ORIGINAL ARTICLE

Estimating Heterogeneous Treatment Effects of Driver Incentives in Ride-hailing Platforms

Nihar Vipulkumar Patel

ABSTRACT

Ride-hailing and delivery platforms rely on a mix of monetary and non-monetary incentives to dynamically balance supply and demand in two-sided marketplaces. However, the average effect of any single incentive can mask substantial heterogeneity in behavioral responses across drivers and riders. This paper provides a methodological roadmap for estimating heterogeneous treatment effects of driver incentives in such an operational environment characterized by dynamic behavior, endogenous matching, and pervasive interference among agents. The roadmap formalizes causal estimands that capture direct effects on treated drivers, spillover effects on other drivers, and two-sided equilibrium effects including rider outcomes. Identification strategies are discussed for settings constrained by limited experimental control with an emphasis on approaches combining randomized variation-such as saturation and clustered randomization designs-with flexible, robust estimators, including meta-learners, causal forests, and doubly robust methods. Building on these heterogeneity estimates, a policy learning layer is proposed that translates HTEs into operationally constrained incentive assignments, subject to budgetary limits, fairness criteria, and service level guardrails. Validation procedures are outlined using held-out geographic regions and simulation-based stress testing to assess policy performance in a wide range of market conditions. The proposed framework is suited for capturing realistic heterogeneity across drivers-for example, full-time, part-time, students, retirees, newly onboarded, or experienced-and riders-for example, commuters, nighttime users, or occasional travelers-and for accounting for time-varying demand states and geographic frictions intrinsic to urban mobility markets.

Keywords: Causal Inference and Policy Learning, Driver Incentives, Heterogeneous Treatment Effects (HTE) , Ride-Hailing Platform Dynamics, Two-Sided Marketplaces

1 | Introduction

The ride-hailing platforms run a complex, two-sided marketplace where they match riders in need of transportation with drivers willing to provide that transportation . Demand for rides surges or drops quickly according to time of day, location, weather, and events. Incentive mechanisms-which include surge pricing multipliers, quest bonuses for completing a series of trips, guaranteed earnings floors, streak bonuses for consecutive rides, and personalized nudges-are what help balance this supply and demand at various times and places. In short, these are key levers for aligning supply with spiky, localized demand. The response to a given incentive can vary dramatically, however. The same incentive program may induce one driver to work additional hours while leaving another driver unmoved. For instance, a high surge multiplier available on a Friday night from 10 PM to 2 AM might be crucial in attracting drivers to nightlife areas, while the same multiplier would be superfluous for the Friday afternoon commute, when many drivers are already out. Or, a guaranteed earnings offer per hour might powerfully incentivize a new or part-time driver to come online-by removing some of their downside risk-but that guarantee might be irrelevant to a seasoned full-time driver who already earns above that threshold. Such examples show how incentives uniform in structure can create starkly different behaviors depending on context and individual. [1]

Traditional A/B testing and average treatment effect analyses mask these disparities. A field experiment might conclude that some incentive has, on average, no significant effect on driver activity. For example, such a “zero average effect” could mask large positive effects for some subpopulations-for instance, side-hustle drivers who respond strongly to modest bonuses-being counterbalanced by negative or negligible effects for others, such as in already saturated areas where additional incentives simply pay drivers for what they would do anyway. Aggregate impacts are a misleading basis for platform decisions because the platform might cancel an incentive that actually is crucial for a subset of drivers or continue an expensive program that only works in highly constrained conditions. It is this recognition that moves the focus from the average treatment effects to heterogeneous treatment effects-quantifying how the impact of an incentive varies across drivers, times, and locations.

In this paper, the aim is to estimate and exploit these heterogeneous treatment effects to design smarter driver incentive policies. This paper develops a comprehensive framework that addresses the unique challenges of the ride-hailing context, wherein driver behavior is dynamic over time, drivers and riders are matched in real time, and any one driver’s decision can influence market outcomes for others. First, the paper formalizes the key causal quantities of interest, including the direct effect of an incentive on an individual driver and the indirect spillover effects that occur when multiple drivers are incentivized simultaneously. This framework is extended to consider equilibrium effects on rider outcomes such as wait times and fares that result from the aggregate supply response and consequent shifts in the marketplace. [2] It next discusses identification strategies for such effects under realistic constraints. Randomized experiments in such marketplaces face challenges with interference-one driver’s treatment affects another’s outcome-and endogenous matching. The paper considers how the use of experimental designs can isolate causal effects with such complications, including cluster randomization with varying saturation levels of treatment, staggered rollouts of incentives across regions or times, and encouragement designs where incentive offers serve as instrumental variables. The analysis also considers how observational or quasi-experimental approaches-for instance, natural experiments based on policy changes at geo-boundaries or discontinuities at incentive eligibility thresholds-can add to the evidence. The study then outlines approaches to estimation for learning HTEs from data: modern machine learning techniques adapted for causal inference-which involve meta-learners that estimate conditional outcomes flexibly for treatment and control (T-learners, S-learners, and X-learners), the R-learner formulation that residualizes outcomes and treatments to remove nuisance variation and then learns the treatment effect function, and tree-based methods such as causal trees and causal forests that directly aim to partition the population by treatment effect differences.

Importantly, these methods have been selected for their ability to handle high-dimensional covariates and to produce estimates that are robust or orthogonal to estimation errors in nuisance parts of the model.

Building on the heterogeneous effect estimates, the paper considers how a platform might seek to optimize incentive allocation. The authors suggest a policy learning approach that takes the estimated HTEs as inputs to a decision of which drivers-or which times and places-receive an incentive offer, at what level, so as to maximize some objective such as platform-wide welfare or efficiency. Importantly, this policy optimization is performed under operational

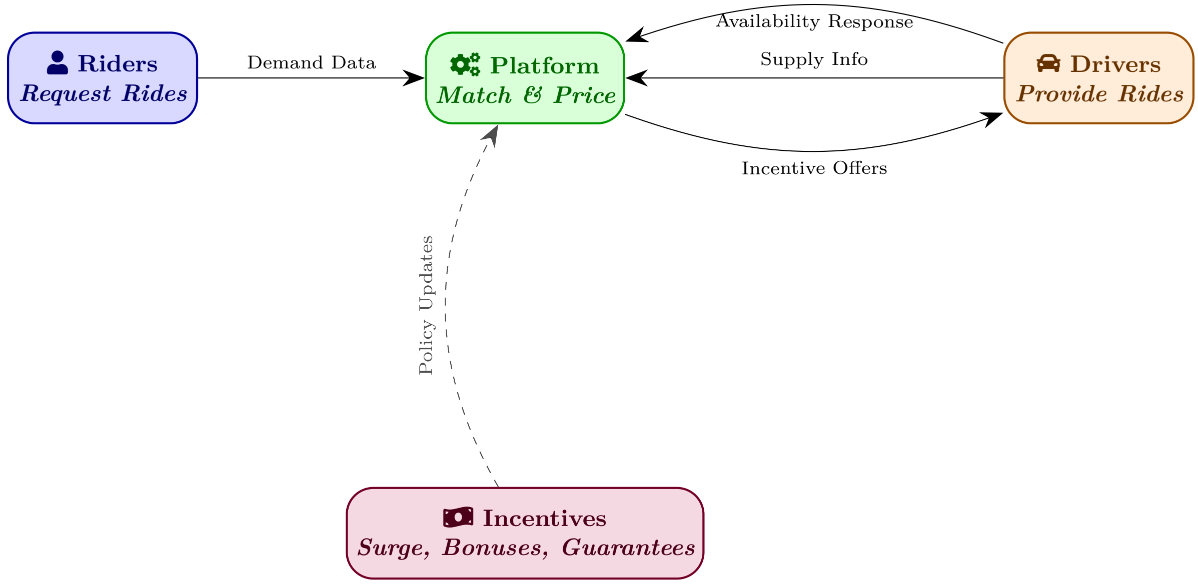

Figure 1: Conceptual overview of the ride-hailing marketplace showing rider demand, driver supply, and platform-mediated incentives.

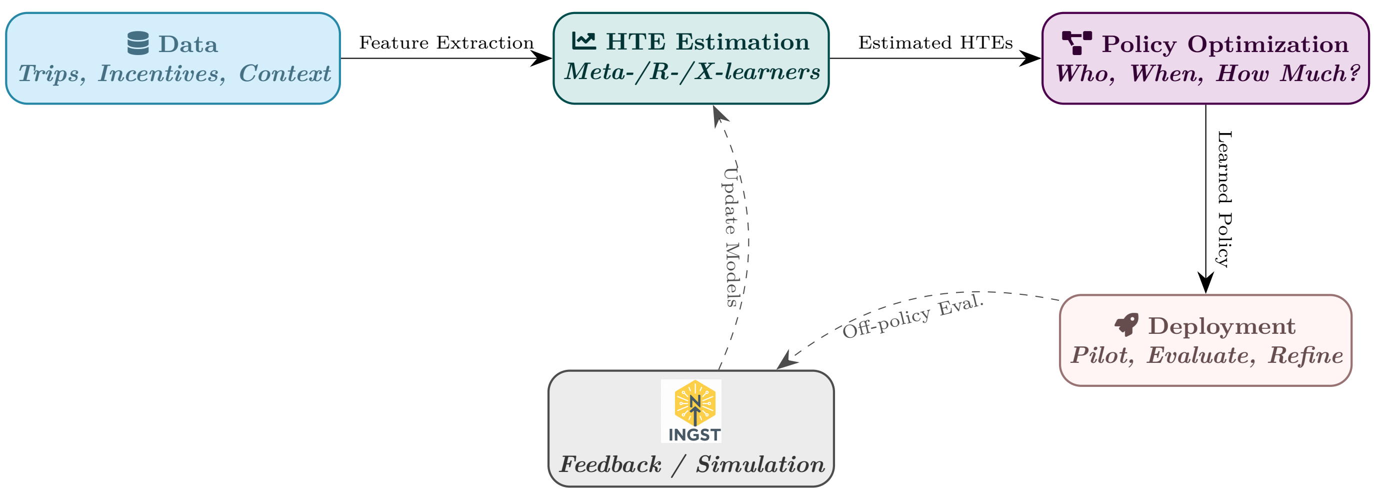

Figure 2: Framework for estimating heterogeneous treatment effects (HTEs) and optimizing incentive policies through iterative learning and deployment.

constraints: the platform has a limited budget for incentives; it needs to keep service quality metrics within desired bounds-e.g., expected wait time for riders below target-and it has to ensure that no group of drivers or riders is unduly disadvantaged by the incentive scheme [3]. The authors discuss how these operational constraints might be integrated into the incentive policy design and what kinds of trade-offs may arise. The study then goes on to detail practical considerations in implementing and validating the proposed approach. An experimental and deployment blueprint-everything from granular data collection to phased rollout of any learned policy by way of shadow testing and small-scale pilots, and continuing with monitoring for unintended consequences-is suggested. Before full deployment, it describes how off-policy evaluation techniques can be used on historical data to estimate the performance of a new incentive policy and how further validation might be provided through discrete-event simulation or other modeling of the marketplace by capturing complex feedback dynamics in a controlled setting. By explicitly accounting for heterogeneity, this work targets incentive strategies that are more efficient, equitable, and robust to ebbs and flows of a real-world ride-hailing market.

2 | Context and Institutional Features

Before describing the methodology in detail, it is useful to consider the environment in which these incentive policies operate. The ride-hailing marketplace is a high-frequency, spatially granular environment; conditions can change from minute to minute and from block to block. A platform has to constantly balance driver supply with rider demand within a city, and incentives are deployed in this environment to influence behaviors in real time. A number of features of this marketplace create scope for heterogeneous effects of incentives: [4]
One critical dimension is the heterogeneity across drivers. The supply of drivers is not homogeneous: it includes full-time drivers who treat ride-hailing as their main source of income, part-time "side hustle" drivers who log on during occasional free hours, students who may drive only during weekends or vacations, retirees seeking supplemental income, and rookies versus veterans of the platform. Drivers vary by vehicle type, which may preclude them from accepting certain types of trips; by their historical reliability and acceptance rates; and by their typical preferred working hours or locations. Owing to these differences, the same incentive can have different relevance and appeal. For example, a quest bonus that rewards completing 20 trips in a week may strongly motivate a part-time driver seeking to make the most of her limited time, but a full-time driver may complete 20 trips with ease as a regular part of her weekly routine and thus find the quest less of a lure. A surge multiplier in a distant neighborhood may be enticing to a driver willing to reposition for higher pay but irrelevant to one who is low on fuel or prefers to stick to familiar areas. In short, driver heterogeneity in availability, preferences, and constraints generates heterogeneity in response to incentives.
Riders are heterogeneous, too, and their heterogeneity indirectly affects the efficiency of driver incentives. Rider demand comes in flavors: there are commuting users who take rides during weekday mornings and evenings with relatively predictable patterns, nightlife or entertainment-oriented users who create late-night and weekend surges in certain districts, and sporadic users who request rides only occasionally [5]. The spatial and temporal distribution of these rider segments implies that at any given time, the composition of demand might vary. At any given moment-for instance, 8 AM on a Tuesday-demand is largely composed of work commuters requesting standard routes, whereas at midnight on a Saturday, the demand may cluster around bars or event venues with riders who may be more willing to wait or pay a premium if supply is short. Driver incentives may need to be stronger to attract drivers into the late-night entertainment district than they are normally during a typical weekday commute rush because the alternative for drivers may be more enticing in the former scenario. Moreover, the sensitivity of the riders to price and wait times can differ; a commuter would not ask for a ride if the prices surge too high, but nightlife riders may be more accepting of the surge price for lack of alternatives. Thus, the rider mix and demand context-which drivers will indirectly respond to through pricing or perceived need-contributes to heterogeneous incentive effects.
Another important factor is the variety of incentive instruments themselves. Platforms use a menu of incentive types, each with its own structure and intended effect. Price-linked incentives like surge pricing or boost multipliers directly raise the payout for trips in certain areas or times, effectively a real-time price adjustment intended to attract drivers to high-demand zones. Quantity-linked incentives such as quest bonuses ("complete X trips to earn $Y extra") or streak bonuses ("complete N trips in a row without rejecting any to earn a reward") encourage drivers to commit to a certain volume or pattern of work [6]. Insurance-style incentives like guaranteed hourly earnings ensure drivers a minimum income rate during a period (useful to assure new drivers or those concerned about low demand that their time will be compensated). Finally, non-monetary and behavioral nudges like notifications ("demand is high in your area, consider going online") or heat maps that highlight busy zones can subtly influence driver behavior without direct payments. Each type of incentive may resonate differently with drivers. For example, a risk-averse driver might value a guaranteed earning offer more than a chance at a high surge, whereas a risk-seeking driver might prefer the opposite. Some drivers might be motivated by game-like quest challenges, while others ignore them. The heterogeneity in driver psychology and circumstances means no single incentive type universally outperforms the others; it depends on the match between incentive design and driver type.
Finally, any incentive scheme must operate subject to practical operational constraints and objectives. The platform has a limited budget for incentives and thus cannot simply offer maximum bonuses to everyone at

