Designing and Sequencing Incentives for Gig Economy Workers

Wichinpong Park Sinchaisri
Haas School of Business, University of California, Berkeley, parksinchaisri@berkeley.edu

Abstract. We study how to design and sequence incentives for platform workers using field evidence from a large on-demand delivery platform. Incentives are expensive and widely used, yet most work evaluates one scheme at a time and overlooks how past contracts shape current responses. In collaboration with the firm, we conducted a two-month switchback experiment in 30 expansion regions and three worker tiers. Each region cycled through weekly contracts from a 2 by 2 family that varies along two dimensions: whether incentives target completed orders versus online hours, and whether they operate at daily versus multi-day streak horizons. Using worker-day and order-level data, we estimate incremental supply hours and incremental cost per supply hour by region cluster, worker tier, incentive type, and incentive sequence. Hour-based guarantees are generally more cost efficient than order-based bonuses, but the ranking varies systematically across markets and worker tiers. We also document history dependence: the same contract yields different responses depending on the incentive used in the previous week. These results provide guidance for platforms on both which incentives to offer and how to phase them over time.

Key words: gig work; incentives; platforms; field experiments; switchback designs; labor supply

1. Related Literature

This paper studies how a platform should design and sequence incentives when workers respond not only to the current contract, but also to the state created by previous contracts. The setting is a food-delivery platform where riders self-schedule, decide how long to remain active, and choose how selectively to accept assignments. The platform's incentive problem is therefore both operational and behavioral. Operationally, incentives are used to acquire capacity in the right market and time window. Behaviorally, incentives can alter routines, earnings expectations, and progress toward goals, making the response to a contract history-dependent. Methodologically, these same dynamics imply that a standard current-treatment comparison may not answer the platform's policy question when the policy object is a sequence of contracts.

Our paper contributes to four streams of research: incentive design and labor supply in flexible work, behavioral responses to nonlinear and intertemporal incentives, learning and adaptation in repeated work environments, and experimentation in platforms with interference and temporal carryover.

1.1. Incentive design and labor supply in flexible platform work

A central challenge in on-demand platforms is that labor supply is not a fixed input. Workers decide when to participate, how long to work, and how much effort or selectivity to exert conditional on being active. This distinguishes platform capacity management from settings where the firm directly schedules labor. In ride-hailing and delivery markets, workers value flexibility and respond to short-run earning opportunities, but these responses are heterogeneous and shaped by institutional features of the platform (Hall and Krueger 2018, Chen et al. 2019, Mas and Pallais 2017). Operations models of self-scheduled capacity similarly emphasize that platforms use prices, pay, and other instruments to bring supply into alignment with time-varying demand (Cachon et al. 2017, Chen and Sheldon 2016, Bimpikis et al. 2019, Ma et al. 2022). Recent empirical work shows that even seemingly operational details, such as payment timing, can change labor supply, underscoring that the relevant contract is not only the expected level of compensation but also how compensation is structured over time (Chen et al. 2025).

The incentive-design literature provides a foundation for why contract form matters. Performance-pay contracts can increase productivity by linking rewards to output, but they also change sorting, effort allocation, and distributional outcomes (Lazear 2000, Shearer 2004, Bandiera et al. 2007). Nonlinear contracts create thresholds and notches that can induce target-hitting behavior and bunching near eligibility cutoffs (Oyer 1998, Saez 2010, Kleven and Waseem 2013). These insights are directly relevant to platform incentives, where bonuses often require riders to complete a minimum number of orders, work a minimum number of hours, or maintain performance across multiple days.

Our contribution to this stream is to separate two contract primitives that are often bundled in practice. The first is the qualification metric: whether the contract rewards availability, such as online hours, or throughput, such as completed orders. The second is the reward horizon: whether the contract resets quickly or creates state dependence through streaks or multi-day requirements. This distinction matters because the same monetary value can activate different behavioral margins. Hours-based incentives primarily purchase presence. Orders-based incentives also change the return to accepting marginal assignments. Streak-based incentives create continuation value because current qualification changes the value of future participation. By organizing common platform incentives around these primitives, we provide a tractable framework for comparing contracts that are operationally familiar but behaviorally distinct.

1.2. Behavioral responses to reference points, goals, and progress

A second stream of work shows why workers may respond to the path of incentives rather than only to contemporaneous marginal pay. Reference-dependent preferences imply that individuals evaluate outcomes relative to an endogenous benchmark, with losses relative to the benchmark weighted more heavily than comparable gains (Kahneman and Tversky 1979, Tversky and Kahneman 1992, Kőszegi and Rabin 2006, Rabin 1998). In labor supply, this mechanism has been used to explain income targeting and state-dependent work decisions among taxi drivers (Camerer et al. 1997, Crawford and Meng 2011, Farber 2015, Goette et al. 2004). Although the empirical debate in that literature concerns the magnitude and interpretation of daily income targeting, the broader implication is important for platform incentives: recent earnings can become a benchmark against which future work opportunities are evaluated.

This mechanism creates a natural reason why incentive sequences can matter. A generous predecessor contract may increase current activity, but it can also raise the earnings reference point for the next contract. If the successor is less generous or more uncertain, workers may evaluate it as a disappointment even if its absolute expected value is positive. We refer to this as a paydown channel: the platform must pay down an elevated reference point created by its own prior incentive. This channel differs from standard incentive effects because it is not a property of the current contract alone. It depends on the comparison between the current contract and the worker's recent experience.

Goal pursuit provides a second behavioral foundation for dynamic responses. Goal-setting theory emphasizes that explicit goals can focus attention and increase persistence (Locke and Latham 2002). The goal-gradient literature further shows that effort can accelerate as agents approach a reward threshold (Hull 1932, Kivetz et al. 2006, Nunes and Drèze 2006), while effort may be weakest when agents are psychologically distant from both the starting point and the goal (Bonezzi et al. 2011). Related work on temporal landmarks shows that calendar boundaries can affect initiation and re-engagement (Dai et al. 2014, 2015). More recently, evidence on streaks shows that consecutive-success structures can increase persistence and that broken streaks can affect subsequent decisions (Silverman and Barasch 2023, Silverman et al. 2023, Mehr et al. 2025).

Our horizon dimension is grounded in this literature. Daily contracts are memoryless within the week: qualification today does not affect tomorrow's eligibility value. Streak contracts are state-dependent: qualification today changes progress toward a later payout. The model formalizes this force through a progress premium, the incremental continuation value of qualifying today. This yields a sharp empirical prediction that is absent from static labor-supply models: under streak contracts, effort and participation should vary with progress toward completion, whereas reset contracts should not generate the same within-week gradient.

1.3. Learning, routines, and adaptation in repeated platform work

Workers on platforms also learn and adapt. The organizational learning literature frames repeated decision-making as a tradeoff between exploration and exploitation (March 1991, Posen and Levinthal 2012). Empirical work shows that performance incentives and goals can shape exploration choices (Lee and Meyer-Doyle 2017, Raveendran et al. 2023), while cognitive-science evidence shows that individuals use both directed and random exploration in sequential choice tasks (Wilson et al. 2014, Schulz and Gershman 2019). In platform work, learning can occur along multiple dimensions: where demand is high, which time windows are attractive, how to evaluate assignments, and how to combine platform work with other obligations.

This stream is relevant because it suggests another source of history dependence. A contract that induces participation today may expose workers to information that changes their productivity or willingness to work tomorrow. Algorithmic decision-support research similarly shows that interventions can shape human sequential decisions beyond the period in which advice or recommendations are provided (Bastani et al. 2025). In gig work, prior evidence indicates that worker behavior reflects both economic incentives and behavioral frictions, including inertia and adaptation (Allon et al. 2023). Related work on multihoming further emphasizes that worker participation depends on outside options and cross-platform opportunities, which can make retention and sequencing central to platform strategy (Allon et al. 2025).

Our baseline theory focuses on two parsimonious inherited states: a participation friction and an earnings reference point. The participation friction captures routines, habit formation, and platform-specific familiarity in reduced form. The earnings reference point captures adaptation in earnings expectations. We treat learning as an important alternative mechanism rather than embedding it as a third primitive in the baseline model. This choice keeps the theory disciplined while preserving an empirical role for learning: if sequence effects are stronger among less experienced workers, persist after recovery periods, or operate through productivity rather than participation and earnings intermediates, learning becomes a plausible complementary explanation.

1.4. Marketplace interference and switchback experimentation

A final challenge is methodological. Marketplace experiments often exhibit interference because treated and untreated users interact through common prices, queues, congestion, and matching mechanisms. In two-sided markets, individual-level randomization can therefore produce biased estimates of the policy effect that would arise if the intervention were deployed market-wide (Blake and Coey 2014). This issue is especially salient in delivery platforms: incentives assigned to one group of riders can change the order pool, wait times, and earnings opportunities faced by other riders. For this reason, platforms often randomize policies at the market-time level rather than the individual level.

Switchback experiments provide a leading framework for such settings. In a switchback design, an aggregate unit, such as a market, is sequentially exposed to different treatments over time. This design is well suited to platform operations because it respects marketplace-level interference while generating repeated experimental variation (Bojinov and Shephard 2019, Bojinov et al. 2021, 2023). Recent work develops optimal designs and inference under carryover, surrogate adjustment, data-driven design, forward-looking demand, and decision-oriented objectives (Chen and Simchi-Levi 2025, Xiong et al. 2024, Wu et al. 2024, Ni 2025). Related work on long-term treatment effects studies how short-term experiments can inform persistent effects when outcomes evolve after treatment exposure (Huang et al. 2023, Tran et al. 2024).

