

Pricing and Demand Management for Integrated Same-Day and Next-Day Delivery Systems

Dipayan Banerjee^{1,2} Alan Erera¹ Alejandro Toriello¹

¹H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology

²Quinlan School of Business, Loyola University Chicago

October 11, 2024

We study a system in which a common delivery fleet provides service to both same-day delivery (SDD) and next-day delivery (NDD) orders placed by e-retail customers who are sensitive to delivery prices. We develop a model of the system and optimize with respect to two separate objectives. First, empirical research suggests that fulfilling e-retail orders ahead of promised delivery days increases a firm’s long-run market share. Motivated by this phenomenon, we optimize for customer satisfaction by maximizing the quantity of NDD orders fulfilled one day early given fixed prices. Next, we optimize for total profit; we optimize for a single SDD price, and we then set SDD prices in a two-level scheme with discounts for early-ordering customers. Our analysis relies on continuous approximation techniques to capture the interplay between NDD and SDD orders, and particularly the effect one day’s operations have on the next, a novel modeling component not present in SDD-only models; a key technical result is establishing the model’s convergence to a steady state using dynamical systems theory. We derive structural insights and efficient algorithms for both objectives. In particular, we show that, under certain conditions, the total profit is a piecewise-convex function with polynomially-many breakpoints that can be efficiently enumerated. In a case study set in metropolitan Denver, approximately 10% of NDD orders can be fulfilled one day early at optimality, and profit is increased by 1-3% in a two-level pricing scheme versus a one-level scheme. We conduct operational simulations for validation of solutions and analysis of initial conditions.

Key words: e-commerce; last-mile delivery; continuous approximations

1. Introduction

Empirical evidence suggests that fulfillment speed is a key factor in e-retail sales (Deshpande and Pendem 2023). In response to increased customer fulfillment speed expectations, e-retail firms now commonly offer same-day delivery (SDD) and next-day delivery (NDD) options. Historically, e-retail giant Amazon has been a leader in providing shorter fulfillment time guarantees (Amazon 2009). In recent years, a wide variety of specialized e-retailers have also expanded their SDD and/or NDD offerings, such as sports merchandise retailer Fanatics (Salgado 2022) and pharmaceutical delivery firm TRxADE (TRxADE Health 2022), among others. In this paper, we consider last-mile delivery design and planning decisions for firms that offer both SDD and NDD fulfillment options.

E-retail systems with SDD and NDD options require intelligent design and management to ensure that target customer service levels are met while controlling delivery costs. A particular operational challenge for rapid e-retail fulfillment is one of information dynamics: to provide the fastest service, it is sometimes

necessary to begin dispatching orders on some delivery routes before all orders and delivery locations are known and before all routes have been planned. This is in contrast to more traditional last-mile delivery settings, in which all of a day's delivery locations are known at the beginning of the day. Additionally, tactical planning of systems that provide both same-day and next-day delivery requires a clear understanding of how decisions made on a particular day impact the following day's operation. For example, how does the time at which we choose to stop accepting orders for SDD today impact the quantity of NDD orders to be fulfilled tomorrow? Similarly, if we choose to fulfill some of today's NDD orders one day early, how are tomorrow's operations affected?

Recent work on system design for e-retail fulfillment (e.g., Snoeck and Winkenbach 2022, Stroh et al. 2022) has focused solely on SDD. Allocating order demand to delivery vehicles in SDD systems is challenging for two primary reasons. First, delivery deadlines create a time capacity for each vehicle, and this capacity is depleted over the course of the operating day whether the vehicle is performing productive delivery work or not (i.e., waiting or returning empty to a stocking location). Second, batching scale economies favor consolidated delivery routes serving more customers; such routes use less time per individual delivery, but vehicles must wait (and deplete capacity) to build a batch of deliveries. In SDD studies to date, the optimization of SDD system design is considered in isolation, assuming that customers are not offered delivery options with longer lead times or that the SDD system operates independently from other systems used to deliver the longer lead-time orders. In practice, however, non-SDD options are usually available to customers, typically at lower cost. Thus, it is important to understand how potential customers choose between fulfillment options and how those choices create demand for transportation resources; this is particularly important when the same resources may be used to serve demand arising from both fulfillment options.

Some prior research has studied multi-day vehicle routing problems (e.g., Angelelli et al. 2007a,b). In general, these studies focus on short-term optimization of operations and do not consider system design questions. Recent work integrating routing and pricing is similarly focused on detailed optimization of operations in a given day (e.g., Afsar et al. 2021, Afsar 2022, Ulmer 2020). Alternatively, in this paper we combine aspects of these areas of research; we study the design of an e-retail delivery system that incorporates price-driven choice of fulfillment options, and that considers concurrent fulfillment of orders with different delivery day guarantees using common transportation resources.

Specifically, we consider a setting where an e-retailer manages a single fleet of delivery vehicles and uses them to fulfill both same-day and next-day delivery orders in a local region. Our focus is on tactical, medium-term decision-making and management. That is, we are not concerned with optimizing long-term physical investments (such as fulfillment center location) nor day-to-day operational issues (such as sequencing customers on vehicle routes). Rather, we focus on aspects of the system that can be revisited periodically every few weeks or months, such as determining prices and order deadlines. To gain transparent

insights and better understand trade-offs, we use continuous approximation and fluid relaxation techniques, which are often applied in logistics system design.

1.1 Contributions

Our specific contributions are as follows.

1. We study an e-retail fulfillment system with price-differentiated SDD and NDD options serviced by a common fleet; to the best of our knowledge, this work is the first to study the design of such delivery systems. To facilitate system design and transparent managerial insights, we propose a deterministic, steady-state model of the system that uses fluid approximations of stochastic customer arrivals and continuous approximations for vehicle tour durations.
2. Empirical evidence shows that exceeding customer delivery time expectations can grow market share (Fisher et al. 2019). Motivated by this research, we use the model to analyze a customer satisfaction objective of maximizing the average daily quantity of NDD orders delivered one day early. We show that, under certain conditions, some orders can be delivered early even when full utilization of vehicles is required to deliver all orders on the day of their fulfillment guarantee. We then derive an optimal temporal structure of the system and an associated solution procedure. We use the model to assess whether a system can feasibly operate in the long run under a given fleet size and pricing scheme.
3. We next use the model to study static pricing schemes that maximize system profit. We first consider the problem of choosing a single SDD price. To incentivize early-arriving customers to choose SDD, thereby improving the efficiency of fulfillment operations, we also propose a two-level SDD pricing scheme in which early-ordering customers receive a discounted price. For this latter scheme, we derive structural insights and an efficient solution approach under mild assumptions.
4. Finally, we test our model on a system with real-world road network data, and we compare our model outputs to simulations. We find that our model's predictions align with operational realities; in particular, the system is insensitive to initial conditions.

This section concludes with a brief review of the relevant literature. Section 2 introduces the e-retail system under consideration, and Section 3 details the deterministic approximation. Section 4 studies the convergence of the model to a steady state. Section 5 analyzes the customer satisfaction objective, while Section 6 analyzes static pricing schemes for profit maximization. Section 7 presents computational studies of both objectives. Section 8 provides concluding remarks. Supplementary appendices contain proofs, technical details, and data omitted from the main body.

1.2 Literature Review

1.2.1 Last-Mile Delivery Optimization. Much of the ongoing research on vehicle routing problems (VRPs) is motivated by the unique challenges of last-mile delivery for e-commerce. The literature in this area is vast; we refer to Archetti and Bertazzi (2021) for a thorough review of recent work and instead focus on models relevant to this paper, specifically those concerning SDD and NDD routing.

To our knowledge, only one SDD-focused study explicitly offers NDD as a fulfillment option in its optimization model: Ulmer (2020) offers NDD to customers for free, but NDD routing is handled separately from SDD by an independent third-party delivery company. In most other SDD studies, there is no explicitly stated NDD option. Among these studies, a customer request for which the delivery system operator cannot (or wishes not to) offer SDD is usually either rejected (e.g., Chen et al. 2022, Klapp et al. 2020) or assigned to an independent third-party delivery company (e.g., Voccia et al. 2019). Other studies provide SDD to all incoming customer requests — without a NDD option — and seek to minimize a temporal objective, such as total customer waiting time (Ulmer and Streng 2019).

In the literature, the term ‘SDD’ is generally used to refer to one of two different types of fulfillment time guarantees. In the first type, every customer delivery request is associated with its own fulfillment time guarantee based on when the order was placed. For example, Ulmer and Thomas (2018) study a setting in which each order must be fulfilled no later than four hours after it is placed. In the second type, every customer who submits a SDD request prior to a daily order deadline faces the same fulfillment time guarantee. For example, Stroh et al. (2022) aim to fulfill every order placed between 9 AM and 2 PM by 6 PM of the same day. Henceforth in this paper, ‘SDD’ refers to guarantees of the latter type. We also differentiate the context of our work from that of attended home delivery (AHD), in which customers are assigned a specific delivery time slot to facilitate physical transfer of goods. We instead assume that orders may be delivered at any time during the service day.

1.2.2 Multi-Period Routing. Outside of the specific SDD context, a class of routing problems related to our work is that of multi-period VRPs. Such problems consider a time horizon of multiple consecutive periods (i.e., days) during which customer delivery requests arrive and must be fulfilled. Some works in this area (Angelelli et al. 2007a,b, Bülbül 2022, Ulmer et al. 2018) allow the delivery service provider to choose which orders are fulfilled on a particular day. In other studies (Albareda-Sambola et al. 2014, Andreatta and Lulli 2008, Angelelli et al. 2009, 2010, Wen et al. 2010), certain requests must be fulfilled on a particular day or within a particular set of days. These studies focus on short-term routing optimization rather than on the design of broader aspects of the system. Additionally, these studies do not leverage pricing to manage the inflow of customer delivery requests; rather, the request arrival processes are exogenous.

Three recent multi-period routing studies share greater similarities with some aspects of our work. Estrada-Moreno et al. (2019) and Yıldız and Savelsbergh (2020) offer discounts to incentivize certain customers to change their preferred delivery day. Avraham and Raviv (2021) seek an optimal assortment optimization policy in a multi-period AHD routing context. The time horizon is assumed indefinite, so the

objective is to maximize the steady-state proportion of accepted orders. The authors take an “approach in which the conditions under which the system is working are assumed stationary over time, and thus, a simpler stationary policy is searched.” This is philosophically similar to our approach in this paper. Our model also assumes that the system is operating in a steady state, and thus our tactical design decisions, such as the choice of SDD price, are stationary over time.

1.2.3 Pricing in Last-Mile Delivery. In addition to Estrada-Moreno et al. (2019), Ulmer (2020), and Yıldız and Savelsbergh (2020) discussed above, several other studies in recent years have analyzed last-mile delivery systems that seek to manage customer demand via pricing or monetary incentives. We next present a representative sample of such studies.

Afsar (2022) studies a profit maximization problem in which a single delivery price is offered to a set of known potential customers, each of which has a maximum price that they are willing to accept; this maximum price is unknown but distributed independently and uniformly at random. A single vehicle tour services all customers who accept the offered price. Afsar et al. (2021) study a similar problem with multiple vehicles in which all customer information is known in advance; in this variant, each geographical subset within the region may have a different delivery price. Klein and Steinhardt (2023) dynamically decide what delivery time guarantees to offer incoming customers and at what prices. All of the guarantees are assumed to be within the same day; the computational studies specifically test delivery guarantees of 90 and 300 minutes after order placement. Galiullina et al. (2024) optimize the operation of a system in which a subset of customers are offered incentives to have their orders delivered to a shared pickup point (i.e., parcel locker) instead of their homes. Other works use pricing to control the demand for delivery time windows in the AHD context (e.g., Köhler et al. 2023). Fleckenstein et al. (2023), Snoeck et al. (2020), and Waßmuth et al. (2023) provide comprehensive reviews of routing problems with pricing decisions. A related stream of literature controls the availability of crowdsourced delivery drivers via the compensation paid to drivers; we refer to Yıldız and Savelsbergh (2019), Behrendt et al. (2024), and the references therein for examples and discussion of such work.

The pricing component of this work addresses decision-making in a different context than the existing literature on routing and pricing. Unlike prior studies, we study pricing decisions at the tactical level by setting static prices that ensure that the system can meet operational targets on average. Instead of explicitly constructing individual vehicle routes, we use continuous approximations, discussed next, to model customer arrivals and route durations. Finally, as mentioned previously, we believe this work to be the first to study the design of systems where price-differentiated SDD and NDD options are fulfilled by the same delivery fleet.

1.2.4 Continuous Approximations. In logistics system design, the use of deterministic continuous approximation (CA) techniques is a common approach for deriving structural insights and understanding parameter relationships. CA models — the logistical counterpart to fluid models in queuing systems — generally incorporate at least one of the following two features: (i) stochastic customer locations and arrivals are replaced with deterministic, continuous fluid accumulations in space and time; and (ii) routing cost or time calculations are replaced with deterministic, continuous expressions characterizing expected routing time as a function of the number of customers on a delivery route. The latter relies on the classical result that the expected traveling salesperson problem (TSP) route duration grows proportionally to the square root of the number of stops on the route (Beardwood et al. 1959), commonly known as the Beardwood-Halton-Hammersely (BHH) Theorem. Recent applications of CA methods to e-commerce systems include expected cost minimization (Stroh et al. 2022), retail channel selection (Ge et al. 2021), and geographical service region design for last-mile delivery (Banerjee et al. 2022, 2023, Carlsson et al. 2024). We refer to Ansari et al. (2018) and Franceschetti et al. (2017) for a broader review of CA methods and applications.

2. System Setting

In this section, we describe the operational details of the system. An e-retailer of non-perishable, physically small products (in contrast to restaurant meals, groceries, or bulky items, such as furniture and appliances) provides NDD and SDD to customers located in a fixed geographical region. We assume that the service region has been geographically partitioned *a priori* into small, independent, fixed delivery zones. We further assume that no additional geographical discrimination occurs within each zone; for example, a customer’s exact location within a zone does not affect the price they must pay for SDD. We focus our attention on one of these zones; henceforth, “system” refers to the management and operation of this zone. The system is served by a single nearby depot from which vehicles dispatch with orders for delivery. We assume all items are available for delivery from the depot when the customer places the corresponding order, and we assume that unsuccessful delivery attempts do not occur.

2.1 Timeline, Order Types, and Vehicles


The system operates on a daily basis in a repeating fashion. Within each full day, a portion is the *service day*. The service day is the fixed daily interval during which the system is fully active: during the service day, vehicles may be active and real-time decisions may need to be made. Each day i begins at the start of its corresponding service day, at time $t = R_i$. The start time of each service day is consistent, e.g., each day begins at 9 AM.

The fixed end of the service day is T_i for a given day i and is also consistent from day to day, $|[R_i, T_i]| = |[R_j, T_j]| = T$ for all days i, j . Let $R_{i+1} - R_i = R$ for all days i . All orders placed after $t = T_i$ on a day i and before $t = R_{i+1}$ are *overnight* orders to be served by $t = T_{i+1}$. On each day i , the system operator offers SDD as a delivery option starting from the beginning of the service day until an SDD *order deadline* C_i . That is,

SDD is available as a potential delivery option alongside NDD during the interval $[R_i, C_i]$, and all customers who select this SDD option are guaranteed to receive their orders by T_i . Because customer order times and locations are uncertain but the available delivery fleet is fixed, this order deadline is adjusted by the system operator on a daily basis to ensure that the service day constraints are satisfied.

Figure 1 illustrates two consecutive days (i and $i + 1$); the overnight intervals are not presented to scale. NDD order availability, marked with a hatched blue line, is always present. On the other hand, SDD order availability, marked with a solid red line, is only present prior to the order deadline on each day. Notice also that the SDD availability window is slightly longer on day $i + 1$ than on day i , illustrating the potential daily variability of the SDD order deadline.

Figure 1 Consecutive service days with concurrent SDD and NDD order windows



The system operates a fleet of m homogeneous vehicles to deliver orders. Since the constraining factors in modern urban last-mile contexts are generally tight delivery deadlines, we assume vehicles are physically uncapacitated; our personal communications indicate that industry practitioners assume vehicles to be uncapacitated when developing models to solve problems in similar contexts. Each vehicle may dispatch from the depot once per service day. The technical results presented in Sections 3–6 are valid for any choice of m . Practically, however, our work is most useful for small delivery fleets; our computational studies use $m \leq 3$. Large delivery fleets would likely see significant benefits from geographical discrimination or further partitioning of the zone.

2.2 Prices and Customer Behavior

As in Ulmer (2020), we assume NDD is the default, free fulfillment option. The SDD price — the cost to a customer and the delivery revenue gained by the firm — is denoted as $p_0 > 0$. The number of overnight orders is distributed as $\text{Poisson}(\mu)$, where $\mu > 0$. During each service day, e-retail customers arrive according to a stochastic process and individually choose to either purchase SDD, default to NDD, or abandon the system entirely. This decision depends both on the characteristics of the system (SDD price and SDD availability) at the time of arrival and on characteristics of the individual customer. There are two distinct *types* of customers placing orders during the service day, and each chooses between two alternatives.

Type A customers purchase products only through the system. That is, a Type A customer either chooses the SDD delivery option or the NDD delivery option. A randomly selected customer k of Type A has a threshold price $Y_k \geq 0$: if $p_0 \leq Y_k$, the customer chooses the SDD option if it is available. If $p_0 > Y_k$ or if

the SDD option is not available, customer k instead selects NDD. Threshold prices Y_1, Y_2, Y_3, \dots of Type A customers are independent and identically distributed (i.i.d.) and are characterized by a known cumulative distribution function F_A satisfying $F_A(p) = 0$ for all $p < 0$ and $F_A(p) < 1$ for all $p \geq 0$. If SDD itself is viewed as the product, threshold prices in our setting are analogous to “willingness-to-pay” values (e.g., Ulmer 2020) or “reservation prices” (e.g., Song and Li 2018) in similar contexts. Type A customers arrive to the system via a potentially non-stationary Poisson process; the arrival rate of Type A customers at time $t \in [0, T]$ is given by a continuous function $\lambda_A(t) > 0$.

