TRANSPORTATION SCIENCE
Vol. 40, No. 3, August 2006, pp. 327–341
ISSN 0041-1655 | EISSN 1526-5447 | 06 | 4003 | 0327

DOI 10.1287/trsc.1050.0136

Incentive Schemes for Attended Home Delivery Services

Ann Melissa Campbell
Department of Management Sciences, Tippie College of Business, University of Iowa, Iowa City, Iowa 52242-1994, ann-campbell@uiowa.edu

Martin Savelsbergh
Department of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, Georgia 30332-0205, martin.savelsbergh@isye.gatech.edu

Many companies with consumer direct service models, especially grocery delivery services, have found that home delivery poses an enormous logistical challenge because of the unpredictability of demand coupled with strict delivery windows and low profit margin products. In this paper, we examine the use of incentives to influence consumer behavior to reduce delivery costs. We propose optimization models for two forms of incentives and demonstrate their value and impact through simulation studies. We conclude with a presentation of insights resulting from our efforts.

Key words: home delivery services; vehicle routing
History: Received: August 2003; revision received: June 2005; accepted: July 2005.

In recent years, many new and existing businesses have adopted a consumer direct (CD) service model that allows customers to purchase goods online and have them delivered directly to their front doors. Crossing this "last mile" provides a huge increase in service for customers but also creates a huge logistics challenge for companies, even with the advantages e-commerce provides. For example, we have seen the rise and subsequent fall of many e-grocers, including Webvan (Farmer and Sandoval 2001) and Shoplink, that have run out of money in the process of finding a distribution model that enables them to stay competitive with local grocery stores. E-grocers that have survived have changed distribution plans and focus, as demonstrated by the closing of the San Francisco-based operations for Peapod (Cox 2001); many continue to enter the arena, including Fresh Direct (Green 2003), with their own ideas on how to succeed. With annual revenue from all goods sold online predicted to be $195 billion by 2006 (Johnson, Delhagen, and Yuen 2003), CD is quickly becoming one of the most important business models, but there are still many open questions about how to run such businesses efficiently and effectively.

There are several issues in developing a successful direct delivery strategy. The fulfillment process for most CD businesses can be divided into three phases: (1) order capture and promise, (2) order sourcing and assembly, and (3) order delivery. Our research effort focuses on the interactions between order promise (deciding on a delivery time) and order delivery (devising efficient delivery schedules). Better integration of these decisions has the potential to substantially improve profitability, especially for those CD businesses offering "attended" deliveries. Attended deliveries are those where the consumers must be present; they may be necessary for security reasons (e.g., expensive computer equipment), because goods are perishable (e.g., milk or flowers), or because goods are being picked up or exchanged (e.g., dry cleaning, videos/DVDs) and are a vital feature of many CD service models. To provide a high service level and to avoid delivery failures, it is customary in attended home delivery services for the company and customer to mutually agree on a narrow delivery window or time slot.

In a previous paper (Campbell and Savelsbergh 2005) we studied and developed methodologies for order-acceptance decisions to maximize overall profit. The key idea underlying these methodologies is to exploit information about potential future orders to evaluate whether it is better to accept a customer's order or to reserve capacity for potential future orders. The techniques for making these decisions are based on modified insertion heuristics. As each order arrives, we compare the value of inserting that particular order versus inserting potential future orders that are properly discounted based on their probability of being realized. Computational results indicate that these order-acceptance strategies can significantly increase profits.

In this paper, we study and develop methodologies for a different aspect of the order-acceptance decision, namely the promise of a delivery window. In practice, it is often the case, (particularly in a struggling industry such as e-groceries where high customer retention is of utmost importance) that a vendor accepts an order unless it is impossible to satisfy the request. In such a scenario, we can potentially make significant improvements in routing costs by influencing customers' choices of delivery windows. If customers select better windows, not only will total distance be less, but a more efficient use of resources may increase the number of orders that can be accepted, thus creating higher revenues. We will look specifically at offering discounts, or *incentives*, to customers to influence window selection.

Home delivery is a fairly new phenomenon; thus few models and algorithms have been proposed and studied that help create an understanding of the complexities and intricacies of these distribution problems. Our goal is to continue to change this by looking at new features and variations of the problem. The main contribution of the work reported in this paper is that we develop incentive schemes and demonstrate that they can significantly increase the profitability of companies providing home delivery. The incentive computations involve the use of insertion heuristics as well as linear programming models. Our computational studies offer greater insight and better understanding of when incentives will be successful and what type of incentives will perform best. Another contribution of this paper, in our view, is that we formally propose two new optimization problems that capture the use of different types of incentives and allow the research community to focus on common problems.

The paper is organized as follows. In §1 we review the relevant literature, and in §2 we introduce the home delivery problem with time slot incentives (HDPTI). This new problem allows the vendor to influence delivery window selection through the use of incentives. We present a model for this problem, discuss important assumptions, propose methods for computing incentives, and present computational results that illustrate their success. In §3 we introduce a second version of the problem, the Home Delivery Problem with Wider Slot Incentives (HDPWI), where vendors offer incentives to customers to accept wider delivery windows. We discuss how to modify the earlier model and incentive computation and illustrate the impact of this change through computational results. We include a side note in §4 on the value of computing incentives for the earliest arriving orders. We conclude in §5 with a summary of the insights obtained and a short discussion of customer behavior modeling in §6.

