts workers to a zone, the resulting congestion reduces per-worker earnings, partially undoing the incentive effect. Under high congestion, the feedback is strong enough that additional spending in a zone can reduce per-worker welfare while simultaneously increasing platform costs. The economic implication is that platforms impose a negative externality on their own workforce through incentive design: a bonus in zone A pulls workers from zone B, increasing congestion in A and reducing it in B. The net welfare effect depends on the demand structure [14]. When demand is concentrated, the externality is manageable because workers move toward genuine coverage gaps. When demand is balanced, bonuses primarily shuffle workers between zones, creating congestion without coverage gains. Standard models that treat each worker's participation decision as independent systematically overestimate bonus effectiveness by ignoring this coupling.

Uniform incentive programs effect an implicit cross-subsidy from high-demand to low-demand zones. When the platform offers uniform bonuses, it pays the same rate per worker in every zone. In high-demand zones with inadequate coverage, this spending generates productive activation. In low-demand zones where supply already meets demand, the spending is a pure transfer. The demand-asymmetry experiment shows that the Stackelberg optimizer eliminates this cross-subsidy by setting low-demand zone bonuses to zero, and the resulting 21% coverage improvement under high asymmetry comes entirely from redirecting wasted spending, not from spending more. From a distributional perspective, a shift from uniform to targeted bonuses would reduce income for workers in low-demand zones while improving delivery reliability in high-demand zones—a tradeoff between worker income support and allocative efficiency that policymakers should consider explicitly.

There exists a regime in which the welfare-maximizing policy is zero incentive spending. Proposition 4 establishes that above a composition threshold, the platform should spend nothing on bonuses. For mature platforms with $\alpha > 0.50$, the windfall share exceeds 0.75, meaning more than three-quarters of every bonus dollar is a pure transfer [40]. Once workers come to expect bonuses, however, removing them risks short-term supply disruption even though the long-run equilibrium without bonuses may be nearly identical (since always-on workers participate regardless). This creates a ratchet effect: bonuses are easy to introduce but costly to withdraw, and over time the workforce composition shifts in a direction that makes them less effective while the expectation of receiving them becomes entrenched. Understanding this dynamic is essential for platforms considering the long-run consequences of their incentive strategies and for regulators evaluating whether platform bonus programs serve a genuine labor-market function or have become self-perpetuating transfers.

8. Conclusion

We developed a Stackelberg framework for incentive design in gig delivery networks, accounting for worker heterogeneity and congestion externalities. The model categorizes workers into always-on, responsive, and dormant types, integrates their decision-making into a logit congestion game, and connects the resulting equilibrium to a platform profit-maximization problem through a bilevel program. A windfall decomposition was introduced to differentiate productive incentive spending from infra-marginal transfers, providing a nuanced understanding of incentive allocation.

The analytical findings reveal several key insights. First, there exists a workforce-composition threshold beyond which no bonus spending is optimal, as demonstrated in Proposition 4. Second, optimal incentive allocation tends to concentrate in the highest-demand zone when demand asymmetry is significant, as shown in Proposition 5. Third, windfall intensity increases monotonically with the share of always-on workers and is further amplified by congestion, as outlined in Proposition 3. Computational experiments on synthetic datasets with individual-level worker heterogeneity validate these theoretical predictions. For instance, in a baseline five-zone scenario, Stackelberg-optimal targeting improved platform profit by approximately 44% compared to a no-bonus baseline, while uniform bonus programs reduced profit by roughly 35%. These results, while instance-dependent, underscore the structural advantages of anticipatory optimization over uniform incentive strategies.

The central conclusion is that the effectiveness of incentive spending in gig delivery networks fundamentally depends on the composition of the workforce rather than the total expenditure. This compositional dependence is overlooked in models that treat the workforce as a homogeneous entity with a single supply elasticity and is underappreciated in current platform practices. Recognizing and leveraging this heterogeneity is critical for designing effective incentive programs.

Several limitations of the current study suggest directions for future research. First, the single-period model does not account for dynamic effects, such as workers learning about the platform's incentive strategies over time or becoming habituated to elevated compensation levels. Extending the model to a multi-period framework that incorporates worker learning would provide a more comprehensive understanding. Second, the analysis focuses on a single platform, whereas real-world markets often feature competing delivery platforms drawing from overlapping worker pools. In such contexts, incentive design becomes a multi-leader game with a more complex equilibrium structure. Third, while the synthetic instances used in this study are parameterized based on published evidence, validation against observed platform outcomes would enhance the empirical robustness of the findings. Establishing partnerships with platforms to enable such validation is a priority for future work.