We build on this literature but ask a different question. In much of the switchback literature, carryover is a threat to estimating a current-treatment effect: past treatments contaminate current outcomes, and the design seeks to control, bound, or exploit that dependence for more efficient estimation. In our setting, carryover is also the object of interest. The behavioral model predicts that prior contracts change the state of the workforce, and the platform's decision is inherently sequential: which contract should be deployed after which predecessor? The relevant estimand is therefore not only the average effect of a current contract, but the sequence-specific mean, namely the outcome under a predecessor-successor pair. Our design contribution is to connect a behavioral model of state formation to a history-aware switchback design that creates support for the predecessor-successor contrasts needed to evaluate incentive sequences.

1.5. Positioning and contribution

Taken together, the prior literature establishes that platforms can shape flexible labor supply through incentives, that nonlinear and progress-based rewards can generate behavioral state dependence, and that marketplace experiments require designs that respect interference and time dynamics. What remains less understood is how these forces interact when platforms rotate a portfolio of incentives over time. In practice, a platform rarely asks only whether one contract is better than another in isolation. It asks whether an hours guarantee should precede an orders bonus, whether a streak contract should be followed by a reset contract, and whether the answer differs across market conditions and worker segments.

This paper addresses that gap in three ways. First, we develop a behavioral operations model that organizes platform incentives along two primitives, qualification metric and reward horizon, and shows how these primitives activate distinct margins of worker behavior. Second, we show that predecessor contracts can affect successor performance through two opposing state variables: routine formation lowers future participation frictions, while reference-point adaptation raises the earnings benchmark that future contracts must satisfy. This yields dynamic complementarity and possible ranking reversals, so the best successor contract depends on the predecessor-induced state. Third, we implement a history-aware switchback experiment in a large delivery platform, rotating a two-by-two family of incentives across 30 regions and multiple worker tiers. The design links theory to practice by estimating sequence-specific responses, not only contemporaneous treatment effects. The result is evidence on both which incentive structures work and how they should be phased over time.

2. Setting: On-Demand Food Delivery in Thailand

We study incentive design on a large on-demand food delivery platform in Thailand. The platform matches consumers and restaurants with a flexible fleet of riders who decide when to work, how long to remain active, and which assignments to accept. These decisions are central to service performance. Demand varies sharply across hours, days, weather conditions, and local markets, while rider supply is self-scheduled and costly to acquire precisely when demand is highest. The platform therefore relies on short-horizon incentive contracts to manage capacity, stabilize service levels, and encourage rider availability during operationally important windows.

Thailand is a useful setting for studying this problem for two reasons. First, the economic structure resembles other large delivery markets. Platforms operate two-sided marketplaces with time-varying demand, algorithmic dispatch, geographically localized matching, and independent riders who retain discretion over their work. Public reports describe the Thai food-delivery sector as a large and competitive market organized around a small number of major platforms, similar in broad structure to other concentrated delivery markets (Thairath Money 2024, Momentum Works 2024). Second, the setting has institutional features that make incentive responses especially informative. Deliveries are predominantly completed by motorbike in dense urban and peri-urban environments, where travel-time uncertainty, weather exposure, and local route knowledge are operationally important. In addition, switching platforms within a shift appears less frictionless than in settings where workers can simply keep multiple apps open and reallocate attention continuously. Platform-branded equipment, onboarding routines, market-specific operating practices, and eligibility rules can make same-day switching costly. These frictions strengthen the link between the focal platform's incentive policy and realized within-platform labor supply.

The setting is therefore well suited to our research question. Incentives are consequential because the platform must acquire capacity from workers it does not directly schedule. At the same time, incentive effects may be dynamic because riders repeatedly face platform contracts, accumulate recent earnings experience, and respond to progress toward thresholds. The platform's problem is not only which incentive contract increases supply in a given week, but how contracts should be sequenced when this week's policy changes the workforce state encountered by next week's policy.

2.1. Platform operations and rider decisions

The platform operates a dispatch-based delivery workflow. When a rider is active, the platform surfaces delivery opportunities through its matching system. Each assignment is associated with pickup and drop-off locations, expected travel and waiting time, and compensation information. The rider then decides whether to accept. After acceptance, the rider completes pickup and delivery, and the platform records timestamps, distance, completion outcomes, and pay components.

This workflow creates three behavioral margins. The first is participation: whether the rider logs on at all during a day or week. The second is availability: how long the rider remains online and whether that availability overlaps with peak demand windows. The third is acceptance selectivity: conditional on being online, how willing the rider is to accept dispatched assignments. These margins are operationally distinct. Increasing participation expands the active labor pool. Increasing online time improves coverage, especially during peak periods. Increasing acceptance can raise completed orders conditional on active supply, but may also expose riders to less attractive assignments and interact with congestion, distance, and waiting time.

This multi-margin structure is the reason contract form matters. A contract that rewards online hours directly targets availability. A contract that rewards completed orders targets realized throughput and therefore also affects acceptance behavior. A contract with a daily reset targets current-day behavior. A contract with multi-day requirements creates continuation value because today's qualification can preserve eligibility for a later payout. The experiment is designed around these distinctions.

2.2. Incentive contracts in the field setting

The platform's incentive menus can be organized along two primitives: the qualification metric and the reward horizon. The qualification metric determines what behavior makes a rider eligible for a bonus. In availability-based contracts, eligibility depends on online hours. In throughput-based contracts, eligibility depends on completed orders. The reward horizon determines whether qualification is rewarded immediately or affects future payout opportunities. Reset contracts determine eligibility within a short window, typically a day. State-dependent contracts condition payout on progress across multiple days or segments within a week.

Crossing these primitives yields the four contract families studied in the experiment. Daily Hours (DH) rewards availability within a day. Daily Orders (DO) rewards completed orders within a day. Streak Hours (SH) rewards repeated satisfaction of hours requirements over a week, sometimes with mandatory high-demand days. Streak Orders (SO) rewards repeated satisfaction of order requirements over segmented windows within a week. A no-incentive condition provides a baseline in which riders receive standard compensation but no additional region-specific incentive.

The operational implementation is richer than this taxonomy but maps cleanly into it. Hours-based incentives typically require both total online hours and peak-window online hours, reflecting the platform's need to acquire capacity when demand is most valuable. Orders-based incentives use tier-specific order thresholds, reflecting the platform's goal of increasing completed service rather than merely online presence. Streak-style incentives create within-week state dependence. Under SH, qualification on one day can preserve progress toward a weekly payout, especially when mandatory days remain. Under SO, the week is divided into segments, and riders earn bonuses by satisfying daily order requirements throughout each segment. In both cases, current qualification changes the value of future participation.

2.3. Rider tiers and target calibration

The platform assigns riders to performance tiers based on recent activity and performance. In the experiment, we refer to these tiers as low, mid, and top. Thresholds and payouts vary by tier so that targets remain attainable for less active riders while remaining meaningful for highly active riders. This tiering is operationally important because the same contract family can imply different effective difficulty across rider segments. A daily order target that is routine for a top-tier rider may be demanding for a lower-tier rider; a streak requirement that creates useful persistence for one segment may be discouraging for another if the target is perceived as unattainable.

For this reason, rider tier is not merely a control variable. It is part of how the platform implements incentive design. In the empirical analysis, we distinguish pre-experiment tier heterogeneity from endogenous changes in activity that may occur after treatment. The tier structure also provides a natural way to test whether contract primitives operate differently when thresholds are more or less attainable.

2.4. From setting to experimental variation

The field setting motivates two design choices. First, the experiment varies contract structure rather than only incentive generosity. The four treatment families isolate the metric and horizon dimensions that are central to platform operations and to the behavioral model. Second, the experiment varies contracts at the region-week level. Riders in the same local market interact through a common pool of orders and matching opportunities, so individual-level randomization would alter the environment faced by both treated and untreated riders. Region-week assignment instead matches the platform's operational decision: which incentive menu to post in a given market and week.

The next sections formalize these ideas. The model shows why metric and horizon activate different behavioral margins, and why predecessor contracts can shape successor responses through inherited worker states. The experimental design then implements these theoretical requirements by rotating the four contract families and a no-incentive baseline across regions and weeks.

3. Model Revised

The setting in §2 highlights three features that a model of platform incentives must capture. First, workers respond along multiple margins: they decide whether to participate, how long to remain online, and how selectively to accept assignments. Second, incentive contracts differ in what they reward. A contract based on hours primarily rewards availability, whereas a contract based on completed orders rewards realized throughput and therefore changes the return to accepting marginal assignments. Third, platform incentives are not one-shot interventions. Workers repeatedly face contracts over time, and prior contracts can change the behavioral state with which they encounter later contracts.

We build a parsimonious model around these features. Contracts are characterized by two primitives: a *qualification metric*, which determines what behavior makes a worker eligible for a bonus, and a *reward horizon*, which determines whether current qualification is rewarded immediately or changes the value of future participation within the cycle. Prior contracts affect future responses through two inherited states: a participation friction K_t and an earnings reference point r_t. The first captures routine formation and reactivation costs; the second captures reference-point adaptation. The model does not attempt to structurally estimate the psychological primitives underlying these states. Instead, it uses them to derive which incentive sequences should be informative, when static contract rankings can fail, and why the experiment must estimate predecessor-successor responses rather than only current-treatment effects.