1. Literature

Research on home delivery strategies is increasing, but most of the initial work has focused primarily on comparing the profits from very different service models rather than optimization in the design or performance of a single model. For example, Saranen and Småros (2001) simulate the delivery costs for two specific models—Streamline.com's unattended delivery policy and Webvan's attended half-hour delivery window policy—and find the more restrictive Webvan model to cost five times more. For unattended home deliveries, where time slots are not of concern, Punakivi (2000) studies the trade-off between the use of fixed routes and the use of optimally sequencing the deliveries on routes as soon as all deliveries are known. Depending on the density of the delivery area, simulation reveals an average savings from using optimal routing of 18%–54%. Yrjölä (2001) compares different strategies for picking the orders but also suggests values to use in evaluating the performance of an online grocer. Lin and Mahmassani (2002) summarize the delivery policies for many online grocers in the United States and use vehicle-routing software to evaluate the impact of some of these policies on a few realistic instances of the problem. Both unattended and attended policies are compared, along with different delivery window widths.

The problem closest to the one addressed in this paper does not involve groceries but scheduling repairmen to visit gas customers. In the problem considered by Madsen, Tosti, and Vælds (1995), requests for service that arrive during one week are scheduled to be serviced during the following week. The request must be scheduled when it arrives, so the challenge is to commit to a particular delivery time window that will lead to efficient routing solutions when all remaining requests for the week have arrived. The proposed solution approach involves the selection of seeds for different areas and the choice of where to insert requests based on insertion costs into routes containing the nearest seeds. Similar to the problem we study here, no probabilistic information about future requests is assumed. The problem is very different, though, because it is the vendor that selects the delivery window.

There are also several related, but also distinctly different, research areas related to the study of incentives for home delivery and good routing practices for the resulting problems. These include the study of revenue management, such as for airline ticket prices, and the study of vehicle-routing problems with stochastic demands and customers. A brief review of the literature in these areas is discussed in our earlier paper on home delivery (Campbell and Savelsbergh 2005). One recent paper not included in this survey is Bent and Van Hentenryck (2004). These authors exploit stochastic information about future requests to schedule requests under consideration, as done in Campbell and Savelsbergh (2005). The objective is to maximize the number of accepted requests, but the authors do not consider the option of rejecting an "expensive" delivery to preserve resources for more future deliveries, as in Campbell and Savelsbergh (2005). The proposed methodology involves maintaining multiple sets of tentative route plans, which can each be examined to see if a particular request can be handled feasibly by the vendor. We will use a related idea in this paper. Bent and Van Hentenryck (2004) test their algorithm on instances with varying degrees of dynamism (Larsen, Madsen, and Solomon 2002), that is, varying ratios of number of requests known in advance to the total number of requests.

2. The Home Delivery Problem with Time Slot Incentives

2.1. Problem Definition and Assumptions

We will now define the first dynamic routing and scheduling problem studied in this paper, which we refer to as the HDPTI. We have to construct a set of delivery routes for a specific day in the not-too-distant future. Requests from a known set of customers for a delivery on that particular day arrive in real time and are considered up to a certain cut-off time T, which precedes the actual execution of the planned delivery routes. We accept each request that arrives if there is available capacity. We assume that each request consumes d_i of the vehicle capacity and results in a revenue of r_i. There is a homogeneous set of m vehicles with capacity Q to serve the accepted orders. To increase the level of service, we guarantee that the actual delivery will take place during a one-hour time slot on the delivery day. These one-hour delivery time slots are nonoverlapping and cover the entire day, for example, 8:00–9:00, 9:00–10:00, …, 19:00–20:00. We assume that for each customer i, we know the probability p_i^t that he or she will choose a delivery in time slot t when an order is placed. When a request for service arrives, the vendor may offer incentives of up to B dollars per time slot. The probability of a customer choosing a particular time slot increases by an amount equal to the incentive offered multiplied by rate x. An increase in the probability of one or more time slots is compensated for by a decrease in the probability of the other time slots with p_i^t values greater than zero. The time slot selection by the customer is based on these modified probabilities. The objective is to maximize the total profit resulting from executing the final set of delivery routes, that is, total revenue minus incentive and delivery costs, where we assume that the delivery costs depend linearly on the travel time. The travel time between two locations i and j is denoted by t_{ij}. For simplicity, we will assume that the cost of one unit of travel time is \$1, but our models can be modified easily to accommodate other cost functions.

The definition above assumes knowledge about the likelihood that a customer selects a particular time slot for delivery (the probabilities p_i^t) and about the effect of incentives on a customer's behavior (the rate x). We believe both are reasonable assumptions, especially given the way businesses are evolving. A variety of industries, such as package delivery, are starting to use historical information about customers to estimate the likelihood of those customers requiring a particular service and are using this information for planning purposes. As technology and computing resources improve, the number of companies tracking and using such information about their customers and their ordering patterns will only increase. Thus, the ability to estimate and use p_i^t values seems a realistic assumption. If there is little or no historical information about a customer placing an order, companies may choose not to offer incentives to this customer, may assume equal probabilities for all time slots, or may use aggregate historical information to construct a prototypical time slot selection profile and use this profile for such customers. A variety of industries, such as online grocery shopping and delivery companies, is similarly starting to collect and use historical information about customers' reaction to incentives. Their experience indicates that even small incentives (a few dollars) can change customers' selection of delivery windows (Thomas Parkinson, personal communication).