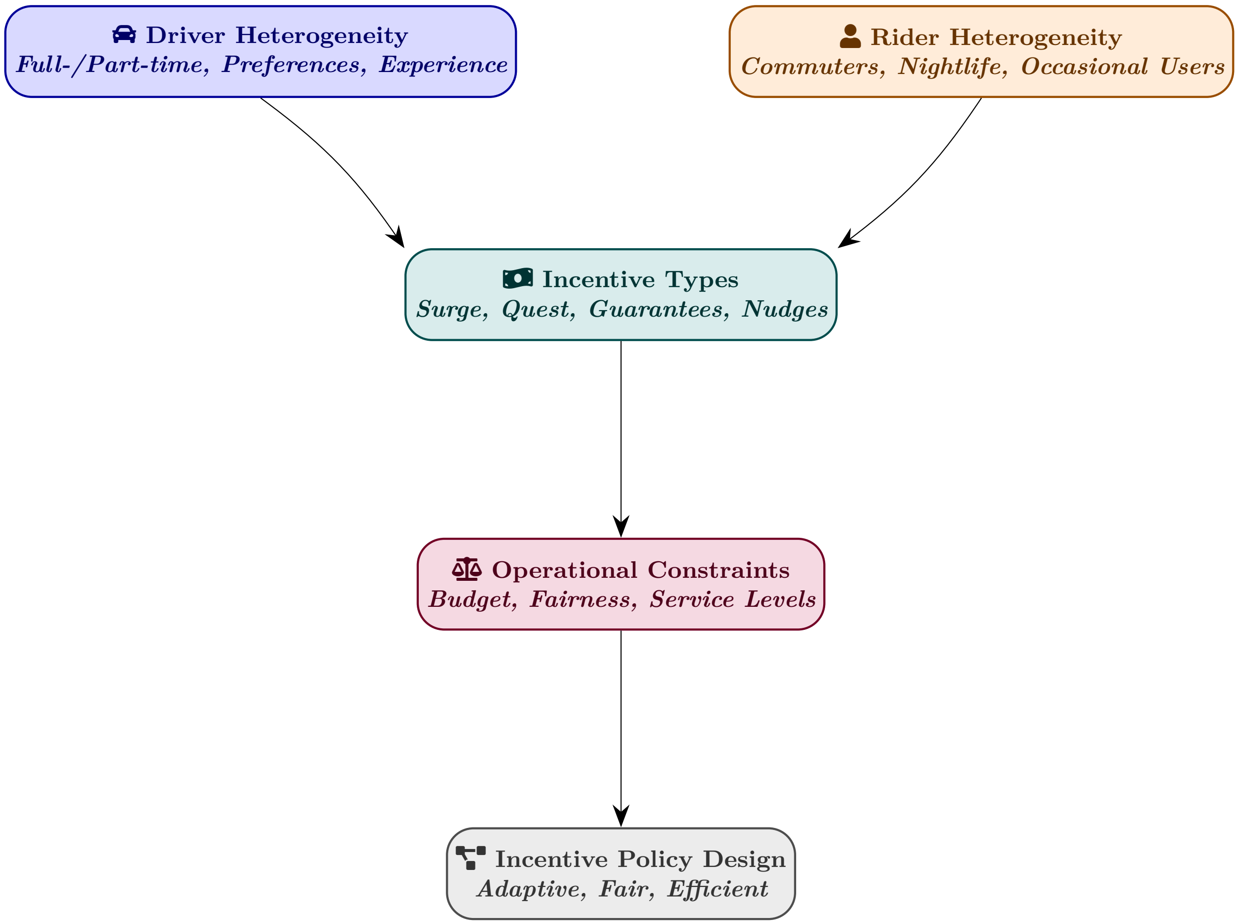

Figure 3: Institutional and contextual sources of heterogeneity in ride-hailing incentive effects. Differences across drivers, riders, incentive types, and platform constraints collectively shape how incentive policies perform.

all times. There are fairness and compliance considerations: regulators or internal policies may require that opportunities-or earnings-are equitably distributed, preventing the platform from concentrating all incentives in one area, or subgroup, for extended periods [7]. Service-level targets also impose constraints: the platform must hit certain key performance metrics-like reasonable estimated time of arrival (ETA) for riders, low cancellation rates, and high fulfillment rates. Over-incentivizing in one region could lead to driver oversupply and idleness there while another region is left underserved, harming overall service reliability. Marketplace health metrics-like driver utilization (fraction of time drivers have passengers) and rider abandonment (riders giving up after waiting too long)-are closely watched. An effective incentive system thus needs to navigate these constraints: spending money efficiently, maintaining fairness or compliance, and achieving service level goals, all while leveraging heterogeneity to improve outcomes. These institutional features set the stage for the causal questions and policy design challenges we address in the rest of the paper.

3 | Conceptual Framework

3.1 Notation and Estimands

This study formalizes the causal framework that underlies the incentive analyses of the marketplace. Let's consider a set of drivers indexed by i, time periods indexed by t (for example, discrete intervals like hours or quarters of an hour), and geographic areas indexed by g - for example, defined zones or grid cells in the city. Let $A_{it} \in \{0, 1\}$ be an indicator of whether driver i is offered or receives a given incentive at time t - with $A_{it} = 1$ meaning the incentive is applied and $A_{it} = 0$ meaning no incentive for that driver at that time. We focus on a binary treatment for simplicity, though in practice incentives may have multiple levels

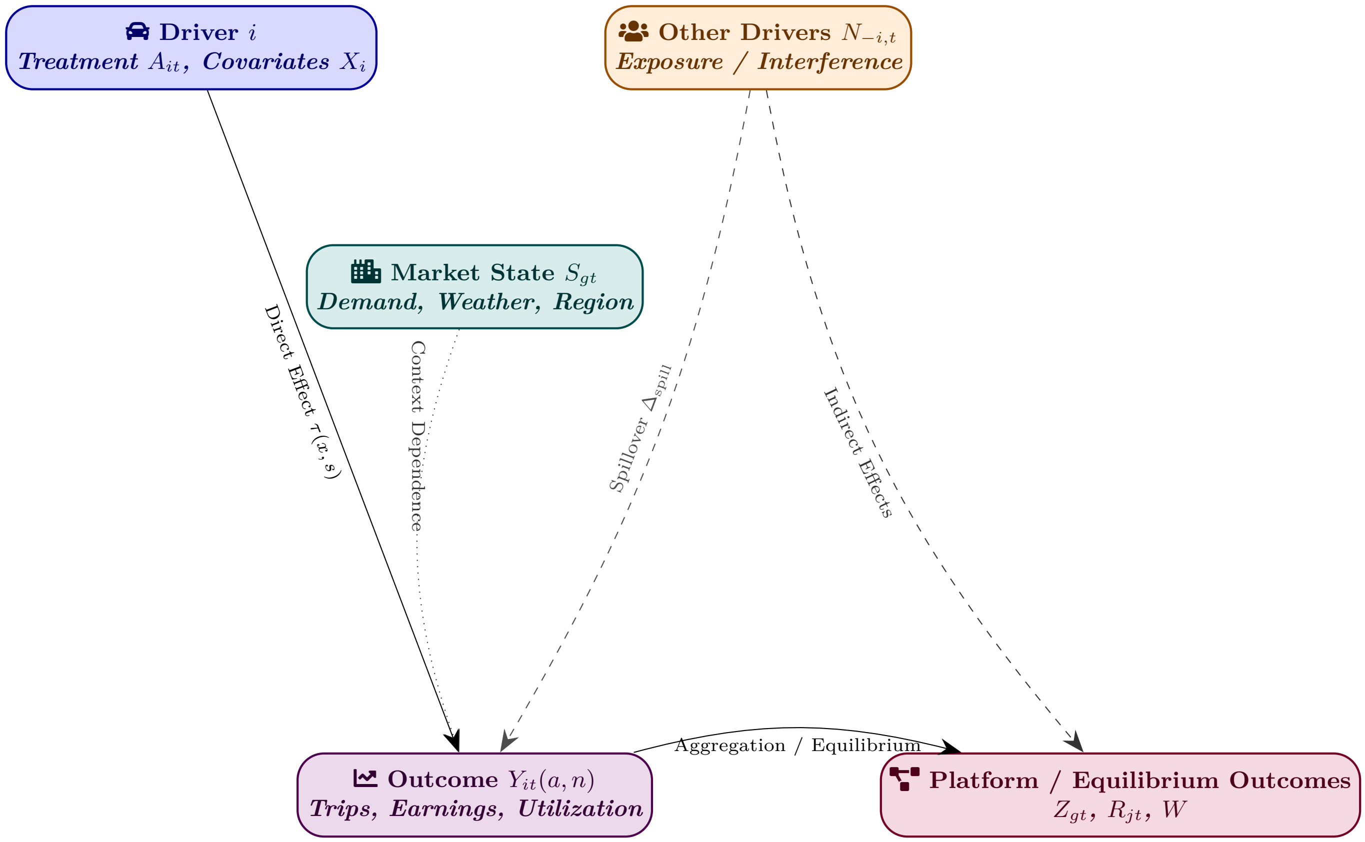

Figure 4: Conceptual causal structure of driver incentive effects. Each driver's outcome $Y_{it}(a,n)$ depends on their own treatment A_{it} (direct effect), the treatment configuration of others $N_{-i,t}$ (spillover effect), and the market state S_{gt} (context). Aggregated across areas and times, these micro-level effects shape platform- and rider-level equilibrium outcomes (Z_{gt}, R_{jt}) and overall welfare W.

or varieties. Let X_i denote the covariates associated with each driver, capturing individual characteristics, such as tenure on the platform, average number of hours driven per week, typical driving locations or times, past responsiveness to incentives, and so on. [8]
We denote by S_{gt} the state of the market in area g at time t. This may include features such as current ride demand intensity, local traffic or weather conditions, the distribution of rider types (commuters vs. leisure travelers), or the current price multiplier in effect.
Let $N_{-i,t}$ denote the vector of incentive assignments to all other drivers at time t (so it collects A_{jt} for all j eqi). This term captures the notion of treatment "exposure" or interference/ externalities: $N_{-i,t}$ describes how many or which other drivers are treated around i, which can impact i's outcomes - for example, if lots of other drivers around them also got incentives, the competitive environment for rides is different than if i alone was treated.
In this potential outcomes framing, we write $Y_{it}(a,n)$ for the outcome of interest for driver i at time t under a scenario where $A_{it} = a$ for that driver and $N_{-i,t} = n$ for the vector of others. The outcome Y_{it} could be a measure of driver i's engagement or success - such as the number of trips completed in that period, total earnings in that period, the fraction of time spent with a passenger (utilization), or even longer-term outcomes like whether driver i is still active on the platform the next week.
By focusing on $Y_{it}(a,n)$, there is recognition that the outcome of driver i could depend not only on his or her own treatment status, a, but also on the entire configuration, n, of others' treatments. First, the individual treatment effect - or direct effect - for a driver with characteristics $X_i = x$ in market state $S_{gt} = s$ is defined as the difference in that driver's expected outcome when they are treated versus not treated, holding the environment at a reference condition. Formally, we define:

$$\tau(x,s) = \mathbb{E}[\, Y_{it}(1,n^\star) - Y_{it}(0,n^\star) \,|\, X_i = x,\ S_{gt} = s]$$

where $n^\star$ represents a reference scenario for the

treatments of other drivers, for example a baseline fraction of other drivers receiving incentives, or perhaps $n^\star$ is a vector of zeros meaning we consider the effect when no one else is treated.
The function $\tau(x, s)$ is the CATE for a driver with profile x in context s. It captures just the direct effect of providing that driver with the incentive, abstracting from general market effects by conditioning on a fixed background n^*: Next comes the concern for spillover effects or externalities: Other drivers being treated can influence a given driver's outcome. We can formalize an "exposure–response" curve. For example, comparing two situations: one in which a proportion of other drivers in the area, π_1, are treated and another in which a larger proportion π_2 are treated, with driver i's own treatment status held fixed in both situations. We define the spillover effect on driver i when the treatment saturation changes from π_1 to π_2 as: [9]

$$\begin{aligned} \Delta_{\text{spill}}(x, s; \pi_1, \pi_2) = \mathbb{E}\Big[Y_{it}\big(a, \quad & (1) \\ N_{-i,t} \sim \pi_2\big) - Y_{it}\big(a, \quad & (2) \\ N_{-i,t} \sim \pi_1\big) \quad & (3) \\ \Big|\, X_i = x,\ S_{gt} = s\Big] \quad & (4) \end{aligned}$$

where notation $N_{-i,t} \sim \pi$ indicates that the other drivers' treatments are assigned according to some scheme with overall rate π – for instance, each other driver independently has probability π of receiving the incentive. That is, Δspill(x,s; 1, 2) provides in words the change in the expected outcome for a driver like x in context s due to increasing the intensity of treatment around them from 1 to 2.
This captures interference effects: if a lot of other drivers receive incentives, for example, this might increase competition for the same rider requests, possibly hurting a particular driver's earnings either because they do not receive the treatment or because they do receive the treatment but share the benefit with many others. On the other hand, broad-based incentives would lead to better overall market efficiency: more drivers online may reduce wait times for riders, leading to more ride requests, benefiting all drivers. Spillover effects may, in general, be positive or negative, and may themselves depend on x, s, and on whether the driver is treated (a) or not. We also broaden our view to outcomes beyond just a single driver. Let Z_{gt} denote platform-level outcomes in area g at time t. These might include metrics such as the average estimated time of arrival (ETA) for riders in that area, the ride fulfillment rate (the fraction of rider requests that get matched with a driver), or aggregate measures such as total platform revenue in that cell and period.
Let R_{jt} be the outcomes for a rider j at time t including, for example, whether rider j got a ride, their waiting time, and price they paid.
Driver incentives can impact these outcomes also through equilibrium adjustments. For example, if many drivers are attracted to an area by an incentive, then Z_{gt} might reflect lower average wait times for riders (a benefit to both riders and the platform), and the R_{jt} for riders in that area might reflect higher likelihood of getting a ride or paying a lower surge price. If incentives result in oversupply in one location, then other locations might get poorer service. [10] We refer to these system-wide effects that propagate through driver-rider matching and pricing as equilibrium effects. With these various outcomes, a platform would normally set an objective function balancing them in designing policies. You can envision a weighted welfare metric that aggregates driver, rider, and platform outcomes. For example, one convenient form is: \$and\$

$$W = \sum_{g,t} \Big(\alpha\, \Delta Z_{gt} + \beta\, \overline{\Delta Y}_{gt} + \gamma\, \overline{\Delta R}_{gt}\Big) \tag{5}$$

In this expression, ΔZ_{gt} denotes the change in a platform-level metric in cell g at time t (compared to a baseline without the incentive). Let $\overline{\Delta Y}_{gt}$ denote the average change in a driver outcome, such as the number of trips completed or earnings, in that cell and time window, and let $\overline{\Delta R}_{gt}$ be the average change in a rider-centric outcome, such as wait time or ride completion probability.
The coefficients α, β, and γ are weights reflecting the platform's priorities. The resulting W aggregates these effects into a single welfare index. Consider an optimal policy design problem where the goal is to maximize W subject to various operational constraints.