3.1. Worker decisions and output

Time is divided into contract cycles $t = 1, 2$, each consisting of D subperiods indexed by $d = 1, \ldots, D$. In the field experiment, a cycle is a week and a subperiod is a day. The platform posts a contract $z_t \in \mathcal{Z}$ at the start of cycle t. The contract applies uniformly to all eligible workers in the market during that cycle.

A worker enters cycle t with inherited state

$$\Omega_t = (K_t, r_t),$$

where K_t is a participation friction and r_t is an earnings reference point. In each subperiod d, the worker chooses participation $x_{td} \in \{0, 1\}$, online hours $h_{td} \geq 0$, and acceptance intensity $a_{td} \in [0, 1]$. The variable a_{td} summarizes the worker's willingness to accept dispatched assignments, with larger values corresponding to lower selectivity. If $x_{td} = 0$, then $h_{td} = 0$.

Conditional on participation, completed orders are stochastic:

$$Q_{td} \sim F(\cdot; \mu_{td}), \qquad \mu_{td} = \theta_{td} h_{td} a_{td}, \tag{1}$$

where $\theta_{td} > 0$ captures market conditions, including demand intensity, matching ease, traffic conditions, and local order availability. The family $F(\cdot; \mu)$ is first-order stochastically increasing in μ. For an order threshold $\bar{Q}$, let

$$A_Q(h, a, \theta; \bar{Q}) = \Pr(Q_{td} \geq \bar{Q} \mid \mu = \theta h a)$$

denote the probability of attaining the threshold. We assume A_Q is increasing in h, a, and θ, and has increasing differences in (a, θ). This condition captures the operational fact that increasing acceptance intensity is more valuable when the market provides enough demand for accepted assignments to translate into completed orders.

Worker earnings in subperiod d are

$$Y_{td} = p_{td} Q_{td} + \pi_z(s_{td}, \chi_{td}, d), \tag{2}$$

where p_{td} is baseline pay per completed order, $\chi_{td} \in \{0, 1\}$ is the contract-specific qualification indicator, s_{td} is the within-cycle progress state, and $\pi_z(\cdot)$ is the incentive payout rule.

Per-subperiod utility is

$$u_{td}(z) = x_{td} \left[\mathbb{E} \left\{ Y_{td} - \lambda (r_t - Y_{td})_+ \right\} - K_t - c(h_{td}) - h_{td} g(a_{td}) \right], \tag{3}$$

where $\lambda \geq 0$ is the loss-aversion parameter, $c(\cdot)$ is the disutility of online time, and $h_{td}g(a_{td})$ is the cost of accepting less attractive work. The functions c and g are continuously differentiable, increasing, and convex, and the feasible choice set is compact.

The reference-dependent term follows the idea that workers evaluate earnings relative to an endogenous benchmark rather than only in levels (Kahneman and Tversky 1979, Kőszegi and Rabin 2006). In the present setting, this benchmark is shaped by prior platform experience. A predecessor contract that raises recent earnings can therefore make a later contract less attractive, even if the later contract is valuable in absolute terms.

3.2. Contract primitives

A contract $z \in \mathcal{Z}$ consists of a qualification metric, a reward horizon, and contract parameters such as thresholds, bonus amounts, peak-hour requirements, mandatory days, and tier-specific eligibility rules. We focus on two qualification metrics and two reward horizons.

Definition 1 (Qualification metric) *A contract is* availability-based *if qualification depends on online hours. In the simplest case,*

$$\chi^H_{td} = \mathbf{1}\{h_{td} \geq \bar{H}\}.$$

A contract is throughput-based *if qualification depends on completed orders:*

$$\chi^Q_{td} = \mathbf{1}\{Q_{td} \geq \bar{Q}\}.$$

Definition 2 (Reward horizon) *A contract has a* reset horizon *if payout depends only on current-subperiod qualification:*

$$\pi_z(s, \chi, d) = b(z)\chi.$$

A contract has a state-dependent horizon *if current qualification changes a progress state,*

$$s_{t,d+1} = T_z(s_{td}, \chi_{td}, d),$$

and thereby changes payout opportunities in later subperiods of the same cycle.

Crossing these two primitives yields the four incentive families used in the field experiment:

	Reset horizon	**State-dependent horizon**
Availability metric	Daily Hours (DH)	Streak Hours (SH)
Throughput metric	Daily Orders (DO)	Streak Orders (SO)

The metric dimension determines which behavioral margin the contract directly activates. Availability-based contracts reward presence. Conditional on being online, the bonus does not directly depend on accepting a larger share of assignments. Throughput-based contracts reward conversion. Raising acceptance intensity increases the probability of meeting an order threshold, and this marginal value is larger when market conditions are favorable. The horizon dimension determines whether qualification is memoryless within the cycle. Reset contracts reward current qualification in isolation. State-dependent contracts create continuation value because qualification today changes the value of future participation.

3.3. Within-cycle progress value

The worker's value from subperiod d onward is

$$V_d(K_t, r_t, s_{td}; z) = \max_{x,h,a} \left\{ u_{td}(z) + \beta \mathbb{E}\left[V_{d+1}(K_t, r_t, s_{t,d+1}; z) \mid s_{td}, h, a \right] \right\}, \tag{4}$$

with terminal value normalized to zero after subperiod D. The incremental expected value of qualifying in subperiod d at progress state s is

$$\begin{aligned} \Delta_d(s; z) = {} & \pi_z(s, 1, d) - \pi_z(s, 0, d) \\ & + \beta \mathbb{E}\left[V_{d+1}(K_t, r_t, T_z(s, 1, d); z) - V_{d+1}(K_t, r_t, T_z(s, 0, d); z) \right]. \end{aligned} \tag{5}$$

We refer to $\Delta_d(s; z)$ as the progress premium. Under a reset contract, the continuation term is zero and the progress premium equals the immediate bonus. Under a state-dependent contract, the continuation term can be positive because qualification advances the worker toward a future payout.

Assumption 1 (Progress premium) *For any state-dependent contract z, the transition rule $T_z(s, \chi, d)$ is weakly increasing in s and χ. In addition, the continuation value gain from qualifying is weakly increasing in current progress:*

$$V_{d+1}(K, r, T_z(s, 1, d); z) - V_{d+1}(K, r, T_z(s, 0, d); z)$$

is weakly increasing in s.

Assumption 1 captures the goal-gradient logic underlying streak contracts. When a worker is close to completing a streak or segment, qualifying today can be especially valuable because it preserves eligibility for a later reward. The assumption generates a within-cycle prediction: under state-dependent contracts, effort and participation should increase with progress toward completion, whereas reset contracts should not exhibit the same progress gradient. The appendix verifies the condition in a canonical streak environment and the empirical analysis tests the corresponding within-week pattern.

3.4. Cross-cycle state transitions

Sequencing enters because current contracts affect states that persist into the next cycle. Let

$$N_t(z) = \mathbb{E}\left[\sum_{d=1}^{D} x^*_{td}(z)\right]$$

denote expected active subperiods under contract z, and let

$$\bar{Y}_t(z) = \frac{1}{D}\mathbb{E}\left[\sum_{d=1}^{D} Y^*_{td}(z)\right]$$

denote expected mean earnings per subperiod. The inherited states evolve according to

$$K_{t+1} = \Phi(K_t, N_t(z_t)), \qquad r_{t+1} = \Psi(r_t, \bar{Y}_t(z_t)). \tag{6}$$

Assumption 2 (State transitions) *The function Φ is weakly increasing in K_t and strictly decreasing in $N_t(z_t)$. The function Ψ is weakly increasing in both arguments and strictly increasing in $\bar{Y}_t(z_t)$.*

The transition for K_t captures a routine channel. Recent participation lowers future activation costs, either because working becomes more habitual or because workers accumulate platform-specific knowledge about routes, demand windows, and operating routines (Wood and Neal 2007, Allon et al. 2023). The transition for r_t captures a paydown channel. Higher recent earnings raise the benchmark against which later contracts are evaluated, so a successor with lower or more uncertain earnings is more likely to fall below the inherited reference point (Kőszegi and Rabin 2006, Camerer et al. 1997, Crawford and Meng 2011).

Assumption 3 (Monotone successor response) *For every successor contract z, the market-level outcome functions $\mathcal{N}(K,r;z)$, $\mathcal{H}(K,r;z)$, and $\mathcal{Q}(K,r;z)$ are weakly decreasing in both K and r.*

Here $\mathcal{N}(K,r;z)$, $\mathcal{H}(K,r;z)$, and $\mathcal{Q}(K,r;z)$ denote aggregate participation, online hours, and completed orders under successor contract z, evaluated at inherited state (K,r). These functions aggregate optimal worker choices over the distribution of riders in a region-week. Assumption 3 states that higher activation frictions and higher earnings reference points weakly reduce the response to any successor contract.

3.5. Platform objective

The platform's one-cycle payoff from posting contract z_t in state (K_t, r_t) is

$$\Pi_t(K_t, r_t; z_t) = R_t\left(\mathcal{H}(K_t, r_t; z_t), Q(K_t, r_t; z_t)\right) - C(K_t, r_t; z_t),$$

where $R_t(\cdot)$ is the gross value of acquired availability and completed service, and $C(\cdot)$ is expected incentive cost. The two-cycle design problem is

$$\max_{z_1, z_2 \in \mathcal{Z}} \left\{\Pi_1(K_1, r_1; z_1) + \delta \mathbb{E}\left[\Pi_2(K_2, r_2; z_2)\right]\right\},$$

where (K_2, r_2) is induced by the predecessor contract through (6).