We assume that incentives (discounts) will only be offered for time slots that have a positive probability of being selected. Because the probability information will be based on historical data, a probability of zero would reflect that a customer has never selected a delivery during that particular slot, so it is likely that this time slot is not feasible or is very undesirable for the customer. On the other hand, if two slots have positive and equal probability, it likely means that a customer is fairly indifferent between two time windows, so a small incentive might be able to influence the probability of choosing one over the other.

We have modeled consumer behavior by stating that the probability of choosing a particular time slot increases by an amount equal to the incentive offered multiplied by a rate x. Because the probability of all time slots must sum to 1, the probability of time slots with positive probability that do not receive an incentive must decrease, and we assume this decrease is in equal portions. That is, the total increase in probability created by incentives is divided up and removed equally from all the time slots with positive probability not receiving an incentive. Consequently, the maximum incentive payout for a set of time slots is not only limited by B, the maximum incentive per time slot, but also by the smallest probability among the time slots that do not receive an incentive (as probabilities cannot go below zero). Other options for decreasing the probabilities of the time slots not receiving incentives exist, such as decreasing these probabilities proportionally. The methodology discussed in this paper can easily be adapted to handle this and other options.

We also have to specify what happens when delivery within at least one of the time slots with positive probability is currently infeasible given all the orders that have already been accepted (and assigned a time slot). We have examined two options. First, if delivering to customer i in a time slot t with positive probability p_i^t of being selected is infeasible (which implies that this time slot is not presented to the customer as a delivery option), we will assume that this probability p_i^t will be redistributed equally among the feasible time slots with positive probability. Second, we will alternately assume that the customer will walk away with probability p_i^t. Note that we do assume that the probability of choosing to walk away, that is, choosing an infeasible time slot, can be reduced by offering incentives to other feasible time slots. In environments where demand is less than capacity (i.e., undersaturated markets), using incentives to prevent customers from walking away can be critical in maximizing profits.

Note that a major difference with the setup considered in our previous paper (Campbell and Savelsbergh 2005) is that we will not assume any stochastic information about future requests. This not only simplifies the incentive computations, but it also eliminates the need to generate accurate and reliable information about future requests. If incentives based only on the set of already accepted requests and the request under consideration prove to be successful in increasing profits, it follows that the results will only improve as more information about the future is included in the incentive computations.

Note also that all orders are received prior to the execution of any delivery schedule. That is, we will not consider "same-day delivery" services, where orders arrive during the execution of a delivery schedule and have to be incorporated immediately. We feel this is justified because in many consumer home delivery environments the vehicles will be loaded with customer-specific orders, so once a vehicle has started its route it cannot be rerouted to a new customer, because the correct inventory will likely not be on board.

For simplicity and ease of presentation, we assume for the remainder of this paper that the size of requests (d_i) is small compared to the vehicle capacity (Q) and that vehicle capacity is not a constraining factor when constructing delivery routes. This is generally true in practice because the constraints imposed by the combination of travel times and delivery windows are far more restrictive.

2.2. Determining Feasibility and Cost

As a first step toward designing a solution approach for the HDPTI, we develop a methodology to dynamically determine whether an order, which is characterized by a size and a delivery address, can be accommodated during a particular time window given the set of already accepted orders. We will use the same methodology to estimate what the order's contribution to profit will be if delivery is made during a given window. A key goal of this study is to determine if using profit estimates based only on currently accepted orders is sufficient to create effective incentives, so these estimates need to be as accurate as possible while quickly computed.

To dynamically determine whether we can accommodate an order in a particular time slot, we have to determine if there exists a set of m routes visiting all previously accepted orders as well as the order under consideration. The order under consideration must be visited during the particular time slot, and all previously accepted orders must be visited in their committed time slots for the resulting schedule to be feasible. It is well known that deciding whether a feasible solution to the vehicle-routing problem with time windows exists is NP-complete (Savelsbergh 1986), so it is natural to consider employing heuristics to answer the question of feasibility quickly.

We evaluate this question in two phases: First we create a set, S, of schedules for the already accepted orders, and second we evaluate the feasibility (and then cost) of inserting the order under consideration into these schedules. Each schedule in S is composed of m delivery routes for the previously accepted orders. By creating and maintaining a set of feasible schedules, rather than just one, we are able to increase the number of insertion options for each order that arrives and are more likely to find a feasible and low-cost insertion point. As mentioned in the literature review, a similar approach was taken in Bent and Van Hentenryck (2004). We use a combination of insertion heuristics and randomization to build the set S of schedules with the previously accepted orders.