3.2 Why Heterogeneity Matters

Having defined these quantities, the study in focus demonstrates the reasons why, within this context, controlling for heterogeneity in treatment effects is essential. [11] From an operational standpoint, targeting incentives can greatly enhance the efficiency of these. The same total of driver bonuses can have a very different outcome depending on where and to whom the money is allocated. For instance, let's say a platform has a fixed pool of funds to give a boost to supply on a Friday night. If those funds are concentrated on drivers who are only marginally interested in driving-say, part-time drivers who otherwise would stay home-in areas and times when rider demand is peaking, such as a downtown entertainment district at midnight, the additional supply brought online can significantly reduce rider shortages and cancellations. On the other hand, if funds were spread thin across drivers or used in areas where supply already exists (or offered to drivers who would drive anyway), the effect might be minimal. In other words, there are diminishing returns in some of the segments and high marginal returns in others. This enables the platform to optimize bang for the buck because identifying where the incentive elasticity-that is, supply's responsiveness to incentives-is highest. Second, heterogeneity is important because of considerations of fairness and compliance. A uniform incentive policy may inadvertently favour particular groups or locales. [12] For example, an incentive structured as a quest that requires many hours of driving may preferentially benefit full-time drivers and exclude others who cannot devote that much time, such as single parents or students. This kind of policy might be subject to equity concerns, or even violate regulations, if it has a disparate impact on a protected class. The understanding of heterogeneity allows the platform to change either incentive structures or targeting so that no particular group is systematically left out or negatively affected. In multi-region platforms, uniform incentives could lead to persistent geographic imbalances, where some neighborhoods always have a surplus of drivers chasing bonuses while others remain underserved.

The heterogeneous approach may help to rotate or modulate incentives so as to maintain more balanced coverage and opportunities. Third, robust performance over different market conditions requires heterogeneity-aware strategies. The effect of an incentive can interact in complex ways with the state of the world S_{gt}. A policy that works well during weekday commuter peaks might fail miserably during weekend nights, and vice versa, because driver motivations and rider behaviors differ. Conversely, a strategy tuned for new drivers may never work for experienced ones. Otherwise, if the platform only looked at average results, it might deploy a policy that seems "okay" on average but actually underperforms in critical scenarios (and perhaps overperforms in trivial ones) [13]. The platform can design incentive programs that remain effective under a range of conditions or can switch tactics depending on the situation by modeling how treatment effects vary with context, including time of day, location, demand level, and driver segment. This results in more resilient operations, and avoids the pitfall of chasing short-term gains at the expense of long-run stability.

4 | Identification in the Presence of Interference and Dynamics

4.1 Randomization with Interference

Designing experiments for driver incentives is complicated because the treatment given to one driver can affect the outcome experienced by another. Technically, the usual Stable Unit Treatment Value Assumption, or SUTVA, does not hold; there is interference between the units. To get unbiased estimates, we need experimental designs that take this interference into account. One powerful strategy involves cluster randomization. Rather than independently randomizing incentives at the individual driver level, the platform can divide up the city or region into clusters-for example, distinct neighborhoods or grid cells-and randomize the incentive conditions at the cluster level. For instance, the platform might assign some cells to a high intensity of the incentive and other cells to a low-or zero-incentive intensity at any given time. Drivers within a cluster would then either mostly all receive the incentive offer or all serve as controls, thereby containing most of the interference within the cluster [14]. We can estimate combined direct and spillover effects with less bias from cross-cluster interference by comparing clusters that received different incentive treatments.

Another approach that builds on clustering is a two-stage randomization or "saturation" design. In the first stage, the experimenter randomizes the overall fraction of drivers to be treated within each cluster-for example, one neighborhood might be assigned a 25% treatment saturation, another 50%, another 0% as a control baseline, and so on. In the second stage, individual drivers within that cluster are randomly assigned to treatment with probability equal to the cluster's target saturation level. The result is that this design allows the estimation of both individual-level effects and spillover effects: by comparing outcomes of treated and control drivers within the same cluster-that is, the same environment-we get an estimate of direct effect, while by comparing clusters with different saturation levels we glean the spillover or indirect effects. It effectively separates the "how many are treated" dimension from the "who is treated" dimension in the experiment.

Another approach is temporal staggering of treatment rollouts. Instead of turning an incentive on in all areas simultaneously, the platform can randomize the starting times (or days) at which different clusters receive a new incentive policy. For instance, one set of cities or regions starts a certain bonus program this week, while others continue without it and only later start it at times that are randomly assigned [15]. This kind of rollout, a form of stepped-wedge or difference-in-differences design, allows identification of the effects while accounting for time trends and reduces the risk of network contagion. Given that drivers across different clusters do interact-for example, by driving across city boundaries-at least the asynchronous timing in staggered rollouts means that not all areas are "treated" at the same time, providing periods even after the program has started where control clusters exist for comparison.

4.2 Encouragement Designs and Partial Compliance

In many experiments, but particularly those with incentives, not everyone offered the treatment actually takes it up. Consider, for example, a driver who is randomly assigned an incentive-meaning the platform sends them a notification of a bonus opportunity-but who does not happen to be online and thus never sees it, or else sees it and ignores it. This can happen because treatment assignment, in this case, involves randomization of an incentive offer rather than some more direct form of treatment administration. The situation described above is called partial compliance or imperfect adherence to the treatment assignment. We handle this distinction by defining two kinds of effects: the intention-to-treat (ITT) effect and the treatment-on-the-treated (TOT) effect. The ITT measures the impact of being assigned to the incentive condition-as opposed to being assigned to the control condition-while the TOT measures the impact of actually receiving or acting on the incentive. To the extent that assignment is random, we can make use of it as an instrumental variable for actual treatment received-meaning, essentially, compare outcomes of those assigned to incentive versus not assigned, and then scale the effect by the take-up rate to infer the effect on compliers-drivers who could and do respond to the incentive-offered [16]. This approach, using an instrumental variable, recovers the causal effect among the subgroup of individuals who can be induced by the offer-compliers-by making use of the randomized assignment as an encouragement. Concretely, this means careful tracking in the experiment design as to who was offered the incentive and who actually acted upon it-for instance, coming online or driving in the target area-so these two layers of effects can be disentangled.

4.3 Observational and Natural Experiments

Randomized experiments are powerful, but sometimes infeasible for all questions. Fortunately, platform data often contain natural experiments or quasi-experiments that can be exploited. One example is to use regression discontinuity (RD) designs around incentive thresholds. Platforms often have rules like "Complete 10 trips to get a bonus"; a driver with 9 trips versus one with 10 trips might be almost identical except that one just earned a bonus and the other did not. If drivers cannot perfectly manipulate being just below or above the cutoff (which can be checked by examining the distribution of drivers around the threshold), comparing their outcomes provides an estimate of the incentive's effect.
There are also geographic boundaries where an incentive program may be rolled out in one area, but not an adjacent one, creating a natural experiment if the areas are otherwise similar. We can use a difference-in-differences approach if we have an incentive policy change at a known time in one region and have a comparable region which can serve as a control; we compare the before-after change in the treated region to the before-after change in the control region [17]. All of these approaches require careful diagnostic checks: we need to check that there are no other confounding changes at the threshold or policy time, and that individuals were not sorting or manipulating around the incentive criteria.

4.4 Accounting for Equilibrium Effects

Broad marketplace changes can be induced by driver incentives; they could change prices - through surge algorithms - and rider behavior - through wait times and availability. These equilibrium effects can be identified only by looking at the level of market outcomes, rather than individual drivers. One approach to capturing such effects is incorporating varied treatment intensity into the experimental design. Rather than simply turning an incentive on or off, the platform can randomize the magnitude of the incentive-or the fraction of drivers receiving it - across comparable units. So for instance, one set of cells might have a high surge multiplier imposed-say 2.0x - while another has a moderate one -1.3x -and another has none -1.0x baseline. We can infer the dose-response relationship of incentive intensity by seeing how outcomes like average wait time, or the rate at which rides get completed, differ across such conditions. This will capture nonlinear effects and thresholds-maybe it takes a certain critical mass of drivers to meaningfully reduce wait times. Another dimension involves analyzing endogenous matching and pricing [18]. If we simply observe that high-incentive periods have lower wait times, it could be because the incentive caused it - or it could be that the algorithm raised prices, which deterred some riders, which then reduced wait times independent of incentive. Randomized assignment of incentives can be used as an instrument for such endogenous variables - like price, or driver availability - in a structural equation. In other words, since the incentive assignment is random, any correlation between that assignment and, say, the surge price can be treated as exogenous; we can then attribute changes in the rider outcomes to the incentive via its effect on price. By instrumenting in this way, we isolate the chain of causality in equilibrium. Finally, the level of analysis should be aligned to be consistent with the level of intervention when measuring equilibrium outcomes. That is, if incentives vary at the 15-minute by neighborhood level, the outcomes need to be aggregated to that spatiotemporal granularity. By creating "market windows" (such as looking at each 15-minute interval in each zone as one observation), we pair the treatment with the relevant outcomes. Aggregation to the window level helps average out some of the noise in individual transactions and matches cause to effect in time. It also makes it possible to include fixed effects or controls at the window level, such as time-of-day or zone-specific trends, which further isolate the effect of interest [19]. In short, careful experimental variation and aligning measurement to that variation allows us to identify not just individual responses but how the entire rider-driver equilibrium shifts under different incentive regimes.

5 | Estimation Approaches for HTE

5.1 Machine Learning Meta-Learners

Estimation of HTEs in a flexible manner requires methodologies that can capture complex functional relationships between incentives, individual characteristics, and contextual factors. An important class of methods involves meta-learners for treatment effects, which extend conventional predictive modeling techniques to estimate conditional causal effects. For instance, a T-learner fits two separate models for outcomes: one that estimates the expected outcome given covariates under treatment, and another estimating the expected outcome under control. Let $\hat{m}_1(x, s)$ be the predicted outcome for a driver with features x in market context s if treated, and $\hat{m}_0(x, s)$ be the corresponding prediction if under control. Then, the estimated treatment effect is given by $\hat{\tau}(x, s) = \hat{m}_1(x, s) - \hat{m}_0(x, s)$, allowing each model to specialize in one treatment condition. On the other hand, the S-learner fits a single model that includes the treatment indicator A as an additional covariate and jointly estimates $\mathbb{E}[Y \mid X = x, S = s, A = a]$ under both treatment states. A related, yet potentially superior, variant called the X-learner has advantages when the size of treatment groups is unequal. The X-learner first imputes the missing potential outcomes for every observation using an initial estimation of outcomes-for example, derived via a T-learner-and subsequently fits a second-stage model on the set of imputed individual effects to directly estimate $\tau(x, s)$. By incorporating knowledge from both the treated and control population, the X-learner has generally greater precision when one group is considerably larger than the other.

5.2 Orthogonalization and Nuisance-Robust Estimation

An alternative methodological strand focuses on enhancing robustness to confounding and high-dimensional nuisance variation by residualization or orthogonalization. The R-learner is a paradigmatic example of this approach [20]. It first estimates two nuisance components: the conditional mean of the outcome as a function of covariates (omitting treatment) and the propensity score representing the probability of receiving treatment given the covariates. Using these estimates, residuals are formed for the outcome (actual minus predicted) and for the treatment indicator (actual minus predicted propensity). Using residuals, the residualized outcome is regressed on the residualized treatment conditional on covariates x and market state s, in order to estimate (x, s). This orthogonalization ensures that small errors in estimating the nuisance functions have only second-order effects on the treatment effect estimator, giving rise to double or debiased machine learning robustness. In practice, cross-fitting is often used to avoid overfitting: data are split in such a way that the training of nuisance models and the training of treatment effect models are always done on different folds. A wide array of flexible learners-random forests, gradient-boosted trees, or neural networks-can play the role of nuisance estimators, while the final stage can use a similarly flexible model or else a more interpretable linear specification.

Tree-based approaches offer a complementary perspective, focusing on interpretability in addition to

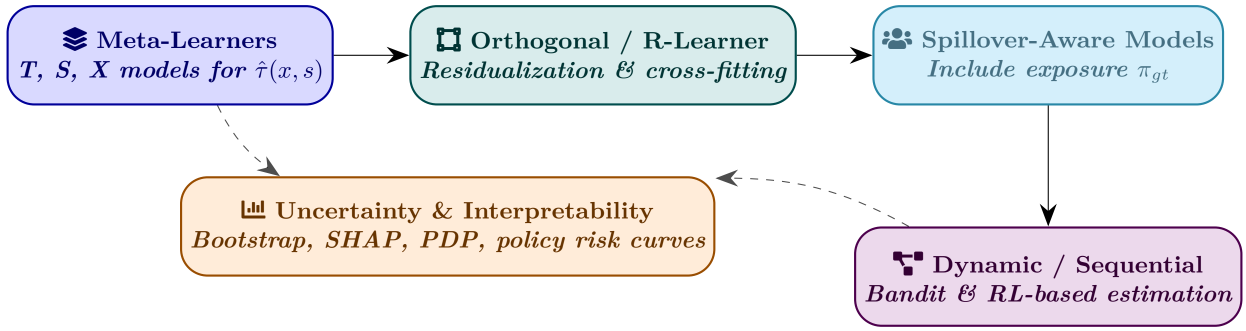

Figure 5: Estimation framework for heterogeneous treatment effects (HTE). Meta-learners (T-, S-, X-learners) estimate conditional effects flexibly; orthogonalized approaches like the R-learner improve robustness to nuisance estimation; spillover-aware models extend the HTE framework to environments with interference; dynamic models capture sequential effects via reinforcement learning; and uncertainty and interpretability tools quantify and explain $\hat{\tau}(x, s)$ estimates.

flexibility. Causal trees are modifications of standard decision tree algorithms to focus on maximizing treatment effect differences between leaves rather than minimizing overall prediction error. The algorithm recursively segments the covariate space (X, S) into subgroups that display maximum heterogeneity in treatment response [21]. In order to reduce overfitting, "honest" estimation procedures are often implemented: one subsample decides on splitting rules, and another estimates treatment effects within the resulting leaves. A causal forest is an ensemble estimator: aggregating across many such trees provides smooth estimates of $\tau(x, s)$ and measures of uncertainty. Such methods pick up nonlinear interactions and produce easy-to-understand segmentations of the population; each leaf represents a subgroup defined by transparent rules such as "drivers with tenure exceeding six months who mainly drive on weekends."

5.3 Spillover-Aware Learning

Traditional HTE estimators typically rely on the Stable Unit Treatment Value Assumption (SUTVA), which precludes any interference between units. In marketplace settings-like ride-hailing-this is rarely valid since a driver's outcome will depend on the treatment assignments of other drivers in his nearby neighborhood. Spillover-aware learning addresses this by augmenting the feature space to explicitly encode treatment exposure. For example, define π_{gt} as the proportion of drivers treated within driver i's spatial cluster g at time t. The treatment effect can then be modeled as a function $\tau(x, s, \pi)$ that conditions not only on individual and contextual features but also on the surrounding treatment saturation. In this formulation, the model estimates $\mathbb{E}[Y_{it}(1, \pi) - Y_{it}(0, \pi) \mid X_i = x, S_{gt} = s]$, capturing how the marginal effect of the incentive varies with environmental exposure. The sensitivity of $\tau(x, s, \pi)$ to π can be analyzed using derivatives or finite differences, yielding estimates of the strength and direction of spillover effects. Here, one might allow for interaction terms with π and key covariates or use nonparametric learners capable of modeling high-order interactions [22]. Resulting estimates can capture, for example, the effectiveness of an incentive for a driver operating in a sparsely treated region versus one surrounded by many incentivized peers, thus informing the optimal saturation level of incentive deployment.