This problem illustrates why a platform cannot generally rank contracts by contemporaneous performance alone. A contract may generate high output in the current cycle but leave the platform with an elevated reference point that weakens the next contract. Another contract may generate more modest current output but lower future activation frictions and improve the effectiveness of a later throughput-oriented contract. The relevant managerial object is therefore not only the effect of contract z, but the effect of contract z after predecessor z'. The next section formalizes this point and derives the implications for experimental design.

4. Theory and Design Implications Revised

The model yields two sets of implications. The first concerns contract form: the qualification metric determines which behavioral margin is directly rewarded, while the reward horizon determines whether current effort creates continuation value within the cycle. The second concerns sequencing: because contracts change inherited participation frictions and earnings reference points, the effect of a current contract depends on the predecessor-induced state in which it is implemented. This section formalizes these implications and shows why they lead naturally to a history-aware switchback design.

4.1. Contract primitives and behavioral margins

We first compare contracts that differ in their primitives but are evaluated at the same inherited state. The metric result is about the margin the contract activates. The horizon result is about how incentives evolve within a cycle.

Proposition 1 (Contract primitive effects) *Fix an inherited state* (K, r).

1. Metric. *On states where the order threshold is attainable but not guaranteed, a throughput-based contract adds a bonus-related return to acceptance intensity relative to an availability-based contract. This incremental return is increasing in market conditions θ.*
2. Horizon. *Under Assumption 1, the progress premium $\Delta_d(s;z)$ is weakly increasing in progress s for any state-dependent contract and is constant in s for any reset contract.*

The metric result clarifies why hours-based and orders-based incentives should not be interpreted as interchangeable increases in compensation. Under an availability-based contract, the bonus depends on time online and does not directly reward accepting marginal assignments. Under a throughput-based contract, the bonus depends on the probability of clearing an order threshold. Since that probability is increasing in acceptance intensity and has increasing differences in acceptance and market conditions, the throughput contract creates the strongest incremental return to acceptance precisely in markets where additional accepted assignments are most likely to convert into completed orders.

The horizon result clarifies why daily and streak incentives are behaviorally distinct even when they reward the same metric. A reset contract rewards current qualification in isolation. A state-dependent contract makes current qualification valuable because it changes the worker's progress toward a later payout. The resulting progress premium generates a within-cycle prediction: effort and participation should rise as workers approach completion under state-dependent contracts, whereas reset contracts should not exhibit the same progress gradient.

4.2. Sequencing through inherited states

We next turn to sequencing. For any predecessor z_1 and successor z_2, define the successor-cycle output as

$$Q_2(z_1,z_2) = Q\left(\Phi(K_1,N_1(z_1)),\Psi(r_1,\bar{Y}_1(z_1));z_2\right).$$

Analogous notation applies to participation $\mathcal{N}_2$, online hours $\mathcal{H}_2$, and incentive cost C_2.

Proposition 2 (Sequencing decomposition) *Within the baseline model, predecessor effects on successor-cycle outcomes operate through the inherited states (K_2,r_2). In particular,*

$$Q_2(z_1,z_2) = Q\left(\Phi(K_1,N_1(z_1)),\Psi(r_1,\bar{Y}_1(z_1));z_2\right),$$

with analogous decompositions for $\mathcal{N}_2$, $\mathcal{H}_2$, and C_2.

This decomposition disciplines the interpretation of history dependence. The predecessor does not enter the successor response as an unrestricted sequence fixed effect. It matters because it changes active participation and earnings in the predecessor cycle, which in turn update the participation friction and the earnings reference point inherited by the successor. This structure yields two channels. A predecessor that increases participation lowers future activation frictions through the routines channel. A predecessor that raises earnings elevates future benchmarks through the paydown channel.

The following definitions organize the main sequencing predictions.

Definition 3 (Sequence classes) *Let $\mathcal{P} \subseteq \mathcal{Z}$ denote participation-building predecessors, which induce relatively high $N_1(z)$. Let $\mathcal{R} \subseteq \mathcal{Z}$ denote earnings-setting predecessors, which induce relatively high $\bar{Y}_1(z)$. Let $\mathcal{T} \subseteq \mathcal{Z}$ denote throughput-intensive successors, for which a reduction in participation friction generates a relatively large increase in completed orders.*

These classes are not defined by the successor-period outcome being tested. In the empirical analysis, participation-building and earnings-setting predecessors are classified using predecessor-period intermediates such as active days and mean earnings, while throughput intensity is tied to the contract primitive that rewards completed orders and to observed sensitivity of output to inherited participation.

Lemma 1 (Routine channel) *If $z^A \in \mathcal{P}$ and $z^B \notin \mathcal{P}$ with $N_1(z^A) > N_1(z^B)$, then, holding the inherited reference point fixed,*

$$Q_2(z^A, z) \geq Q_2(z^B, z) \qquad \text{for every successor } z.$$

Lemma 2 (Paydown channel) *If $z^A \in \mathcal{R}$, $z^B \notin \mathcal{R}$, and $\bar{Y}_1(z^A) > \bar{Y}_1(z^B)$, then, holding inherited participation friction fixed,*

$$Q_2(z^A, z) \leq Q_2(z^B, z) \qquad \text{for every successor } z,$$

with strict utility loss whenever $\lambda > 0$ and the successor exposes workers to positive downside risk, $\Pr(Y_{td} < r_2(z^A)) > 0$.

The two channels run in opposite directions. The routine channel is dynamically beneficial: a predecessor that gets workers to show up lowers the future cost of activation and improves the state in which the successor is implemented. The paydown channel is dynamically costly: a predecessor that raises earnings also raises the benchmark future contracts must satisfy. Hence a contract can be attractive in isolation but unattractive as a predecessor if it leaves workers with expectations that subsequent contracts cannot sustain.

For later interpretation, note that the gain-loss component implies

$$\frac{\partial}{\partial r}\mathbb{E}\left[Y - \lambda(r - Y)_+\right] = -\lambda \Pr(Y < r). \tag{7}$$

Thus the paydown channel is strongest when the successor's earnings distribution places substantial mass below the inherited reference point.

4.3. Sequencing complementarity

The routine channel is especially valuable when the successor can convert a larger active workforce into completed output. This complementarity is formalized through an increasing-differences condition.

Assumption 4 (Increasing differences in successor output) *The function $Q(K, r; z)$ has increasing differences in $(-K, z)$ with respect to the throughput-intensity ordering. Equivalently, the output gain from reducing K is weakly larger for successors in $\mathcal{T}$ than for successors outside $\mathcal{T}$.*

The assumption captures a simple operational idea. A throughput-intensive successor is useful only if workers are active enough for the acceptance and conversion margins to matter. When participation frictions are high, few workers are available to respond to the successor. When a participation-building predecessor has lowered those frictions, a throughput-intensive successor can convert the active worker base into completed orders.

Theorem 1 (Sequencing complementarity) *Let $z^A \in \mathcal{P}$ and $z^B \notin \mathcal{P}$ with $N_1(z^A) > N_1(z^B)$. Let $z^H \in \mathcal{T}$ and $z^L \notin \mathcal{T}$. Under Assumptions 2, 3, and 4, holding the inherited reference point fixed,*

$$Q_2(z^A, z^H) - Q_2(z^B, z^H) \geq Q_2(z^A, z^L) - Q_2(z^B, z^L).$$

Theorem 1 is the core sequencing result. It says that the value of a participation-building predecessor is larger when the successor is throughput-intensive. The implication is not merely that history matters. It is that the value of a predecessor depends on what follows, and the value of a successor depends on the state created by what came before. Contract performance is therefore a property of sequences, not only of individual contracts.

4.4. Ranking reversal

The complementarity result implies that static contract rankings can fail. Define

$$G(K,r) = Q(K,r;z^H) - Q(K,r;z^L),$$

where $z^H \in \mathcal{T}$ is throughput-intensive and $z^L \notin \mathcal{T}$ is less throughput-intensive.

Theorem 2 (Ranking reversal) *Suppose the conditions of Theorem 1 hold and Assumption 4 holds strictly. Let*

$$K^{hi} = K_2(z^B), \qquad K^{lo} = K_2(z^A),$$

with $K^{hi} > K^{lo}$. If

$$G(K^{hi},r) < 0 < G(K^{lo},r), \tag{8}$$

then

$$Q_2(z^B,z^L) > Q_2(z^B,z^H) \qquad \text{and} \qquad Q_2(z^A,z^H) > Q_2(z^A,z^L).$$

Thus, the ranking of successors reverses across predecessor-induced states.

The reversal condition is substantive and testable. It requires the participation-building predecessor to reduce K enough that the throughput-intensive successor changes from inferior to superior. When this condition holds, a platform that ranks contracts by average contemporaneous performance may choose the wrong contract for the state it faces. The same successor can be ineffective after a weak predecessor and effective after a participation-building predecessor.

4.5. History-mixture representation

The ranking reversal changes the experimental object. Let $Y_{ct}(z',z)$ denote the potential outcome in region c and cycle t when the predecessor contract is z' and the current contract is z. Define the sequence-specific mean

$$\mu_b(z',z) = \mathbb{E}\left[Y_{ct}(z',z) \mid B_{ct} = b\right], \tag{9}$$

where $B_{ct} = b$ denotes a design block, such as a calendar week, market stratum, or their interaction.