Each time an order is accepted, we keep the least cost schedule that contains all previously accepted orders plus the newly accepted one (in the time slot selected by the customer). This schedule becomes the first one in the set S used to evaluate the feasibility and cost of the next order that arrives. Thus, we always have at least one feasible schedule for the already accepted orders in the set S. The other schedules in S are created one at a time. To build each one, we iteratively insert orders into the (partial) schedule until all orders are included or an infeasibility occurs. At each iteration of the construction, we evaluate the feasibility and the cost of inserting each of the as-yet-unscheduled orders at each point in the partially built schedule. Because all previously accepted orders have specific guaranteed time slots, an insertion of a previously accepted order is feasible only if it can occur during its guaranteed time slot. For each inserted order i in the schedule under construction, we maintain values e_i and l_i representing the earliest and latest times delivery can begin given the committed time slot and position on the partially built route. An order j can feasibly be inserted between already inserted orders $i-1$ and i in time slot t ($begin^t, end^t$) if and only if

$$e_j = \max(e_{i-1} + \text{service time} + t_{i-1,j}, begin^t), \quad (1)$$

$$l_j = \min(l_i - t_{j,i} - \text{service time}, end^t), \quad (2)$$

and

$$e_j \leq l_j. \quad (3)$$

Service time represents the time required to deliver the product once the vehicle has arrived. We assume the service time is the same for all customers. If the insertion is feasible, then we can evaluate the increase in cost (travel time) as

$$t_{i-1,j} + t_{j,i} - t_{i-1,i}. \quad (4)$$

We maintain a restricted candidate list of unscheduled orders with the smallest insertion costs at each iteration (and their corresponding insertion points in the current schedule) and randomly choose from among these to determine which order is inserted next in the partial schedule. This is a typical greedy randomized adaptive search procedure (Feo and Resende 1995). After each insertion, we update the e_i values for the inserted order as well as all orders following it and the l_i value for the inserted order and all orders preceding it. (See Campbell and Savelsbergh 2004 for a survey of efficient implementations of insertion heuristics.) We iterate until all orders are scheduled or one of the orders cannot feasibly be inserted. If the construction process completes with all orders scheduled, we add the schedule s to the set S and record its total cost $C(s)$. We try to generate a schedule a fixed number, say n, of times. Therefore, the set S will have cardinality less than or equal to $n+1$. Let $C(*) = \min_{s \in S} C(s)$. Due to the randomization, the routes making up the schedules may be quite different, but if the size of the restricted candidate list is small, say two or three, the constructed schedules usually have fairly similar costs.

We evaluate the cost of accepting the order currently under consideration, say j, in a particular time slot t by computing the cost of including it in each of the schedules in S. For each schedule $s \in S$, we try to insert j during time slot t at each possible insertion point. Feasibility is quickly evaluated using the same e_j and l_j calculations as described above. If feasible, we compute

$$\text{cost} = (t_{i-1,j} + t_{j,i} - t_{i-1,i}) + C(s) - C(*). \quad (5)$$

This value represents the "true" added cost associated with making this delivery to j. Without the $C(*)$ and $C(s)$ terms, the computation would only represent the increase in cost with respect to schedule s. This could disguise the fact that a cheap insertion can only occur if an otherwise much more expensive schedule is used. We maintain the lowest insertion cost for each time slot t and represent its final value, after evaluating all $|S|$ schedules, with C^t. (We also maintain the associated schedule s and point of insertion into s.) These C^t values form the basis for our incentive calculations, as they indicate the minimum cost of making a delivery to j in time slot t.

2.3. Modeling the Home Delivery Problem with Time Slot Incentives

In the previous section, we described how to determine quickly whether it is feasible to insert an order in a time slot t and how to compute an associated value C^t for that insertion. If we find that the C^t values vary widely for different time slots, then we may want to offer an incentive to the customer for choosing a time slot with a higher profit. Offering incentives raises many challenging questions, such as—

- How do we decide which time slot(s) receive an incentive?
- How do we decide on the size of the incentive(s)?

We can start to answer this second question by modeling the relationships described in the problem definition for the HDPTI. Recall that for each customer i, delivery in time slot t will be selected with probability p_i^t if no incentives are offered. Furthermore, the probability of choosing a particular time slot increases by an amount equal to the incentive offered multiplied by rate x, and the probability of all slots that do not receive an incentive will decrease by equal amounts.

Next, note that we cannot offer an incentive to all time slots with $p_i^t > 0$, because the increased probability resulting from the incentives must be discounted from other slots. Thus, to model this problem, we must divide the set of time slots with positive probability of being selected into two groups. Let

- O = set of time slots with $p_i^t > 0$

- $U =$ subset of O that may receive an incentive
- $V =$ subset of O not receiving an incentive.

We want to find

- $I^t =$ the incentive for time slot t
- $z =$ the reduction in probability for all time slots in V to maximize expected profitability.

Given the above and our basic assumption that insertion costs are a good reflection of future costs, the incentive decision for customer i can be represented by the following incentive optimization problem:

$$\max \sum_{t\in U}(r_i - C^t - I^t)(p_i^t + xI^t) + \sum_{t\in V}(r_i - C^t)(p_i^t - z) \quad (6)$$

subject to:

$$z \le p_i^t \quad \forall\, t \in V \quad (7)$$

$$\sum_{t\in U} xI^t = z|V| \quad (8)$$

$$0 \le I^t \le B \quad \forall\, t \in U. \quad (9)$$

In the objective, the first portion represents the product of the adjusted profit and adjusted probability associated with awarding an incentive I^t to time slot t in U. This product is the expected profitability from time slots where incentives are offered. Likewise, the second portion represents the expected profits from the slots with no incentives, with profits and probabilities adjusted accordingly. The first constraint in Equation (7) limits z such that the adjusted probability of each slot not receiving an incentive cannot fall below zero. The second constraint, Equation (8), sets z equal to the increase in probability created by incentives divided by the number of time slots in V, so the sum of all probabilities will remain equal to one. Finally, Equation (9) restricts each incentive to be less than the specified limit B.