5.4 Dynamic and Sequential Models

Repeated or continuous incentive offerings introduce dynamic dependencies that might not be captured by static treatment effect models. In such contexts, a driver's current response may depend on past exposure to incentives, accumulated fatigue, or expectations of future offers. A simplified approach treats each decision epoch independently, framing the problem within the paradigm of contextual bandits: the platform observes the current state, comprising driver features x_t and market context s_t, and chooses whether to offer an incentive with the aim of maximizing immediate expected outcomes. One-step reward signals are incorporated in this approach through the term $\tau(x_t, s_t)$, under the assumption that intertemporal dependencies are either negligible or accounted for via periodic re-estimation of the model. However, as soon as the intertemporal effect starts to be important, one needs a more general framework. The problem can be cast as an MDP with a state that summarises both the current market conditions and driver histories, such as recent receipt of incentives or cumulative fatigue. The actions correspond to incentive decisions, while rewards can incorporate both immediate and longer-term performance metrics. Different value functions, which represent the cumulative expected utility of different sequences of incentives, can then be estimated using reinforcement learning methods such as Q-learning and policy gradients [23]. Such approaches balance immediate supply mobilization with longer-run considerations aimed at driver retention and well-being. The implementation typically relies on off-policy evaluation methods in order to guarantee the validity of learning from observational or quasi-experimental data collected under previous incentive regimes.

5.5 Uncertainty and Interpretability

In view of the complexity of machine learning estimators, quantification and communication of uncertainty are indispensable for credible inference and policy guidance. Confidence intervals for can be obtained either by nonparametric bootstrap procedures-resampling of drivers or driver clusters-or via analytical variance approximations, when available. Providing uncertainty bands for subgroup-specific estimates-say, treatment effects for new drivers during peak weekend periods-avoids overconfidence in point estimates. Interpretability techniques further enhance transparency and stakeholder trust. Partial dependence plots illustrate the marginal relationship between a selected covariate and the predicted treatment effect, highlighting nonlinearities or threshold effects. If monotonic relationships are indicated by theoretical or operational considerations, interpretability constrained or smoothed versions of such plots may be applied. Feature attribution methods, such as SHAP (Shapley Additive Explanations), extended to treatment effect models, decompose individual-level predictions into contributions from constituent features, and can be used for diagnostic assessment of model behavior [24]. Beyond estimation, policy selection under uncertainty can be guided using policy risk curves, plotting expected welfare outcomes against measures of uncertainty-for instance, the lower confidence bound of expected welfare at a given confidence level. These curves visualize the trade-off between maximizing estimated gain and ensuring reliability, enabling decision-makers to select policies that achieve a favorable balance between efficiency and robustness.

6 | Outcomes and Measurement

The impact of driver incentives should be evaluated on several levels since the influence that this has on the marketplace is multifaceted. From the driver level, we are interested in metrics that reflect driver engagement and productivity and satisfaction. This includes obvious measures around the number of minutes or hours the driver is online-that is, active on the app and available to accept rides-during the incentive period and the number of trips the driver completes. Acceptance rate can also be informative as incentives might encourage drivers to accept more rides, particularly marginal or long-distance ones they would reject otherwise. Driver earnings is an important outcome in itself, but also its stability or dispersion-for example, we can look at variance in earnings across drivers or across time to see if incentives concentrate earnings or smooth them out. Moreover, as discussed above, overuse of incentives could lead to fatigue or burnout; we can include metrics such as the length of the driver's session-that is, how long they stay online or log off sooner-or a proxy for driver fatigue-such as a drop in acceptance rate or an increase in break frequency toward the end of a session. A critical long-term driver outcome is retention: whether the driver continues to be active on the platform in subsequent weeks or months [25]. There is a trade-off when incentives boost short-term activity but at the cost of long-term participation or vice versa; therefore, measuring retention or comeback rates after the incentive is valuable.

At the rider level, important outcomes surround the rider experience in terms of convenience, cost, and reliability. A key metric is the wait time or ETA of a driver for the rider. Effective incentives that entice more drivers into an area should shrink wait times for riders. We can measure the distribution of wait times-not just the average-percentiles (such as the 90th percentile wait time) are useful to make sure even the worst case delays are improving. Another rider-centric outcome is the trip fulfillment rate: when a rider requests a ride, what is the probability that a driver actually picks them up within a reasonable time? If incentives work, we would expect fewer requests by riders to go unfulfilled or get canceled because no driver was available. The price paid by the rider-which could be dynamic because of surge pricing-is also an outcome of interest; from the rider's perspective, a good outcome might be that incentives enable lower surge multipliers and thus more affordable rides, though from the platform's perspective, lower prices may mean transferring more surplus to riders. We

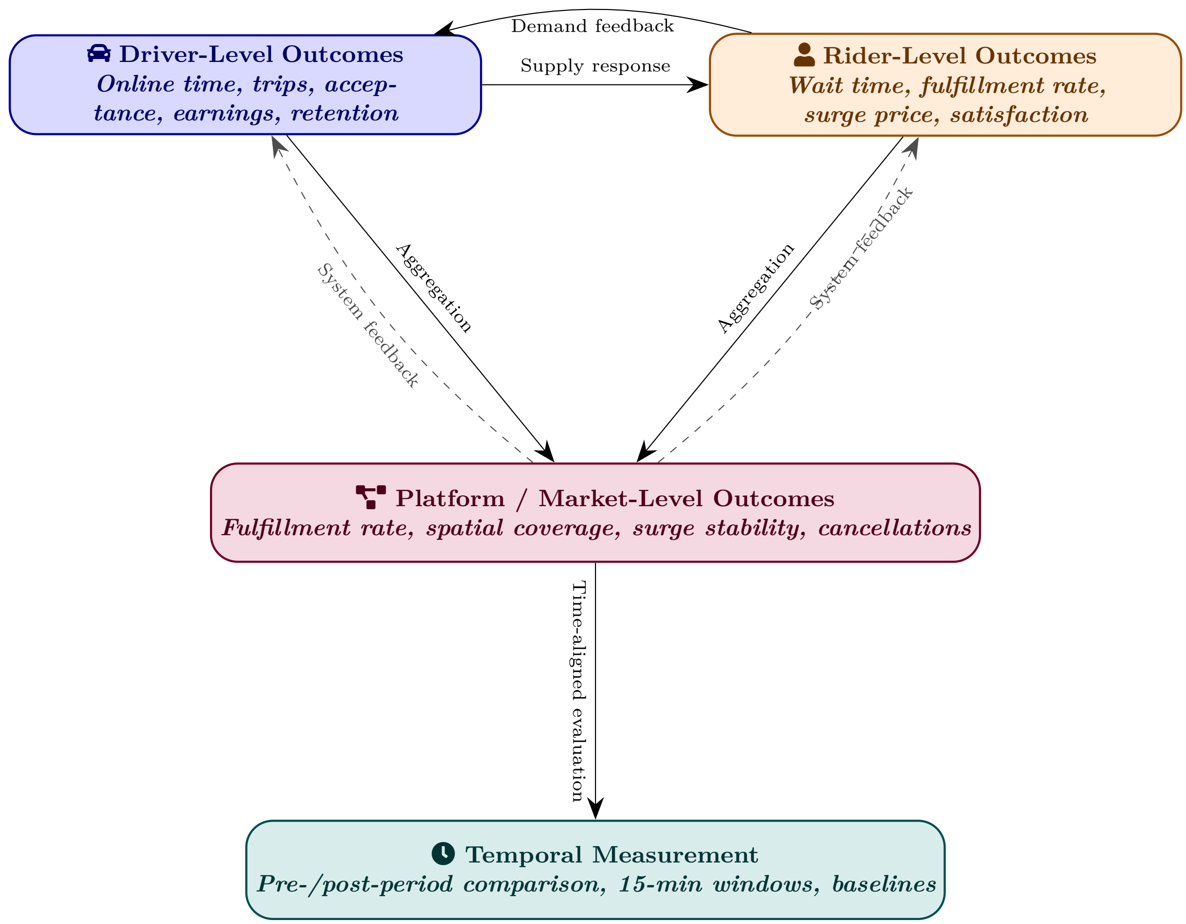

Figure 6: Measurement framework linking outcomes across system levels. Driver-level metrics capture engagement, productivity, and retention; rider-level metrics measure wait times, fulfillment, and satisfaction; platform-level metrics reflect overall market balance and service reliability. Temporal baselines and pre/post windows ensure causal alignment between incentives and observed changes.

might also consider rider satisfaction, or quality of service, indirectly through metrics such as rider ratings of trips, or frequency with which riders use the platform-for example, do riders take more trips because service improved? Some such metrics are more difficult to measure directly, but one could track proxies-such as whether riders who faced long waits come back and request again soon, or whether they drop off the platform.

We aggregate the effects at the platform or market level to see overall marketplace health [26]. There is the overall fulfillment rate, which is the fraction of ride requests that are successfully completed; this metric encapsulates both the supply and demand sides well and is closely tied to revenue. There is the spatial distribution of service: incentives might improve coverage in typically underserved areas. We track the average and peak surge levels. If incentives are properly balancing the market, we should find fewer extreme surge events, indicating a more stable supply-demand matching. The platform may also monitor externalities such as how many rides were canceled either by drivers or riders, since miscoordination can sometimes increase with lots of drivers relocating or chasing incentives, leading to mismatches or rider cancellations. Beyond these direct metrics, we often incorporate composite indices or use multiple metrics on a dashboard to ensure that improvement in one metric-like wait time-does not inadvertently worsen another, like driver idle time or rider cancellation rate. It is important to measure outcomes on consistent time scales relative to the intervention, typically looking at outcomes in the same time window during which the incentive was active. We would want to look at driver and rider outcomes in that region for that time window of 15 minutes and perhaps a short period after. We often include a baseline measurement before the incentive period, to control for any pre-existing trend [27]. By differencing or controlling for the pre-period, we get a cleaner view of the change attributable to the incentive. All these measurements together allow us to build a comprehensive picture of how the incentive affected individual behavior and overall market performance.

7 | Off-Policy Evaluation of Incentive Policies

Even after heterogeneous effect estimates, a platform would often want to try out new incentive allocation rules (policies) offline, before actually deploying them to drivers. Off-policy evaluation means estimating the performance of a hypothetical policy using existing logged data generated by a different policy (often this is an experiment, or the status quo operation). If one wants to reliably make such an estimate, then one must correct for the fact that the data were collected under a different assignment mechanism. One common technique is inverse propensity weighting: each observed outcome is weighted by the ratio of the probability that the new policy would have chosen that action to the probability that the old policy (which generated the data) did choose it. For example, if the new policy tends to give an incentive in situations similar to a particular data point, that data point gets a higher weight in evaluating the new policy; if the new policy rarely would do what was done in the data, that data point gets a low weight. To make this concrete, suppose in the historical data, a given hour had no incentive (control), but a new policy would give an incentive in that hour nearly always; any good results that occurred in that hour under control are not very relevant to the new policy, and IPW will downweight that data, or in a sense realize that we lack information for that scenario. There is one critical requirement for IPW: the logging (old) policy's probabilities must be known for each decision (this is usually true in experiments where randomization probabilities are set, or can be approximated from policy rules) [28]. To keep the variance manageable weights are often clipped or normalized: one common approach is self-normalizing: scale all weights so that they sum to 1 over the data set.
A more robust approach is the doubly robust estimator, combining propensity weighting with outcome modeling. In a DR estimator, one uses a predictive model to estimate what the outcome would be under each action-incentive or not-for each data point-and also uses weights that account for the policy difference, then takes some combination such that if either the model is correct or the propensities are correct, the evaluation is unbiased. This often has much lower variance than pure IPW and guards against mistakes in either component of the estimator.
When implementing OPE in practice, it's smart to stratify evaluation across relevant segments: for example, compute the policy's performance separately for peak hours versus off-peak, or for different regions or driver demographics, just to make sure the policy works well everywhere and doesn't just average out. The results can be aggregated (with appropriate weights reflecting frequency or importance of each segment) to get an overall performance estimate.
Finally, one should build safety tests into OPE. One should check that the new policy never violates the budget or other hard constraints in the data that we have-if it does, that is a red flag that the policy might not be implementable without issues. We can also simulate what would happen to service levels under the new policy by looking at metrics such as predicted average ETA or fulfillment rates; if any of these fall outside acceptable ranges-for example, if in some simulations the ETA would exceed a threshold-the policy might be adjusted or discarded [29]. Fairness considerations can be included by examining how different groups fare under the new policy's evaluation-if one group is systematically worse off, the policy might be refined. Using these OPE techniques, the platform can virtually "test" a new incentive scheme on past data, catch potential problems, and estimate effectiveness before trying live with real drivers and riders.