Assumption 5 (First-order behavioral exposure) *For the outcomes studied here, treatment history is summarized by the predecessor-successor pair:*

$$Y_{ct}(\mathbf{z}_{1:t}) = Y_{ct}(Z_{c,t-1}, Z_{ct}).$$

The exposure assumption does not rule out longer-run persistence. Rather, it defines the first-order behavioral exposure targeted by the experiment. Recovery weeks in the field design reduce sensitivity to earlier histories and allow empirical checks for longer carryover.

Theorem 3 (History-mixture representation) *Let*

$$w_{z'|z,b} = \Pr(Z_{c,t-1} = z' \mid Z_{ct} = z, B_{ct} = b)$$

be the predecessor weights induced by the assignment schedule within block b. Under random assignment and Assumption 5,

$$\mathbb{E}[Y_{ct}^{obs} \mid Z_{ct} = z, B_{ct} = b] = \sum_{z' \in \mathcal{Z}} w_{z'|z,b}\, \mu_b(z', z).$$

Theorem 3 shows that a current-treatment mean is a design-weighted mixture of sequence-specific means. The weights are not structural parameters of the market. They are induced by the assignment schedule. If $\mu_b(z', z)$ varies with z', then changing the predecessor distribution changes the estimand recovered by a current-treatment comparison.

Corollary 1 (Incomplete static ranking) *Under the reversal condition in Theorem 2, there exist predecessor weight distributions for which*

$$\mathbb{E}[Y \mid Z = z^H] > \mathbb{E}[Y \mid Z = z^L],$$

and other predecessor weight distributions for which

$$\mathbb{E}[Y \mid Z = z^H] < \mathbb{E}[Y \mid Z = z^L].$$

A design that varies only current treatments identifies an incomplete policy object when the platform's decision is which sequence to deploy.

This result complements the switchback literature. Existing designs provide rigorous tools for estimating causal effects when treatment is assigned at a market-time level and outcomes may depend on recent treatments (Bojinov and Shephard 2019, Bojinov et al. 2021, 2023, Chen and Simchi-Levi 2025, Xiong et al. 2024, Ni 2025). Our point is substantive rather than corrective. In incentive sequencing, predecessor history is not only a source of contamination. It is part of the treatment exposure because prior contracts shape worker states. The policy-relevant object is therefore $\mu_b(z', z)$, the performance of a current contract conditional on the predecessor-induced state.

4.6. Implications for switchback design

The model implies three requirements for an experiment on incentive sequencing.

First, the treatment menu must vary the contract primitives that theory identifies as behaviorally distinct. Varying only payout generosity would not separate the availability and throughput margins, nor would it separate reset incentives from state-dependent incentives. The field experiment therefore crosses qualification metric with reward horizon.

Second, the assignment schedule must generate support for predecessor-successor pairs. A sequence-specific mean $\mu_b(z', z)$ is not identified for histories that are never assigned. Full balance over every ordered pair may be infeasible in a field experiment with limited market-weeks, but the design should create support for the histories theory predicts to be most informative: participation-building predecessors followed by throughput-intensive successors, and earnings-setting predecessors followed by modest successors.

Third, assignment should occur at the market-time level. Riders in the same local market compete for a common pool of orders and interact through dispatch and congestion. Individual-level randomization would alter the environment faced by both treated and untreated riders. Region-week assignment instead matches the platform's operating decision and estimates the market-level policy effect of posting an incentive menu in a given market and week.

The field experiment implements these requirements through a history-aware switchback design. The design rotates four incentive families, Daily Hours, Daily Orders, Streak Hours, and Streak Orders, plus a no-incentive baseline, across region-weeks. Each monthly cycle includes a recovery week followed by two incentive weeks, so the Week-2 contract serves as the predecessor for Week-3 outcomes. Within each incentive week, assignments are balanced across win, swing, and lose market strata. The design is history-aware rather than fully history-balanced: it is constructed to create the predecessor-successor contrasts implied by the model, while preserving operational feasibility.

4.7. Empirical predictions

The theory yields five empirical implications.

1. **Metric.** Throughput-based contracts should generate larger changes in acceptance rates than availability-based contracts, with the gap increasing in market strength.
2. **Horizon.** State-dependent contracts should generate within-cycle effort acceleration as workers approach completion; reset contracts should not.
3. **Routines.** For $z^A \in \mathcal{P}$, $z^B \notin \mathcal{P}$, and $z^H \in \mathcal{T}$,

$$\mu(z^A, z^H) > \mu(z^B, z^H).$$

4. **Ranking reversal.** The ranking of a throughput-intensive successor z^H and a less throughput-intensive successor z^L can reverse across predecessor states:

$$z^L \succ z^H \text{ after weak participation predecessors,}$$

while

$$z^H \succ z^L \text{ after participation-building predecessors.}$$

5. **Paydown.** For $z^A \in \mathcal{R}$, $z^B \notin \mathcal{R}$, and a modest successor z,

$$\mu(z^A, z) < \mu(z^B, z),$$

especially when the successor exposes workers to substantial downside risk relative to the inherited reference point.

The empirical analysis estimates sequence-specific outcome matrices for online hours, completed orders, acceptance rates, and incentive cost. It then uses predecessor-week active days and earnings to assess the routine and paydown channels, and examines learning as an alternative mechanism using worker experience and persistence through recovery weeks.

5. Experimental Design

The theory in §3 and §4 implies that an informative experiment must do more than compare one incentive against another. First, the treatment menu must vary the two contract primitives that shape worker behavior: the qualification metric and the reward horizon. Second, the assignment schedule must generate predecessor-successor variation, because the effect of a current contract may depend on the state created by the previous contract. Third, assignment must occur at a market-time level rather than at the individual-worker level, because riders operating in the same local market interact through a common order pool, dispatch system, and congestion environment. The field experiment was designed to satisfy these requirements while remaining feasible within the platform's operating constraints.

5.1. Experimental unit and market sample

The experiment was implemented in 30 expansion regions outside Bangkok. The unit of assignment is a region-week. In each treated region-week, the platform posted one incentive menu that applied to all eligible riders operating in that region during that week. This assignment level matches the platform's actual policy decision: which incentive menu to post in a market for a given planning period. It also avoids the most direct form of within-market interference that would arise under rider-level randomization, where treated and untreated riders would compete for the same orders while facing different incentives.

Before the experiment, the platform classified the 30 regions into three market strata based on historical performance and competitive conditions: ten win markets, ten swing markets, and ten lose markets. These strata were operationally meaningful to the platform because they captured differences in market maturity, baseline supply conditions, and competitive pressure. They also improved the experimental design by allowing treatment assignment to be balanced within broad market types.

5.2. Treatment menu

The treatment menu follows directly from the contract space in the model. The first dimension is the qualification metric. Availability-based incentives reward online hours, while throughput-based incentives reward completed orders. The second dimension is the reward horizon. Reset incentives determine qualification over a daily horizon, while state-dependent incentives condition payout on progress across multiple days or segments within the week.

Crossing these dimensions yields four incentive families:

$$\{\mathrm{DH}, \mathrm{DO}, \mathrm{SH}, \mathrm{SO}\},$$

where DH denotes Daily Hours, DO denotes Daily Orders, SH denotes Streak Hours, and SO denotes Streak Orders. The experiment also includes a no-incentive baseline, denoted by 0, in which riders received the platform's standard compensation but no additional region-specific weekly incentive.

The treatment menu is therefore not an arbitrary collection of contracts. DH and DO isolate the metric dimension under a reset horizon. SH and SO isolate the metric dimension under a state-dependent horizon. DH and SH compare reset versus state-dependent incentives holding the availability metric fixed. DO and SO compare reset versus state-dependent incentives holding the throughput metric fixed. These contrasts map directly to the primitive effects in Proposition 1.

5.3. Timing and assignment schedule

The experiment ran over two consecutive months, July 2025 and August 2025. Each month followed a three-week cadence. Week 1 was a recovery week with no region-specific weekly incentive. Weeks 2 and 3 were incentive weeks, during which each region was assigned to one of the four incentive families or to the no-incentive baseline.

This cadence serves both operational and inferential purposes. Operationally, it matched the platform's planning rhythm for regional incentives. Inferentially, the recovery week attenuates behavioral states inherited from the prior month, while the two consecutive incentive weeks generate the predecessor-successor variation central to the theory. In each month, the Week-2 assignment serves as the predecessor for Week-3 outcomes. The primary sequence-specific contrasts therefore compare outcomes under the same Week-3 contract after different Week-2 predecessors.

Assignment was balanced within each incentive week and market stratum. Let C_g denote the set of regions in stratum $g \in \{\text{win}, \text{swing}, \text{lose}\}$. For every incentive week t and every condition $z \in \{\text{DH}, \text{DO}, \text{SH}, \text{SO}, 0\}$, the design satisfies

$$\sum_{c \in C_g} \mathbf{1}\{Z_{ct} = z\} = 2 \qquad \text{for each } g,$$

and therefore

$$\sum_{c=1}^{30} \mathbf{1}\{Z_{ct} = z\} = 6.$$

Thus, every incentive week contains all five conditions, and each condition appears in two win markets, two swing markets, and two lose markets. This balance reduces reliance on functional-form assumptions about common week shocks and broad market-type differences. Identification of contemporaneous treatment effects comes from comparisons across regions within the same calendar week and stratum, while identification of sequence effects comes from the predecessor-successor variation generated between Weeks 2 and 3.