Type B customers require their products urgently. If the SDD price is too high, an arriving Type B customer will abandon and purchase their desired products via a different channel, for example from a different SDD provider or by visiting a retail store. In contrast to Type A customers, who are flexible with respect to fulfillment time guarantees, Type B customers are relatively impatient. The formal characterization of Type B customers is similar to that of Type A customers. A randomly selected customer ℓ of Type B has a threshold price $Z_\ell \geq 0$: if $p_0 \leq Z_\ell$, the customer will choose the SDD option if it is available. If $p_0 > Z_\ell$ or if the SDD option is not available, customer ℓ will instead abandon without a purchase. Threshold prices Z_1, Z_2, Z_3, \dots of Type B customers are i.i.d. and are characterized by a known cumulative distribution function F_B satisfying $F_B(p) = 0$ for all $p < 0$ and $F_B(p) < 1$ for all $p \geq 0$. Type B customers arrive to the system via a potentially non-stationary Poisson process; the arrival rate of Type B customers at time $t \in [0, T]$ is given by a continuous function $\lambda_B(t) \geq 0$. With respect to delivery locations, arriving customers of both types are assumed to have a common geographic distribution, and customer locations are i.i.d.

We define *willing proportion* (WP) functions w_A, w_B for each customer type as $w_A(p) = 1 - F_A(p)$ and $w_B(p) = 1 - F_B(p)$. For a given SDD price p_0 , these functions indicate the proportion of customers of each type who are willing to pay p_0 for SDD. That is, the probability that a random customer of Type A chooses SDD over NDD is $w_A(p_0)$, and the probability that a random customer of Type B chooses SDD over the abandonment option is $w_B(p_0)$. Note that our descriptions are general and do not restrict customer behavior to a specific class of binomial choice model (e.g., logit). When SDD is available, the arrival of all SDD-purchasing customers is given by a Poisson process with rate $w_A(p_0)\lambda_A(t) + w_B(p_0)\lambda_B(t)$.

From the perspective of discrete choice theory, this characterization of customer purchasing behavior represents a mixture of two binomial choice models. This behavior can be equivalently characterized by a single random utility model without initially separating customers into two types. Specifically, such a model would assume that each incoming customer has utilities denoted by the random variables U_0 and U_1 for SDD and NDD, respectively, that satisfy $U_0 > U_1$. Customers for whom $U_0 > U_1 \geq 0$ are Type A customers with threshold price $U_0 - U_1$. Customers for whom $U_0 \geq 0 > U_1$ are Type B customers with threshold price U_0 . Therefore, the analysis in this paper does not depend on the segmentation of customers into two types. Given that NDD is free, our segmentation is mathematically equivalent to a single customer stream that chooses between three options (SDD, NDD, and abandonment) under the mild assumption that

customers always choose SDD if it is available *and* free. We henceforth use our original characterization of separate Type A and Type B customer streams for clarity in exposition; see Gallego and Topaloglu (2019) and Strauss et al. (2018) for further discussion on discrete choice models in similar settings.

3. Modeling Details

In this section, we motivate and detail an approximate model of the system. We then consider the fundamental problem of long-run system feasibility.

3.1 Vehicle Dispatching in an Asymptotic Regime

Traditional applications of CA techniques generally assume a known number of random points to be visited on each vehicle tour. Corollary 1 slightly extends the usual statement of the BHH Theorem to a setting in which the number of points in the TSP tour is itself a random variable. This result motivates the use of the square-root functional form approximation of vehicle tour durations even when the number of stops on the tour is unknown.

COROLLARY 1. Let \mathcal{R} be a compact planar region equipped with a distance metric \mathbf{d} . Let N_1, N_2, \dots be non-negative integer-valued i.i.d. random variables with mean $\nu < \infty$, and let $N = \sum_{i=1}^k N_i$. Note that N has mean $k\nu$. Let X_1, X_2, \dots denote points chosen i.i.d. from an absolutely continuous distribution having support \mathcal{R} . Suppose that the speed of travel in \mathcal{R} is \sqrt{k} ; i.e., the travel time between $x, y \in \mathcal{R}$ is $\mathbf{d}(x, y)/\sqrt{k}$. Then, there exists a finite $\beta > 0$ such that the duration of the TSP tour through $\{X_1, X_2, \dots, X_N\}$ converges almost surely to $\beta\sqrt{\nu}$ as $k \rightarrow \infty$.

The proof, formalized briefly for completeness in Appendix A.1, applies the Strong Law of Large Numbers (SLLN) to the BHH Theorem. In particular, Corollary 1 implies almost sure convergence to $\beta\sqrt{\nu}$ when $N \sim \text{Poisson}(k\nu)$. This specific Poisson case is in fact an intermediate result within classical proofs of the BHH Theorem that use Poissonization techniques (e.g., Beardwood et al. 1959, Karp and Steele 1985, Steele 1997).

The asymptotic regime of Corollary 1 is analogous to commonly studied regimes in which arrival and service rates are scaled by the same factor (e.g., Varma et al. 2023). Here, our scaling of the expected spatial point density by k and the speed by \sqrt{k} is due to the dimensions of the corresponding units of measure (e.g., customers per square mile vs. miles per hour). Additionally, if the tour includes a setup time of $\alpha \geq 0$ and a per-stop service time of $\gamma/k \geq 0$, the same result holds with a limit of $\alpha + \beta\sqrt{\nu} + \gamma\nu$ by the SLLN.

3.2 Approximated System Model

In order to understand parameter relationships and gain managerial insights, we consider average-case system behavior via CA. There are two facets to our CA approach: (i) replacing TSP tours over random customer locations with expected tour durations, and (ii) replacing stochastic, discrete customer requests

with continuous (i.e., fluid) order accumulations. Corollary 1 motivates a functional approximation of tour durations when order accumulations are continuous.

In our model, the time taken for a single vehicle to dispatch from the depot, serve $n \in \mathbb{R}_{\geq 0}$ orders, and return to the depot is given by a deterministic *dispatch duration function* denoted $f(n)$. Our technical results require f to be strictly concave, increasing, non-negative, and continuous. In light of Corollary 1, we particularly focus on a dispatch duration function of the form $f(n) = \alpha + \beta\sqrt{n} + \gamma n$. The parameters $\alpha \geq 0$, $\beta > 0$, and $\gamma \geq 0$ can be estimated *a priori* by simulating TSP tours through the depot over customer sets generated randomly via the known geographic distribution. In the function, α represents any potential setup time and linehaul travel time to and from the depot, γ is a service time per order, and $\beta\sqrt{n}$ is the routing time approximation of Corollary 1.

Let C denote the target order deadline; i.e., the target average of C_i values over time. Then, we can focus on one average-case day in the system by replacing stochastic order arrivals with continuous accumulations: the day (and service day) begins at $t = 0$, SDD orders are accepted until a time $C \geq 0$, the service day ends at a given time $t = T > C$ (equal to $T_i - R_i$ for all i), and the day ends at a given time $t = R \geq T$. To avoid trivially infeasible systems, we assume $C < T - \alpha$. Figure 2 illustrates this timeline. Since T and R are fixed system parameters and μ orders accumulate during $[T, R]$ in the deterministic model regardless of SDD price, we henceforth disregard the dashed portion of the timeline and focus on the service day $[0, T]$.





Figure 2 Single-day timeline, deterministic model

Assume C and p_0 are fixed. Orders accumulate at the rates implied by the customer arrival rates and SDD price. Suppose for now that all NDD orders are served on the day of their delivery deadline. Then, on our average-case day, there are two types of orders that we need to serve. At the beginning of the day, there is an available *time-zero accumulation* of NDD orders (from the previous day) and overnight orders ready to be delivered. The total size of this time-zero accumulation is equal to the sum of (i) all NDD orders placed by Type A customers who arrived to the system during $[0, C]$ on the previous day and found the SDD price too high, (ii) all orders placed by Type A customers who arrived during $[C, T]$ on the previous day when SDD was unavailable, and (iii) all overnight orders. During today's interval $[0, C]$, SDD orders accumulate continuously and deterministically at a rate $w_A(p_0)\lambda_A(t) + w_B(p_0)\lambda_B(t)$. Figure 3 illustrates these orders, both those available at $t = 0$ (in blue, hatched) and those accumulating during the service day (in red, solid). The time-zero accumulation is labeled with an expression for the total number of orders available at $t = 0$; the first integral computes NDD orders placed during $[0, C]$ on the previous day, the second integral computes NDD orders placed during $[C, T]$ on the previous day, and μ is the quantity of overnight orders. Observe that increasing p_0 or decreasing C both increase the size of the time-zero accumulation.

Figure 3 Time-zero accumulation quantity and SDD order rate

3.3 Vehicle Dispatch Policies and Feasibility

Our initial goal is to determine whether the system can serve these orders with m vehicles, each dispatching once per day. We define an m -vehicle *dispatching policy* as an m -tuple of ordered pairs $((t_1, q_1), \dots, (t_m, q_m))$ with $t_1 \leq t_2 \leq \dots \leq t_m \leq C$ and $q_1, \dots, q_m \geq 0$. For each vehicle $d \in [m]$, t_d denotes the vehicle's time of departure and q_d denotes the order quantity served by the vehicle. A dispatch policy is feasible if the following conditions hold,

$$q_d \leq \eta + \int_0^{t_d} (w_A(p_0)\lambda_A(t) + w_B(p_0)\lambda_B(t))dt - \sum_{j=1}^{d-1} q_j \quad \forall d \in [m], \quad (1a)$$

$$t_d + f(q_d) \leq T \quad \forall d \in [m], \quad (1b)$$

$$\sum_{d=1}^m q_d = \eta + \int_0^C (w_A(p_0)\lambda_A(t) + w_B(p_0)\lambda_B(t))dt, \quad (1c)$$

where $\eta := \int_0^C ((1 - w_A(p_0))\lambda_A(t)dt + \int_C^T \lambda_A(t)dt + \mu)$ is the size of the time-zero accumulation. Constraints (1a) ensure that vehicles do not dispatch to serve orders which have not yet accumulated. Constraints (1b) ensure that all vehicles return to the depot by the end of the service day. Constraint (1c) ensures that all SDD and NDD orders are served. As implied by Figure 3, these conditions assume all orders are served on the day of their delivery deadline – no NDD orders are served a day early.

For a setting defined by $C, T, m, p_0, \lambda_A(t), \lambda_B(t)$, and the WP functions, does a vehicle dispatching policy exist that can feasibly fulfill all orders on the day of their delivery guarantee? If the answer is “yes,” then a steady state exists, in the sense that this deterministic model of the system can operate feasibly in the same manner every day. If the answer is “no,” then additional vehicles are required. Since feasibility is not immediately clear by inspecting constraints (1a)–(1c), we now focus on constructing a vehicle dispatching policy sufficient for establishing feasibility.


3.4 System Feasibility

For $t \leq T - f(0)$, define $Q(t) = q$ such that $t + f(q) = T$. In other words, $Q(t) = f^{-1}(T - t)$ is the maximum number of orders a vehicle can serve when departing at t . We define the m -vehicle *greedy policy* as follows. First, dispatch as many vehicles as possible at $t = 0$ with $Q(0)$ orders each; these vehicles return at $t = T$. Then, dispatch all but one of the remaining vehicles sequentially carrying as many orders as possible such

that each of these vehicles (i) takes all unserved accumulated SDD orders at its time of departure, and (ii) returns no later than time T . The last (m -th) vehicle dispatches at $t = C$ with all remaining unserved orders. By definition, the greedy policy is feasible if and only if the m -th vehicle returns by time T . Additionally, if some vehicle $\ell < m$ returns to the depot prior to T , then all subsequent vehicles dispatch with zero orders (or need not be dispatched); this implies that no more than ℓ vehicles are required.

Figure 4 illustrates the structure of the greedy policy. Suppose that $m = 3$. Arrows represent vehicle dispatch departure and return times. The first vehicle serves all of the time-zero accumulation as well as the SDD orders placed before its departure time. The second vehicle serves all SDD orders placed since the departure time of the first vehicle. The third and final vehicle serves all of the remaining SDD orders. Observe that the final vehicle returns to the depot prior to $t = T$, indicating that the system is feasible for the chosen parameters.

Figure 4 Greedy policy illustration (order accumulations not to scale)



Algorithm 1 in Appendix A.2 formalizes the computation of the greedy policy. The greedy policy can be computed efficiently given the value of $Q(0)$, requiring at most m calls to a univariate root-finding routine. As we illustrate next, efficiency in computation is only one of several reasons why the greedy policy is fundamental to the analysis of our system.

Let the greedy policy be denoted by $P = ((t_1, q_1), \dots, (t_m, q_m))$. We say that our system is m -saturated (or simply *saturated*) if, for all $d \in [m]$, $t_d + f(q_d) = T$. That is, the system is m -saturated if all vehicles return exactly at the end of the service day in the m -vehicle greedy policy. The following lemma, proved in Appendix A, demonstrates the usefulness of the greedy policy when the system is saturated.

LEMMA 1. *If a system is m -saturated, the m -vehicle greedy policy is the only m -vehicle policy that can feasibly serve the system.*

A natural question is whether there exist non-saturated scenarios in which the greedy policy is infeasible, but there exists a different policy that can still serve the system. The answer to this question is negative, and the proof (in Appendix A) is a direct application of the preceding lemma.

THEOREM 1. *If the m -vehicle greedy policy is infeasible, no other m -vehicle policy can feasibly serve the system.*

This result implies that checking whether the system is feasible is as simple as computing the greedy policy. However, when the system is not saturated, there may still be many feasible policies to choose from. Among these feasible policies, it can be shown that the greedy policy is also optimal from a cost-minimization perspective, as we formalize in Theorem 2 and prove in Appendix A.6.

THEOREM 2. *Suppose no m' -vehicle policy is feasible for $m' < m$, but the m -vehicle greedy policy is feasible. Then the m -vehicle greedy policy minimizes the total routing time across all feasible policies. More generally, if the cost associated with a dispatch is a continuous, non-negative, strictly increasing, strictly concave function of the dispatch quantity, then the m -vehicle greedy policy minimizes the total cost across all feasible policies.*

3.4.1 Vehicle Utilization and Fulfillment Difficulty. Even if m is the minimum fleet size required for feasibility, if the system is not saturated the final vehicle may return to the depot prior to T in the greedy policy, as in Figure 4. Intuitively, having the same quantity of orders accumulate earlier in the day tends to make the system “easier” to serve, whereas having the same quantity of orders accumulate later in the day tends to make the system more “difficult” to serve. This is because earlier realization of customer demand improves the potential for batching on delivery vehicles (i.e., leveraging the economies of scale associated with the concave function f). As an example, consider a one-vehicle system in which 100 customers demand SDD. If all of these customers place their orders by 11 AM, there is a greater likelihood that the vehicle can feasibly deliver these orders, compared to a situation in which 50 of these customers delayed their orders until 4 PM. In the extreme case, this principle is illustrated by the fact that NDD requests are significantly easier to fulfill than SDD requests, since NDD requests can be batched into quantities of $Q(0)$ the day after the order is placed.

From a managerial perspective, a feasible system that is “too easy” to serve indicates potential unrealized gains or underutilized delivery capacity. However, we have multiple levers at our disposal — such as increasing C and/or decreasing p_0 — that we can use in order to remedy these inefficiencies and optimize various chosen system objectives while maintaining steady state feasibility. We next apply our model to the optimization of two such objectives. We begin by leveraging the greedy policy to revisit the concept of the system steady state.

4. Steady State Existence and Convergence

Our analysis until this point has assumed that we are operating in a steady state, yet we have disregarded two subtle but important questions. First, does our system always reach a steady state? Second, are there multiple steady states associated with the same set of system parameters? For the purposes of this analysis, we assume a fixed SDD price and fleet size with our model’s usual deterministic order accumulations and vehicle tour durations. On the other hand, we no longer assume that we are currently in a steady state (i.e., we may vary the dispatching policy and SDD order deadline each day).

Assume our system has the capacity to serve some SDD orders; that is, $\mu + \int_0^T \lambda_A(t)dt < mQ(0)$. Given p_0 and m , suppose we choose a target SDD order deadline C so that the system is saturated; such a $C \in (0, T)$ exists by the continuity of f . Although we know how vehicles should optimally be dispatched in the steady state, this steady state is associated with a time-zero accumulation of η orders at the beginning of each day. If there are instead no orders requiring immediate/same-day fulfillment at $t = 0$ (time R_0), will the system converge such that η time-zero orders arise at each future time R_i for some sufficiently large i ?

For a day with q orders accumulated at time zero, we want the SDD cutoff to be as late as possible while maintaining feasibility with respect to the end-of-day deadline T . Since this SDD order deadline depends on q , we use the function $G(q)$ to denote the q -dependent deadline. To compute $G(q)$, we solve the greedy policy with a time-zero accumulation value of q and extend the SDD order deadline until vehicle m returns exactly at $t = T$. By definition, $G(\eta) = C$. In the scenario faced by the system manager, the order cutoff on day 0 will be $G(0)$, and $G(q)$ decreases as q increases.

Delaying the SDD order deadline induces some later-arriving Type A customers to purchase SDD when they otherwise would have been obligated to select the NDD option. Therefore, a higher value of $G(\cdot)$ today corresponds to a lower time-zero accumulation quantity tomorrow. Specifically, we define the function $H(q)$ as the quantity of *tomorrow's* time-zero accumulation given a time-zero accumulation of q units *today*. The value of $H(q)$ is determined by $G(q)$:

$$H(q) = \mu + \int_0^{G(q)} ((1 - w_A(p_0))\lambda_A(t)dt + \int_{G(q)}^T \lambda_A(t)dt). \quad (2)$$

Because G is a decreasing function, $H(q)$ increases as q increases. Just as $G(\eta) = C$, we also have $H(\eta) = \eta$. In this scenario, the order cutoff on day 0 will be $G(0)$.

Beginning at day 0, the system therefore evolves based on the initial accumulated quantity of orders. For example, when there are no initial orders to be served, the system on day 3 has a time-zero accumulation of $H(H(H(0)))$ units and a SDD order deadline of $G(H(H(H(0))))$. For notational purposes, we denote the time-zero accumulation on day j given q initial orders as $H^{(j)}(q)$, the j -th composition of H . The system's SDD order deadline on day j is then $G(H^{(j)}(q))$, which we denote $G^{(j)}(q)$. Note, however, that $G^{(j)}$ does not represent a function composition in the same manner as $H^{(j)}$.

As measured by the daily time-zero accumulation $H^{(j)}$ and SDD order deadline $G^{(j)}$, the system indeed converges to the chosen steady state when there are zero orders requiring delivery at $t = 0$ on day 0. More generally, we can show that the system converges to the chosen steady state for any initial accumulated quantity satisfying $mQ(0) \geq q$, i.e., the vehicles are able to serve the q orders on day 0. Theorem 3, proved in Appendix A.8, formalizes this result. From a dynamical system perspective, the saturated steady state represents a Lyapunov-stable attracting fixed point.