9. Proofs

9.1. Proof of Proposition 1

Define the mapping $\Phi: [0, N]^{ZT} \to [0, N]^{ZT}$ by $\Phi(\boldsymbol{s})_{zt} = \sum_{k \in \mathcal{K}} N_k \, P_{zt}^k(U^k(b, s))$. Under Assumption 1, Γ_{zt} is continuous, so U_{zt}^k is continuous in $\boldsymbol{s}$. The logit probability (2) is a smooth function of utilities, so Φ is continuous. The domain $[0, N]^{ZT}$ is compact and convex. By Brouwer's fixed-point theorem, Φ has at least one fixed point, which constitutes an equilibrium.

9.2. Proof of Proposition 2

Consider an increase in b_{zt} holding other components fixed. In the logit model, U_{zt}^R increases, which raises $\exp(\mu U_{zt}^R)$ and thus P_{zt}^R, while the outside-option utility is unchanged. The congestion effect partially offsets this: more workers in zone z raises Γ_{zt}, reducing utilities. However, the type- R participation rate $1 - P_0^R = 1 - (1 + \sum_{z',t'} \exp(\mu U_{z't'}^R))^{-1}$ depends on the *best* available alternative. Since b_{zt} has increased, the maximum attainable utility for type R has weakly increased, so total participation weakly increases even after congestion adjustment.

9.3. Proof of Proposition 3

Part (a). Increasing α (holding N fixed) raises N_A and reduces N_R. In the numerator of (5), x_{zt}^A increases because the always-on population grows. In the denominator, $s_{zt} = x_{zt}^A + x_{zt}^R$; the always-on increase exceeds the responsive decrease because type-A workers have lower reservation utility and are less congestion-sensitive. Therefore the ratio increases.

Part (b). Higher γ_{zt} reduces the net utility for all workers, but disproportionately affects type R (whose participation is marginal). At equilibrium, higher congestion reduces x_{zt}^R more than x_{zt}^A, shifting supply composition toward always-on workers and raising the windfall share.

9.4. Proof of Proposition 4

We establish global optimality of $b^\star = 0$ in the two-zone specialization. Define $\Pi(b_1, b_2)$ as the platform objective with equilibrium supply substituted: $\Pi(b) = \sum_z [R_z \lambda_z(s_z(b), d_z) - \bar{w}_z \lambda_z - b_z s_z(b) - C^u(d_z - \lambda_z)^+]$.

Step 1: Marginal analysis at $b = 0$. Consider increasing b_1 by $\varepsilon > 0$. The marginal benefit is additional coverage revenue: $\partial \Pi / \partial b_1|_{b=0} = (R_1 - \bar{w}_1 + C^u) \cdot \partial s_1 / \partial b_1 - s_1$, where the first term is the marginal value of additional supply (positive only when $s_1 <$

d_1) and the second is the bonus cost to all workers already in zone 1. When α is large, s_1 is large and $\partial s_1/\partial b_1$ is small (the responsive pool is thin), so $\partial\Pi/\partial b_1|_{b=0} < 0$.

Step 2: Concavity in the two-zone case. Under linear congestion ($\Gamma_z = \gamma s_z$) and the logit specification, the composite objective $\Pi(b_1, b_2)$ is concave in b for the two-zone, two-type specialization. To see this, note that $s_z(b)$ is a smooth, concave function of b_z (the logit response saturates as bonuses increase, and congestion feedback further dampens the supply response). The revenue component $R_z \min(s_z, d_z)$ is concave in s_z (it is piecewise linear with a non-positive slope change at $s_z = d_z$). The bonus cost $b_z s_z(b)$ is convex in b_z (product of b_z and the increasing, concave function s_z). Therefore Π is the sum of concave terms minus convex terms, yielding concavity.

Step 3: Global optimality. Concavity of Π ensures that any local maximum is global. Since $\partial\Pi/\partial b_z|_{b=0} < 0$ for both zones when $\alpha > \alpha^\star$, and the feasible set $b \geq 0$ is convex, the point $b = 0$ is a global maximum. Setting $\partial\Pi/\partial b_1|_{b=0} = 0$ and solving for α yields the threshold (7).