Table 1 reports the full region-by-week assignment schedule. The schedule is history-aware rather than fully history-balanced. It creates support for the theoretically relevant predecessor-successor contrasts, while preserving weekly treatment balance and operational feasibility. In the empirical section, we report the realized transition-count matrix to make clear which ordered pairs are identified with direct experimental support.

5.4. Implementation of incentive families

All incentive menus were implemented through the platform's standard posting and payout systems. Within each incentive family, thresholds and payouts varied by rider performance tier. Riders were assigned to three tiers, low, mid, and top, based on recent activity and performance. The platform used tier-specific requirements to keep targets attainable for less active riders while maintaining meaningful goals for highly active riders. In the empirical analysis, we distinguish pre-experiment tier heterogeneity from endogenous changes in activity after treatment.

Table 1 **Region-by-week incentive assignment schedule in the 30-city experiment.**

Region	Cluster	July 2025 W2 scheme	July 2025 W3 scheme	August 2025 W2 scheme	August 2025 W3 scheme
Saraburi (Phraphutthabat)	Win	○ Daily Hours	■ Streak Orders	● Streak Hours	— No Incentive
Chumphon	Win	○ Daily Hours	● Streak Hours	○ Daily Hours	— No Incentive
Nakhonnayok	Win	■ Streak Orders	□ Daily Orders	□ Daily Orders	● Streak Hours
Sakonnakhon	Win	■ Streak Orders	— No Incentive	■ Streak Orders	● Streak Hours
Saraburi (Nongkhae)	Win	□ Daily Orders	○ Daily Hours	— No Incentive	□ Daily Orders
Sisaket	Win	□ Daily Orders	— No Incentive	○ Daily Hours	□ Daily Orders
Trat	Win	● Streak Hours	■ Streak Orders	— No Incentive	■ Streak Orders
Sakaeo (Aranyaprathet)	Win	● Streak Hours	□ Daily Orders	□ Daily Orders	■ Streak Orders
Krabi	Win	— No Incentive	○ Daily Hours	● Streak Hours	○ Daily Hours
Yasothon	Win	— No Incentive	● Streak Hours	■ Streak Orders	○ Daily Hours
Nakhonsrithammarat	Swing	○ Daily Hours	■ Streak Orders	● Streak Hours	— No Incentive
Nongkhai	Swing	○ Daily Hours	● Streak Hours	○ Daily Hours	— No Incentive
Surin	Swing	■ Streak Orders	□ Daily Orders	□ Daily Orders	● Streak Hours
Tak (Maesot)	Swing	■ Streak Orders	— No Incentive	■ Streak Orders	● Streak Hours
Phayao	Swing	□ Daily Orders	○ Daily Hours	— No Incentive	□ Daily Orders
Lampang	Swing	□ Daily Orders	— No Incentive	○ Daily Hours	□ Daily Orders
Phetchaburi (Chaam)	Swing	● Streak Hours	■ Streak Orders	— No Incentive	■ Streak Orders
Prachinburi	Swing	● Streak Hours	□ Daily Orders	□ Daily Orders	■ Streak Orders
Kanchanaburi	Swing	— No Incentive	○ Daily Hours	● Streak Hours	○ Daily Hours
Ratchaburi (Banpong)	Swing	— No Incentive	● Streak Hours	■ Streak Orders	○ Daily Hours
Chanthaburi	Lose	○ Daily Hours	■ Streak Orders	● Streak Hours	— No Incentive
Suratthani (City)	Lose	○ Daily Hours	● Streak Hours	○ Daily Hours	— No Incentive
Nakhonsawan	Lose	■ Streak Orders	□ Daily Orders	□ Daily Orders	● Streak Hours
Khonkaen	Lose	■ Streak Orders	— No Incentive	■ Streak Orders	● Streak Hours
Udonthani	Lose	□ Daily Orders	○ Daily Hours	— No Incentive	□ Daily Orders
Songkhla (City)	Lose	□ Daily Orders	— No Incentive	○ Daily Hours	□ Daily Orders
Hat Yai	Lose	● Streak Hours	■ Streak Orders	— No Incentive	■ Streak Orders
Uttaradit	Lose	● Streak Hours	□ Daily Orders	□ Daily Orders	■ Streak Orders
Chiangrai	Lose	— No Incentive	○ Daily Hours	● Streak Hours	○ Daily Hours
Nakhonpathom	Lose	— No Incentive	● Streak Hours	■ Streak Orders	○ Daily Hours

Legend:

○ Daily Hours	(single-day, hours-based)	● Streak Hours	(multi-day, hours-based)
□ Daily Orders	(single-day, orders-based)	■ Streak Orders	(multi-day, orders-based)
— No Incentive			(control period)

Let H_{icdt} denote rider i's total online hours in region c on day d of week t, and let H_{icdt}^{peak} denote online hours during platform-defined peak windows. Let Q_{icdt} denote completed orders. Let $g(i,c,t) \in \{\text{low}, \text{mid}, \text{top}\}$ denote rider i's tier for the relevant region-week.

Under Daily Hours, a rider qualified for a daily bonus by satisfying both a total-hours requirement and a peak-hours requirement:

$$\mathbf{1}\{\text{Qualify}_{icdt}^{DH} = 1\} = \mathbf{1}\{H_{icdt} \geq \bar{H}_{g(i,c,t)}\} \cdot \mathbf{1}\{H_{icdt}^{\text{peak}} \geq \bar{H}_{g(i,c,t)}^{\text{peak}}\}.$$

The peak-hours requirement reflects the platform's objective of acquiring availability when it is operationally most valuable.

Under Daily Orders, a rider qualified for a daily bonus by completing at least a tier-specific number of orders:

$$\mathbf{1}\{\text{Qualify}_{icdt}^{DO} = 1\} = \mathbf{1}\{Q_{icdt} \geq \bar{Q}_{g(i,c,t)}\}.$$

Unlike DH, this contract rewards realized throughput and therefore directly interacts with acceptance behavior and market conditions.

The two state-dependent contracts preserve the same metric distinction while adding within-week continuation value. Under Streak Hours, riders had to satisfy the hours requirement on a specified number of days within the week and, in some implementations, on mandatory high-demand days. Let $\mathcal{M}$ denote the mandatory-day set and let K_g denote the required number of qualifying days for tier g. Weekly qualification is

$$\mathbf{1}\{\text{Earn}_{ict}^{SH} = 1\} = \mathbf{1}\left\{\sum_{d=1}^{D} \text{Qualify}_{icdt}^{DH} \geq K_{g(i,c,t)}\right\} \cdot \prod_{d \in \mathcal{M}} \text{Qualify}_{icdt}^{DH}.$$

This structure creates continuation value because a rider who has already accumulated qualifying days has more to lose from missing a later requirement.

Under Streak Orders, the week was divided into several segments, and riders could earn bonuses by satisfying the daily order requirement throughout each segment. Let $\{\mathcal{S}_r\}_{r=1}^{R}$ denote the segment day sets. Segment-r qualification is

$$\mathbf{1}\{\text{Earn}_{ict}^{SO,r} = 1\} = \prod_{d \in \mathcal{S}_r} \text{Qualify}_{icdt}^{DO},$$

and the rider's total weekly SO bonus is the sum across completed segments. This segmented structure creates multiple progress windows within the week and allows the platform to target persistence across operationally important subperiods.

In the no-incentive baseline, riders received the platform's standard compensation without any additional region-specific weekly incentive. Platform-wide business-as-usual policies that applied to all regions were not part of the experimental treatment.

5.5. Design estimands

The design supports two types of estimands. The first is the contemporaneous effect of a current contract family. For outcome Y_{ct}, this object compares Y_{ct} across region-weeks assigned to different current contracts within the same week and market stratum. These comparisons estimate how the four contract families differ in their average current-period effects.

The second, and more central, object is the sequence-specific mean introduced in §4:

$$\mu_b(z', z) = \mathbb{E}[Y_{ct}(z', z) \mid B_{ct} = b],$$

where z' is the predecessor contract, z is the current contract, and B_{ct} denotes the relevant design block. This object answers the platform's sequencing question: how does contract z perform when the market enters the week after predecessor z'? The design identifies these effects by comparing Week-3 outcomes across regions assigned to the same Week-3 contract but different Week-2 predecessors, within the support created by the assignment schedule.

The sequence-specific estimands are especially important for testing the theory. The routines channel predicts that successors should perform better after participation-building predecessors, particularly when the successor is throughput-intensive. The paydown channel predicts weaker responses when high-earning predecessors are followed by more modest successors. The design therefore allows us to estimate not only which contracts work, but also which predecessor-successor transitions are beneficial or costly.

5.6. Outcome measurement

The primary outcomes are chosen to map directly to the behavioral margins in the model. Participation is measured by the number of active riders and the probability that a rider is active in a region-day or region-week. Availability is measured by total online hours and peak-window online hours. Throughput is measured by completed orders and orders per online hour. Acceptance behavior is measured by acceptance rates and related selectivity measures, where available. Cost outcomes include total incentive payout, incentive cost per incremental online hour, and incentive cost per incremental completed order.

We also construct mechanism outcomes. To study progress incentives, we measure within-week changes in participation, hours, and orders as riders approach or lose eligibility for state-dependent bonuses. To study routines, we measure predecessor-week active days and online hours. To study paydown, we measure predecessor-week earnings and the gap between current earnings opportunities and prior earnings benchmarks. To assess learning as an alternative mechanism, we examine heterogeneity by pre-experiment rider tier, tenure, and prior activity, and test whether sequence effects persist after recovery weeks.