THEOREM 3. *Let η and C denote the time-zero accumulation and SDD order deadline, respectively, associated with the saturated steady state. For any starting quantity $q \in [0, mQ(0)]$ of orders at time $t = 0$ on day 0, $\lim_{j \rightarrow \infty} G^{(j)}(q) = C$ and $\lim_{j \rightarrow \infty} H^{(j)}(q) = \eta$. Additionally, the convergence is monotonic.*


5. Early Next-Day Delivery Maximization

Empirical evidence suggests that fulfilling e-commerce orders earlier than the promised delivery day leads to increased sales in the long run (Fisher et al. 2019). Driven by these findings, we introduce the concept of NDD^+ orders: NDD orders that are fulfilled on the same day they are placed (i.e., one day early). In other words, NDD^+ orders receive the same treatment as SDD orders without incurring the higher delivery fee. Assuming a fixed SDD price p_0 and deadline C , we view the quantity of NDD^+ orders as a measurement of customer satisfaction; thus motivated, we seek to maximize the quantity of NDD orders served as NDD^+ orders with the intent of capturing a greater market share by exceeding customer expectations.

5.1 Structural Analysis


We again assume that the system is operating in steady state, and that C has been chosen so that the system is feasible when all orders are fulfilled on their due date, but not necessarily saturated. We seek to convert NDD orders to NDD^+ orders while maintaining system feasibility. Note that every order served as NDD^+ today entails one fewer NDD order to be served tomorrow. Thus, in steady state, treating some quantity of NDD orders during the day as NDD^+ orders leads to a corresponding reduction in the size of the time-zero accumulation. This idea is illustrated in Figure 5: some NDD orders in the middle of the day are converted to NDD^+ (depicted in striped gold), so the size of the time-zero accumulation is reduced by the same amount (not to scale).

Figure 5 Conversion of NDD orders to NDD^+ orders




Assuming system feasibility, let $q^+ \in \mathbb{R}_{\geq 0}$ denote the maximum number of NDD orders that we can convert to NDD^+ orders. In order to gain insight into the optimization problem at hand, we begin by considering a simple question: when can we serve at least *some* NDD^+ orders? In other words, when is $q^+ > 0$? It is clear that, if the final vehicle in the greedy policy returns before $t = T$ as in Figure 4, the “slack” allows us to serve a positive quantity of NDD orders as NDD^+ . This scenario is illustrated in Figure 6 with some NDD orders placed just before C being converted to NDD^+ orders; the first vehicle now returns before $t = T$ due to some of its quantity being reassigned to the final vehicle.

Perhaps counterintuitively, however, this is not a necessary condition to be able to fulfill some NDD^+ orders. Consider a saturated system in which the last vehicle to serve any orders from the time-zero accumulation — let us denote this as the k -th vehicle — also serves some SDD orders (i.e., $t_k > 0$); this occurs

Figure 6 Feasibly fulfilling NDD^+ orders when $t_m + f(q_m) < T$, cf. Figure 4

if the size of the time-zero accumulation is not an integer multiple of $Q(0)$. Then, we can convert a small quantity of the NDD orders placed prior to t_k into NDD^+ orders. This conversion will slightly reduce the size of the time-zero accumulation by the same small quantity. As illustrated in Figure 7, the original greedy policy can still feasibly serve this new system that features a small positive quantity of NDD^+ orders at the beginning of the day. Therefore, even when the system is saturated we can convert some NDD orders into NDD^+ orders without changing the dispatching policy at all.

Figure 7 Feasibly fulfilling NDD^+ orders when $t_k > 0$ ($k = 1$)

We can show that these two conditions represent the only two scenarios in which $q^+ > 0$. Theorem 4, proved in Appendix A.9, formalizes this result.


THEOREM 4. *Assume that the system is feasible when all orders are fulfilled on their due date; denote the associated greedy policy as $P = ((t_1, q_1), \dots, (t_m, q_m))$ and denote the associated total quantity of orders available at $t = 0$ as η . The system can serve a positive number of NDD orders a day early ($q^+ > 0$) if and only if at least one of the following conditions hold: (i) $t_m + f(q_m) < T$, or (ii) $\eta/Q(0) \notin \mathbb{Z}$.*

This result informs whether we have the ability to provide any NDD^+ service, but it does not indicate *which* NDD orders to fulfill early. We must still determine the time interval(s) during which NDD orders should be treated as NDD^+ .

Consider a solution in which some amount of NDD orders placed during time interval $[t, t + \varepsilon]$ are converted to NDD^+ , while some NDD orders placed earlier than t are not converted. Without loss of feasibility,

the same solution can convert earlier orders instead of the later ones. It follows that there is no benefit to converting NDD orders placed during some time interval unless *all* NDD orders placed earlier in the day are also designated as NDD⁺ orders. Therefore, the optimal solution to our problem must adhere to the following structure: for some “early deadline” \hat{C} , all NDD orders placed during the interval $[0, \hat{C}]$ are designated as NDD⁺ orders, while remaining NDD orders are fulfilled normally, on the day of their delivery guarantee. Figure 8 illustrates this structure; the labeled NDD⁺ accumulation rate prior to \hat{C} is in addition to the usual SDD accumulation rate labeled separately.

Figure 8 Time-zero accumulation quantity, NDD⁺ order rate, and SDD order rate



5.2 Computational and Managerial Implications

The key consequence of this structure is that the problem of optimizing the quantity of NDD⁺ orders is equivalent to maximizing \hat{C} while maintaining feasibility of the greedy policy (and thereby maintaining feasibility of the system). We also know that if a particular value of \hat{C} is feasible, then all earlier values of \hat{C} are also feasible. Therefore, assuming that the system is feasible for $\hat{C} = 0$, we can solve for the optimal NDD⁺ deadline \hat{C}^* (within a given tolerance) by bisection search over the interval $[0, T]$.

However, because we do not immediately gain any additional revenue from fulfilling NDD orders early, in a typical system we should expect \hat{C}^* to be much earlier than C . If the system is feasible even with $\hat{C} \geq C$, as an extreme case, every customer who places an order prior to C receives SDD. This indicates the system is operating inefficiently, and other parameters should be reevaluated. For example, the system’s fleet size should be reduced, or C should be delayed.

To improve on the efficiency of naïve bisection, we introduce an auxiliary function $F(\hat{C})$, which we define as the return time of the final vehicle in the greedy policy when the NDD⁺ deadline is set to \hat{C} . Because we seek a value of \hat{C} that induces saturation, the problem of maximizing \hat{C} is reduced to solving $\max\{\hat{C} | F(\hat{C}) = T\}$. Fortunately, we can show that this set usually contains only a single point and the problem is thus amenable to univariate root-finding routines. Proposition 1, proved in Appendix A.10, formalizes this property.

PROPOSITION 1. *The function $F(\hat{C})$ is non-decreasing for $\hat{C} \in [0, C)$. Additionally, if $F(0) < T$ and $F(C) > T$, then $F(\hat{C}) = T$ has a unique solution \hat{C}^* in the interval $(0, C)$.*

This result allows us to leverage sophisticated root-finding routines, such as Brent’s method, that use the value of $F(\hat{C}) - T$ to calculate \hat{C}^* more efficiently than bisection search. In our experience, using this approach (instead of bisection search) reduces computation time by 30–40%. While this improvement may not be necessary for a single instance, the effect is noticeable when optimizing hundreds or thousands of instances sequentially, perhaps in order to compare various combinations of parameter settings.

The analysis in this section motivates a simple set of guidelines for managing the system: (i) all NDD orders placed before the optimal NDD⁺ deadline \hat{C}^* are fulfilled on the same day, (ii) all SDD orders are fulfilled on the same day, and (iii) all other orders are fulfilled on the following day. Our analysis assumes the chosen NDD⁺ deadline is not advertised, to prevent early-arriving customers from acting strategically and selecting the free NDD option when they otherwise would have paid for the SDD guarantee.

5.3 Mitigating Customers’ Strategic Behavior


Despite the lack of advertisement, a few customers may slowly learn from experience about the NDD⁺ deadline, particularly if some customers are placing orders frequently. This may potentially lead to a small loss in the firm’s profit over long time scales because such customers would eventually be less willing to pay for SDD prior to \hat{C}^* . An e-retail firm may wish to proactively prevent such strategic behavior while still occasionally exceeding customers’ expectations.

We propose one potential approach to mitigate customer learning: introduce randomization when choosing which customers to upgrade to NDD⁺. That is, instead of providing NDD⁺ to all customers within a particular time window, thereby risking some customers learning the structure of the e-retailer’s management guidelines, we upgrade only a random subset of the incoming NDD⁺ orders during a longer time window. The relative size of this random subset should be chosen to be low enough that customers cannot effectively learn strategic behavior quickly enough to affect the e-retailer’s tactical planning.

Formally, suppose that the e-retailer determines that upgrading only ξ proportion of NDD orders during $[0, C]$ is sufficient to effectively nullify customer learning. In the stochastic system, this means that every incoming NDD order during $[0, C]$ is randomly upgraded to NDD⁺ with probability ξ (independently of any other upgrade decision). In the deterministic approximation, this corresponds to an effective NDD⁺ arrival rate of $\xi \lambda_A(t)$ for $t \in [0, C]$. Figure 9 illustrates this structure.

It remains to find the value of C that maximizes the number of NDD⁺ orders while maintaining system feasibility. As before, if a particular value of \hat{C} is feasible, then all earlier values of \hat{C} are also feasible. Therefore, bisection search can again be employed to solve for \hat{C}^* in this context.

The underlying hypothesis of this approach is that, by introducing randomness, the customers can no longer rely on guaranteed upgrades during an early time window, thereby making it difficult for them to learn and adapt to the firm’s strategy. Unfortunately, we are not aware of any prior research that studies this hypothesis in the context of e-retail delivery. In a different environment, however, Sriram et al. (2015) find

Figure 9 Time-zero accumulation quantity, effective NDD⁺ order rate, and SDD order rate

“a slower rate of learning among households that experience high variability” in the service quality of a video-on-demand company. Future empirical work may seek to extend these findings to the e-retail delivery sector and inform a data-driven method for choosing an appropriate value of ξ for our proposed strategy.

To conclude this section, we stress again that the percentage of all NDD orders during $[0, C]$ given an upgrade should be small. If a significant number of NDD⁺ orders are being fulfilled daily, it may indicate that the vehicle fleet is not being utilized to its full earning potential. In such a case, it may be prudent to forgo these NDD⁺ strategies and instead re-optimize the SDD order deadline and price, as we do next.

6. Static Pricing and Temporal Discounting

The customer satisfaction objective studied in the preceding section is useful for firms attempting to better establish themselves in a market at the expense of optimizing short-term profit. A well-established firm, on the other hand, may wish to instead design the system purely for profit maximization; we now apply our model to this objective. Instead of serving NDD⁺ orders, we choose the SDD price p_0 and order deadline C — previously assumed to be fixed — in order to maximize daily profit in the system.

6.1 Single SDD Price

We wish to design our system by selecting a static SDD price p_0 and SDD order deadline C . The price must be selected from a given discrete, fixed price menu (e.g., $\{1, 2, \dots, 10\}$), while the SDD order deadline may take any value in the interval $(0, T)$. Because the current focus is the system’s profit, we also consider potential fulfillment costs under two different cost structures.

6.1.1 Flat per-vehicle costs. Suppose that we may use at most m_{\max} vehicles, but the daily cost per vehicle is $\varphi > 0$. Therefore, the total daily profit in our model is $p_0 \int_0^C (\lambda_A(t) w_A(p_0) + \lambda_B(t) w_B(p_0)) dt - \varphi m$. Our goal is to maximize this profit.

Fix a particular SDD price p_0 and fleet size m . Because the costs are fixed at φm , we seek the revenue-maximizing SDD order deadline. Observe that we may assume full utilization of vehicles without loss of optimality. Thus, the revenue-maximizing SDD order deadline is the unique value C for which all vehicles return at T in the associated feasible greedy policy. Applying a similar approach as in Section 5 allows us

to calculate this value of C , because increasing the time of the SDD order deadline progressively increases the fleet's utilization. Repeating this procedure for every p_0 in the price menu produces the optimal solution for a given m . Finally, optimizing the revenue for each $m \leq m_{\max}$ and comparing the associated profits gives the optimal fleet size.

6.1.2 Dispatch-duration costs. Suppose instead that costs are incurred proportional to the fleet's total dispatch duration. Under this alternate cost structure, the total daily profit in the model is $p_0 \int_0^C (\lambda_A(t)w_A(p_0) + \lambda_B(t)w_B(p_0))dt - \psi \sum_{i=1}^m f(q_i)$ for any $m \leq m_{\max}$ and a positive scaling constant ψ .

Recall that the greedy policy minimizes total dispatch duration for any feasible combination of m , C , and p_0 . The profit's dependence on total dispatch duration implies that it is no longer necessarily optimal to choose the value of C that fully utilizes the fleet's temporal capacity for a given choice of m and p_0 . When arrival rates are time-homogeneous, however, we can show that the profit-maximizing solution admits a convenient structure.


PROPOSITION 2. *Assuming costs proportional to total dispatch duration, constant customer arrival rates, and a fixed SDD price p_0 , the total profit is a piecewise-convex function of C . Additionally, the breakpoints of this function correspond to values of C for which either (i) all utilized vehicles in the associated greedy policy return exactly at time $t = T$, or (ii) the quantity of the associated time-zero accumulation is an integer multiple of $Q(0)$.*

The result is proved in Appendix A.11. In other words, condition (i) means that any vehicle that is responsible for a positive quantity of orders must be fully utilized at optimality. Condition (ii) means that any vehicle that serves NDD orders must dispatch at $t = 0$.

Finding the optimal solution simply requires finding and comparing these candidate optimal values of C (i.e., the breakpoints). For a given value of m and for each vehicle $i \leq m$, there exists at most one value of C such that vehicle i is the last vehicle to be utilized *and* vehicle i returns at $t = T$. Similarly, there exists at most one value of C such that vehicle i is the last vehicle to serve any NDD orders *and* vehicle i departs at $t = 0$. Thus, there are $\mathcal{O}(m)$ candidate optimal values of C that can be enumerated by bisection search over $C \in [0, T)$. Repeating this process for all prices in the menu and all $m \leq m_{\max}$ gives the profit-maximizing combination of C , p_0 , and m .

6.2 Temporal Discounting

Because orders placed earlier in the day can be served more efficiently than orders placed later in the day, we now consider a slightly more sophisticated pricing scheme. We choose a base SDD price p_0 and a discounted SDD price $\hat{p}_0 \leq p_0$ from the menu. The discounted SDD price \hat{p}_0 is offered to all customers prior to a *discounted SDD order deadline* \hat{C} . The base SDD price is offered to all customers after \hat{C} but prior to the *final SDD order deadline* C , where $C \geq \hat{C}$. Figure 10 illustrates this system design; compare this to the simpler design in Figure 3.

Figure 10 Two-level static SDD pricing

Analogous to our approach in the single-price problem with per-vehicle costs, consider maximizing revenue given a fixed price combination $\{\hat{p}_0, p_0\}$ and fleet size m . In the single-price problem, there were no degrees of freedom given a fixed SDD price; any chosen p_0 directly determined the associated SDD order deadline due to the necessity for vehicles to be fully utilized at optimality. In this two-level pricing scheme, however, fixing the prices still requires us to make one decision. We must choose the discounted deadline \hat{C} , which then implies a single value of C that ensures all vehicles are fully utilized.


The resulting combination of \hat{C} and C generates a total revenue of $\hat{p}_0 \int_0^{\hat{C}} (\lambda_A(t)w_A(\hat{p}_0) + \lambda_B(t)w_B(\hat{p}_0))dt + p_0 \int_{\hat{C}}^C (\lambda_A(t)w_A(p_0) + \lambda_B(t)w_B(p_0))dt$. In general, even this univariate problem of maximizing revenue as a function of \hat{C} displays no particular structural properties. However, if the arrival rates of each type of customer are modeled as constant over time (but potentially distinct), we can derive a structural insight that is important for optimization under the per-vehicle cost structure.

THEOREM 5. *Suppose that the arrival rates $\lambda_A(t)$ and $\lambda_B(t)$ are constant over time. For a fixed price combination $\{\hat{p}_0, p_0\}$ and fleet size m , the total revenue is a continuous, piecewise convex function of \hat{C} . The breakpoints of this function correspond to values of \hat{C} that satisfy at least one of the following conditions: (i) a vehicle departs exactly at \hat{C} in the associated greedy policy, or (ii) the quantity of the associated time-zero accumulation is an integer multiple of $Q(0)$.*

The proof of this result, deferred to Appendix A.12, is lengthy and requires analyzing several cases. Figure 11 illustrates a solution that satisfies both conditions.

With respect to optimization, this result significantly reduces the search space for potential global optimal solutions to these breakpoints. In two practical cases, the number of breakpoints for each price combination is bounded by a polynomial function of the fleet size m : (1) when all overnight and potential NDD orders can be served by a single vehicle, or (2) when all customers are of Type A. The former case tends to occur when the zone is chosen to be sufficiently small and Type B customers represent the majority of the customer base. The latter represents a reasonable modeling assumption when the product(s) sold by our e-retail firm cannot be purchased nor substituted via a different e-retail firm or brick-and-mortar retailer. In both cases, the potential optimal solutions can be efficiently enumerated; we provide details in Appendix A.13.

Figure 11 Potential optimal two-level SDD pricing solution with associated dispatching policy



7. Computational Study

In this section, we apply our models to a case study in the Denver metropolitan area. We first describe geographical aspects of the system and build the corresponding CA model. Then, we optimize the system with respect to the objectives from Sections 5 and 6. Finally, we perform operational simulations in order to validate our model predictions in an operational setting.

We query population values via the open-access Population Estimation Service (NASA SEDAC and Columbia University CIESIN 2018), which uses United States census data. We generate customer locations and travel time matrices with the open-source tools Openrouteservice (HeiGIT 2022) and VeRoViz (Peng and Murray 2022). Optimization, including the computation of TSP tours, is implemented in Python 3.7.3 via Gurobi and SciPy.

7.1 System Details and Model Parameters

7.1.1 Geography and Vehicle Routing. Our study is set in the suburbs of Denver, Colorado, USA. Specifically, we are focused on the management and operation of a roughly rectangular delivery zone measuring approximately 4 miles east-to-west and 3 miles north-to-south. The total population of this zone is roughly 81,000. In order to capture the geographical inhomogeneity of customer locations in practice, we discretize the zone into 12 one mile-by-one mile tracts, with populations ranging from roughly 4,100 to 11,600. We assume the probability that a random customer is located within a tract is proportional to the tract's population. Within a selected tract, the location of the customer is chosen uniformly at random along the road network. The zone is served by a depot located to the northwest; the driving distance between the zone and depot is approximately five miles, corresponding to a driving time of approximately ten minutes. Figure 12 depicts the location of the depot, marked with a flag, and the zone relative to downtown Denver. Each service day begins at 9 AM, and deliveries must be completed and all vehicles must return to the depot by 6 PM.