We can use the above model to compute a set of incentives that maximizes expected profits given the partitioning of time slots into sets U and V. This still leaves the question, though, of how to decide which time slots should be assigned to sets U and V. This decision is not as straightforward as it may seem because of the interaction between the insertion costs and the probabilities. The following are some observations concerning the selection of set U.

Observation 1. If a single time slot is considered for an incentive ($|U| = 1$), it is possible that the optimal incentive is zero even if that time slot has the uniquely lowest insertion cost.

Example. Let $r_j = 20$, $C^1 = 10$, $C^2 = 12$, $C^3 = 16$, $p_i^1 = 0.5$, $p_i^2 = 0.3$, $p_i^3 = 0.2$, and $x = 0.1$. Consider time slot 1 for an incentive; that is, $U = \{1\}$ and $V = \{2, 3\}$. If $I^1 = 0$, then the objective function value is

$$(20-10)(0.5) + (20-12)(0.3) + (20-16)(0.2)$$
$$= 5 + 2.4 + 0.8 = 8.2.$$

If $I^1 = 2\epsilon$, then the objective function value decreases to

$$(20-10-2\epsilon)(0.5+0.2\epsilon) + (20-12)(0.3-0.1\epsilon)$$
$$+ (20-16)(0.2-0.1\epsilon) = 8.2 - 0.2\epsilon - 0.4\epsilon^2 < 8.2.$$

In fact, the optimal value for any single incentive I^t is zero whenever the following holds:

$$r_i - C^t - \frac{p_i^t}{x} - \frac{\sum_{t\in V}(r_i - C^t)}{|V|} \le 0. \quad (10)$$

This can be derived by manipulating the objective function given the observation that if there is a single time slot t that can receive an incentive, then $z = xI^t/|V|$. Equation (10) shows that if we consider offering an incentive for a single slot, the optimal value will be zero unless the profit from this time slot is greater than the average profit from the other slots by at least p_i^t/x. In the example, the average profit for time slots other than time slot 1 is 6 and $p_i^1/x = 0.50/0.10 = 5$. The profit from slot 1 is 10, which is less than $6 + 5 = 11$, so it is better to offer no incentive to time slot 1. The result does not preclude it being profitable to offer an incentive for a time slot with a higher insertion cost.

The above can be generalized to situations where we consider providing incentives for more than one time slot. With more than one time slot, substituting $z = (\sum_{t\in U} xI^t)/|V|$ into the objective function yields an expression that contains no terms involving products of incentive values. Thus, we can extend our result for each time slot under consideration.

Observation 2. The optimal value for any incentive I^t for $t \in U$ is zero whenever the following holds:

$$r_i - C^t - \frac{p_i^t}{x} - \frac{\sum_{t\in V}(r_i - C^t)}{|V|} \le 0. \quad (11)$$

Observation 3. If a single time slot is considered for an incentive ($|U| = 1$) and the optimal incentive is zero, it is possible that it will receive a positive incentive when considered in conjunction with another time slot.

Example. Consider the instance presented with Observation 1, but let $U = \{1, 2\}$ and $V = \{3\}$. If $I^1 = 0.5$ and $I^2 = 0.333$, then the expected profit is

$$(20-10-0.5)(0.5+0.1(0.5)) + (20-12-0.333)$$
$$\cdot (0.3+0.1(0.333)) + (20-16)(0.2-0.1(0.5+0.333))$$
$$= (9.5)(0.55) + (7.666)(0.3333) + (4)(0.1167)$$
$$= 5.225 + 2.555 + 0.467 = 8.247 > 8.2.$$

Observe that here an incentive for a (single) time slot decreases the probabilities associated with other (high-profit or low-cost) time slots, resulting in lower overall expected profits, whereas when the time slot is considered together with another time slot, an incentive may increase expected profits. In the above example, *both* time slots 1 and 2 would not receive an incentive if considered by themselves, but when considered together, it is beneficial for both to receive a positive incentive. This result is significant because it shows that even if a time slot does not receive an incentive when considered by itself, we cannot stop considering it when trying to find the best set U.

The above also demonstrates that it may be necessary to offer incentives for more than one time slot to increase expected profits. This can be a result of the quadratic nature of the expected profit function as well as the limit B.

These observations demonstrate that selecting a set U of time slots to consider for incentives so as to maximize the expected profit is nontrivial. The likely candidates for the set U are the time slots with the cheapest insertion costs because of the profits they offer, but customer behavior, represented through the selection probabilities, also impacts the choice.