8 | Policy Learning Under Constraints

Having estimated who is most responsive to incentives and what the effects are, the platform's next challenge is to decide on an optimal incentive allocation policy that respects practical constraints. We can frame this as an optimization problem. For instance, suppose a_i represents the incentive (if any) given to driver i (this could encode both whether and how much incentive, but think of it simply as a choice variable per driver or driver-category). One might write an objective like:

$$\max_{\{a_i\}} \quad \mathbb{E}[W(a)] \tag{6}$$

$$\text{subject to} \quad \sum_i \mathrm{Cost}(a_i) \leq B, \tag{7}$$

$$\mathrm{ETA}_{gt}(a) \leq \bar{\eta}, \tag{8}$$

$$\mathrm{Fairness}(a) \geq \phi\,. \tag{9}$$

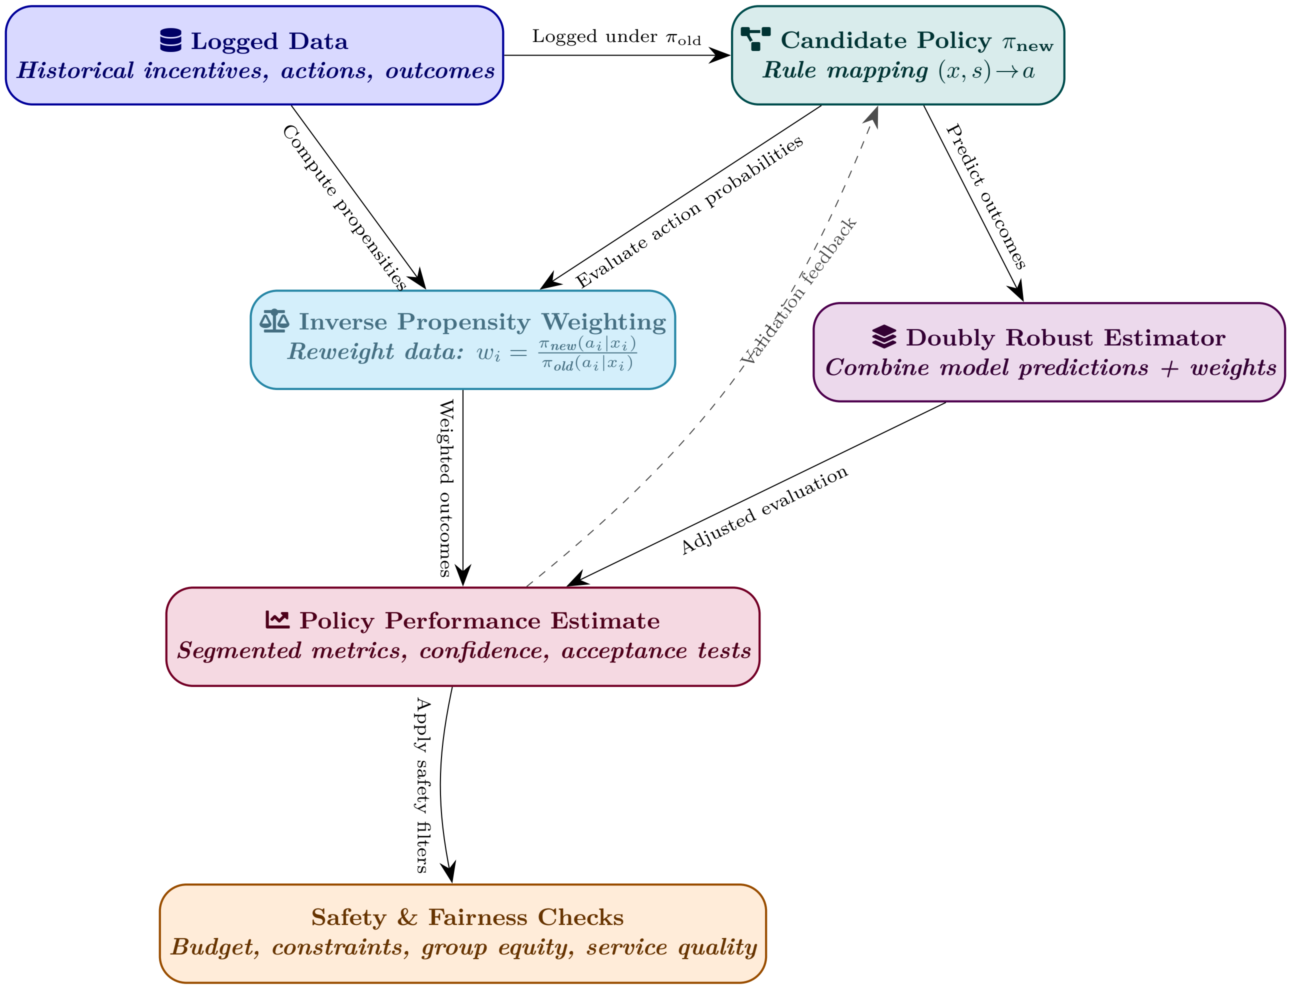

Figure 7: Off-policy evaluation (OPE) pipeline for testing incentive policies using historical data. Logged actions and outcomes under the old policy are reweighted via inverse propensity weighting (IPW) and refined using doubly robust (DR) estimators to assess a new policy's expected performance. Safety and fairness diagnostics ensure that evaluated policies respect budgetary, service, and equity constraints before live deployment.

where $W(a)$ is the welfare objective discussed earlier as a function of the assignment $a = \{a_i\}$ to all drivers, the cost constraint ensures the total spending on incentives does not exceed a budget B, the service-level constraint $\text{ETA}_{gt}(a) \leq \bar{\eta}$ means that for each area g and time t, the expected rider wait time under the policy a stays below a threshold $\bar{\eta}$ (this is one example of a guardrail), and $\text{Fairness}(a) \geq \phi$ denotes a requirement that some fairness metric remains above a minimum value ϕ (for example, that the difference in outcomes between any two demographic groups of drivers is small, or that every region gets at least a certain level of incentive support). In reality, the formulation might be more complex (and possibly solved approximately rather than exactly), but this expresses the idea that the policy should maximize overall benefit while obeying constraints on budget, service quality, and fairness. The budget constraint is straightforward: the platform usually has an earmarked amount it can spend on incentives (per day, week, or event), and the policy must not exceed this. Thus, even if many drivers have high estimated $\tau(x, s)$, the algorithm has to be selective in whom to actually give the incentive to [30]. This often turns the problem into a form of knapsack or allocation problem: among many possible driver-times where incentives could yield benefit, pick the set that maximizes W without over-spending. Solutions might involve ranking drivers by some score related to $\tau(x, s)$ per cost and picking the top ones until the budget runs out, possibly with some adjustments for the other constraints.

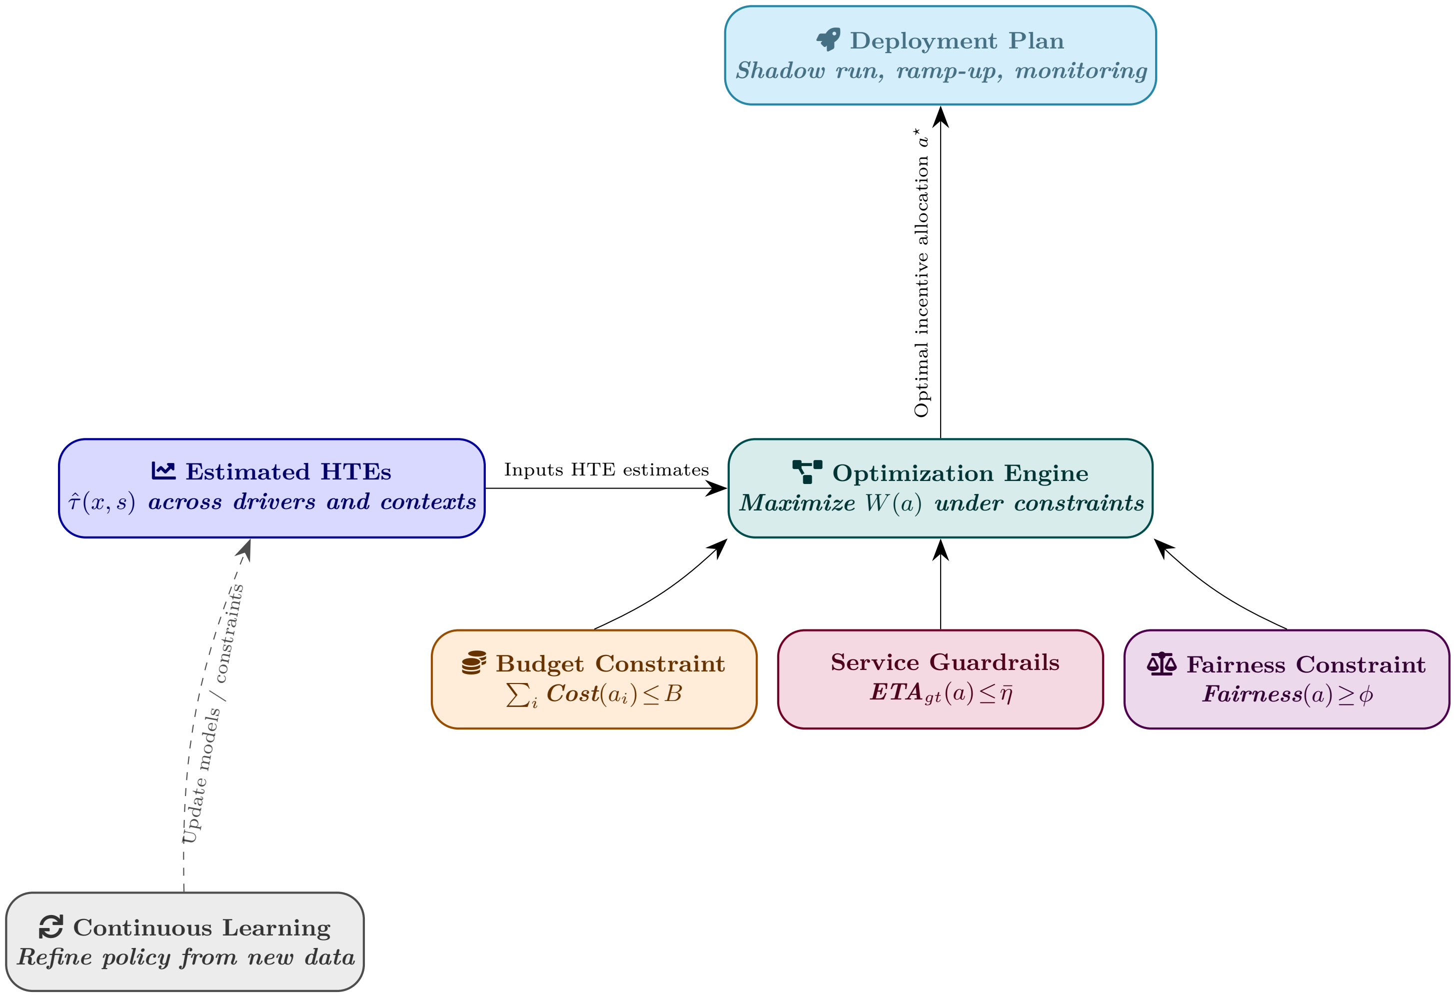

Figure 8: Constraint-aware policy learning pipeline. Estimated heterogeneous treatment effects (HTEs) feed into an optimization that maximizes expected welfare $W(a)$ subject to budget, service-level (guardrail), and fairness constraints. The resulting allocation policy is validated through shadow testing, phased deployment, and continuous feedback to ensure stability, compliance, and learning over time.

By guardrails, we refer to hard constraints that ensure the platform continues to operate within bounds. In the context of incentives, one guardrail might be: do not let rider-side metrics degrade beyond a threshold. For example, even with changes in driver distribution, ensure that the 95th percentile of rider wait times does not exceed some number of minutes or that the driver cancellation rate-which covers drivers rejecting rides or canceling after acceptance-does not spike above a set percentage. Some of these can be encoded as explicit constraints in the optimization-like the ETA constraint in the example objective. Others might be implicit: the platform may simply not consider policies that, based on simulations or past data, are likely to cause a violation. Guardrails often come from business requirements or regulatory obligations-for example, contractual service levels or safety requirements-and act as nonnegotiables in policy design.

Fairness constraints make sure that the incentive policy does not create or reinforce inequities [31]. There is a range of ways to formalize fairness in this context. One approach is group fairness: for example, requiring that drivers in different neighborhoods, or drivers of different demographic groups, have roughly equal opportunity or treatment in terms of incentive distribution or earnings outcomes. Another approach is individual fairness: making sure similar drivers in similar situations are treated similarly by the policy. Imposing fairness could mean adding terms into the objective which penalize disparity - for instance, subtracting a penalty if one area consistently gets much more incentive spend per driver than another after controlling for demand - or adding explicit constraints, such as each region getting at least X incentives per hour, or the ratio of incentives given between any two groups lies between certain bounds. There is often a trade-off: a purely welfare-maximizing policy might concentrate incentives where they are most effective, which could be in a subset of areas or drivers, whereas a fair policy might spread incentives more broadly at some cost in efficiency. The platform has to decide on these trade-offs and encode them in the policy learning process.

It is also prudent to impose a smoothness and stability requirement on the learned policy. Smoothness in this case means that the incentive allocation should avoid erratic or highly discontinuous behavior. For example, we would not want a policy that gives an incentive in one neighborhood but then withdraws it in the neighboring area in such a way as to encourage drivers to repeatedly drive back and forth across the boundary to chase incentives [32]. A policy that oscillates rapidly in time, turning incentives on and off every few minutes in the same area, is similarly likely to confuse drivers and result in gaming or loss of trust. We can penalize large differences in a_i between similar contexts, such as neighboring locations or consecutive time intervals, by incorporating terms that regularize the policy to vary more smoothly in space and time. In other words, if two adjacent areas have very similar demand and driver conditions, this algorithm shouldn't give a high incentive in one but none in the other unless there is a good reason to. Even after the offline calculations yield an optimal or near-optimal policy, there still needs to be care in deployment. A suggested approach is to start in a shadow mode-incentive policy runs in the background, making decisions on real-time data (a copy), but it does not actually send those incentives out to the drivers. This allows collecting data on what the policy would have done and how the marketplace would have looked, counterfactually, if those actions were taken, which can then be compared with what actually happened. It is a sort of trial run sans consequences to catch any glaring issues. From there, the platform can do an incremental ramp-up to enable the incentive policy for a small random subset of drivers or regions initially and, following the results closely, gradually expand its scope. For instance, 10% of drivers get targeted by the new policy in week 1, and in week 2, 30% of drivers are targeted and so on, so that when some side effects (maybe an unforeseen way that drivers respond, or an imbalance the algorithm did not account for) start to surface, it can halt or adjust prior to full deployment [33]. Even after deployment, continuous learning can be incorporated: the policy keeps observing data and updating its decisions, effectively treating the initial deployment as yet another experiment to refine the model. In short, careful constraint-aware optimization and cautious phased deployment manage this leap from analysis to action.

9 | Experimental Design Recommendations

Designing field experiments to learn about heterogeneity requires great foresight and methodological precision. First, the estimands of interest need to be clearly defined before the experiment is launched. The specific direct effects, spillover effects and equilibrium outcomes that are to be measured should be determined a priori along with the units of analysis. For example, "exposure" to others' treatment can be quantified by defining a spatial neighborhood radius, or using the fraction of drivers treated within a given zone, and the experimental design can be structured to vary this level of exposure. Having the relevant causal effects formally specified, including any hypotheses articulated in Section 10, guides both the randomization scheme and the data collection strategy. The estimands should be precisely defined in order to ensure interpretability of the resulting estimates. This involves identifying the causal quantities of interest-direct, spillover, and equilibrium effects-and describing how they will be operationalized in the experimental data. [34]

Explicit definitions of exposure and treatment intensity ensure that randomization and measurement processes align with the intended causal interpretation. Cluster-based randomization facilitates identification of both individual and interference effects. In practice, geographic or social clusters might be randomly assigned to varying incentive intensities. For example, neighborhoods might be assigned to one of a set of conditions: no incentives (control group) and low, medium, and high treatment groups, who have the specified percentage of drivers eligible for a bonus.