In order to build the system model, we empirically estimate a routing time approximation suitable for our particular system's road network, driving speed limits, depot location, and customers' geographical distribution. For each $n \in \{15, 30, 45, \dots, 105\}$, we randomly generate 50 sets of n customers within the

Figure 12 System geography for computational study (generated in Leaflet via VeRoViz)




zone. For each set, we solve for the fastest TSP tour through the n points and the depot, where the travel time between two locations depends on the road network and may be asymmetric. We then fit the functional form $\alpha + \beta\sqrt{n} + \gamma n$ to these 350 tour durations. We approximate the duration of a TSP tour visiting n randomly generated customer locations in the zone from the depot as $15.83 + 13.61\sqrt{n} + 0.33n$ minutes.

In a recent case study in London, Allen et al. (2018) empirically estimate 4.1 minutes of non-driving service time per customer in a last-mile delivery context. This includes unloading packages, walking between the delivery vehicle and the customers' residence, and "gaining proof-of-delivery" (e.g., taking a picture of the delivered package or obtaining a signature). Because finding parking close to customer residences is easier in our suburban setting relative to the urban environment of London, we assume a slightly lower service time of 3 minutes per delivery in our study. Additionally, we include 5 minutes per dispatch to account for any setup time, such as traveling from the depot's loading bay to the actual road. Combining all of these components produces the dispatch time function $f(n) = 20.83 + 13.61\sqrt{n} + 3.33n$, which we use in our model.

7.1.2 Prices and Customer Behavior. We model behavior of each customer type with a logit-type sigmoid function, rescaled so that all customers choose SDD when $p_0 = 0$:

$$w_A(p_0) = \frac{1 - (1 + e^{3-0.7p_0})^{-1}}{1 - (1 + e^3)^{-1}} \quad \text{and} \quad w_B(p_0) = \frac{1 - (1 + e^{5-0.5p_0})^{-1}}{1 - (1 + e^5)^{-1}}. \quad (3)$$

Figure 13 illustrates the WP functions for each customer type. Type A customers are less willing to pay for SDD than Type B customers at every price point, as evidenced by w_B dominating w_A in Figure 13. For

Figure 13 Proportion of customers choosing SDD as a function of SDD price

example, at a SDD price of $p_0 = 5$ dollars, approximately 40% of arriving Type A customers and 93% of arriving Type B customers choose the SDD option if available.

We assume a total customer arrival rate that corresponds to 7.5% of the residents of our zone attempting to place an order once per month (28 days) on average,


$$\mu + \int_0^T (\lambda_A(t) + \lambda_B(t)) dt = (0.075 \times 81116) / 28 \approx 217.28. \quad (4)$$

The daily quantity of orders will be lower than this value due to some Type B customers abandoning the system upon encountering either an excessive SDD price or a closed SDD ordering window. The actual arrival rates $\lambda_A(t)$ and $\lambda_B(t)$ we use are different for each objective and are specified in the following subsection.

7.2 System Optimization

7.2.1 Customer Satisfaction. We first consider the objective of maximizing the quantity of NDD⁺ orders. Customer arrivals are assumed to be time-varying; Type A customers' arrivals peak around noon, while Type B customers' arrivals increase over the course of the day. Figure 14 illustrates the specific arrival rate of each type of customer. We assume that the quantity of overnight orders is $\mu = 30$. This corresponds to an average of two Type A customer arrivals per hour between 6 PM and 9 AM, approximately 67% of the Type A arrival rate at 6 PM. In this experiment, the SDD price is $p_0 = 5$ dollars and the SDD order deadline is 3 PM ($C = 360$ minutes). There are $m = 3$ vehicles assigned to the zone, each dispatching once daily.

Applying the greedy policy to the deterministic model produces the following solution. Vehicle #1 departs from the depot at 9:56 AM ($t = 55.88$) to serve 98.52 orders (55.54 NDD, 30 overnight, and 12.99 SDD) and returns at 6 PM ($t = 540$). Vehicle #2 departs from the depot at 1:17 PM ($t = 256.85$) to serve 49.89 orders (all SDD) and returns at 6 PM ($t = 540$). Finally, vehicle #3 departs from the depot at 3 PM ($t = 360$) to serve 26.02 orders (all SDD) and returns at 5:57 PM ($t = 536.92$). It is immediately clear that both conditions of Theorem 4 are satisfied, since the system is in an unsaturated steady state ($t_3 + f(q_3) < T$), and no vehicles depart at $t = 0$. Therefore, we expect that it is feasible to convert some NDD orders into NDD⁺ orders.


Figure 14 Customer arrival rates, NDD⁺ objective

This is indeed the case; the optimal NDD⁺ quantity in this system is $q^+ = 8.28$ orders per day. The optimal solution is to fulfill all NDD orders placed by 10:25 AM ($\hat{C} = 85.25$) on the same day. When compared with the minuscule temporal slack in the original system — the final vehicle returns to the depot just three minutes prior to the end of the service day — it may initially seem that this optimal value of \hat{C} is too high. However, recall that $t_k = t_1 = 55.88 > 0$ implies that we can increase the NDD⁺ deadline to some extent without needing to modify the dispatching policy at all. Indeed, in this case, we can increase the NDD⁺ deadline until $t = 55.88$ without changing any departure times, return times, or dispatch quantities in the original greedy policy. Increasing the deadline by an additional half-hour to $\hat{C} = 85.25$ saturates the system. The concave nature of f implies that economies of scale can be leveraged more effectively earlier in the day; thus, it is unsurprising that this half-hour delay towards the beginning of the day entails a three-minute delay in the return time of the final vehicle at the end of the day.

The associated optimal vehicle dispatching policy is as follows. Vehicle #1 departs from the depot at 10:01 AM ($t = 60.98$) to serve 97.25 orders (47.25 NDD, 30 overnight, 5.81 NDD⁺, and 12.99 SDD), vehicle #2 departs from the depot at 1:14 PM ($t = 254.25$) to serve 50.50 orders (2.48 NDD⁺ and 48.02 SDD), and vehicle #3 departs from the depot at 3 PM ($t = 360$) to serve 26.68 orders (all SDD). All vehicles return at 6 PM ($t = 540$).

A benefit of CA methods is computationally efficient sensitivity analysis. As an example, system managers may be interested in the effect of varying the upgrade proportion ξ in the randomized NDD⁺ extension studied in Section 5.3. The solution detailed in the preceding paragraphs is equivalent to the randomized policy with $\xi = 1$, but how many fewer NDD customers can we upgrade to NDD⁺ if we choose a lower value of ξ ? To answer this question, we maximize the NDD⁺ deadline, thereby maximizing the quantity of NDD⁺ orders fulfilled, for each $\xi \in \{0.05, 0.051, \dots, 0.999, 1\}$. Figure 15a plots the optimal NDD⁺ deadline for each value of ξ , and Figure 15b plots the corresponding quantity of NDD⁺ orders.

Both plots depict a continuous, piecewise relationship with a single breakpoint. The breakpoint occurs at the value of ξ for which the types of orders fulfilled by the third vehicle changes: the third vehicle fulfills some NDD⁺ orders for values of ξ to the left of the breakpoint, and the third vehicle fulfills no NDD⁺ orders

Figure 15 Sensitivity analysis of randomized NDD⁺ strategy

for values of ξ to the left of the breakpoint. Each of the components of the piecewise graph in Figure 15b is slightly nonlinear despite the curve's appearance in the plot. These figures provide managerial insights for a potential implementation of the randomized NDD⁺ policy. For example, it is interesting to observe that upgrading only half of the incoming NDD orders during the designated time window still allows for 5.38 NDD⁺ orders fulfilled per day – just 35% lower than the 8.28 NDD⁺ orders fulfilled under the original optimal policy with $\xi = 1$.

7.2.2 Profit Maximization. We next turn our focus to maximizing delivery profit via our system model under the two-level SDD pricing scheme with per-vehicle costs. We assume here that the overnight accumulation remains at $\mu = 30$ orders, but the arrival rates of Type A and Type B customers are constant over time. Specifically, $\lambda_A(t) \approx 0.09$ and $\lambda_B(t) \approx 0.26$ arrivals per minute for all $t \in [0, T]$.

We optimize over the discrete price menu $\{0.5, 1, 1.5, \dots, 14.5, 15\}$ dollars; thus, there are 465 potential price combinations to analyze. Up to three vehicles may be used. For each $m \in \{1, 2, 3\}$, recall that we seek the combination of base (late) SDD price p_0 , discounted (early) price \hat{p}_0 , discounted SDD order deadline \hat{C} , and final SDD order deadline C that entails the maximum total daily delivery revenue. The total delivery profit is equal to the delivery revenue minus $m\phi$, where ϕ represents the daily cost per vehicle. For sensitivity purposes, we are also interested in understanding how the optimal fleet size varies with ϕ .

Table 1 Optimal revenues, SDD prices, and SDD deadlines

Fleet size m	Discounted SDD price \hat{p}_0	Discounted SDD order deadline \hat{C}	Base SDD price p_0	Final SDD order deadline C	Total Revenue
1	-	-	\$9	10:27 AM	\$129.44
2	\$8	10:23 AM	\$9	2:10 PM	\$470.50
3	\$8	1:54 PM	\$9	3:42 PM	\$626.30

For each price combination, we observe that at most one vehicle is required to service all NDD and overnight orders. Thus, for each price combination, the optimal choice of deadlines is one of polynomially


many options. For each $m \in \{1, 2, 3\}$, the computational time required to enumerate and check the revenue of each of the potential optimal solutions (see Appendix A.13 for computational details) for all price combinations is less than three minutes in total. Table 1 summarizes the optimal solution and associated optimal delivery revenue for each potential fleet size; there is no benefit to using a two-level pricing scheme for a fleet size of $m = 1$ because the associated optimal solution sets the discounted deadline to $t = 0$. Based on these results, we observe that using $m = 3$ vehicles is optimal for $\varphi \leq \$156.80$, using $m = 2$ vehicles is optimal for $\$156.80 \leq \varphi \leq \341.06 , and using $m = 1$ vehicle is optimal for $\varphi \geq 341.06$. However, we operate at a loss for any $\varphi > \$235.25$.

Examining the list of feasible solutions, we observe that price combinations similar to the optimal combination entail near-optimal total revenues, while less-similar price combinations entail clearly suboptimal total revenues. As an example for $m = 3$, the discounted/base SDD price combination of \$7.50 and \$8.50 entails a total daily revenue of \$624.08, while the discounted/base SDD price combination of \$3 and \$12 entails a total daily revenue of \$403.77. This implies that the total revenue function is fairly “flat” around the optimum, so managers can slightly adjust the optimal prices without fear of significant lost revenue. We include the five best solutions for each m in Appendix B to further illustrate this point.

As with the NDD⁺ objective, we can also perform a sensitivity analysis for the two-level pricing model. To illustrate, suppose we seek to understand how the composition of the customer base (i.e., the relative proportion of overnight, Type A, and Type B orders) affects the optimal revenue and pricing solutions. The original solutions assumed that Type B customers make up approximately 64.6% of all arrivals. Assuming that the expected total number of daily customer arrivals remains fixed, how does varying this percentage affect the optimal solution? We analyze this question by re-solving for the optimal prices and deadlines while varying the proportion of Type B orders between 50% and 80% by 0.5% increments.

Figure 16a plots the maximum revenue across this range for $m \in \{1, 2, 3\}$. We observe that revenue increases nonlinearly with the Type B proportion. The positive correlation is as expected because relatively impatient Type B customers are willing to pay higher prices for SDD. Next, we observe that there is no value of φ for which using $m = 1$ vehicle is optimal *and* profitable. Thus, we must choose between using $m = 2$ or $m = 3$ vehicles; the decision depends on both the proportion of Type B customers and the value of φ . The additional revenue associated with using the third vehicle increases as we increase the proportion of Type B orders, from \$143.47 at one end of the range to \$159.06 at the other end.

Figure 16b plots the early and late SDD deadlines as a function of the proportion of Type B customers for $m = 2$. We observe a piecewise function with four apparent breakpoints. Unsurprisingly, these four breakpoints (marked with dashed vertical lines) are also where the optimal price combination changes. Specifically, an increase in Type B customers tends to increase the prices in the optimal combination. From left to right in Figure 16b, the four optimal price combinations are $(\hat{p}_0, p_0) = (\$7.50, \$8.50)$, $(\$8, \$8.50)$, $(\$8, \$9)$, and $(\$8.50, \$9)$. However, we again observe that the total revenue is fairly flat around the optimum, so

Figure 16 Sensitivity analysis of customer type composition

multiple price combinations are almost optimal. Specifically, the best solution with $(\hat{p}_0, p_0) = (\$8, \$9)$ — the original price combinations from Table 1 — achieves at least 99.5% of the optimal revenue in every case analyzed here. In Appendix B, we conduct a similar sensitivity analysis of customers' threshold prices.

Finally, it is worthwhile for system managers to quantify the benefit of a two-level SDD pricing scheme versus a simpler design with a single SDD price. In our zone, the two-level scheme entails an additional \$2.02 in daily revenue for $m = 2$ and an additional \$3.91 for $m = 3$ when compared to the best solution with a single SDD price, assuming our original parameter settings. The associated relative increase in profit is 1.2% for $\varphi = \$100$ and 2.9% for $\varphi = \$200$; this is a significant margin if a firm operates dozens of zones daily in a metropolitan region. It is ultimately the decision of system managers to determine on a zone-by-zone basis whether the monetary benefits of the two-level pricing scheme outweigh the slight increase in the complexity of managing the system.

7.3 Operational Simulations and Convergence Analysis

As a final step, we assess the convergence of a stochastic operational model and compare its performance to our predictions. Specifically, over a 12-week horizon, we simulate discrete customer arrivals and fulfill orders by exactly computing optimal TSP tours over customers' delivery locations. We compare the results of our simulations with the system behavior predicted by the deterministic model and analyze the effect of varying initial conditions. Specifically, this simulation study analyzes the optimal solution to the early delivery maximization problem studied earlier in this section with NDD⁺ and SDD order deadlines of 10:25 AM ($\hat{C} = 85.25$) and 3 PM ($C = 360$), respectively.

We generate daily stochastic customer arrivals as follows. Customers arrive via non-stationary Poisson processes at rates depicted in Figure 14. A customer arriving into the system at a time $t \in (0, T)$ is associated

with a type and threshold price. If the customer’s threshold price does not exceed the SDD price ($p_0 = \$5$), the customer places a SDD order. Otherwise, the customer places a NDD order (Type A) or abandons the system (Type B). Customers’ threshold prices are i.i.d. for each type, with distributions given by the WP functions (3). Excluding the first day of the horizon, the quantity of overnight orders prior to each day is distributed as Poisson(30). The delivery location associated with each customer is randomly generated by the same procedure used to empirically estimate the function f .

The daily operation broadly proceeds as follows, beginning at $t = 0$ and assuming some quantity of time-zero accumulation. As orders arrive into the system, they are classified based on their fulfillment date: if a customer places a SDD order at any time *or* a NDD order prior to the NDD⁺ deadline, it is designated for delivery today; otherwise, it is relegated for fulfillment tomorrow. When an order is designated for delivery today, it is added to the set of unserved orders; this set may also contain yesterday’s NDD orders and overnight orders. At all times, the system operator maintains the *current dispatch duration*, calculated as the fastest TSP tour — including setup and per-order service times — through the depot and the delivery locations associated with all of the orders in the unserved set. The first vehicle departs to fulfill all orders when the current time plus the current dispatch duration equals T . At this time, the unserved set is reset and the same procedure repeats for the second and third vehicles. Among all orders designated for fulfillment on a particular day, orders are batched into dispatches based on their arrival time; no geographic discrimination occurs. Occasionally, an arriving order may push a vehicle’s return time past T . In such cases, that order is not served by the current vehicle, and the current vehicle dispatches immediately. The departure time of the third vehicle also constitutes that day’s recorded SDD order deadline. The NDD⁺ deadline of 10:25 AM remains unchanged from day to day.


For each of 25 trials, we simulate the system for a 12-week (84-day) time horizon. In our model’s steady state solution, the daily time-zero accumulation is $\eta \approx 77.25$. Therefore, for each trial, the time-zero accumulation on day 1 is chosen as Poisson(77.25). For each trial and day, we record representative data about the system, including the SDD order deadline, the fulfilled quantity of each type of order, and information about each of the three dispatches. In aggregate, we expect this data to align closely with the values predicted by our deterministic model. Table 2 lists the averages and standard deviations (over all trials and days) of some representative values. For comparison, the table also lists the corresponding values associated with the optimal solution to the deterministic model (labeled ‘predicted’).

Table 2 Representative daily system values: stochastic simulations vs. deterministic model predictions

	SDD order deadline (min. after 9 AM)	Total dispatch duration (min.)	SDD orders fulfilled	NDD ⁺ orders fulfilled
Simulated, mean (\pm st. dev.)	359.66 (\pm 15.66)	941.60 (\pm 44.03)	88.13 (\pm 6.38)	8.31 (\pm 2.95)
Predicted	360	944.77	88.90	8.28

All of the simulated averages are within 1% of the values predicted by the deterministic model. This is in line with existing studies that use CA methods in the last-mile delivery context. Because this data is aggregated across days, it is natural to ask whether the behavior of the system changes as the system evolves. Specifically, does the system exhibit greater deviation from the steady state as time goes on? Figure 17 records the mean, minimum, and maximum SDD order deadline across all trials for each day in the horizon. The chart also displays the mean plus/minus one standard deviation across all trials for each day. Observe that no trend is evident over time for any of the summary statistics. Additionally, an analysis of each individual simulation trial reveals that SDD order deadlines never stray outside of the 2:30–3:30 PM time window for more than two consecutive days. This data shows that the system tends to stay near the steady state over time and does not deviate in either direction.

Figure 17 Daily SDD order deadline (minutes after 9 AM) summary statistics across trials



In these simulations, we intentionally set up our system’s initial condition to mimic the steady state. We showed previously that the deterministic model converges to the saturated steady state. It is natural to ask whether the stochastic system exhibits similar behavior when the quantity of orders at $t = 0$ on the first day differs from the expected steady-state daily time-zero accumulation η . To answer this question, we re-simulate each of the 25 trials over the entire 84-day horizon twice more. In the first case, we begin the simulation with zero orders accumulated at $t = 0$ on day 1. In the second case, we begin the simulation with $\text{Poisson}(3\eta)$ orders accumulated at $t = 0$ on day 1; the value of 3η is chosen to ensure all initial orders can be served on day 1. In both cases, every other customer arrival remains unchanged (with respect to customer type, location, arrival time, and threshold price) from the initial simulations.