In our computational experiments, we will compare choosing the l time slots with the cheapest insertion costs for the set U with an exhaustive search over all possible combinations of l time slots for the set U. The latter option may be impractical in terms of run time, especially for larger values of l, but will give a good measure of how well simply choosing the time slots with the cheapest insertion cost performs. We could, alternatively, incorporate the choice of whether or not a slot receives an incentive in the model itself (using binary selection variables), but this would lead to a mixed integer program that would be much more difficult and time consuming to solve.

2.4. Solving the Home Delivery Problem with Time Slot Incentives

Before we can evaluate the impact of incentives on the actual profits, we must figure out how to actually solve the proposed model quickly. After removing constant terms and with a little rewriting, the objective function in the model becomes

$$\max \sum_{t\in U}(x(r_i-C^t)-p_i^t)I^t-\sum_{t\in U}x(I^t)^2-\sum_{t\in V}(r_i-C^t)z. \quad (12)$$

This highlights how profitability is gained and where it is lost as a result of incentives. We observe that the objective function has a quadratic term for each time slot for which we offer an incentive. Because solving quadratic programs can be (too) time consuming for an online algorithm, we will use a linear approximation of the problem.

We approximate each quadratic term $(I^t)^2$ with a piecewise linear function over $f-1$ intervals between $I^t=0$ and $I^t=u$ where

$$u=\min\left(B, \frac{\min_{t\in V} p_i^t}{x}|V|\right). \quad (13)$$

The first term defining the upper bound u of I^t is the upper limit available for individual incentives, and the second term is the amount of incentive that can be spent on a single time slot before the time slot $t\in V$ with the smallest positive probability becomes zero. This requires f (additional) variables $y_1^t, \ldots, y_f^t$ per time slot. Normally, piecewise linear approximations require the introduction of integer variables. However, because $-(I^t)^2$ is nonincreasing and convex on the interval of interest and we are maximizing, the integer variables are not needed. Thus, we can solve the approximation as a linear program. (For a general discussion of how to approximate a continuous function of one variable with a piecewise linear function, see Nemhauser and Wolsey 1988.)

The resulting incentive optimization linear program is as follows:

$$\max \sum_{t\in U}(x(r_i-C^t)-p_i^t)(I^t)-\sum_{t\in U}x\left(\left(\frac{u}{f-1}\right)^2 y_2^t + \left(\frac{2u}{f-1}\right)^2 y_3^t+\cdots+(u)^2 y_f^t\right)-\sum_{t\in V}(r_i-C^t)z \quad (14)$$

subject to:

$$\sum_{i=1}^{f} y_i^t = 1 \quad \forall\, t\in U \quad (15)$$

$$I^t=\frac{u}{f-1}y_2^t+\frac{2u}{f-1}y_3^t+\cdots+uy_f^t \quad (16)$$

$$z\le p_i^t \quad \forall\, t\in V \quad (17)$$

$$\sum_{t\in U} xI^t = z|V| \quad (18)$$

$$0\le I^t\le B \quad \forall\, t\in U. \quad (19)$$

Equations (17) and (18) are the same as before, but Equations (15) and (16) are added because of the linearization. Both I^t and $(I^t)^2$ terms are now based on the y values.

2.5. Computational Experiments

Our primary goal in this section is to conduct computational experiments to determine the impact on total profit of using incentives to influence customer behavior. Furthermore, we want to study and compare methods for choosing the value of these incentives. Finally, we want to analyze the impact of instance characteristics on the performance of our proposed methodology.

In our testing, we want to compare the total profit, that is, total revenue − total costs − total incentives paid, associated with using incentive schemes of different forms. The following is a list of the five methods for which results are included in the tables. These five were selected to illustrate the impact of certain characteristics of the problem and solution methodology. Each of the five methods, except for BSTLP, completes within a couple of seconds.

• *NOINC* (no incentives): To evaluate the impact of incentives, we determine the profits when no incentives are provided. This means that each customer's time slot selection is based on initial p_i^t values. (Appropriate adjustments are made if one or more of these time slots are infeasible. Recall that if delivering to customer i in a time slot t with positive probability p_i^t of being selected is infeasible, we will redistribute this probability equally among the remaining feasible time slots with positive probability.)

• *CHPFLT* (flat incentive to cheapest slots): To evaluate the impact of sophisticated incentive computations (e.g., our linear programming-based incentive optimization techniques), we determine the profits when a simple incentive scheme is used. A simple and straightforward incentive scheme awards equal incentives to the l time slots with cheapest insertion costs. (No incentives are awarded if all time slots have equal insertion costs, because there is clearly no advantage to offering incentives in that situation.)

• *CHPLP* (LP-based incentive to cheapest slots): We want to evaluate the impact of different rules for choosing the sets U and V used in the incentive optimization problem. A natural rule is to choose the l time slots with the cheapest insertion costs to be in the set U. Note that this does not mean that all l time slots will necessarily receive an incentive, but only these time slots *may* receive an incentive. In fact, when the optimal incentive for a time slot in U is zero, it is moved to the set V, and the incentive optimization problem is resolved for the updated sets U and V. This process repeats until all optimal incentives are positive. (Otherwise, we have a situation in which we reduce the probabilities of the time slots in V, but, at the same time, have time slots in U that do not receive an incentive, but their probability is not reduced.)