In each treated cluster, a random subset of drivers is offered the incentive corresponding to the intensity of treatment assigned. The randomization at two levels allows the estimation of both within-cluster contrasts-treated versus untreated drivers in the same environment-and between-cluster contrasts-areas with differing levels of treatment saturation-providing estimates of both direct and spillover effects. Temporal staggering of the introduction of treatment across

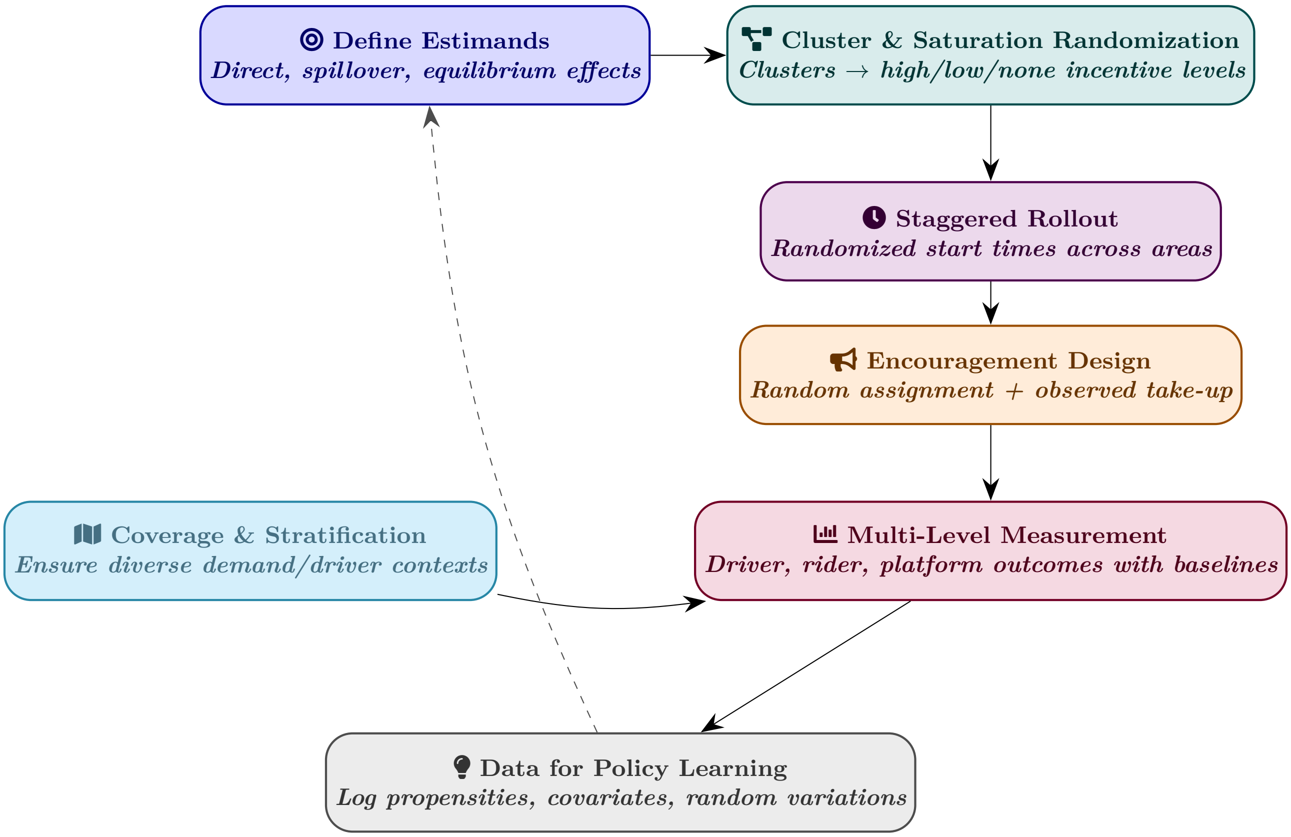

Figure 9: Recommended experimental design for learning heterogeneous treatment effects. Researchers should begin by defining causal estimands, then use cluster-level randomization with varied saturation and staggered rollouts to capture both direct and spillover effects. Encouragement designs address partial take-up. Measuring outcomes across driver, rider, and platform levels, with proper baselines and stratified coverage, enables rich post-experiment policy learning.

entities provides an additional source of identification that makes estimates more robust to global shocks. This is often done by randomly assigning clusters, or cities, to start treatment at various points in time, creating a series of temporally staggered rollouts. A design like this allows for the identification of dynamic effects and enables difference-in-differences across clusters that enter treatment at different times. [35] This, in turn operationally reduces risk by not exposing all the markets to new incentive conditions simultaneously.

Partial compliance to incentive offers is a fact in many real-world situations, and this heterogeneity in uptake needs to be explicitly incorporated into the design. Recording should be done for both assignment, the targeted drivers, and actual uptake, the responding drivers.

The random assignment is a legitimate instrumental variable for estimating treatment-on-the-treated effects in addition to intent-to-treat effects. This structure permits a decomposition of policy impacts into compliance behavior and treatment responsiveness components. Measurement of outcomes should occur at several levels: driver, rider, and platform; and involve appropriate baselines.

For each experimental unit (e.g., cluster-hour), key outcomes should be recorded: driver supply, trip counts, and earnings; rider wait times, fulfillment rates, and prices; platform-level utilization or revenue. Pre-treatment baselines, like the corresponding hour on a previous day or week, serve as the reference values for difference-in-differences estimation and, thus, improve the precision and internal validity. The randomization should include a representative range of market conditions and driver segments. [36] Both peak and off-peak, commuting hours on weekdays and nightlife on weekends, and core and peripheral parts of the city should be represented across treatment and control conditions. If responsiveness is to vary systematically with driver characteristics (tenure,

experience, etc.), then stratified, or blocked, randomization provides adequate representation of each subgroup.
The resulting balance will increase the ability to detect heterogeneous treatment effects and enhance the generalizability of the findings. Data logging should be implemented comprehensively from the outset to support subsequent causal modeling and policy optimization. The probability of treatment assignment for each experimental unit should be recorded to facilitate computation of the propensity weights for off-policy evaluation. Detailed covariate data on driver attributes and contextual features should be likewise collected, to enable the estimation of heterogeneous treatment functions (x, s), as well as the simulation of counterfactual incentive policies. The experimental design should be viewed not as a way to test an isolated incentive but rather as a method of generating data toward future policy learning. Random variation in incentive magnitudes or message framing can further extend the support of the data, enabling richer causal inference. Moreover, structured post-treatment feedback from drivers can provide qualitative context for interpreting quantitative results [37]. A well-designed field experiment on learning heterogeneity combines formal specification of estimands, structured randomization, comprehensive measurement, and anticipatory data collection. A strategic approach like that changes isolated experiments into systematic platforms for cumulative causal learning and adaptive incentive policy development.

10 | Anticipated Patterns of Heterogeneity

H1 (Tenure): New drivers versus experienced drivers. Newly onboarded drivers are expected to exhibit greater short-run responsiveness to certain incentives, particularly guarantees or minimum-earnings offers, relative to experienced drivers. New drivers typically face higher uncertainty about earnings and may be more influenced by income assurances; a guaranteed hourly rate can determine whether they decide to log in at all. Experienced full-time drivers, in contrast, may already earn above the guarantee threshold and maintain established driving routines. However, experienced drivers may respond more strongly to incentives such as surge multipliers in well-known hotspots, leveraging their city knowledge to strategically position themselves for high-paying trips. Tenure thus conditions the relative effectiveness of incentive types: incentives that reduce downside risk are likely to be more effective for new drivers, whereas those that enhance upside potential may better motivate experienced drivers.
H2 (Engagement mode): Part-time side-hustlers versus full-time drivers [38]. Drivers who treat ride-hailing as a supplementary income source are expected to respond differently from those who drive as their primary occupation. Part-time drivers may be especially motivated by short-term, limited-scope bonuses that align with their available hours, such as weekend quests that offer a fixed reward for completing a small number of trips, because their participation is time-constrained and goal-oriented. These drivers may also value incentives that enable rapid attainment of modest earning targets. Full-time drivers, by contrast, are likely to find small bonuses trivial and to value incentives that smooth income fluctuations or reduce idle time. For instance, incentives targeting typically slow periods may be more appealing to full-time drivers than brief, high-intensity offers. Hence, the impact of an incentive is expected to interact with engagement mode: short-duration bonuses for part-timers and stabilization-oriented incentives for full-timers.
H3 (Rider mix): Commuter-heavy versus nightlife-heavy demand. The effectiveness of incentives is anticipated to vary with the temporal and spatial composition of rider demand. In areas and periods characterized by volatile, surge-prone demand, driver incentives are likely to yield greater marginal benefits. This is especially true in nightlife districts during late-night hours [39]. Under such conditions, small increases in active driver supply can substantially reduce rider wait times and cancellations. In contrast, during stable commuter periods (for example, weekday mornings), the incremental effect of additional incentives may be limited, as many drivers already plan to work and riders possess alternative travel options. Accordingly, heterogeneity by rider mix is expected to manifest as larger positive effects on rider-side outcomes in leisure-oriented contexts relative to routine commuter settings. These outcomes include reduced cancellations and lower surge prices.
H4 (Dynamics): Effects of repeated incentives and potential fatigue. Repetitive exposure to the same incentive may attenuate its effectiveness over time. A diminishing marginal effect is hypothesized, whereby the initial instances of an incentive strongly influence driver behavior, but continued repetition leads to habituation or expectation. For example, a nightly streak bonus for consecutive trips may initially increase activity but eventually become normalized,

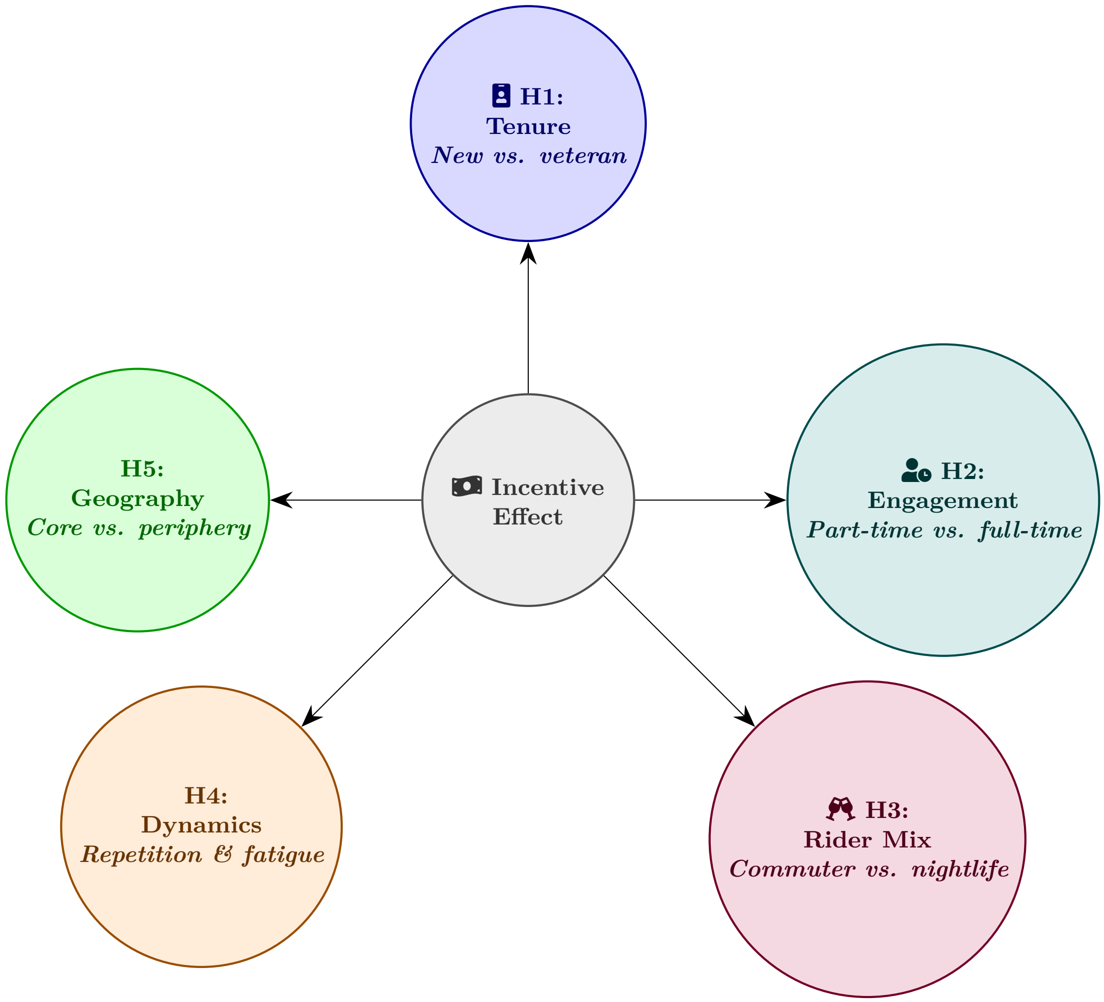

Figure 10: Map of hypothesized heterogeneity patterns. Five axes capture anticipated differences in incentive responsiveness: (1) driver tenure (new vs. experienced), (2) engagement mode (side-hustle vs. full-time), (3) rider mix (commuter vs. nightlife demand), (4) dynamic fatigue or repetition effects, and (5) geographic variation (core vs. peripheral areas). Each factor may alter both the magnitude and direction of incentive impacts across contexts.

such that its removal would reduce engagement, while its continued presence adds little incremental motivation. Rotating incentive types or incorporating cooldown intervals may mitigate this fatigue and preserve responsiveness. Empirical testing of this hypothesis would examine whether treatment effects decline in later experimental periods relative to earlier ones [40]. Additionally, excessive incentive use may produce burnout, where drivers overexert themselves during promotion periods and subsequently reduce activity, offsetting intended gains.

H5 (Geography): Core city versus peripheral areas. Geographic context introduces structural differences in both driver supply and rider demand. In central urban zones, baseline trip volume is high, and incentives may primarily reallocate drivers across busy locations rather than generate new supply. In peripheral or low-density regions, however, baseline driver availability is limited due to concerns about low demand and high deadhead time. Incentives in these areas are expected to exhibit a convex or threshold effect: modest incentives may have negligible influence, but sufficiently strong incentives can attract enough drivers to meaningfully improve service quality. A small increase in active drivers in such zones can sharply reduce rider wait times and cancellation rates. Consequently, incentive elasticity in peripheral areas is hypothesized to be higher once a critical mass of drivers is reached, producing stronger effects on fulfillment rates and rider experience per unit of incentive expenditure compared to central zones.

11 | Robustness and Validity Checks

Whatever the rigour of study design, it is essential to investigate empirical results for possible sources of bias or fragility [41]. One important diagnostic involves interference leakage: even under cluster randomization, drivers and riders can cross cluster boundaries-for example, a driver from a treated area provides a trip that ends in a control area. Sensitivity analyses can further be performed by excluding a buffer zone around cluster borders, or by redefining clusters (e.g., merging adjacent smaller clusters) in order to assess whether estimated effects remain stable. If interference extends across clusters, treatment-control contrasts would be expected to be attenuated near boundaries, a pattern that can be detected through diagnostic plots or models incorporating distance-based exposure measures.

If a large proportion of the drivers in the treatment group never saw or responded to the incentive, the ITT estimate will be a lower bound on the true effect for compliers. This concern can be dealt with by calculating compliance rates-namely, the proportion of treated drivers exhibiting behavior consistent with the incentive-and by reporting both ITT and treatment-on-the-treated estimates.

A big difference between ITT and TOT would mean that the effectiveness of the policy is focused on a subgroup of drivers. Temporal confounders, such as seasonality or exogenous events, must also be considered. If the experimental period coincides with holidays, major sporting events, or atypical weather, these factors may affect the outcome variable independently of the incentive intervention.