The results suggest that the system rapidly converges towards the steady state. In each of the 25 trials and for both of the new initial conditions, the evolution of the system coincides with that of the original simulation no later than day 4. That is, for any chosen trial, the dispatching and fulfillment data are identical to those of the original simulation from day 4 onward (and usually earlier) for both new initial conditions.

Table 3 Simulated effects of initial conditions

Initial time-zero accumulation	Mean SDD order deadline (minutes after 9 AM)				
	Day 1	Day 2	Day 3	Day 4	Day 5
0	405.56	361.16	355.25	355.52	357.62
Poisson(η)	356.34	359.67	355.25	355.52	357.62
Poisson(3η)	182.32	353.25	355.08	355.52	357.62

For each of the three initial conditions, Table 3 records the average simulated SDD order deadline for the first five days across all trials.

The largest changes take place within the first two days, suggesting that the system rapidly adjusts itself — by way of varying the SDD order deadline as necessary — to return to the steady state. These results illustrate that the idea of a system steady state, which we assumed to exist when building our deterministic model, is indeed reflected in practice. Our analysis of initial conditions can be generalized to an additional managerial insight: the system is robust in the sense that it can recover fairly quickly from unexpected, brief demand shocks — positive or negative — to return to steady-state operation.

8. Concluding Remarks

We studied an e-retail system in which price-sensitive customers choose between price-differentiated SDD and NDD fulfillment options; both types of orders are fulfilled concurrently by the same delivery fleet. We proposed a deterministic model, in which stochastic order arrivals are replaced with fluid order accumulations, in order to facilitate system design and analysis. We began by showing that the feasibility of serving all orders in the model directly corresponds to the feasibility of a specific greedy vehicle dispatching policy. We then showed that the system model converges to a steady state over time.

We then analyzed two design objectives. First, we sought to improve customer satisfaction by maximizing the daily quantity of orders fulfilled a day early. Second, we aimed to maximize profit via a two-level static SDD pricing scheme in which early-ordering customers face a potentially-discounted SDD price compared to late-ordering customers. For both objectives, we derived structural managerial insights and efficient solution approaches. Finally, we applied our model to a case study set in suburban Denver and validated our deterministic approach via stochastic simulations with exactly computed vehicle routes. Our simulations suggest that the system designs produced by our deterministic model are reflected in the day-to-day stochastic realities of the system, regardless of initial conditions.

The effectiveness of the greedy policy was shown to be a consequence of the constraints of the system (m vehicles dispatching once daily). The results related to feasibility, convergence, and pricing in this paper therefore depend directly on these constraints. Future work may derive analogous structural results under different constraints on the vehicle fleet (e.g., m vehicles dispatching twice daily, one vehicle dispatching m times daily, etc.) for which greedy-type policies may or may not be optimal.

For tactical planning purposes, an underlying assumption of the model is that no geographical discrimination or geographical batching occurs within the zone. As a result, delivery locations are not taken into

account when deciding which NDD customers to upgrade to NDD⁺ or which SDD customers pay the cheaper price \hat{p}_0 . This allows for fair customer service across the zone and streamlined decision-making at the operational level. However, future work may seek to slightly improve system performance by relaxing this assumption to allow for geographical discrimination. For example, potential NDD⁺ upgrades could be restricted to customers in certain locations closer to the depot. Similarly, the operational work of Afsar et al. (2021), in which different parts of a service area are offered different delivery prices, could be applied to tactical decision-making. Such an approach may seek a set of location-dependent SDD prices that induce a feasible steady-state system model. As discussed earlier, these strategies for geographical discrimination would be especially valuable for delivery fleets larger than the ones considered in this paper.

We anticipate other potential extensions of the modeling framework proposed in this paper. For example, we could consider incorporating customer arrival rates that are dependent on the day of the week or other seasonal variations, dynamic pricing guided by optimal static prices, and integrating other price-differentiated fulfillment options (like two-day delivery). Future empirical research on customer behavior in rapid e-retail systems would be valuable for motivating and analyzing extensions to our model. In particular, understanding how customers' ordering strategies change in the presence of random free delivery upgrades would be instrumental for empirically validating our randomized NDD⁺ strategy proposed in Section 5.3.

Acknowledgments

Dipayan Banerjee's work is supported by the NSF Graduate Research Fellowship (DGE-1650044). The authors thank the editorial team and three anonymous referees for their constructive feedback.

References

- Afsar HM (2022) Traveling salesperson problem with unique pricing and stochastic thresholds. *Computers & Industrial Engineering* 173:108696.
- Afsar HM, Afsar S, Palacios JJ (2021) Vehicle routing problem with zone-based pricing. *Transportation Research Part E: Logistics and Transportation Review* 152:102383.
- Albareda-Sambola M, Fernández E, Laporte G (2014) The dynamic multiperiod vehicle routing problem with probabilistic information. *Computers & Operations Research* 48:31–39.
- Allen J, Piecyk M, Piotrowska M, McLeod F, Cherrett T, Ghali K, Nguyen T, Bektas T, Bates O, Friday A, Wise S, Austwick M (2018) Understanding the impact of e-commerce on last-mile light goods vehicle activity in urban areas: the case of London. *Transportation Research Part D: Transport and Environment* 61(B):325–338.
- Amazon (2009) Amazon launches same day delivery in seven major cities and expands Saturday delivery options. URL <https://press.aboutamazon.com/news-releases/news-release-details/amazon-launches-same-day-delivery-seven-major-cities-and-expands>.
- Andreatta G, Lulli G (2008) A multi-period TSP with stochastic regular and urgent demands. *European Journal of Operational Research* 185(1):122–132.

- Angelelli E, Bianchessi N, Mansini R, Speranza MG (2010) Comparison of policies in dynamic routing problems. *Journal of the Operational Research Society* 61(4):686–695.
- Angelelli E, Savelsbergh M, Speranza MG (2007a) Competitive analysis for dynamic multiperiod uncapacitated routing problems. *Networks* 49(4):308–317.
- Angelelli E, Savelsbergh M, Speranza MG (2007b) Competitive analysis of a dispatch policy for a dynamic multi-period routing problem. *Operations Research Letters* 35(6):713–721.
- Angelelli E, Savelsbergh M, Speranza MG (2009) Short term strategies for a dynamic multi-period routing problem. *Transportation Research Part C: Emerging Technologies* 17(2):106–119.
- Ansari S, Başdere M, Li X, Ouyang Y, Smilowitz K (2018) Advancements in continuous approximation models for logistics and transportation systems: 1996–2016. *Transportation Research Part B: Methodological* 107:229–252.
- Archetti C, Bertazzi L (2021) Recent challenges in routing and inventory routing: e-commerce and last-mile delivery. *Networks* 77(2):255–268.
- Avraham E, Raviv T (2021) The steady-state mobile personnel booking problem. *Transportation Research Part B: Methodological* 154:266–288.
- Banerjee D, Erera A, Stroh A, Toriello A (2023) Who has access to e-commerce and when? Time-varying service regions in same-day delivery. *Transportation Research Part B: Methodological* 170:148–168.
- Banerjee D, Erera A, Toriello A (2022) Fleet sizing and service-region partitioning for same-day delivery systems. *Transportation Science* 56(5):1327–1347.
- Beardwood J, Halton J, Hammersley J (1959) The shortest path through many points. *Mathematical Proceedings of the Cambridge Philosophical Society* 55(4):299–327.
- Behrendt A, Savelsbergh M, Wang H (2024) Task assignment, pricing, and capacity planning for a hybrid fleet of centralized and decentralized couriers. *Transportation Research Part C: Emerging Technologies* 160:104533.
- Bülbül Z (2022) *Dynamic Route Planning for Last-Mile Delivery*. Ph.D. thesis, University of Waterloo.
- Carlsson JG, Liu S, Salari N, Yu H (2024) Provably good region partitioning for on-time last-mile delivery. *Operations Research* 72(1):91–109.
- Chen X, Ulmer M, Thomas B (2022) Deep Q-learning for same-day delivery with vehicles and drones. *European Journal of Operational Research* 298:939–952.
- Deshpande V, Pendem P (2023) Logistics performance, ratings, and its impact on customer purchasing behavior and sales in e-commerce platforms. *Manufacturing & Service Operations Management* 25(3):827–845.
- Estrada-Moreno A, Savelsbergh M, Juan A, Panadero J (2019) Biased-randomized iterated local search for a multi-period vehicle routing problem with price discounts for delivery flexibility. *International Transactions in Operational Research* 26(4):1293–1314.
- Fisher M, Gallino S, Xu JJ (2019) The value of rapid delivery in omnichannel retailing. *Journal of Marketing Research* 56(5):732–748.

- Fleckenstein D, Klein R, Steinhardt C (2023) Recent advances in integrating demand management and vehicle routing: A methodological review. *European Journal of Operational Research* 306(2):499–518.
- Franceschetti A, Honhon D, Laporte G, van Woensel T, Fransoo J (2017) Strategic fleet planning for city logistics. *Transportation Research Part B: Methodological* 95:19–40.
- Galiullina A, Mutlu N, Kinable J, Van Woensel T (2024) Demand steering in a last-mile delivery problem with home and pickup point delivery options. *Transportation Science* 58(2):454–473.
- Gallego G, Topaloglu H (2019) *Revenue Management and Pricing Analytics* (Springer).
- Ge J, Honhon D, Fransoo J, Zhao L (2021) Supplying to Mom and Pop: Traditional retail channel selection in megacities. *Manufacturing & Service Operations Management* 23(1):19–35.
- HeiGIT (2022) Openrouteservice. URL <https://openrouteservice.org/>.
- Karp RM, Steele JM (1985) Probabilistic analysis of heuristics. Lawler EL, Lenstra JK, Rinnooy Kan AHG, Shmoys DB, eds., *The Traveling Salesman Problem*, chapter 6 (Wiley).
- Klapp M, Erera A, Toriello A (2020) Request acceptance in same-day delivery. *Transportation Research Part E: Logistics and Transportation Review* 143:102083.
- Klein V, Steinhardt C (2023) Dynamic demand management and online tour planning for same-day delivery. *European Journal of Operational Research* 307(2):860–886.
- Köhler C, Ehmke JF, Campbell AM, Cleophas C (2023) Evaluating pricing strategies for premium delivery time windows. *EURO Journal on Transportation and Logistics* 12:100108.
- NASA SEDAC, Columbia University CIESIN (2018) Population Estimation Service, Version 3. URL <https://doi.org/10.7927/H4DR2SK5>.
- Peng L, Murray C (2022) VeRoViz: A vehicle routing visualization toolkit. *INFORMS Journal of Computing* 34(4):1842–1848.
- Salgado A (2022) Quiet Platforms will offer faster Fanatics deliveries in more markets. Supply Chain Dive, URL <https://bit.ly/3QtBV1M>.
- Snoeck A, Merchán D, Winkenbach M (2020) Revenue management in last-mile delivery: state-of-the-art and future research directions. *Transportation Research Procedia* 46:109–116.
- Snoeck A, Winkenbach M (2022) A discrete simulation-based optimization algorithm for the design of highly responsive last-mile distribution networks. *Transportation Science* 56(1):201–222.
- Song B, Li M (2018) Dynamic pricing with service unbundling. *Production and Operations Management* 27(7):1334–1354.
- Sriram S, Chintagunta P, Manchanda P (2015) Service quality variability and termination behavior. *Management Science* 61(11):2739–2759.
- Steele JM (1997) *Probability Theory and Combinatorial Optimization* (Society for Industrial and Applied Mathematics).

- Strauss A, Klein R, Steinhardt C (2018) A review of choice-based revenue management: theory and methods. *European Journal of Operational Research* 271(2):375–387.
- Stroh A, Erera A, Toriello A (2022) Tactical design of same-day delivery systems. *Management Science* 68(5):3444–3463.
- TRxADE Health (2022) TRxADE Health, Inc., launches next day delivery services via TRxADE Prime Plus program. URL <https://finance.yahoo.com/news/trxade-health-inc-launches-next-123000666.html>.
- Ulmer M (2020) Dynamic pricing and routing for same-day delivery. *Transportation Science* 54(4):1016–1033.
- Ulmer M, Soeffker N, Mattfeld D (2018) Value function approximation for dynamic multi-period vehicle routing. *European Journal of Operational Research* 269(3):883–899.
- Ulmer M, Streng S (2019) Same-day delivery with pickup stations and autonomous vehicles. *Computers & Operations Research* 108:1–19.
- Ulmer M, Thomas B (2018) Same-day delivery with heterogeneous fleets of drones and vehicles. *Networks* 72(4):475–505.
- Varma SM, Bumpensanti P, Maguluri ST, Wang H (2023) Dynamic pricing and matching for two-sided queues. *Operations Research* 71(1):83–100.
- Voccia S, Campbell AM, Thomas B (2019) The same-day delivery problem for online purchases. *Transportation Science* 53(1):167–184.
- Waßmuth K, Köhler C, Agatz N, Fleischmann M (2023) Demand management for attended home delivery – a literature review. *European Journal of Operational Research* 311(3):801–815.
- Wen M, Cordeau JF, Laporte G, Larsen J (2010) The dynamic multi-period vehicle routing problem. *Computers & Operations Research* 37(9):1615–1623.
- Yıldız B, Savelsbergh M (2019) Service and capacity planning in crowd-sourced delivery. *Transportation Research Part C: Emerging Technologies* 100:177–199.
- Yıldız B, Savelsbergh M (2020) Pricing for delivery time flexibility. *Transportation Research Part B: Methodological* 133:230–256.

Appendix A Omitted Technical Details

A.1 Proof of Corollary 1

Let $P_i = X_{1+\sum_{j=1}^{i-1} N_j}, \dots, X_{\sum_{j=1}^i N_j}$ for all i ; note that $\{X_1, X_2, \dots, X_N\} = \bigcup_{i=1}^k P_i$. Let $\varepsilon > 0$. By the SLLN, it holds with probability (w.p.) 1 that, for all sufficiently large k ,

$$\left| \sum_{i=1}^k N_i - vk \right| \leq vk \implies \sum_{i=1}^k N_i \leq 2vk \implies \sqrt{\sum_{i=1}^k N_i} \leq \sqrt{2vk} \quad (\text{A1})$$

and

$$\left| \sum_{i=1}^k N_i - vk \right| \leq \frac{vk}{2} \implies \sum_{i=1}^k N_i \geq \frac{vk}{2}. \quad (\text{A2})$$

Next, the BHH (1959) Theorem and (A2) imply the existence of a finite $\beta > 0$ such that, w.p. 1,

$$\left| \text{TSP} \left(\bigcup_{i=1}^k P_i \right) - \beta \sqrt{\sum_{i=1}^k N_i} \right| \leq \frac{\varepsilon}{2\sqrt{2v}} \sqrt{\sum_{i=1}^k N_i}$$

for all sufficiently large k . By the SLLN, it then holds w.p. 1 that, for all sufficiently large k , $|\sum_{i=1}^k N_i - vk| \leq \varepsilon^2 k / 4\beta^2$. It then holds w.p. 1 that, for all sufficiently large k ,

$$\begin{aligned} \left| \text{TSP} \left(\bigcup_{i=1}^k P_i \right) - \beta \sqrt{vk} \right| &\leq \left| \text{TSP} \left(\bigcup_{i=1}^k P_i \right) - \beta \sqrt{\sum_{i=1}^k N_i} \right| + \left| \beta \sqrt{\sum_{i=1}^k N_i} - \beta \sqrt{vk} \right| \\ &\leq \left| \text{TSP} \left(\bigcup_{i=1}^k P_i \right) - \beta \sqrt{\sum_{i=1}^k N_i} \right| + \beta \sqrt{\left| \sum_{i=1}^k N_i - vk \right|} \\ &\leq \frac{\varepsilon}{2\sqrt{2v}} \sqrt{\sum_{i=1}^k N_i} + \beta \sqrt{\frac{\varepsilon^2 k}{4\beta^2}} \\ &\leq \frac{\varepsilon \sqrt{2vk}}{2\sqrt{2v}} + \beta \sqrt{\frac{\varepsilon^2 k}{4\beta^2}} \\ &= \varepsilon \sqrt{k}. \end{aligned}$$

Rearranging gives $\left| \frac{1}{\sqrt{k}} \text{TSP} \left(\bigcup_{i=1}^k P_i \right) - \beta \sqrt{v} \right| \leq \varepsilon$ as desired. \square

In the following proofs and algorithms, we deal solely with the deterministic approximation of the system. Therefore, in the remainder of this appendix, the term ‘‘system’’ is used to mean ‘‘approximated model of the system’’ unless stated otherwise.

A.2 Algorithmic Details of Greedy Policy

For all $t \in [0, T]$, define $\lambda_0(t)$ as the accumulation rate of orders that are to be served today; specifically, $\lambda_0(t) = w_A(p_0)\lambda_A(t) + w_B(p_0)\lambda_B(t) > 0$ for all $t \in [0, C]$ and $\lambda_0(t) = 0$ for all $t \in (C, T]$.

Algorithm 1: Computing the greedy policy

Input : parameters C, T, m, η and functions $\lambda_0(t), f(n), Q(t)$

Output : m -vehicle greedy policy $((t_1, q_1), \dots, (t_m, q_m))$

```

1 set  $k = \min \{ \lfloor \eta / Q(0) \rfloor, m - 1 \}$ 
2 for  $d = 1, 2, \dots, k$  do
3    $\lfloor$  set  $t_d \leftarrow 0, q_d \leftarrow Q(0)$ 
4 set  $d \leftarrow k + 1$ 
5 while  $d < m$  do
6   if  $d = k + 1$  then
7     set  $t_d$  s.t.  $t_d + f(\eta - \sum_{j=1}^k q_j + \int_0^{t_d} \lambda_0(t) dt) = T$ 
8     set  $q_d \leftarrow \eta - \sum_{j=1}^k q_j + \int_0^{t_d} \lambda_0(t) dt$ 
9   else
10    set  $t_d$  s.t.  $t_d + f(\int_{t_{d-1}}^{t_d} \lambda_0(t) dt) = T$ 
11    set  $q_d \leftarrow \int_{t_{d-1}}^{t_d} \lambda_0(t) dt$ 
12   if  $t_d > C$  then
13      $\lfloor$  set  $t_d \leftarrow C$ 
14   set  $d \leftarrow d + 1$ 
15 set  $t_m \leftarrow C$  and  $q_m \leftarrow \eta + \int_0^C \lambda_0(t) dt - \sum_{j=1}^{m-1} q_j$ 
16 return  $((t_1, q_1), \dots, (t_m, q_m))$ 

```

A.3 Preliminaries

LEMMA 2. $Q(t)$ is continuous, convex, and decreasing.