5.7. Implementation fidelity and threats to interpretation

The assignment schedule was pre-specified and implemented through the platform's standard incentive-posting mechanism. We verify implementation using administrative records on posted menus, eligibility flags, rider tier assignments, and realized payouts. Deviations due to operational exceptions are documented and addressed in robustness checks.

Three threats are particularly relevant. The first is cross-region mobility. If riders frequently work across region boundaries, treatment may spill over across assignment units. We measure cross-region work patterns and report robustness analyses restricting the sample to riders whose activity is concentrated in a single region during the relevant window. The second is anticipation. If riders learn future assignments before they are posted, predecessor and successor effects may be confounded by expected future incentives. The platform's normal posting process limits such anticipation, and the empirical analysis checks for pre-period changes before incentive weeks. The third is longer-run carryover beyond the predecessor week. The recovery week reduces but need not eliminate longer histories. We therefore test whether outcomes remain sensitive to second-order histories after conditioning on the predecessor contract and observed predecessor-week state variables.

These threats are not peripheral to the design. They reflect the same dynamic and marketplace features that motivate a region-week switchback. The empirical strategy therefore pairs fixed-effects estimates with design-based checks that respect the assignment schedule and reports sequence-specific contrasts within the support of the experiment.

6. Concluding Remarks

References

Allon G, Cohen MC, Moon K, Sinchaisri WP (2025) Managing multihoming workers in the gig economy. *Working Paper* .

Allon G, Cohen MC, Sinchaisri WP (2023) The impact of behavioral and economic drivers on gig economy workers. *Manufacturing & Service Operations Management* 25(4):1376–1393.

Bandiera O, Barankay I, Rasul I (2007) Incentives for managers and inequality among workers: Evidence from a firm-level experiment. *The Quarterly Journal of Economics* 122(2):729–773.

Bastani H, Bastani O, Sinchaisri WP (2025) Improving human sequential decision making with reinforcement learning. *Management Science* .

Bimpikis K, Candogan O, Saban D (2019) Spatial pricing in ride-sharing networks. *Operations Research* 67(3):744–769.

Blake T, Coey D (2014) Why marketplace experimentation is harder than it seems: The role of test-control interference. *Proceedings of the Fifteenth ACM Conference on Economics and Computation (EC '14)*, 567–582 (ACM).

Bojinov I, Rambachan A, Shephard N (2021) Panel experiments and dynamic causal effects: A finite population perspective. *Quantitative Economics* 12(4):1171–1196.

Bojinov I, Shephard N (2019) Time series experiments and causal estimands: exact randomization tests and trading. *Journal of the American Statistical Association* 114(528):1665–1682.

Bojinov I, Simchi-Levi D, Zhao J (2023) Design and analysis of switchback experiments. *Management Science* 69(7):3759–3777.

Bonezzi A, Brendl CM, De Angelis M (2011) Stuck in the middle: The psychophysics of goal pursuit. *Psychological science* 22(5):607–612.

Cachon GP, Daniels KM, Lobel R (2017) The role of surge pricing on a service platform with self-scheduling capacity. *Manufacturing & Service Operations Management* 19(3):368–384.

Camerer C, Babcock L, Loewenstein G, Thaler R (1997) Labor supply of new york city cabdrivers: One day at a time. *The Quarterly Journal of Economics* 112(2):407–441.

Chen H, Simchi-Levi D (2025) Efficient switchback experiments with surrogate variables: Estimation and experimental design. *Management Science* .

Chen K, Feinerman K, Haggag K (2025) Flexible pay and labor supply: Evidence from uber's instant pay. *Management Science* .

Chen MK, Rossi PE, Chevalier JA, Oehlsen E (2019) The value of flexible work: Evidence from uber drivers. *Journal of political economy* 127(6):2735–2794.

Chen MK, Sheldon M (2016) Dynamic pricing in a labor market: Surge pricing and flexible work on the uber platform. *EC'16: Proceedings of the 2016 ACM Conference on Economics and Computation* .

Crawford VP, Meng J (2011) New york city cab drivers' labor supply revisited: Reference-dependent preferences with rational-expectations targets for hours and income. *American Economic Review* 101(5):1912–1932.

Dai H, Milkman KL, Riis J (2014) The fresh start effect: Temporal landmarks motivate aspirational behavior. *Management Science* 60(10):2563–2582.

Dai H, Milkman KL, Riis J (2015) Put your imperfections behind you: Temporal landmarks spur goal initiation when they signal new beginnings. *Psychological science* 26(12):1927–1936.

Farber HS (2015) Why you can't find a taxi in the rain and other labor supply lessons from cab drivers. *The Quarterly Journal of Economics* 130(4):1975–2026.

Goette L, Huffman D, Fehr E (2004) Loss aversion and labor supply. *Journal of the European Economic Association* 2(2–3):216–228.

Hall JV, Krueger AB (2018) An analysis of the labor market for uber's driver-partners in the united states. *Ilr Review* 71(3):705–732.

Huang S, Wang C, Yuan Y, Zhao J, Zhang J (2023) Estimating effects of long-term treatments. *arXiv preprint arXiv:2308.08152* .

Hull CL (1932) The goal-gradient hypothesis and maze learning. *Psychological Review* 39(1):25–43.

Kahneman D, Tversky A (1979) Prospect theory: An analysis of decision under risk. *Econometrica* 47(2):263–291.

Kivetz R, Urminsky O, Zheng Y (2006) The goal-gradient hypothesis resurrected: Purchase acceleration, illusionary goal progress, and customer retention. *Journal of marketing research* 43(1):39–58.

Kleven HJ, Waseem M (2013) Using notches to uncover optimization frictions and structural elasticities: Theory and evidence from pakistan. *The quarterly journal of economics* 128(2):669–723.

Kőszegi B, Rabin M (2006) A model of reference-dependent preferences. *The Quarterly Journal of Economics* 121(4):1133–1165.

Kőszegi B, Rabin M (2006) A model of reference-dependent preferences. *The Quarterly Journal of Economics* 121(4):1133–1165.

Lazear EP (2000) Performance pay and productivity. *American Economic Review* 90(5):1346–1361.

Lee S, Meyer-Doyle P (2017) How performance incentives shape individual exploration and exploitation: Evidence from microdata. *Organization Science* 28(1):19–38.

Locke EA, Latham GP (2002) Building a practically useful theory of goal setting and task motivation: A 35-year odyssey. *American psychologist* 57(9):705.

Ma H, Fang F, Parkes DC (2022) Spatio-temporal pricing for ridesharing platforms. *Operations Research* 70(2):1025–1041.

March JG (1991) Exploration and exploitation in organizational learning. *Organization science* 2(1):71–87.

Mas A, Pallais A (2017) Valuing alternative work arrangements. *American Economic Review* 107(12):3722–3759.

Mehr KS, Silverman J, Sharif MA, Barasch A, Milkman KL (2025) The motivating power of streaks: Increasing persistence is as easy as 1, 2, 3. *Organizational Behavior and Human Decision Processes* 187:104391.

Momentum Works (2024) Food delivery platforms in southeast asia 2024. URL https://momentum.asia/insights/detail/food-delivery-platforms-in-southeast-asia-2024.

Ni T (2025) Decision analytics of switchback experiments: A robust optimization approach. *Available at SSRN 5245482* .

Nunes JC, Drèze X (2006) The endowed progress effect: How artificial advancement increases effort. *Journal of Consumer Research* 32(4):504–512.

Oyer P (1998) Fiscal year ends and nonlinear incentive contracts: The effect on business seasonality. *The Quarterly Journal of Economics* 113(1):149–185.

Posen HE, Levinthal DA (2012) Chasing a moving target: Exploitation and exploration in dynamic environments. *Management science* 58(3):587–601.

Rabin M (1998) Psychology and economics. *Journal of economic literature* 36(1):11–46.

Raveendran M, Srikanth K, Ungureanu T, Zheng GL (2023) How do performance goals influence exploration-exploitation choices? *Organization Science* 34(6):2464–2486.

Saez E (2010) Do taxpayers bunch at kink points? *American economic Journal: economic policy* 2(3):180–212.

Schulz E, Gershman SJ (2019) The algorithmic architecture of exploration in the human brain. *Current opinion in neurobiology* 55:7–14.

Shearer B (2004) Piece rates, fixed wages and incentives: Evidence from a field experiment. *The Review of Economic Studies* 71(2):513–534.

Silverman J, Barasch A (2023) On or off track: How (broken) streaks affect consumer decisions. *Journal of Consumer Research* 49(6):1095–1117.

Silverman J, Barasch AP, Small DA (2023) Hot streak! inferences and predictions about goal adherence. *Organizational Behavior and Human Decision Processes* 179:104281.

Thairath Money (2024) Line man 44% grab 1 . URL https://www.thairath.co.th/money/tech_innovation/tech_companies/2813800.

Tran A, Bibaut A, Kallus N (2024) Inferring the long-term causal effects of long-term treatments from short-term experiments. *Proceedings of Machine Learning Research* .

Tversky A, Kahneman D (1992) Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty* 5(4):297–323.

Wilson RC, Geana A, White JM, Ludvig EA, Cohen JD (2014) Humans use directed and random exploration to solve the explore–exploit dilemma. *Journal of experimental psychology: General* 143(6):2074.