• *BSTLP* (LP-based incentive to best set U): Another rule, though computationally intensive, is to enumerate all possible sets U of size 1 up to l, solve the incentive optimization problem for each set U, and select the one that maximizes expected profits. This allows us to evaluate how much we give up by greedily choosing the l cheapest time slots to comprise the set U.

• *BST* (best case): To serve as an "upper bound," we also evaluate what happens when customers always select the time slot preferred by the vendor (ideal customer behavior). In this setup, the time slot with the cheapest insertion cost is always selected and no incentives are paid. (Our computational experiments will show that this is not a true upper bound, as the associated profit can sometimes be exceeded by other approaches. This does not represent erroneous behavior, but is a reflection of the true dynamic nature of the problem.)

Such analysis can only be performed by means of simulation. Simulation is used to generate a stream of delivery requests at different points in time, uniformly distributed between zero and the cut-off time T. Given this stream of arrivals of delivery requests, we can evaluate the behavior of the different methods listed above by using them to decide incentives and then simulate consumer response to these offers.

In all experiments, unless specifically stated otherwise, 30 requests are generated for customers uniformly distributed over an area of dimension 60 units by 60 units, where each vehicle can travel one unit per minute. The service time to complete a delivery to a customer is 20 minutes. There are 12 time slots of 60 minutes available for making deliveries. For each customer, eight consecutive time slots will have a positive probability (possibly wrapping around at the end of the day, e.g., the first two time slots at the beginning of the day and the last six time slots at the end of the day). This is to reflect that few customers are willing or able to accept delivery during all possible time slots. The probabilities for the acceptable time slots follow one of three patterns, differing in terms of how much one time slot is preferred to the others. In the first pattern ($PROBPAT = 1$), all time slots (the ones with positive probability) are equally likely to be selected by the customer. In the second pattern ($PROBPAT = 2$), there is one time slot that is two times more likely to be selected than the others. In the third pattern ($PROBPAT = 3$), there is one slot that is three times more likely to be selected than the others. The slot that is "preferred" is randomly selected. Other patterns are possible, of course, but these provide a good starting point for an initial analysis. The probability of selecting a time slot increases at a rate of 0.2 per dollar of incentive. The (implied) limit on each incentive is \$5. Revenue for each delivery is \$100. For each computational experiment, 25 instances are generated, and the results are averaged to create the associated table.

During schedule construction, we try to generate 50 schedules; that is, $|S| \leq 51$, using a restricted candidate list with three orders. The quadratic term in the incentive optimization problem is approximated with five linear pieces.

We start by presenting the results for the 25 instances of the basic data set described above. The five methods are compared not only for different probability patterns, but also for different sizes of the set U. Besides the difference in profit associated with different sizes of U, there may be other reasons why a certain number of members of U is preferable. For example, experience may indicate that it is less

Table 1 Base Results

PROBPAT	$\lvert U\rvert$	*NOINC*	*CHPFLT*	*CHPLP*	*BSTLP*	*BST*
1	1	1,224.72	1,373.20	1,354.78	1,362.24	1,349.29
	2	1,224.72	1,304.28	1,405.99	1,379.40	1,349.29
	3	1,224.72	1,300.77	1,436.41	1,416.87	1,349.29
	4	1,224.72	1,219.82	1,370.68	1,446.67	1,349.29
Average		1,224.72	1,299.52	1,391.97	1,401.30	1,349.29
% improvement		0.00	6.11	13.66	14.42	10.17
2	1	1,235.63	1,375.60	1,380.03	1,311.60	1,349.29
	2	1,235.63	1,365.10	1,391.44	1,398.97	1,349.29
	3	1,235.63	1,261.17	1,389.83	1,375.15	1,349.29
	4	1,235.63	1,247.64	1,377.88	1,437.28	1,349.29
Average		1,235.63	1,312.38	1,384.80	1,380.75	1,349.29
% improvement		0.00	6.21	12.07	11.74	9.20
3	1	1,228.93	1,311.89	1,370.46	1,333.48	1,349.29
	2	1,228.93	1,336.17	1,372.35	1,407.83	1,349.29
	3	1,228.93	1,361.99	1,378.48	1,400.10	1,349.29
	4	1,228.93	1,294.06	1,346.86	1,402.10	1,349.29
Average		1,228.93	1,326.03	1,367.04	1,385.88	1,349.29
% improvement		0.00	7.90	11.24	12.77	9.79

confusing for a customer to be offered one incentive than to negotiate many varied incentives. Thus, we are interested in evaluating how many incentives are needed to obtain the majority of benefits that incentives can offer or finding out if it is true that more is "always better." In each experiment, we report the average profits from each method over all choices of U (*average*) and compare this to the results when no incentives are used (% *improvement*). These serve as a quick summary of the results.

The results presented in Table 1 clearly demonstrate the value of incentives. With a simple incentive scheme (*CHPFLT*), profits increase by 6%–8%, and with more sophisticated incentive schemes (*CHPLP* and *BSTLP*), profits increase by 12%–15%. It is interesting to observe that the sophisticated incentive schemes consistently outperform the "upper bound." Recall that when computing the upper bound, we assume that the customer always does exactly what we want him or her to do without having to pay incentives. The fact that incentive based approaches are able to achieve higher profits reflects the dynamic nature of the problem. Optimal decisions based on the orders that have arrived so far may not be optimal after additional orders arrive.