Proof. Because $Q = f^{-1}(T - t)$, Q is continuous on account of being the composition of continuous functions. Because f is increasing and $f(Q(t)) = T - t$ by definition for all t , Q is decreasing. Since f is increasing and concave, f^{-1} is increasing and convex. Additionally, $T - t$ is a convex function of t , so $Q(t) = f^{-1}(T - t)$ is convex. \square

A.4 Proof of Lemma 1

Proof. Denote the greedy policy by $P = ((t_1, q_1), \dots, (t_m, q_m))$. Suppose that the system is saturated, and there exists some other m -vehicle policy $\hat{P} = ((\hat{t}_1, \hat{q}_1), \dots, (\hat{t}_m, \hat{q}_m))$ feasible for the system. Define $d' = \max\{d \mid \hat{q}_d > q_d\}$. Such a d' must exist because P and \hat{P} are different policies, so at least one element must differ. Since the instance is saturated, it must hold that $t_{d'} + f(q_{d'}) = T$. Observe that $d' \neq m$ because the policy \hat{P} would be infeasible otherwise, since $\hat{q}_m > q_m$ and $C + f(q_m) = T$ together imply $\hat{t}_m + f(\hat{q}_m) = C + f(\hat{q}_m) > T$.

Together, $t_{d'} + f(q_{d'}) = T$ and $\hat{q}_{d'} > q_{d'}$ imply $\hat{t}_{d'} < t_{d'}$ because f is an increasing function (equivalently, because Q is a decreasing function). If $\hat{t}_{d'} < 0$, then we have a trivial contradiction to the feasibility of \hat{P} ; suppose that this is not the case. Because dispatches in the greedy policy depart immediately upon their corresponding quantity being accumulated, $\hat{t}_{d'} < t_{d'}$ further implies

$$\sum_{d=1}^{d'} \hat{q}_d < \sum_{d=1}^{d'} q_d. \quad (\text{A3})$$

As a result,

$$\sum_{d=d'+1}^m \hat{q}_d > \sum_{d=d'+1}^m q_d. \quad (\text{A4})$$

Then, there must exist some $d \in \{d' + 1, \dots, m\}$ such that $\hat{q}_d > q_d$, which contradicts the definition of d' . Therefore, P is the only policy which can feasibly serve the system. \square

A.5 Proof of Theorem 1

Proof. Denote the greedy policy by $P = ((t_1, q_1), \dots, (t_m, q_m))$. To avoid triviality, assume $t_m > 0$. Suppose that the greedy policy is infeasible, and there exists some other m -vehicle policy $\hat{P} = ((\hat{t}_1, \hat{q}_1), \dots, (\hat{t}_m, \hat{q}_m))$ feasible for the system. Consider a modified system in which all of the parameters are identical, except that C is replaced with $t_m < C$ and $\lambda_0(t) = 0$ for all $t > t_m$. This modified system is saturated by definition. However, \hat{P} is still feasible for the modified system. This contradicts Lemma 1; thus, no m -vehicle policy can serve the original system. \square

A.6 Proof of Theorem 2

We first show that the greedy policy minimizes total routing time. This result and its proof are generalizations of Theorem 1 of Stroh et al. (2022).

Proof. If the system is saturated, then the greedy policy is the only feasible policy, so it must necessarily be optimal. Assume instead that the system is not saturated, and there exists an optimal non-greedy policy $P^1 = ((t_1^1, q_1^1), \dots, (t_m^1, q_m^1))$. We aim to show that P^1 can be transformed into either the greedy policy with the same total dispatch time, or into a policy with a strictly lower total dispatch time.

First, shift the dispatch departure times of P^1 earlier such that each vehicle dispatches as soon as its corresponding quantity is accumulated. Denote this new policy $P^2 = ((t_1^2, q_1^2), \dots, (t_m^2, q_m^2))$. Observe that the total dispatch times associated with P^1 and P^2 are equal since each dispatch in both policies serves the same quantity.

Now, considering P^2 , we wish to modify the dispatches such that the dispatch quantities are non-increasing. If the dispatch quantities of P^2 are already non-increasing, set $P^3 = P^2$; if not, let $d \in [m - 1]$ be the smallest dispatch index such that $q_d^2 < q_{d+1}^2$. Swap the quantities of dispatches d and $d + 1$ and modify their departure times so that they depart as soon as their quantity is accumulated. This preserves feasibility for the following reasons.

- $t_{d+1}^2 + f(q_{d+1}^2) \leq T$ implies that the new d -th dispatch returns to the depot by T , since the new departure time of the new d -th dispatch serves q_{d+1}^2 and departs no later than t_{d+1}^2 by construction.
- $t_{d+1}^2 + f(q_{d+1}^2) \leq T$ implies that the new $(d+1)$ -th dispatch returns to the depot by T , since the new $(d+1)$ -th dispatch serves a quantity less than q_{d+1}^2 and departs no later than t_{d+1}^2 by construction.

Observe that this swapping does not change the total dispatch duration. Continue to perform this pairwise swapping procedure between consecutive dispatches until dispatch quantities are non-increasing. Denote the resulting policy P^3 . (★)

In the policy P^3 , each dispatch departs as soon as its corresponding quantity is accumulated, and dispatch quantities are non-increasing. Additionally, all dispatch quantities are strictly positive, as otherwise we could remove all zero-quantity dispatches and serve I with fewer than m vehicles. If P^3 is the greedy policy, the greedy policy is optimal and we are done. Suppose instead that P^3 is not the greedy policy. Then, there is some dispatch $d \in [m-1]$ such that $t_d^3 + f(q_d^3) < T$. There must exist some sufficiently small $\delta > 0$ and $\varepsilon \in [0, t_{d+1}^3 - t_d^3]$ such that the following dispatch updates preserve the policy's feasibility:

$$\begin{aligned} (t_d^3, q_d^3) &\rightarrow (t_d^3 + \varepsilon, q_d^3 + \delta), \\ (t_{d+1}^3, q_{d+1}^3) &\rightarrow (t_{d+1}^3, q_{d+1}^3 - \delta). \end{aligned}$$

Because $q_d^3 \geq q_{d+1}^3$ and f is strictly concave and increasing, $f(q_d^3 + \delta) + f(q_{d+1}^3 - \delta) < f(q_d^3) + f(q_{d+1}^3)$. Therefore, this updated policy has a strictly lower total dispatch duration than P^1 , a contradiction. As a result, the greedy policy is optimal with respect to total dispatch duration. □

Next, suppose that the cost associated with a dispatch (with quantity q) is not necessarily $f(q)$ but rather a continuous, non-negative, strictly increasing, strictly concave function $\kappa(q)$. We prove that the greedy policy remains cost-optimal in this case.

Proof. The proof is identical up to (★), since moving dispatch departure times and swapping dispatch quantities does not affect the total cost over all dispatches.

In the policy P^3 , each dispatch departs as soon as its corresponding quantity is accumulated, and dispatch quantities are non-increasing. Additionally, all dispatch quantities are strictly positive, as otherwise we could remove all zero-quantity dispatches and serve I with fewer than m vehicles. If P^3 is the greedy policy, the greedy policy is optimal and we are done. Suppose instead that P^3 is not the greedy policy. Then, there is some dispatch $d \in [m-1]$ such that $t_d^3 + f(q_d^3) < T$. There must exist some sufficiently small $\delta > 0$ and $\varepsilon \in [0, t_{d+1}^3 - t_d^3]$ such that the following dispatch updates preserve the policy's feasibility:

$$\begin{aligned} (t_d^3, q_d^3) &\rightarrow (t_d^3 + \varepsilon, q_d^3 + \delta), \\ (t_{d+1}^3, q_{d+1}^3) &\rightarrow (t_{d+1}^3, q_{d+1}^3 - \delta). \end{aligned}$$

Because $q_d^3 \geq q_{d+1}^3$ and κ is strictly concave and increasing, $\kappa(q_d^3 + \delta) + \kappa(q_{d+1}^3 - \delta) < \kappa(q_d^3) + \kappa(q_{d+1}^3)$. Therefore, this updated policy has a strictly lower total cost than P^1 , a contradiction. As a result, the greedy policy is optimal with respect to total cost as defined via $\kappa(\cdot)$. □

A.7 Technical Details: Partial Ordering and Induced Subsystems

Let the function $\lambda(t)$ denote the total accumulation rate of orders that must be served by the end of the service day; this includes both SDD orders and any NDD^+ orders (see Section 5). Modifying prices and cutoffs in our CA system model changes the size of the time-zero accumulation η as well as $\lambda(t)$. Assuming a fixed region, m , f , and T , we consider two systems I (defined by the parameters η^I, C^I and function $\lambda^I(t)$) and J (defined by the parameters η^J, C^J and function $\lambda^J(t)$). We seek to formally compare the difficulty of feasibly serving system I and system J .

We say that I is *at least as difficult* as J (denoted as $I \succeq J$) if both of the following conditions hold:

- (i) $C^J \leq C^I$,
- (ii) $\eta^J + \int_0^T \lambda^J(t) dt \leq \eta^I + \int_0^T \lambda^I(t) dt$, and
- (iii) $\eta^J + \int_0^t \lambda^J(s) ds \geq \eta^I + \int_0^t \lambda^I(s) ds$ for all $t \in [0, C^J]$.

Condition (i) states that SDD and NDD^+ orders continue to accumulate at least as late in system I versus system J . Condition (ii) states that the total quantity of orders in system I is at least that of system J . Condition (iii) implies that orders are available earlier in system J than in system I . We also say that $I \succeq J$ if, instead of the above conditions, I and J satisfy

- (iv) $\eta^J \leq \eta^I$, and
- (v) $\lambda^J(t) \leq \lambda^I(t)$ for all $t \in [0, T]$.

These conditions state that the accumulation in system I always “dominates” that of J . Note that “ \succeq ” represents a partial order and not a total order, since two systems are not necessarily comparable via these conditions. This formalism implies that if I can feasibly be served by a policy $P^0 = ((t_1^0, q_1^0), \dots, (t_m^0, q_m^0))$ and $I \succeq J$, then J can be feasibly served by a policy $P^1 = ((t_1^0, q_1^1), \dots, (t_m^0, q_m^1))$ with departure times identical to those of P^0 and with $q_i^1 \leq q_i^0$ for all $i \in [m]$.

In later proofs, it will occasionally be useful to consider a *subsystem* with a reduced total quantity relative to the entire system. Given a base system I and a quantity χ that is no greater than the total quantity in I (not including Type B customers who abandon the system), we say that J is a χ -induced subsystem of I if

$$\eta^J + \int_0^{C^J} \lambda^J(t) dt = \chi, \quad (\text{A5})$$

where $\eta^J = \eta^I$ and $\lambda^J = \lambda^I$, assuming $\eta^I \leq \chi$. If instead $\eta^I > \chi$, then $C^J = 0$ and $\eta^J = \chi$. Given a policy $P = ((t_1, q_1), \dots, (t_m, q_m))$ applied to I , we occasionally refer to the $(\sum_{i=1}^d q_i)$ -induced subsystem as “the subsystem induced by the first d dispatches of policy P .” These ideas, which formalize the concept of system difficulty, are used extensively in subsequent proofs.

A.8 Proof of Theorem 3

A.8.1 Supporting Results

LEMMA 3. *The function $G(q)$ is decreasing and continuous for all $q \geq 0$.*

Proof. The decreasing nature of G follows directly from the prior discussion on partial orders. We show that G is Lipschitz continuous, which implies that G is continuous. Let $q \geq 0$.

Let $1/K = \min_{t \in [0, T]} \lambda_0(t)$; such a value exists by the continuity of λ_0 and the Extreme Value Theorem. Let $\delta > 0$ be small. Consider two independent service days: day i with q orders available at $t = 0$ and day j with $q + \delta$ orders available at $t = 0$; assume that we serve both days with the greedy policy and extend the SDD cutoff so that all vehicles return at $t = T$ on both days. Now, suppose that the SDD cutoff on day j (i.e., $G(q + \delta)$) is earlier than $G(q) - K\delta$. Then, the total quantity served on day j is strictly less than the total quantity served on day i . By the structure of the greedy policy and the fact that $q + \delta > q$ orders are available at $t = 0$ on day j , this implies that the final vehicle returns before $t = T$ on day j . This is a contradiction, so we must have that $|G(q + \delta) - G(q)| \leq K\delta$. Therefore, G is Lipschitz continuous. \square

LEMMA 4. *The function $H(q)$ is increasing and continuous for all $q \geq 0$.*

Proof. By definition,

$$H(q) = \mu + \int_0^{G(q)} ((1 - w_A(p_0))\lambda_A(t))dt + \int_{G(q)}^T \lambda_A(t)dt. \quad (\text{A6})$$

Because $G(q)$ is decreasing in q and $1 - w_A(p_0) < 1$, this expression implies $H(q)$ is increasing in q . Additionally, since $G(q)$ is continuous, this expression also implies that $H(q)$ is continuous. \square

We also make use of the following known result.

LEMMA 5. *If a function g is continuous at a point x , then for any sequence (x_0, x_1, x_2, \dots) with $\lim_{n \rightarrow \infty} x_n = x$, it must also hold that $\lim_{n \rightarrow \infty} g(x_n) = g(x)$.*

A.8.2 Proof of Theorem 3

Proof. We first aim to show that, for any non-negative $q < \eta$, $\lim_{j \rightarrow \infty} H^{(j)}(q) = \eta$, and the convergence is monotonically increasing.

Because H is an increasing function, $H(0) < H(\eta) = \eta$. Also, $H(0) \geq \mu > 0$; in other words, $H^{(1)}(0) > H^{(0)}(0)$. The increasing nature of H implies $H(H^{(1)}(0)) > H(H^{(0)}(0))$; equivalently, $H^{(2)}(0) > H^{(1)}(0)$. Repeating this process indefinitely implies

$$H^{(0)}(0) < H^{(1)}(0) < H^{(2)}(0) < H^{(3)}(0) < \dots \quad (\text{A7})$$

Additionally, induction implies that, for every $j \in \mathbb{N}$, $H^{(j)}(0) < H(\eta) = \eta$. By the Monotone Convergence Theorem, the sequence $(H^{(0)}(0), H^{(1)}(0), H^{(2)}(0), H^{(3)}(0), \dots)$ must therefore converge to some $\hat{q} \leq \eta$.

Suppose for the purposes of contradiction that $\hat{q} < \eta$. Given this assumption, further suppose for the purposes of contradiction that $H(\hat{q}) = \hat{q}$; then, there is a saturated system steady state associated with a daily time-zero accumulation of quantity \hat{q} and a daily SDD order deadline of $G(\hat{q}) > G(\eta) = C$. Note, however, that this $G(\hat{q})$ -cutoff saturated system is more difficult and serves a greater total daily quantity than the

C -cutoff system. This implies that the greedy policy used to serve the $G(\hat{q})$ -cutoff system is also feasible to serve the C -cutoff system. This is a contradiction, since C was chosen so that the system is saturated. Thus, if $\hat{q} < \eta$, then $H(\hat{q}) \neq \hat{q}$.

Furthermore,

$$(H^{(0)}(0), H^{(1)}(0), H^{(2)}(0), H^{(3)}(0), \dots) \rightarrow \hat{q} \quad (\text{A8})$$

implies

$$(H(H^{(0)}(0)), H(H^{(1)}(0)), H(H^{(2)}(0)), H(H^{(3)}(0)), \dots) \rightarrow H(\hat{q}) \quad (\text{A9})$$

by Lemma 5. However, the latter sequence is simply $(H^{(1)}(0), H^{(2)}(0), H^{(3)}(0), H^{(4)}(0), \dots)$, and hence it converges to \hat{q} . This implies $H(\hat{q}) = \hat{q}$, a further contradiction. Therefore, $\hat{q} = \eta$; equivalently,

$$(H^{(0)}(0), H^{(1)}(0), H^{(2)}(0), H^{(3)}(0), \dots) \rightarrow \eta \quad (\text{A10})$$

in a monotonically increasing fashion. By using the fact that H is increasing, we can generalize this result to say

$$(H^{(0)}(q), H^{(1)}(q), H^{(2)}(q), H^{(3)}(q), \dots) \rightarrow \eta \quad (\text{A11})$$

in a monotonically increasing fashion for any $q \in [0, \eta]$.

Let us denote $q^{\max} = mQ(0)$ for clarity. We next aim to show that, for any $q \in (\eta, q^{\max}]$, $\lim_{j \rightarrow \infty} H^{(j)}(q) = \eta$, and the convergence is monotonically decreasing. Because H is an increasing function, $H(q^{\max}) > H(\eta) = \eta$. Also,

$$H(q^{\max}) \leq \mu + \int_0^T ((1 - w_A(p_0))\lambda_A(t)dt + \int_0^T \lambda_A(t)dt < q^{\max}; \quad (\text{A12})$$

in other words, $H^{(1)}(q^{\max}) < H^{(0)}(q^{\max})$. The increasing nature of H implies $H(H^{(1)}(q^{\max})) < H(H^{(0)}(q^{\max}))$; equivalently, $H^{(2)}(q^{\max}) < H^{(1)}(q^{\max})$. Repeating this process indefinitely implies

$$H^{(0)}(q^{\max}) > H^{(1)}(q^{\max}) > H^{(2)}(q^{\max}) > H^{(3)}(q^{\max}) > \dots \quad (\text{A13})$$

Additionally, induction implies that, for every $j \in \mathbb{N}$, $H^{(j)}(q^{\max}) > H(\eta) = \eta$. By the Monotone Convergence Theorem, the sequence $(H^{(0)}(0), H^{(1)}(q^{\max}), H^{(2)}(q^{\max}), H^{(3)}(q^{\max}), \dots)$ must therefore converge to some $\hat{q} \geq \eta$. An analogous argument as above, omitted to avoid redundancy, implies that $\hat{q} = \eta$. Therefore,

$$(H^{(0)}(q^{\max}), H^{(1)}(q^{\max}), H^{(2)}(q^{\max}), H^{(3)}(q^{\max}), \dots) \rightarrow \eta \quad (\text{A14})$$

in a monotonically decreasing fashion. By using the fact that H is increasing, we can generalize this result to say

$$(H^{(0)}(q), H^{(1)}(q), H^{(2)}(q), H^{(3)}(q), \dots) \rightarrow \eta \quad (\text{A15})$$

in a monotonically decreasing fashion for any $q \in (\eta, q^{\max}]$. Therefore, $\lim_{j \rightarrow \infty} H^{(j)}(q) = \eta$ for all $q \in [0, q^{\max}]$ and the convergence is monotonic. Because G is continuously decreasing and $G(\eta) = C$, the relationship between G and H implies that $\lim_{j \rightarrow \infty} G^{(j)}(q) = C$ for all $q \in [0, q^{\max}]$ and the convergence is monotonic. \square

A.9 Proof of Theorem 4

Proof. Denote the associated greedy policy as $P = ((t_1, q_1), \dots, (t_m, q_m))$ and denote the total quantity of orders available at $t = 0$ as η . First, suppose that $t_m + f(q_m) < T$. By definition, $t_m = C$. Then, by the continuity of f , there must exist some small quantity $\delta \in (0, q_1)$ and duration $\varepsilon \in (0, C)$ such that $t_m + f(q_m + \delta) \leq T$ and $\delta = \int_{C-\varepsilon}^C ((1 - w_A(p_0))\lambda_A(t)dt)$. Therefore, we can convert all δ of the NDD orders placed during the time interval $[C - \varepsilon, C]$ as NDD⁺ orders to be feasibly served by the final vehicle in addition to q_m . Because this conversion does not increase the total quantity of orders to be served across all dispatches, the new system with some NDD⁺ fulfillment can be feasibly served by the policy $((t_1, q_1 - \delta), \dots, (t_m, q_m + \delta))$. Thus, condition (i) of Theorem 4 is sufficient.