Including time fixed effects or event-specific controls could mitigate the risk of misattributing these kinds of shocks to treatment effects. [42]

A further diagnostic is given by the visual inspection of outcome trends for treatment and control groups before intervention. Parallel pre-treatment trajectories are supportive of causal attribution of post-treatment divergence to the intervention. One potential issue is manipulation around incentive thresholds. If drivers anticipate a rule, such as receiving a bonus after one more completed trip, and adjust their behavior to cross that threshold, then causal estimates based on local comparisons, such as those from regression discontinuity designs, will be biased.

The distribution of the running variable should therefore be examined for evidence of bunching at incentive thresholds-a sign that manipulation may have taken place.

Complex estimators used for either heterogeneous treatment effect estimation or policy learning may overfit or capture spurious correlations. The risk of overfitting can be guarded against by both comparing results across model classes-e.g., checking whether a simple linear specification yields qualitatively similar estimates of high-$\tau(x,s)$ segments as more flexible forest-based methods-and using out-of-sample validation to assess predictive consistency. When possible, held-out data or time periods can be used to assess whether the estimated heterogeneity patterns generalize outside of the estimation sample. Policy risk curves and related diagnostics can further convey uncertainty and robustness of policy recommendations to decision-makers. Estimates of treatment effects that vary across different market conditions-for instance, weekday rush hours versus weekend nightlife periods-should be compared for stability [43]. A policy that is performing adequately in one regime but poorly in another may need to be segmented or conditionally deployed. Empirical summaries disaggregating treatment effects by both time and location, as tables or figures, bring clarity over contextual performance. Comprehensive robustness and diagnostic analyses build confidence that conclusions on incentive effects and optimized policy designs hold good in the face of real-world operational variability.

12 | Implementation

Turning the proposed framework into an operational system on a live platform requires a carefully coordinated and technically sophisticated implementation plan. The first and foundational element is data engineering and infrastructure. The platform must construct a comprehensive and reliable data pipeline that captures and aggregates all relevant variables at the appropriate level of spatiotemporal granularity. For instance, the platform may discretize the city into small geographic units, such as hexagonal cells a few hundred meters across, and divide time into discrete intervals, such as 15-minute or hourly slots. For each cell–time unit, the data system should compute and store key features including the number of ride requests (as a measure of demand intensity), the number of available drivers (as a proxy for supply), average rider wait times, the prevailing surge multiplier, contemporaneous weather conditions, and indicators for special events such as concerts, sports games, or holidays. In parallel, the system should integrate information about the active driver pool within each cell–time observation, including counts of drivers by tenure, rating tier, and typical engagement level (full-time versus part-time), as well as the

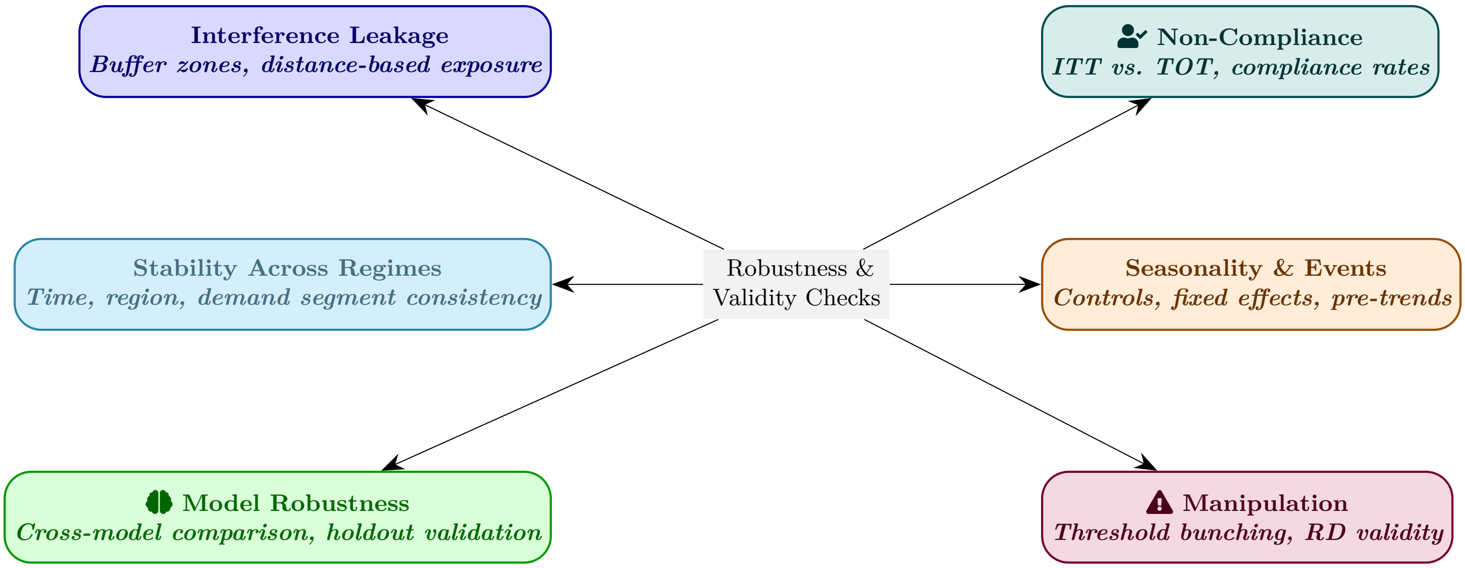

Figure 11: Framework for robustness and validity checks. Each diagnostic addresses a key potential threat: (1) spatial interference and spillover leakage, (2) non-compliance between assignment and uptake, (3) confounding from seasonality or external events, (4) manipulation around experimental thresholds, (5) model risk and overfitting, and (6) cross-regime stability of effects. Together these checks ensure the reliability and external validity of the estimated incentive impacts.

proportion of drivers with past responsiveness to incentive campaigns [44]. This engineered dataset serves as the backbone for experimentation and modeling, allowing outcomes to be analyzed by both region and time and enabling the mapping of treatment assignments to contextual conditions.

The second stage involves curating a comprehensive feature set for the estimation of heterogeneous treatment effects. The model that predicts $\tau(x, s)$ can only be as effective as the covariates provided to distinguish driver types and market contexts. Accordingly, driver-specific features should capture dimensions such as tenure categories (for example, newly onboarded, 1–3 months, 3–12 months, or more than one year on the platform), typical weekly driving hours to separate full-time from part-time drivers, historical trip ratings and cancellation rates as indicators of reliability and service quality, and prior responsiveness to incentive offers if available (for instance, whether the driver typically participates in promotions). Contextual features should include temporal indicators such as hour of day, day of week, and whether the period falls on a public holiday; spatial indicators describing the zone type (for instance, downtown commercial, residential, or airport areas); and measures of historical demand volatility, since some regions exhibit highly bursty demand while others are relatively stable. Real-time operational metrics such as current surge levels, rider queue lengths, and ongoing event flags should also be incorporated. Geographic attributes such as distance from the city center, proximity to transit hubs, and local competition intensity (measured by the number of nearby active drivers) provide further explanatory richness. By continuously updating and maintaining this feature repository, the platform enables the learning algorithms to detect and exploit nuanced, context-dependent patterns in incentive responsiveness.

The third component concerns the execution of the experimental design in the live environment [45]. Implementing the cluster-saturation randomization requires close coordination between the data science and engineering teams. The back-end systems must include logic that, for each cluster and time window, determines whether an incentive is active and, if so, which subset of drivers within the cluster is eligible to receive it, in accordance with the randomized assignment schedule. A dedicated monitoring dashboard should track the proportion of treated drivers in real time, verifying adherence to the planned treatment saturation levels and enabling rapid identification of any deviations. It is also essential that randomization assignments are persistently logged within the data infrastructure. Each trip or session-level record should include explicit fields indicating the incentive offered, if any, and the driver's

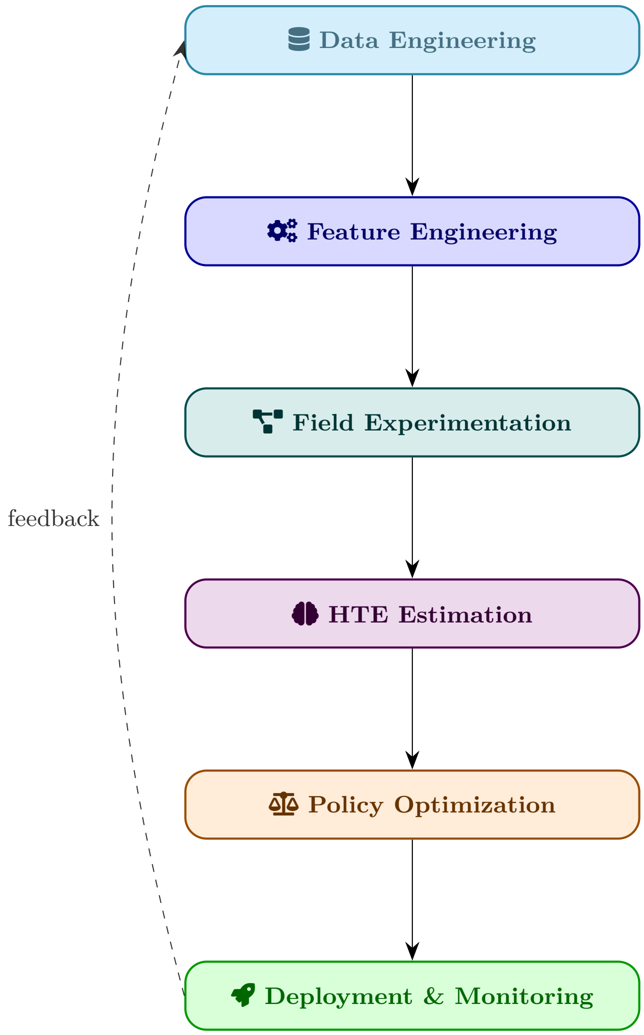

Figure 12: Implementation pipeline. Data flows downward through engineering, experimentation, heterogeneity estimation, and policy optimization, then loops back via feedback from deployment monitoring.

experimental group assignment. Absence of this metadata would severely limit the validity and interpretability of subsequent causal analyses. The experimental rollout may proceed in a phased manner to ensure operational stability, but ultimately the full experimental design should be executed to obtain sufficient data coverage for inference.

The fourth step entails the development of the learning algorithms used for heterogeneous treatment effect estimation and policy optimization. The uplift modeling pipeline should begin with the selection of appropriate model architectures, such as an X-learner implemented with gradient-boosted regression trees, complemented by a causal forest as a nonparametric benchmark [46]. Cross-validation and hold-out testing should be used to tune hyperparameters and evaluate generalization performance. In some cases, it may be useful to reserve entire clusters for final validation. Orthogonalization procedures should be applied to mitigate confounding biases: this involves estimating propensity scores, residualizing both outcomes and treatment indicators, and using the residualized components for causal estimation. Each training instance should also incorporate spillover-related features, such as the fraction of drivers treated within the same cluster or within a certain radius, to enable the model to capture local interference effects. After training, a comprehensive suite of diagnostics should be generated, including variable importance metrics, partial dependence plots, and assessments of extrapolation risk to confirm that the learned model operates within the empirical support of the observed data.

The fifth stage introduces the policy optimization layer, which converts the estimated $\tau(x, s)$ into operational incentive allocation decisions. In the simplest implementation, the model's predicted treatment effect for each eligible driver–context pair can be compared against a cost threshold, such that an incentive is allocated only when the expected return exceeds the marginal cost. More sophisticated implementations may involve formulating the allocation as a constrained optimization problem, akin to a knapsack or budgeted assignment model, where a fixed number of incentive slots per time period must be distributed to maximize overall expected welfare or revenue. In this formulation, the platform's budget, fairness, and service-quality constraints are encoded explicitly. Fairness constraints may, for example, ensure that each geographic region receives a minimum number of incentive opportunities, or that the variance of driver earnings does not increase beyond a specified tolerance [47]. Approximate or heuristic algorithms can often achieve near-optimal allocations with substantially lower computational overhead. One example is greedy selection based on τ/cost ratios. Simulated stress testing of the policy should also be conducted to verify that the resulting allocations do not inadvertently degrade service levels in any region or subpopulation.

The sixth component of implementation involves validation and testing prior to large-scale deployment. Off-policy evaluation using logged experimental data provides an initial means of estimating the performance of proposed incentive policies without affecting live operations. When sufficient data are available, multiple policy candidates can be compared through replay-based evaluation or counterfactual simulation. Controlled pilot rollouts should follow, beginning with a limited number of cities or driver cohorts, to evaluate real-world performance under operational constraints. Complementary discrete-event simulations can further assess system behavior under extreme or atypical scenarios, such as sudden demand surges or temporary supply shortages. These simulations can help identify corner cases and ensure the robustness of the incentive policy. Qualitative validation should also be incorporated when feasible, including structured feedback from drivers or focus groups to confirm that the incentive scheme is understandable and not producing unintended behavioral distortions.

Finally, operational dashboards should continuously track key performance indicators across time, geography, and driver subgroups. These indicators include average estimated time of arrival (ETA), fulfillment rates, driver earnings distributions, and incentive budget utilization. [48]. Automated safeguards should be embedded within the system: if regional wait times exceed acceptable thresholds or service levels deteriorate, the platform should be able to trigger automatic adjustments, such as temporarily increasing incentives in affected areas or reverting to a fallback policy. Real-time monitoring of budget expenditure is equally critical to ensure financial discipline. The system should also track shifts in the underlying data distribution. Examples include changes in driver composition or rider demand patterns that might signal model drift and necessitate retraining. Periodic reevaluation of the heterogeneity models and policy performance should be scheduled, for example on a monthly or quarterly basis, to update $\tau(x, s)$ estimates and recalibrate the optimization algorithm as the market evolves. Implementation is not a one-time exercise but a continuous cycle of experimentation, model refinement, policy deployment, and monitoring, forming a feedback loop that embeds adaptive learning into the platform's operational workflow.