9.5. Proof of Proposition 5

At an interior solution, first-order conditions require equal marginal returns per bonus dollar across zones. The marginal return in zone z depends on the gap $d_z - s_z$: positive gap yields coverage revenue; zero or negative gap yields zero revenue but full bonus cost. When $|d_1 - d_2|$ exceeds $\delta^\star$, the low-demand zone has $s_z \geq d_z$ at the optimum, so its marginal return is negative. The optimizer sets $b_2 = 0$, concentrating all spending in zone 1.

References

1. M. K. Chen, P. E. Rossi, J. A. Chevalier, and E. Oehlsen, "The value of flexible work: Evidence from Uber drivers," *Journal of Political Economy*, vol. 127, no. 6, pp. 2735–2794, 2019.
2. G. P. Cachon, K. M. Daniels, and R. Lobel, "The role of surge pricing on a service platform with self-scheduling capacity," *Manufacturing & Service Operations Management*, vol. 19, no. 3, pp. 368–384, 2017.
3. M. Hu and Y. Zhou, "Dynamic type matching," *Manufacturing & Service Operations Management*, vol. 24, no. 1, pp. 125–142, 2022.
4. J. D. Angrist, S. Caldwell, and J. V. Hall, "Uber versus taxi: A driver's eye view," *American Economic Journal: Applied Economics*, vol. 13, no. 3, pp. 272–308, 2021.
5. H. Guda and U. Subramanian, "Your uber is arriving: Managing on-demand workers through surge pricing, forecast communication, and worker incentives," *Management Science*, vol. 65, no. 5, pp. 1995–2014, 2019.
6. A. Alnaggar, F. Gzara, and J. H. Bookbinder, "Compensation guarantees in crowdsourced delivery: Impact on platform and driver welfare," *Omega*, vol. 122, p. 102965, 2024.
7. W. Sutherland and M. H. Jarrahi, "The sharing economy and digital platforms: A review and research agenda," *International Journal of Information Management*, vol. 43, pp. 328–341, 2018.
8. O. Candogan, K. Bimpikis, and A. Ozdaglar, "Optimal pricing in networks with externalities," *Operations Research*, vol. 60, no. 4, pp. 883–905, 2012.
9. J. V. Hall and A. B. Krueger, "An analysis of the labor market for Uber's driver-partners in the United States," *ILR Review*, vol. 71, no. 3, pp. 705–732, 2018.
10. Y. Dai, S. He, and Z. Ma, "Pricing decisions for an on-demand crowdsourced delivery platform with differentiated service levels," *IEEE Transactions on Engineering Management*, vol. 71, pp. 14350–14364, 2024.
11. T. A. Taylor, "On-demand service platforms," *Manufacturing & Service Operations Management*, vol. 20, no. 4, pp. 704–720, 2018.
12. H. S. Farber, "Why you can't find a taxi in the rain and other labor supply lessons from cab drivers," *The Quarterly Journal of Economics*, vol. 130, no. 4, pp. 1975–2026, 2015.
13. C. Cook, R. Diamond, J. V. Hall, J. A. List, and P. Oyer, "The gender earnings gap in the gig economy: Evidence from over a million rideshare drivers," *The Review of Economic Studies*, vol. 88, no. 5, pp. 2210–2238, 2021.
14. M. Labbé, P. Marcotte, and G. Savard, "A bilevel model of taxation and its application to optimal highway pricing," *Management Science*, vol. 44, no. 12-part-1, pp. 1608–1622, 1998.
15. M. W. Ulmer and M. Savelsbergh, "Workforce scheduling in the era of crowdsourced delivery," *Transportation Science*, vol. 54, no. 4, pp. 1113–1133, 2020.
16. L. Brotcorne, M. Labbé, P. Marcotte, and G. Savard, "A bilevel model for toll optimization on a multicommodity transportation network," *Transportation Science*, vol. 35, no. 4, pp. 345–358, 2001.
17. K. Baltz, *Behavioral Game Theory. Experiments in Strategic Interaction*, 2004.
18. Z. Q. Luo, J. S. Pang, and D. Ralph, *Mathematical Programs with Equilibrium Constraints*, Cambridge University Press, 1996.