Next, suppose that $\eta/Q(0) \notin \mathbb{Z}$ (equivalently, $\eta \bmod Q(0) > 0$). Let k denote the index of the latest dispatch that serves some orders from the time-zero accumulation; the quantity of orders from the time-zero accumulation that dispatch k serves is $\eta \bmod Q(0)$. Define

$$\delta = \min \left\{ \eta \bmod Q(0), \int_0^{t_k} ((1 - w_A(p_0))\lambda_A(t)dt) \right\} \quad (\text{A16})$$

and $\varepsilon > 0$ satisfying $\int_0^\varepsilon ((1 - w_A(p_0))\lambda_A(t)dt) = \delta$. If we convert all δ of the NDD orders placed during the time interval $[0, \varepsilon]$ to NDD⁺ orders to be served by the k -th vehicle (and make no other changes to the system), then the unchanged policy P remains feasible. Thus, condition (ii) of Theorem 4 is also sufficient.

Finally, suppose for the purposes of contradiction that neither condition holds; that is, $t_m + f(q_m) = T$ and $\eta/Q(0) \in \mathbb{Z}$. If we were to convert $q^+ > 0$ NDD orders to NDD⁺ orders, then the resulting system is at least as difficult as the original system (let us denote this as $J \succeq I$). However, because neither condition holds, at least one of the first k vehicles (where k is defined as above) departs after $t = 0$ in the greedy policy used to serve J . Because $t_1 = t_2 = \dots = t_k = 0$, the greedy policies used to serve I and J are different. Therefore, the greedy policy used to serve J cannot be feasible, as otherwise it would also be feasible for the original system I , a contradiction. Thus, at least one of the two conditions is necessary for $q^+ > 0$ to hold. \square

A.10 Proof of Proposition 1

Proof. Let $P^0 = ((t_1^0, q_1^0), \dots, (t_m^0, q_m^0))$ denote the greedy policy applied to the system with no NDD⁺ orders (that is, $\hat{C} = 0$). Let dispatch k denote the last dispatch that serves any orders from the time-zero accumulation, i.e., $k = \max_{j \in [m]} \{j \mid \sum_{i=1}^{j-1} q_i^0 < \eta\}$. As discussed previously, we know that the greedy policy remains unchanged if we increase \hat{C} until a value $\bar{C} \in [0, T]$. Thus, $F(\hat{C}) = F(0)$ for all $\hat{C} \in [0, \bar{C}]$.

Next, we aim to show that $F(\cdot)$ is increasing if we increase \hat{C} beyond \bar{C} . Let $\hat{C} \in [\bar{C}, C)$ and $\varepsilon \in (0, C - \hat{C})$. Since $\hat{C} \geq \bar{C}$, we know that $k < m$. Let $P^1 = ((t_1^1, q_1^1), \dots, (t_m^1, q_m^1))$ denote the greedy policy applied to the system with a NDD⁺ deadline of \hat{C} . By definition, the subsystem J^1 (relative to the system with a NDD⁺ deadline of \hat{C}) induced by the first $m - 1$ dispatches of P^1 is $(m - 1)$ -saturated. Similarly, let $P^2 = ((t_1^2, q_1^2), \dots, (t_m^2, q_m^2))$ denote the greedy policy applied to the system with a NDD⁺ deadline of

$\hat{C} + \varepsilon$. Note that $t_k^1 < t_k^2$ because (i) the time-zero accumulation shrinks as the NDD⁺ deadline increases and (ii) $\hat{C} \geq \bar{C}$.

For the purposes of contradiction, suppose that $F(\bar{C}) \geq F(\hat{C})$. Then, $q_m^1 \geq q_m^2$, implying that $\sum_{i=1}^{m-1} q_i^1 \leq \sum_{i=1}^{m-1} q_i^2$. Let J^2 denote the subsystem (relative to the system with a NDD⁺ deadline of $\hat{C} + \varepsilon$) induced by the quantity $\sum_{i=1}^{m-1} q_i^2$. By definition, $J^2 \succeq J^1$ and J^2 is also $(m-1)$ -saturated. This implies that the first $m-1$ dispatches of P^2 can feasibly serve both J^2 and J^1 . However, since $t_k^1 < t_k^2$, the first $m-1$ dispatches of P^2 constitute a different policy than the first $m-1$ dispatches of P^1 . Therefore, the $(m-1)$ -dispatch greedy policy is not the only $(m-1)$ -dispatch policy that can serve J^2 . This is a contradiction, since J^2 is saturated. Therefore, $F(\bar{C}) < F(\hat{C})$. It follows that if $F(\bar{C}) = F(0) < T$ and $F(C) > T$, then $F(\hat{C}) = T$ has a unique solution in the interval $(0, C)$. \square

A.11 Proof of Proposition 2

Throughout this section, assume that λ_A and λ_B represent fixed scalar quantities and not functions.

Proof. Let $(C_{\min}, C_{\max}]$ denote the interval for which $C \in (C_{\min}, C_{\max}]$ implies that, in the associated feasible greedy policy, k vehicles serve a positive quantity of NDD orders and ℓ vehicles serve a positive quantity of orders. Without loss of generality, assume that $k = 1$ and $\ell = m$. Because the total revenue is a linear function of C , it suffices to show that the greedy policy's total dispatch duration as a function of C is concave over the interval $(C_{\min}, C_{\max}]$.

Let $\lambda_0 = w_A(p_0)\lambda_A + w_B(p_0)\lambda_B$ for notational convenience. Define the function $R(y)$ to take the value of t that satisfies $t + f(\lambda_0 t - y) = T$. That is, $R(y)$ is the time t at which a vehicle serving $\lambda_0 t - y$ quantity should dispatch in order to return exactly at time T . It is clear that $R(\cdot)$ is continuous and strictly increasing in y . Fix some $\hat{y} \geq 0$ and let $R(\hat{y}) = \hat{t}$. Let $\delta > 0$ be small, and let $\varepsilon > 0$ be such that $R(\hat{y} + \delta) = \hat{t} + \varepsilon$; that is, $(\hat{t} + \varepsilon) + f((\hat{t} + \varepsilon)\lambda_0 - (\hat{y} + \delta)) = T$. Because $f(\cdot)$ is strictly concave and increasing, we then have that $(\hat{t} + 2\varepsilon) + f((\hat{t} + 2\varepsilon)\lambda_0 - (\hat{y} + 2\delta)) < T$. This implies that $R(\cdot)$ is a midpoint-convex function of y and thus a convex function of y .

Returning to the feasible greedy policy, the first dispatch time t_1 solves

$$t_1 + f(\lambda_0 t_1 + \mu + \lambda_A T - \lambda_A w_A(p_0)C) = T. \quad (\text{A17})$$

By the strictly increasing and convex nature of $R(\cdot)$, we can conclude that t_1 is a strictly increasing and convex function of C . Similarly, the second dispatch time (assuming $m \geq 3$) t_2 solves $t_2 + f(\lambda_0 t_2 - \lambda_0 t_1) = T$. Therefore, t_2 is a strictly convex, increasing function of t_1 , further implying that t_2 is a strictly convex, increasing function of C . Continuing this argument shows that the i -th dispatch time t_i is a convex, increasing function of C for all $i < m$. By the structure of the greedy policy, the dispatch durations $f(q_i) = T - t_i$ are strictly decreasing, concave functions of C for all $i < m$.

The penultimate dispatch time t_{m-1} is a convex, increasing function of C , and $q_m = (C - t_{m-1})\lambda_0$. Thus, q_m is a concave function of C . It follows that the final dispatch duration $f(q_m)$ is also a concave function of C . Thus, the greedy policy's total dispatch duration $\sum_{i=1}^m f(q_i)$ is a concave function of C , as desired. \square

A.12 Proof of Theorem 5

Throughout this section, assume that λ_A and λ_B represent fixed scalar quantities and not functions.

LEMMA 6. *Let the early SDD price \hat{p}_0 , late SDD price p_0 , and number of vehicles m be fixed. With some abuse of notation, let the function $C(\hat{C})$ denote the time of the final SDD order deadline, given the discounted SDD order deadline \hat{C} , that induces saturation. Choose \hat{C}_{\min} and \hat{C}_{\max} satisfying $\hat{C}_{\min} < \hat{C}_{\max} < C$ such that for all $\hat{C} \in [\hat{C}_{\min}, \hat{C}_{\max}]$, the greedy policy applied to the associated saturated system (induced by appropriately varying C) has: (i) k dispatches serving a positive quantity of NDD orders and (ii) ℓ vehicles departing by \hat{C} . Then, the function $C(\hat{C})$ is convex over the interval $\hat{C} \in [\hat{C}_{\min}, \hat{C}_{\max}]$.*

Proof. The function $C(\hat{C})$ is decreasing and continuous. The former property follows from the discussion on partial orders (Appendix A.7), and the latter property is implied by the continuity of Q and the customer arrival rates. It follows that the function is bijective and invertible. Assume that $k = 1$ without loss of generality. We will show that $C(\hat{C})$ is midpoint convex for all $\hat{C} \in [\hat{C}_{\min}, \hat{C}_{\max}]$.

Fix some $\hat{C}^1 \in [\hat{C}_{\min}, \hat{C}_{\max}]$, and choose some $\hat{C}^3 \in (\hat{C}, \hat{C}_{\max})$. Define $\hat{C}^2 = (\hat{C}^1 + \hat{C}^3)/2$. Let $\tau = \hat{C}^2 - \hat{C}^1 = \hat{C}^3 - \hat{C}^2 > 0$. Next, define $C^1 = C(\hat{C}^1)$, $C^2 = C(\hat{C}^2)$, and $\varepsilon = C^1 - C^2$. Set $C^3 = C^2 - \varepsilon = C^1 - 2\varepsilon$. For each $i \in \{1, 2, 3\}$, denote the instance induced by \hat{C}^i and C^i by I^i . For each $i \in \{1, 2, 3\}$, let the policy $P^i = ((t_1^i, q_1^i), \dots, (t_m^i, q_m^i))$ denote the feasible greedy policy applied to I^i . We know that I^1 and I^2 are saturated by construction. For the purposes of contradiction, let us assume that $q_m^3 > Q(C^3)$.

Let η^1, η^2, η^3 denote the corresponding quantities of the time-zero accumulations. By construction, $\eta^3 - \eta^2 = \eta^2 - \eta^1$; let the value of this expression be $\delta \in \mathbb{R}$.

Because the final vehicle always departs at the final SDD order deadline, observe that $t_m^3 < t_m^2 < t_m^1$ and $t_m^2 - t_m^3 = t_m^1 - t_m^2 = \varepsilon$. Assume for now that $m > \ell + 1$. Because $Q(t)$ is decreasing and convex, $Q(t_m^3) > Q(t_m^2) > Q(t_m^1)$ with $Q(t_m^3) - Q(t_m^2) > Q(t_m^2) - Q(t_m^1)$. Therefore, because the customer arrival rates are constant and $m - 1 > \ell$, it follows that $t_{m-1}^3 < t_{m-1}^2 < t_{m-1}^1$ and $t_{m-1}^2 - t_{m-1}^3 > t_{m-1}^1 - t_{m-1}^2$. Repeating this argument demonstrates that this effect “propagates” sequentially to earlier dispatches, implying that $t_{\ell+1}^3 < t_{\ell+1}^2 < t_{\ell+1}^1$ and $t_{\ell+1}^2 - t_{\ell+1}^3 > t_{\ell+1}^1 - t_{\ell+1}^2$. More generally, we have that $t_{\ell+1}^3 < t_{\ell+1}^2 < t_{\ell+1}^1$ and $t_{\ell+1}^2 - t_{\ell+1}^3 \geq t_{\ell+1}^1 - t_{\ell+1}^2$ once we relax the assumption that $m > \ell + 1$. We proceed next by cases, each beginning at this point in the proof, labeled (†) for convenience. (†)

Case I: $m > 1$ and $\ell + 1 > k$. By definition, the discounted SDD order deadline falls between the ℓ -th and $(\ell + 1)$ -th vehicle departure times in all three of the policies being considered. However, recall that this discounted deadline is delayed by τ when moving from \hat{C}^1 to \hat{C}^2 , and then delayed again by τ when moving from \hat{C}^2 to \hat{C}^3 .

Because $Q(t)$ is decreasing and convex, $Q(t_{\ell+1}^3) > Q(t_{\ell+1}^2) > Q(t_{\ell+1}^1)$ with $Q(t_{\ell+1}^3) - Q(t_{\ell+1}^2) > Q(t_{\ell+1}^2) - Q(t_{\ell+1}^1)$. Therefore, $t_{\ell}^2 - t_{\ell}^3 > t_{\ell}^1 - t_{\ell}^2$. However, because the discounted deadline is varying, it is not guaranteed that $t_{\ell}^3 < t_{\ell}^2 < t_{\ell}^1$. Consider two sub-cases.

Case Ia: $\delta \geq 0$. In other words, this condition states that the time-zero accumulation quantity increases by the same amount (or remains unchanged) as the discounted deadline is delayed by τ twice. By the decreasing structure of $Q(t)$, this condition implies that $t_1^3 \leq t_1^2 \leq t_1^1$. Following the greedy policy, this inequality implies $t_\ell^3 \leq t_\ell^2 \leq t_\ell^1$. Similarly, by tracking the greedy policy in the reverse order, we see that $t_\ell^3 \leq t_\ell^2 \leq t_\ell^1$ implies $t_1^3 \leq t_1^2 \leq t_1^1$ and therefore $\delta \geq 0$. Thus, the condition “ $\delta \geq 0$ ” must be equivalent to the condition “ $t_\ell^3 \leq t_\ell^2 \leq t_\ell^1$ ”.

Beginning at dispatch ℓ and following the propagation argument used prior to (\dagger) , the relationships $t_\ell^3 \leq t_\ell^2 \leq t_\ell^1$ and $t_\ell^2 - t_\ell^3 > t_\ell^1 - t_\ell^2$ imply $t_1^2 - t_1^3 > t_1^1 - t_1^2$. However, the decreasing and convex nature of $Q(t)$ instead implies that $t_1^2 - t_1^3 \leq t_1^1 - t_1^2$, a contradiction. Thus, in Case Ia, we must have $q_m^3 \leq Q(C^3)$.

Case Ib: $\delta < 0$. In other words, this condition states that the time-zero accumulation quantity decreases by the same amount (or remains unchanged) as the discounted deadline is delayed by τ twice. By the decreasing structure of $Q(t)$, this condition implies that $t_1^3 > t_1^2 > t_1^1$. Following the greedy policy, this inequality implies $t_\ell^3 > t_\ell^2 > t_\ell^1$. Similarly, by tracking the greedy policy in the reverse order, we see that $t_\ell^3 > t_\ell^2 > t_\ell^1$ implies $t_1^3 > t_1^2 > t_1^1$ and therefore $\delta < 0$. Thus, the condition “ $\delta < 0$ ” must be equivalent to the condition “ $t_\ell^3 > t_\ell^2 > t_\ell^1$ ”.

Therefore, we can write $0 < t_\ell^3 - t_\ell^2 < t_\ell^2 - t_\ell^1$. Beginning at dispatch ℓ and following the propagation argument used prior to (\dagger) , it follows that $t_1^2 - t_1^1 > t_1^3 - t_1^2$. However, the decreasing and convex nature of $Q(t)$ instead implies that $t_1^1 - t_1^2 < t_1^2 - t_1^3$, a contradiction. Thus, in Case Ib, we must have $q_m^3 \leq Q(C^3)$.

Case II: $\ell + 1 = k$. Returning to (\dagger) , recall that $t_{\ell+1}^3 < t_{\ell+1}^2 < t_{\ell+1}^1$ and $t_{\ell+1}^2 - t_{\ell+1}^3 \geq t_{\ell+1}^1 - t_{\ell+1}^2$. Hence, it must hold that $Q(t_{\ell+1}^3) > Q(t_{\ell+1}^2) > Q(t_{\ell+1}^1)$ and $Q(t_{\ell+1}^3) - Q(t_{\ell+1}^2) > Q(t_{\ell+1}^2) - Q(t_{\ell+1}^1)$.

By direct calculation of order accumulations, we have

$$\begin{aligned} Q(t_{\ell+1}^3) - Q(t_{\ell+1}^2) &= \delta + \tau \times (w_A(\hat{p}_0) - w_A(p_0)) \times \lambda_A + \tau \times (w_B(\hat{p}_0) - w_B(p_0)) \times \lambda_B \\ &\quad - (t_{\ell+1}^2 - t_{\ell+1}^3) \times w_A(p_0) \times \lambda_A \end{aligned} \quad (\text{A18})$$

and

$$\begin{aligned} Q(t_{\ell+1}^2) - Q(t_{\ell+1}^1) &= \delta + \tau \times (w_A(\hat{p}_0) - w_A(p_0)) \times \lambda_A + \tau \times (w_B(\hat{p}_0) - w_B(p_0)) \times \lambda_B \\ &\quad - (t_{\ell+1}^1 - t_{\ell+1}^2) \times w_A(p_0) \times \lambda_A. \end{aligned} \quad (\text{A19})$$

Then,

$$[Q(t_{\ell+1}^3) - Q(t_{\ell+1}^2)] - [Q(t_{\ell+1}^2) - Q(t_{\ell+1}^1)] = \left[(t_{\ell+1}^3 - t_{\ell+1}^2) - (t_{\ell+1}^2 - t_{\ell+1}^1) \right] \times w_A(p_0) \times \lambda_A \leq 0. \quad (\text{A20})$$

This implies $Q(t_{\ell+1}^3) - Q(t_{\ell+1}^2) \leq Q(t_{\ell+1}^2) - Q(t_{\ell+1}^1)$, contradicting the previously derived inequality $Q(t_{\ell+1}^3) - Q(t_{\ell+1}^2) > Q(t_{\ell+1}^2) - Q(t_{\ell+1}^1)$. Thus, in Case II, we must also have $q_m^3 \leq Q(C^3)$.