Wood W, Neal DT (2007) Habit and intention in everyday life: The multiple processes by which past behavior predicts future behavior. *Psychological Review* 114(4):843–863.

Wu Y, Johari R, Syrgkanis V, Weintraub GY (2024) Switchback price experiments with forward-looking demand. *arXiv preprint arXiv:2410.14904* .

Xiong R, Chin A, Taylor SJ (2024) Data-driven switchback experiments: Theoretical tradeoffs and empirical bayes designs. *arXiv preprint arXiv:2406.06768* .

E-Companion to *"Designing and Sequencing Incentives for Gig Economy Workers"*

Appendix A: Proofs

Proof of Proposition 1

Metric. Under an availability-based contract, the bonus term depends on h through the qualification indicator $\mathbf{1}\{h \geq \bar{H}\}$ and does not directly depend on acceptance intensity a. Hence the bonus component has no increasing-differences term in (a, θ). Under a throughput-based contract, the expected bonus-relevant component is

$$\Delta_d(s; z) A_Q(h, a, \theta; \bar{Q}),$$

where $\Delta_d(s; z) > 0$ whenever qualification has positive value. By assumption, A_Q has increasing differences in (a, θ). Multiplication by the positive scalar $\Delta_d(s; z)$ preserves increasing differences. Therefore, on states where the order threshold is attainable but not guaranteed, the throughput-based contract adds an incremental return to a, and this incremental return is increasing in θ. This proves the metric claim.

Horizon. Under a reset contract, the within-cycle state is payoff-irrelevant and current qualification does not change future payout opportunities. The continuation term in (5) is therefore zero, so $\Delta_d(s; z) = b(z)$, which is constant in s. Under a state-dependent contract, Assumption 1 states that the continuation value gain from qualifying is weakly increasing in s. Since the current payout difference $\pi_z(s, 1, d) - \pi_z(s, 0, d)$ is nondecreasing in progress for the contract class studied here, the sum of the current payout difference and the continuation value gain is weakly increasing in s. Hence $\Delta_d(s; z)$ is weakly increasing in progress.

To see why Assumption 1 is natural for streak-style contracts, consider a canonical contract in which the worker receives a bonus B if at least k qualifying days are achieved by the end of a D-day cycle. Suppose, for this verification only, that each remaining day has a fixed qualification probability $q \in (0, 1)$. With n days remaining and j additional qualifying days required, the continuation value is

$$B \Pr(\text{Binomial}(n, q) \geq j).$$

Advancing progress by one reduces the remaining requirement from j to $j - 1$. The marginal value of this progress is

$$B \left[\Pr(\text{Binomial}(n, q) \geq j - 1) - \Pr(\text{Binomial}(n, q) \geq j)\right] = B \Pr(\text{Binomial}(n, q) = j - 1).$$

Near completion, as the remaining requirement becomes small relative to the number of remaining opportunities, this marginal value rises as the worker approaches the threshold. This is the canonical goal-gradient region captured by Assumption 1. □

Proof of Proposition 2

By (6), the predecessor contract z_1 affects the cycle-2 inherited states only through

$$K_2 = \Phi(K_1, N_1(z_1)), \qquad r_2 = \Psi(r_1, \bar{Y}_1(z_1)).$$

Given successor z_2, the cycle-2 output function is $Q(K_2, r_2; z_2)$. Substituting the state transitions gives

$$Q_2(z_1, z_2) = Q\left(\Phi(K_1, N_1(z_1)), \Psi(r_1, \bar{Y}_1(z_1)); z_2\right).$$

The same substitution yields the analogous expressions for $\mathcal{N}_2$, $\mathcal{H}_2$, and C_2. □

Proof of Lemma 1

Since $N_1(z^A) > N_1(z^B)$ and Φ is strictly decreasing in participation by Assumption 2,

$$K_2(z^A) < K_2(z^B).$$

Holding the inherited reference point fixed at r, Assumption 3 implies that $Q(K,r;z)$ is weakly decreasing in K. Therefore,

$$Q(K_2(z^A),r;z) \geq Q(K_2(z^B),r;z)$$

for every successor z. This is exactly

$$Q_2(z^A,z) \geq Q_2(z^B,z).$$

□

Proof of Lemma 2

Since $\bar{Y}_1(z^A) > \bar{Y}_1(z^B)$ and Ψ is strictly increasing in prior earnings by Assumption 2,

$$r_2(z^A) > r_2(z^B).$$

For a fixed successor contract and fixed choices, the reference-dependent component satisfies

$$\frac{\partial}{\partial r}\mathbb{E}\left[Y - \lambda(r-Y)_+\right] = -\lambda \Pr(Y<r).$$

Thus, when $\lambda > 0$ and $\Pr(Y < r_2(z^A)) > 0$, increasing the inherited reference point strictly lowers expected utility. Holding K fixed, Assumption 3 implies that the market-level output response is weakly decreasing in r. Therefore,

$$Q(K,r_2(z^A);z) \leq Q(K,r_2(z^B);z),$$

which yields

$$Q_2(z^A,z) \leq Q_2(z^B,z).$$

□

Proof of Theorem 1

By Lemma 1,

$$K^{lo} = K_2(z^A) \leq K^{hi} = K_2(z^B).$$

Assumption 4 states that $Q(K,r;z)$ has increasing differences in $(-K,z)$ with respect to the throughput-intensity ordering. Since $z^H \in \mathcal{T}$ and $z^L \notin \mathcal{T}$, this implies

$$Q(K^{lo},r;z^H) - Q(K^{hi},r;z^H) \geq Q(K^{lo},r;z^L) - Q(K^{hi},r;z^L).$$

Substituting $K^{lo} = K_2(z^A)$ and $K^{hi} = K_2(z^B)$ gives

$$Q_2(z^A,z^H) - Q_2(z^B,z^H) \geq Q_2(z^A,z^L) - Q_2(z^B,z^L).$$

□

Proof of Theorem 2

The first inequality in the theorem follows from $G(K^{hi}, r) < 0$. By definition of G,

$$G(K^{hi}, r) = Q(K^{hi}, r; z^H) - Q(K^{hi}, r; z^L) < 0,$$

which is equivalent to

$$Q(K^{hi}, r; z^L) > Q(K^{hi}, r; z^H).$$

Since $K^{hi} = K_2(z^B)$, this gives

$$Q_2(z^B, z^L) > Q_2(z^B, z^H).$$

The second inequality follows from $G(K^{lo}, r) > 0$. Again by definition,

$$G(K^{lo}, r) = Q(K^{lo}, r; z^H) - Q(K^{lo}, r; z^L) > 0,$$

which is equivalent to

$$Q(K^{lo}, r; z^H) > Q(K^{lo}, r; z^L).$$

Since $K^{lo} = K_2(z^A)$, this gives

$$Q_2(z^A, z^H) > Q_2(z^A, z^L).$$

Strict increasing differences ensure that $G(K, r)$ is strictly decreasing in K. Hence a sufficiently large reduction in K induced by a participation-building predecessor can move $G(K, r)$ from negative to positive. □

Proof of Theorem 3

Under Assumption 5, the observed outcome satisfies

$$Y_{ct}^{obs} = Y_{ct}(Z_{c,t-1}, Z_{ct}).$$

Fix a design block $B_{ct} = b$ and current contract $Z_{ct} = z$. By the law of total expectation over the predecessor contract,

$$\mathbb{E}[Y_{ct}^{obs} \mid Z_{ct} = z, B_{ct} = b] = \sum_{z' \in \mathcal{Z}} \Pr(Z_{c,t-1} = z' \mid Z_{ct} = z, B_{ct} = b) \mathbb{E}[Y_{ct}^{obs} \mid Z_{c,t-1} = z', Z_{ct} = z, B_{ct} = b].$$

Under random assignment and the exposure mapping,

$$\mathbb{E}[Y_{ct}^{obs} \mid Z_{c,t-1} = z', Z_{ct} = z, B_{ct} = b] = \mathbb{E}[Y_{ct}(z', z) \mid B_{ct} = b] = \mu_b(z', z).$$

Substituting this expression and defining

$$w_{z'|z,b} = \Pr(Z_{c,t-1} = z' \mid Z_{ct} = z, B_{ct} = b)$$

yields

$$\mathbb{E}[Y_{ct}^{obs} \mid Z_{ct} = z, B_{ct} = b] = \sum_{z' \in \mathcal{Z}} w_{z'|z,b} \mu_b(z', z).$$

□

Proof of Corollary 1

By Theorem 2, there exist predecessor states u and v such that

$$\mu(u, z^H) - \mu(u, z^L) > 0$$

and

$$\mu(v, z^H) - \mu(v, z^L) < 0.$$

Consider a design that places weight w on predecessor u and weight $1 - w$ on predecessor v. The current-treatment contrast between z^H and z^L is

$$D(w) = w[\mu(u, z^H) - \mu(u, z^L)] + (1 - w)[\mu(v, z^H) - \mu(v, z^L)].$$

At $w = 1$, $D(w) > 0$. At $w = 0$, $D(w) < 0$. Since $D(w)$ is continuous in w, there exists an interior weight at which the sign changes. Thus, the direction of the current-treatment ranking depends on the predecessor weights induced by the design. A design that estimates only current-treatment means therefore does not identify a complete policy ranking when the platform's decision is sequence-specific. □

Appendix B: Additional Detail