19. K. Bimpikis, O. Candogan, and D. Saban, "Spatial pricing in ride-sharing networks," *Operations Research*, vol. 67, no. 3, pp. 744–769, 2019.
20. S. Lawphongpanich and D. W. Hearn, "An MPEC approach to second-best toll pricing," *Mathematical Programming*, vol. 101, no. 1, pp. 33–55, 2004.
21. H. Xiao, M. Xu, and S. Wang, "A game-theoretic model for crowd-shipping operations with profit improvement strategies," *International Journal of Production Economics*, vol. 262, p. 108914, 2023.
22. S. Dempe, *Foundations of Bilevel Programming*, Springer US, Boston, MA, 2002.
23. T. L. Magnanti and R. T. Wong, "Network design and transportation planning: Models and algorithms," *Transportation Science*, vol. 18, no. 1, pp. 1–55, 1984.
24. B. Colson, P. Marcotte, and G. Savard, "An overview of bilevel optimization," *Annals of Operations Research*, vol. 153, no. 1, pp. 235–256, 2007.
25. J. Luy, G. Hiermann, and M. Schiffer, "Strategic workforce planning in crowdsourced delivery with hybrid driver fleets," *Production and Operations Management*, vol. 33, no. 11, pp. 2177–2200, 2024.
26. Z. S. Qian and H. M. Zhang, "Modeling multi-modal morning commute in a one-to-one corridor network," *Transportation Research Part C: Emerging Technologies*, vol. 19, no. 2, pp. 254–269, 2011.
27. F. He, Y. Yin, and S. Lawphongpanich, "Network equilibrium models with battery electric vehicles," *Transportation Research Part B: Methodological*, vol. 67, pp. 306–319, 2014.
28. A. Behrendt, M. Savelsbergh, and H. Wang, "A prescriptive machine learning method for courier scheduling on crowdsourced delivery platforms," *Transportation Science*, vol. 57, no. 4, pp. 889–907, 2023.
29. I. Constantin and M. Florian, "Optimizing frequencies in a transit network: A nonlinear bi-level programming approach," *International Transactions in Operational Research*, vol. 2, no. 2, pp. 149–164, 1995.
30. A. Migdalas, P. M. Pardalos, and P. Värbrand, Eds., *Multilevel Optimization: Algorithms and Applications*, vol. 20, Springer Science & Business Media, 2013.
31. A. Behrendt, M. Savelsbergh, and H. Wang, "Task assignment, pricing, and capacity planning for a hybrid fleet of centralized and decentralized couriers," *Transportation Research Part C: Emerging Technologies*, vol. 160, p. 104533, 2024.
32. N. Kafle, B. Zou, and J. Lin, "Design and modeling of a crowdsource-enabled system for urban parcel relay and delivery," *Transportation Research Part B: Methodological*, vol. 99, pp. 62–82, 2017.
33. C. Cheng, M. Sim, and Y. Zhao, "Robust workforce management with crowdsourced delivery," *Operations Research*, vol. 73, no. 2, pp. 595–612, 2025.
34. I. Dayarian and M. Savelsbergh, "Crowdshipping and same-day delivery: Employing in-store customers to deliver online orders," *Production and Operations Management*, vol. 29, no. 9, pp. 2153–2174, 2020.
35. C. Archetti, M. Savelsbergh, and M. G. Speranza, "The vehicle routing problem with occasional drivers," *European Journal of Operational Research*, vol. 254, no. 2, pp. 472–480, 2016.
36. R. W. Rosenthal, "A class of games possessing pure-strategy Nash equilibria," *International Journal of Game Theory*, vol. 2, no. 1, pp. 65–67, 1973.
37. M. Savelsbergh and M. W. Ulmer, "Challenges and opportunities in crowdsourced delivery planning and operations—an update," *Annals of Operations Research*, vol. 343, no. 2, pp. 639–661, 2024.
38. D. Merchán, J. Arora, J. Pachon, K. Konduri, M. Winkenbach, S. Parks, and J. Noszek, "2021 Amazon last mile routing research challenge: Data set," *Transportation Science*, vol. 58, no. 1, pp. 8–11, 2024.
39. D. Monderer and L. S. Shapley, "Potential games," *Games and Economic Behavior*, vol. 14, no. 1, pp. 124–143, 1996.
40. R. D. McKelvey and T. R. Palfrey, "Quantal response equilibria for normal form games," *Games and Economic Behavior*, vol. 10, no. 1, pp. 6–38, 1995.