These cases are exhaustive; therefore, we can conclude that $q_m^3 \leq Q(C^3)$. In other words, P^3 is feasible. Because extending the final SDD order deadline (without modifying the discounted deadline) can only make the system more difficult to serve, it follows that $C(\hat{C}^3) \geq C^3$. By generalization, the function $C(\hat{C})$ is midpoint convex and thus convex on the chosen interval. \square

LEMMA 7. *Suppose that the conditions and assumptions from the statement of Lemma 6 hold. Then, the total profit is a convex function of \hat{C} over the interval $\hat{C} \in [\hat{C}_{\min}, \hat{C}_{\max}]$.*

Proof. The total revenue as a function of \hat{C} can be written as

$$\left[\hat{p}_0 \times \hat{C} \times (\lambda_A \times w_A(\hat{p}_0) + \lambda_B \times w_B(\hat{p}_0)) \right] + \left[p_0 \times (C(\hat{C}) - \hat{C}) \times (\lambda_A \times w_A(p_0) + \lambda_B \times w_B(p_0)) \right].$$

Because the prices are fixed, the first bracketed expression is a linear function of \hat{C} , and the second bracketed expression is a convex function of \hat{C} over the interval (by Lemma 6). Thus, the total revenue is a convex function of \hat{C} over the interval. \square

Because dispatch times are continuous functions of \hat{C} , this result directly implies Theorem 5.

A.13 Computational Details for Pricing Optimization

In this section, we detail the efficient optimization of SDD deadlines given a fixed price combination $\{\hat{p}_0, p_0\}$. We discuss two cases in which the number of breakpoints is bounded by a polynomial function of the fleet size m : (i) when the total time-zero accumulation quantity η (overnight plus NDD) can be served by a single vehicle, or (ii) when all customers are of Type A. Throughout this section, assume that λ_A and λ_B represent positive scalars and not functions. For notational convenience, let the “early” SDD order accumulation rate (prior to \hat{C}) be denoted as $\lambda_E = w_A(\hat{p}_0) \times \lambda_A + w_B(\hat{p}_0) \times \lambda_B$, and let the “late” SDD order accumulation rate (between \hat{C} and C) be denoted as $\lambda_L = w_A(p_0) \times \lambda_A + w_B(p_0) \times \lambda_B$.

A.13.1 Case 1. We first assume that the order arrivals satisfy $\mu + \lambda_A \times (1 - w_A(p_0)) \times T < Q(0)$. That is, even if the higher SDD price is offered for the entire service day, the entire time-zero accumulation can be served by a single vehicle dispatching at $t = 0$ with some slack remaining. This condition implies that — regardless of the values of \hat{C} and C chosen (as long as saturation is maintained) — no vehicles depart at $t = 0$. Therefore, the optimal solution to the profit-maximization problem must correspond to a feasible solution wherein a vehicle departs exactly at \hat{C} . We show that there are at most m such feasible solutions.

PROPOSITION 3. *Assume the conditions stated above. For each $\ell \in \{1, 2, \dots, m\}$, there is at most one feasible solution in which the ℓ -th vehicle departs at \hat{C} .*

Proof. Fix $\ell \in \{1, 2, \dots, m\}$. As usual, we assume the usage of the greedy policy without loss of optimality. Suppose that there exists at least one discounted (early) deadline that induces a saturated system wherein the dispatch time of the ℓ -th vehicle is equal to the early deadline. Select one such discounted deadline arbitrarily; label it \hat{C}' . Denote the associated greedy policy as $P' = ((t'_1, q'_1), \dots, (t'_m, q'_m))$. The total quantity

of Type A customers served by the policy is trivially equal to the total daily number of Type A customers in the system $\lambda_A \times T$. Let $\varepsilon > 0$.

We will now perform the following adjustment. Consider increasing t'_ℓ and \hat{C}' by ε to t''_ℓ and \hat{C}'' while adjusting the dispatches' departure times and quantities to ensure that all dispatches still return at T . For example, $t''_{\ell-1} = t''_\ell - \frac{Q(t''_\ell)}{\lambda_E} > t'_\ell - \frac{Q(t'_\ell)}{\lambda_E} = t'_{\ell-1}$. Similarly, $t''_{\ell+1}$ is the unique time that satisfies $\lambda_L \times (t''_{\ell+1} - t''_\ell) = Q(t''_{\ell+1})$; thus, $t''_{\ell+1} > t'_{\ell+1}$ with $q''_{\ell+1} < q'_{\ell+1}$.

Now, observe that in this resulting policy, all departure times are shifted forward. Therefore, all dispatch quantities have been reduced. Because λ_A is constant over time, this implies that the quantity of Type A customers served by each dispatch is also reduced, so the total quantity of Type A customers served by the new policy is less than $\lambda_A \times T$. Similarly, performing the analogous adjustment for some $\varepsilon < 0$ would entail a policy in which the total quantity of Type A customers served is greater than $\lambda_A \times T$.

It is clear that the total quantity of Type A customers served is a continuous function of ε . It follows that, in some neighborhood around 0, the total quantity of Type A customers served is a decreasing function of ε . In a saturated solution, however, it must hold that the total quantity of Type A customers served is $\lambda_A \times T$. Because \hat{C}' was chosen arbitrarily among all feasible choices of discounted deadlines for which the ℓ -th vehicle departs at the discounted deadline in the associated greedy policy, it must then be true that \hat{C}' is the only such discounted deadline. \square

Fix $\ell \in \{1, 2, \dots, m\}$. This result and its proof imply a straightforward procedure for calculating the unique value of \hat{C}' associated with the solution wherein the ℓ -th vehicle dispatches exactly at \hat{C}' . We will define an auxiliary function $g_\ell(\hat{C})$ as follows.

- Define $t_\ell = \hat{C}$. Define $t_{\ell-1} = t_\ell - \frac{Q(t_\ell)}{\lambda_E}$, then $t_{\ell-2} = t_{\ell-1} - \frac{Q(t_{\ell-1})}{\lambda_E}$, and so on until t_1 ; assume that \hat{C} was chosen so that $t_1 > 0$. For all $d \in \{1, 2, \dots, \ell\}$, let $q_d = Q(t_d)$.
- Define $t_{\ell+1}$ as the unique time that satisfies $\lambda_L \times (t_{\ell+1} - t_\ell) = Q(t_{\ell+1})$. Then, define $t_{\ell+2}$ as the unique time that satisfies $\lambda_L \times (t_{\ell+2} - t_{\ell+1}) = Q(t_{\ell+2})$, and so on until t_m . For all $d \in \{\ell+1, \dots, m\}$, let $q_d = Q(t_d)$.
- Define the total quantity of Type A customers served by the first dispatch as $A_1 = q_1 - \mu - t_1 \times (\lambda_B \times w_B(\hat{p}_0))$. For all other $d \leq \ell$, define the total quantity of Type A customers served by dispatch d as $A_d = q_d - (t_d - t_{d-1}) \times (\lambda_B \times w_B(\hat{p}_0))$. Similarly, for all $d > \ell$, define the total quantity of Type A customers served by dispatch d as $A_d = q_d - (t_d - t_{d-1}) \times (\lambda_B \times w_B(p_0))$.

Let the total quantity of Type A customers served by this policy be $g_\ell(\hat{C}) = \sum_{d=1}^m A_d$. By our proof above, the unique solution to the equation $g_\ell(\hat{C}) = \lambda_A \times T$ corresponds to the desired value of \hat{C}' for which the ℓ -th vehicle dispatches at \hat{C}' and the system is saturated. We know that g_ℓ is continuous, $g_\ell(\hat{C}) > \lambda_A \times T$ for all $\hat{C} < \hat{C}'$, and $g_\ell(\hat{C}) < \lambda_A \times T$ for all $\hat{C} > \hat{C}'$. Thus, any general-purpose root finding method (e.g., bisection search) can solve the equation if a solution exists. This can be repeated for all $\ell \in \{1, 2, \dots, m\}$ to find all desired breakpoints. The total profit can be calculated for the solution corresponding to each breakpoint; the

breakpoint with the highest profit is chosen as the best solution for the given prices. Finally, this procedure can be repeated for every price combination to calculate the greatest profit across all price combinations.

A.13.2 Case 2. Suppose now that the number of vehicles required to serve the time-zero accumulation is arbitrary, but that all customers in the system are of Type A. In this case, we will show that there are at most m solutions for which η is an integer. We will do this by showing that η is a decreasing function of \hat{C} when C is adjusted to maintain saturation.

PROPOSITION 4. *Assume that all customers in the system are of Type A. With some abuse of notation, let $\eta(\hat{C})$ represent the quantity of the time-zero accumulation as a (continuous) function of \hat{C} , assuming that C is adjusted to maintain saturation. Then, $\eta(\hat{C})$ is a decreasing function.*

Proof. Recall that $C(\hat{C})$ is a decreasing function. Suppose instead for the purposes of contradiction that $\eta(\hat{C})$ is a non-decreasing function. Then, when \hat{C} is increased, every order is available no later than it was originally; because all customers are of Type A, the total daily quantity of orders does not change, so the new system is no more difficult to solve than it was previously. However, because $C(\hat{C})$ is a decreasing function, the final vehicle's departure time is earlier than it was previously. Thus, the system is no longer saturated, a contradiction. Therefore, η is a decreasing function of \hat{C} . \square

This result implies that, for each $k \in \{1, 2, \dots, m\}$, the equation $\eta(\hat{C}) = kQ(0)$ can be solved efficiently by any general-purpose root-finding method. The solution to this equation, when it exists, is therefore the unique discounted deadline associated with the dispatching policy in which k vehicles dispatching at $t = 0$ together serve all of the time-zero accumulation. Via this result, we can also identify the interval of values of \hat{C} for which k vehicles are required to serve the time-zero accumulation. Within this interval, we can then disregard the first $k - 1$ vehicles and apply the method from Case 1 to identify all of the values of \hat{C} associated with a vehicle dispatching exactly at \hat{C} . Hence, the total number of breakpoints for a fixed price combination is no more than m^2 , all of which can be efficiently computed via univariate root-finding.

Appendix B Additional Computational Data

B.1 Omitted Solution Details


Table B1 Feasible pricing solutions with highest revenue for $m \in \{1, 2, 3\}$

Fleet size m	Discounted SDD price \hat{p}_0	Discounted SDD order deadline \hat{C}	Base SDD price p_0	Final SDD order deadline C	Total Revenue
1	-	-	\$9	10:27 AM	\$129.44
	-	-	\$8.50	10:23 AM	\$129.44
	-	-	\$8	10:21 AM	\$127.99
	-	-	\$9.50	10:30 AM	\$127.81
	-	-	\$7.50	10:19 AM	\$125.33
2	\$8	10:23 AM	\$9	2:10 PM	\$470.50
	\$7.50	10:20 AM	\$9	2:09 PM	\$469.70
	\$8.50	10:25 AM	\$9	2:11 PM	\$469.70
	\$8	10:23 AM	\$8.50	2:01 PM	\$469.28
	-	-	\$8.50	2:02 PM	\$468.48
3	\$8	1:54 PM	\$9	3:42 PM	\$626.30
	\$8	1:54 PM	\$9.50	3:47 PM	\$625.36
	\$8	1:54 PM	\$8.50	3:37 PM	\$625.23
	\$7.50	10:22 AM	\$8.50	3:40 PM	\$624.08
	\$7.50	1:46 PM	\$9	3:38 PM	\$623.43

B.2 Omitted Sensitivity Analysis

This sensitivity analysis studies the effect of customers' threshold prices on the optimal pricing solutions. In relation to the original threshold prices given in Figure 13, we change *all* customers' threshold prices by a certain percentage, ranging from -15% to +15% by 0.5% increments. Figure B1 plots the results in the same manner as Figure 16. Of note in Figure B1b are the frequent changes in the optimal pricing combinations. From left to right, these combinations increase from $(\hat{p}_0, p_0) = (\$7, \$7.50)$ to $(\$9, \$10)$. However, the best solution with $(\hat{p}_0, p_0) = (\$8, \$9)$ achieves at least 98% of the optimal revenue in every case analyzed here.

Figure B1 Sensitivity analysis of threshold prices



(a) Delivery revenue vs. threshold prices

(b) SDD deadlines vs. threshold prices, $m = 2$

B.3 Impact of Depot Location

The depot location was located outside the zone in the original computational study. We now place the depot at an intersection near the center of the zone and briefly summarize the impact of this new depot location.

Assuming the same setup time of five minutes and per-order service time of three minutes, we re-estimate the dispatch duration function and re-solve the NDD⁺ maximization problem. With $m = 3$, $p_0 = 5$, and $C = 360$ (3 PM), the optimal early cutoff time is $\hat{C} = 324.44$ (2:24 PM). From the managerial perspective discussed previously, this value of \hat{C} is quite late, indicating that increasing C in conjunction with using the NDD⁺ strategy may be a better use of resources. For example, a late deadline of $C = 375$ (3:15 PM) entails a more practical NDD⁺ deadline of $\hat{C} = 162.69$ (11:43 AM). In this $(\hat{C}, C) = (162.69, 375)$ optimal NDD⁺ solution, the order quantities fulfilled by each of the three vehicles are 95.97, 53.21, and 28.12 orders, respectively. We validate this solution via simulation in the same manner described in Section 7. As before, 25 trials are each simulated over 84 days. Table B2 summarizes the results of these simulations.

	SDD order deadline (min. after 9 AM)	Total dispatch duration (min.)	SDD orders fulfilled	NDD ⁺ orders fulfilled
Simulated, mean (\pm st. dev.)	374.79 (\pm 15.88)	892.41 (\pm 44.39)	91.91 (\pm 6.59)	16.56 (\pm 4.07)
Predicted	375	897.74	92.63	16.57

The simulation results relative to the predicted values are very similar to those displayed in Table 2 in the paper, suggesting that the placement of the depot inside or outside the zone does not affect the validity of the tactical model in this operational setting. Of note is that the average simulated SDD deadline is 374.79 minutes after 9 AM. That is, on average, we stop accepting SDD orders just 13 seconds before our target deadline of 3:15 PM — a very slight improvement over the analogous statistic in Table 2.

B.4 Impact of Zone Size and Customer Density

As previously discussed, our model is intended to be used for tactical planning of small delivery fleets in small geographical areas. This section briefly summarizes the impact of increasing the size of the zone. First, we dilate the dimensions of the original zone by a factor of two, thereby increasing the zone's area by a factor of four from 12 square miles to 48 square miles. We reduce customers' geographical density by a factor of four to maintain the same total expected customer arrivals. As in the previous section, the depot is located centrally, the SDD price is $p_0 = 5$, the setup time is five minutes, and each delivery incurs a per-order service time of three minutes.

With the resulting increase in driving distances, the previous fleet size of $m = 3$ is infeasible for $C = 360$ (even when no NDD orders are upgraded to NDD⁺). If the fleet size is increased to $m = 4$, the final vehicle returns ten minutes before the end of the service day when no NDD orders are upgraded, indicating that

some NDD⁺ service is possible. Indeed, the optimal NDD⁺ solution is $(\hat{C}, C) = (185.27, 360)$, implying that all NDD orders placed by 12:05 PM should be upgraded for free. The order quantities fulfilled by each of the four vehicles are 75.62, 49.66, 30.36, and 18.79 orders, respectively, in this optimal NDD⁺ solution. For validation, 25 trials are again simulated over 84 days. Table B3 summarizes these results.

Table B3 Stochastic simulations vs. deterministic model predictions, centrally located depot and larger zone

	SDD order deadline (min. after 9 AM)	Total dispatch duration (min.)	SDD orders fulfilled	NDD ⁺ orders fulfilled
Simulated, mean (\pm st. dev.)	358.68 (\pm 17.59)	1306.74 (\pm 60.64)	87.78 (\pm 6.32)	19.04 (\pm 4.40)
Predicted	360	1314.99	88.90	19.04

In this experiment, the average simulated SDD order deadline is approximately 1.32 minutes before the predicted deadline of $C = 360$ (3 PM), a difference of only 0.37%. This gap, although very small relative to the length of the day, is slightly larger than the analogous gaps between predicted and simulated SDD deadlines in Tables 2 and B2. Continuous approximation methods are known to perform better as the number of stops on a route grows larger; the slightly larger gap in this setting is likely therefore caused by the lower quantity fulfilled by the final dispatch (18.79 orders).

As the area grows further, the zone may be too large to be efficiently served and may benefit instead by a partitioning. A large service zone will require a larger vehicle fleet where some vehicles may make very few deliveries. To illustrate with an extreme case, we dilate the dimensions of the original zone by a factor of four, resulting in a sixteen-fold increase of its area to 192 square miles. We reduce customers' geographical density by a factor of sixteen to maintain the same total expected quantity of customer arrivals. With the other parameters remaining the same, this new setting requires $m = 8$ vehicles for feasibility. The optimal NDD⁺ solution of $(\hat{C}, C) = (351, 360)$ specifies a final dispatch requiring three hours to deliver only 8.21 orders. Although such a solution seems likely to be impractical, we simulate it for completeness. Table B4 summarizes the results; note the 1.33% gap between the simulated and predicted SDD deadlines. Despite the vehicle routes having fewer stops, these results suggest that the model remains a useful approximation even in this extreme setting. However, we reiterate that our model is most useful for small delivery fleets, and large fleets serving large zones would benefit from smarter allocation of orders to vehicles based on geographic proximity.

Table B4 Stochastic simulations vs. deterministic model predictions, centrally located depot and much larger zone

	SDD order deadline (min. after 9 AM)	Total dispatch duration (min.)	SDD orders fulfilled	NDD ⁺ orders fulfilled
Simulated, mean (\pm st. dev.)	355.23 (\pm 17.99)	2559.91 (\pm 110.22)	86.92 (\pm 7.08)	35.55 (\pm 5.83)
Predicted	360	2577.90	88.90	36.11