13 | Conclusion and Discussion

Driver incentives are a potent tool for managing the dynamics of a ride-hailing marketplace, but their effects are inherently context-dependent. An exclusive concentration on average treatment effects risks missing substantial heterogeneity across segments of drivers, locations, and temporal conditions. This paper emphasizes the need for an explicit modeling and estimation of HTEs in order to understand the conditions under which incentives have meaningful effects. The proposed framework decomposes total effects into direct impacts on individual drivers, spillover effects among drivers operating in shared environments, and equilibrium feedbacks on the rider side of the platform. [49]

The methodological approach embeds the following complementary elements: experimental designs that allow for interference and complex market interactions, machine learning methods that flexibly capture heterogeneity in treatment effects, and policy optimization procedures that embed practical constraints such as budget limits, fairness, and service quality targets. Collectively, these elements establish a path for translating causal insight into actionable operational decisions. For example, saturation or cluster-based randomized experiments can generate variation necessary to identify causal effects, which may subsequently be analyzed using orthogonalized estimators to show where and for whom incentives are most effective. These empirical findings can subsequently be incorporated into optimization frameworks that determine how to allocate a finite incentive budget to maximize overall welfare during high-demand periods. The resulting system is inherently adaptive, designed to evolve as new data are collected and the marketplace itself changes over time. This points toward an evidence-based incentive strategy that is sensitive to heterogeneity: the platform can then deploy its resources more effectively without distorting driver behavior or market balance. The incentives can be focused on contexts where they have the greatest marginal benefit, such as offering guarantees to new or part-time drivers during critical supply shortages, or giving localized bonuses to mitigate emerging demand spikes. Fairness is ensured, and service quality remains consistent. To conclude, explicit modeling and exploitation of heterogeneous treatment effects allow the design of incentive mechanisms that are more nuanced, more effective, and fairer, hence closing the gap between causal inference and real-world operational decision-making in ride-hailing markets. [50]

Appendix A: Example Estimation Recipe

To help further concretize how one might go about estimating heterogeneous treatment effects and translating those into actionable policy recommendations, consider a simplified, step-by-step workflow for estimating heterogeneous treatment effects using an end-to-end example. First, one would set up a cluster-based randomized experiment with variation in the level of saturation across clusters. For example, each geographic area might be randomly assigned a target treatment rate, such as 0%, 25%, 50%, 75%, or 100% of drivers receiving the incentive. Within each such area, individual drivers are then randomly assigned to treatment in a manner conforming to that area's target rate, inducing natural variation across both clusters and saturation intensities. Experimental data are aggregated into area-time blocks, recording outcomes of interest, such as number of trips completed, wait time, and earnings, along with realized fraction of treated drivers and known assignment probabilities that reflect the experimental design.

The second step is the estimation of the necessary nuisance components. Models are fitted to predict outcomes under no treatment-the so-called baseline model-and to estimate treatment propensities as a function of observable covariates; in a randomized experiment, the latter may be determined by design. Flexible estimators such as gradient boosting machines or neural networks may be employed where appropriate. Step three is orthogonalization, where residuals are computed for both outcomes and treatment assignments by subtracting the model predictions [51]. This transformation removes baseline variation and prepares the data for the main causal estimation. The orthogonalized data are used to train a second-stage model-such as a random forest, X-learner, or linear model-to estimate the conditional treatment effect function. Exposure variables or saturation levels can be included to capture spillover effects where relevant.

Step four in the process involves heterogeneity assessment and interpretation of the estimated treatment effects. The distribution of (x; s) can be investigated across drivers, periods, and areas to come up with subpopulations where the incentives are most effective. Visualization tools such as partial dependence plots can be used to depict how estimated effects vary with covariates like driver tenure or local demand intensity. Confidence intervals obtained through resampling or bootstrap methods provide insight into the precision and robustness of these estimates across different segments. The fifth step translates the estimated treatment effects into operational policies. Using the $\tau(x, s)$ estimates, a policy optimization algorithm selects incentive allocations that maximize an objective such as welfare W subject to budgetary and fairness constraints [52]. In a simple approach, the top-ranked driver-context pairs with the highest estimated treatment effects above a well-chosen cost threshold can be selected until the budget is exhausted. More sophisticated approaches may model the decision as an optimization problem that jointly balances costs, benefits, and fairness metrics. Off-policy evaluation in the experimental dataset can then be used to simulate expected performance for alternative policies, comparing heterogeneous versus uniform incentive strategies to assess efficiency gains. Finally, the procedure should be validated on hold-out samples or through follow-up experiments to confirm generalizability. While simplified, this recipe gives a complete pipeline-from experimental design via causal estimation to optimized decision-making-that can guide iterative learning and deployment of incentive policies in live marketplaces.

Appendix B: Illustrative Outcomes and Weights

In formulating the welfare objective W used for incentive policy optimization, the platform must explicitly define and weight the outcomes of interest for each stakeholder group. Platform-level metrics may include overall fulfillment rate (positively weighted), average estimated time of arrival (negatively weighted), and cancellation rate (negatively weighted). Revenue or profit may also enter the objective, although these typically correlate with fulfillment and surge dynamics. Driver-level metrics may include total trips completed (positively weighted to reflect productivity), earnings variability (negatively weighted if stability is preferred), and retention or continued engagement (positively weighted to promote long-term supply sustainability) [53]. Additional metrics such as driver satisfaction, safety incidents, or quality ratings can be incorporated where available. Rider-side metrics might encompass ride request conversion rates (positively weighted), average wait times (negatively weighted), and fare fairness or affordability (which may receive non-linear treatment depending on strategic priorities). Customer satisfaction scores or in-app ratings can serve as further indicators of rider welfare. The coefficients α, β, γ in the welfare function W represent the platform's prioritization among these objectives. For example, a platform focused on market growth may assign greater weight to reducing rider wait times, while one prioritizing long-term driver retention may emphasize earnings stability. The precise weights can be determined through strategic decision-making, stakeholder consultation, or empirical calibration using historical data that reveal implicit trade-offs previously made by the platform. Sensitivity analysis should accompany any chosen weighting scheme to examine how variations in these parameters affect the resulting optimal policy. If the policy outcomes are highly sensitive to specific weights, additional data collection or analysis may be warranted before deployment. Explicitly defining and analyzing these weights ensures transparency in how incentive policies align with broader platform objectives and provides a principled basis for balancing efficiency, fairness, and sustainability within the ride-hailing ecosystem. [54]

References

[1] A. D. Smith, "Exploring the supply chain management aspects of uber in saudi arabia," 9 2016.

[2] A. Rubel, C. Castro, and A. Pham, "Agency laundering and information technologies," *Ethical Theory and Moral Practice*, vol. 22, pp. 1017–1041, 10 2019.

[3] D. Coyle, J. Adams-Prassl, and A. Adams-Prassl, "Uber & beyond: Policy implications for the uk," *SSRN Electronic Journal*, 1 2021.

[4] J. Pinsof, "A new take on an old problem: Employee misclassification in the modern gig-economy," *Michigan Telecommunications and Technology Law Review*, vol. 22, pp. 341–373, 7 2016.

[5] J. Berg and H. Johnston, "Too good to be true? a comment on hall and krueger's analysis of the labor market for uber's driver-partners:," *ILR Review*, vol. 72, pp. 39–68, 10 2018.

[6] N. F., C. W. Yuan, M. Ghafurian, and B. V. Hanrahan, "Chi - using stakeholder theory to examine drivers' stake in uber," in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, pp. 83–12, ACM, 4 2018.

[7] G. D. Morton, "Neoliberal eclipse: Donald trump, corporate monopolism, and the changing face of work," *Dialectical Anthropology*, vol. 42, pp. 207–225, 8 2017.

[8] S. Lee, B. Hubert-Wallander, M. Stevens, and J. M. Carroll, "Chi - understanding and designing for deaf or hard of hearing drivers on uber," in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pp. 529–12, ACM, 5 2019.

[9] M. J. Loewenstein, "Agency law and the new economy," 11 2017.

[10] S. Kim, E. B. Marquis, R. Alahmad, C. Pierce, and L. P. Robert, "Cscw companion - the impacts of platform quality on gig workers' autonomy and job satisfaction," in *Companion of the 2018 ACM Conference on Computer Supported Cooperative Work and Social Computing*, pp. 181–184, ACM, 10 2018.

[11] C.-W. You, C. W. Yuan, N. Bi, M.-W. Hung, P.-C. Huang, and H.-C. Wang, "Chi - go gig or go home: Enabling social sensing to share personal data with intimate partner for the health and wellbeing of long-hour workers," in *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pp. 1–16, ACM, 5 2021.

[12] M. Amaxopoulou, M. Durovic, and F. Lech, *Regulation of Uber in the UK*, pp. 101–135. Springer Singapore, 9 2020.

[13] S. Tashiro and S. Choi, "Labor market outcomes under digital platform business models in the sharing economy: the case of the taxi services industry.," *Business economics (Cleveland, Ohio)*, vol. 56, pp. 1–12, 8 2021.

[14] R. Cox, "Gender, work, non-work and the invisible migrant: au pairs in contemporary britain," *Palgrave Communications*, vol. 4, pp. 1–4, 10 2018.

[15] C. Cook, R. Diamond, and P. Oyer, "Older workers and the gig economy," *AEA Papers and Proceedings*, vol. 109, pp. 372–376, 5 2019.

[16] O. S. Malgonde, H. Zhang, B. Padmanabhan, and M. Limayem, "Managing digital platforms with robust multi-sided recommender systems," *Journal of Management Information Systems*, vol. 39, pp. 938–968, 10 2022.

[17] G. Iazzolino, "Going karura: colliding subjectivities and labour struggle in nairobi's gig economy," 7 2021.

[18] R. D. Smith and K. Lee, "Global health governance: we need innovation not renovation.," *BMJ global health*, vol. 2, pp. e000275–, 3 2017.

[19] R. L. Redfearn, "Sharing economy misclassification: Employees and independent contractors in transportation network companies," 4 2016.

[20] E. Moriuchi, ““social credit effect” in a sharing economy: A theory of mind and prisoner’s dilemma game theory perspective on the two-way review and rating system,” *Psychology & Marketing*, vol. 37, pp. 641–662, 11 2019.

[21] C. Knittel, “Racial and gender discrimination in transportation network companies,” 10 2016.

[22] D. J. Doorey, “Legal oshidashi, uber, labor law’s boundaries, and law’s response,” *SSRN Electronic Journal*, 1 2020.

[23] J. Hall and A. B. Krueger, “An analysis of the labor market for uber’s driver-partners in the united states,” *ILR Review*, vol. 71, pp. 705–732, 6 2017.

[24] N. K. Chan, “The rating game: The discipline of uber’s user-generated ratings,” *Surveillance & Society*, vol. 17, pp. 183–190, 3 2019.

[25] J. Ticona and R. Tsapatsaris, “Platform counterpublics: Gossip & contested knowledge about onlinelabor platforms,” *AoIR Selected Papers of Internet Research*, 10 2020.

[26] G. Iazzolino, “‘going karura’: colliding subjectivities and labour struggle in nairobi’s gig economy:,” *Environment and Planning A: Economy and Space*, vol. 55, pp. 1114–1130, 7 2021.

[27] C. P. Woo and R. A. Bales, “The uber million dollar question: Are uber drivers employees or independent contractors?,” 4 2016.

[28] K. Daniels and D. Turcic, “Matching technology and competition in ride-hailing marketplaces,” *SSRN Electronic Journal*, 1 2021.

[29] J. Hall and A. B. Krueger, “An analysis of the labor market for uber’s driver-partners in the united states,” 11 2016.

[30] J. Allen-Robertson, “The uber game: Exploring algorithmic management and resistance,” 5 2017.

[31] E. MacEachen, E. R. Musson, E. Bartel, J. Carriere, S. B. Meyer, S. Varatharajan, A. Kosny, P. Bigelow, and R. Saunders, “284 safety management systems: peer-to-peer ratings in the sharing economy,” *Vibration and Noise*, vol. 75, pp. A516.2–A516, 4 2018.

[32] S. M. T. Haque, R. Rashed, M. B. Morshed, M. U. Rony, N. Hassan, and S. I. Ahmed, “Exploring the tensions between the owners and the drivers of uber cars in urban bangladesh,” *Proceedings of the ACM on Human-Computer Interaction*, vol. 5, pp. 1–25, 4 2021.

[33] K. Phung, S. Buchanan, M. Toubiana, T. Ruebottom, and L. Turchick-Hakak, “When stigma doesn’t transfer: Stigma deflection and occupational stratification in the sharing economy,” *Journal of Management Studies*, vol. 58, pp. 1107–1139, 4 2020.

[34] V. P. Sisiopiku, “Stride project i2,” 8 2022.

[35] R. Grohmann, G. Pereira, A. Guerra, L. C. Abílio, B. Moreschi, and A. Jurno, “Platform scams: Brazilian workers’ experiences of dishonest and uncertain algorithmic management,” 10 2021.

[36] N. K. Chan and L. Humphreys, “Mediatization of social space and the case of uber drivers,” *Media and Communication*, vol. 6, pp. 29–38, 5 2018.

[37] A. Davidson and D. A. Theriault, “How consumer experience is shaped by the political orientation of service providers,” *Journal of Consumer Psychology*, vol. 31, pp. 792–800, 3 2021.

[38] J. Meakin, “Labour movements and the effectiveness of legal strategy: Three tenets,” *International Journal of Comparative Labour Law and Industrial Relations*, vol. 38, pp. 187–210, 6 2022.

[39] D. Mrvos, “Illusioned and alienated: Can gig workers organise collectively?,” *tripleC: Communication, Capitalism & Critique. Open Access Journal for a Global Sustainable Information Society*, vol. 19, pp. 262–276, 4 2021.

[40] R. Hamilton and L. L. Price, “Consumer journeys: developing consumer-based strategy,” *Journal of the Academy of Marketing Science*, vol. 47, pp. 187–191, 2 2019.

[41] A. Sandhu and P. Fussey, “The ’uberization of policing’? how police negotiate and operationalise predictive policing technology,” 7 2023.

[42] J. D. Angrist, S. Caldwell, and J. Hall, “Uber vs. taxi: A driver’s eye view,” 9 2017.

[43] C. Knittel, “Racial and gender discrimination in transportation network companies,” 10 2016.

[44] R. Sprague, “Are airbnb hosts employees misclassified as independent contractors,” 8 2020.

[45] E. B. Marquis, S. Kim, R. Alahmad, C. Pierce, and L. P. Robert, "Cscw companion - impacts of perceived behavior control and emotional labor on gig workers," in *Companion of the 2018 ACM Conference on Computer Supported Cooperative Work and Social Computing*, pp. 241–244, ACM, 10 2018.

[46] V. P. Sisiopiku, "Stride project i2," 8 2022.

[47] V. Kameswaran, L. Cameron, and T. R. Dillahunt, "Chi - support for social and cultural capital development in real-time ridesharing services," in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, pp. 342–12, ACM, 4 2018.

[48] T. Berger, C. B. Frey, G. Levin, and S. R. Danda, "Uber happy? work and well-being in the "gig economy"," *Economic Policy*, vol. 34, pp. 429–477, 7 2019.

[49] C. Denis, "Aventuras con facturas: Notes on some of mexico's self-inflicted fiscal wounds," *Latin American Policy*, vol. 14, pp. 638–648, 11 2023.

[50] Y. Li and W. Huang, *GeoAI@SIGSPATIAL - Imitation Learning from Human-Generated Spatial-Temporal Data*. ACM, 11 2019.

[51] P. Geldsetzer, "Using rapid online surveys to assess perceptions during infectious disease outbreaks: a cross-sectional survey on covid-19 among the general public in the united states and united kingdom," 3 2020.

[52] M. Tarafdar, X. Page, and M. Marabelli, "Algorithms as co-workers: Human algorithm role interactions in algorithmic work," *Information Systems Journal*, vol. 33, pp. 232–267, 5 2022.

[53] A. J. Schmitz, "Considering uber technologies inc v heller under us law," 4 2021.

[54] E. MacEachen, E. R. Musson, E. Bartel, J. Carriere, S. B. Meyer, S. Varatharajan, A. Kosny, P. Bigelow, and R. Saunders, "285 the sharing economy: hazards of being an uber driver," *Small Scale Enterprises and Informal Sector*, vol. 75, pp. A494.3–A495, 4 2018.