Next, we will focus on the results for probability pattern 1, that is, each customer has no preference among the delivery windows that are acceptable. We see that when the number of time slots eligible for an incentive increases, the performance of the simple incentive scheme (*CHPFLT*) degrades (from 1,373.20 to 1,219.82), as it has no mechanism to differentiate between different time slots. In fact, for $|U| = 4$ the profit is less than when no incentives are offered. This indicates that the increased profits are insufficient to recoup the incentive payouts. On the other hand, we see that the performance of the linear programming-based incentive schemes improves (from 1,362.24 to 1,446.67 for *BSTLP*). Although the same behavior can be observed with other probability patterns, it is not as pronounced. This is understandable because the stronger the preference is for a particular delivery window, the harder it will be to influence the customer's behavior with incentives. This shows that a thorough analysis of customer behavior is important when considering the use of incentive schemes.

In the above instances, even though 30 orders arrive, only approximately 14 can be served. Next, we investigated whether there is a noticeable change when only the first 15 orders are considered. With fewer customers, each customer can significantly impact cost and revenue.

The results presented in Table 2 show that incentives are still able to increase profits, but that the increases are smaller (about 4% with the simple incentive scheme (*CHPFLT*) and 6%–8% with more sophisticated incentive schemes (*CHPLP* and *BSTLP*)). The most striking difference, when compared with the earlier results, is that the upper bound now dominates the other approaches. This indicates that some of the differences we observed in the earlier results were caused by being able to accept a single order late in the simulation when the schedule was nearly full (timewise).

For the remaining experiments in this section, we revert back to 30 orders arriving but consider only two time slots for incentives ($|U| = 2$). The base results indicated that considering two time slots for incentives leads to substantial increases in profits, and

Table 2 Base Results with Fewer Orders

PROBPAT	$\lvert U\rvert$	*NOINC*	*CHPFLT*	*CHPLP*	*BSTLP*	*BST*
1	1	1,022.69	1,095.78	1,100.65	1,085.79	1,129.35
	2	1,022.69	1,075.67	1,104.85	1,090.83	1,129.35
	3	1,022.69	1,050.81	1,120.48	1,116.88	1,129.35
	4	1,022.69	1,031.35	1,090.85	1,133.04	1,129.35
Average		1,022.69	1,063.40	1,104.21	1,106.64	1,129.35
% improvement		0.00	3.98	7.97	8.21	10.43
2	1	1,034.43	1,094.82	1,089.96	1,063.02	1,129.35
	2	1,034.43	1,096.18	1,099.15	1,115.78	1,129.35
	3	1,034.43	1,074.24	1,127.31	1,101.43	1,129.35
	4	1,034.43	1,039.00	1,083.81	1,119.70	1,129.35
Average		1,034.43	1,076.06	1,100.06	1,099.98	1,129.35
% improvement		0.00	4.02	6.34	6.34	9.18
3	1	1,033.82	1,075.48	1,096.90	1,095.26	1,129.35
	2	1,033.82	1,076.72	1,084.56	1,117.06	1,129.35
	3	1,033.82	1,097.31	1,112.58	1,116.89	1,129.35
	4	1,033.82	1,059.10	1,086.46	1,126.63	1,129.35
Average		1,033.82	1,077.15	1,095.13	1,113.96	1,129.35
% improvement		0.00	4.19	5.93	7.75	9.24

Table 3 Base Results for Varying Incentive Upper Limit

B	*PROBPAT*	*NOINC*	*CHPFLT*	*CHPLP*	*BSTLP*	*BST*
2	1	1,224.72	1,304.28	1,391.58	1,414.88	1,349.29
	2	1,235.63	1,365.10	1,330.80	1,384.19	1,349.29
	3	1,228.93	1,336.17	1,366.02	1,342.97	1,349.29
Average		1,229.76	1,335.18	1,362.80	1,380.68	1,349.29
% improvement		0.00	8.57	10.82	12.27	9.72
5	1	1,224.72	1,304.28	1,405.99	1,379.40	1,349.29
	2	1,235.63	1,365.10	1,391.44	1,398.97	1,349.29
	3	1,228.93	1,336.17	1,372.35	1,407.83	1,349.29
Average		1,229.76	1,335.18	1,389.93	1,395.40	1,349.29
% improvement		0.00	8.57	13.02	13.47	9.72

considering two time slots for incentives simplifies the analysis of the results.

Table 3 reflects what happens when we vary the limit on the incentive for a time slot (B). The results show that although increasing the limit on possible incentive payout does lead (on average) to larger profits, the differences are relatively small. This may be caused, to a large extent, by the fact that the incentives awarded are bounded not only by the limit B, but also by the smallest positive probability of time slots not receiving an incentive. We have observed that it is often the latter that is constraining.

We have modelled incentive impact as a linear relation: For each dollar of incentive the probability that the customer selects a delivery window increases by x; that is, $\bar{p}_i^t = p_i^t + xI^t$. Next, we investigate what happens for various values of x. These results can be found in Table 4. As expected, we see that our ability to increase profits by awarding incentives increases with the impact the incentives have on customer behavior. The increase is more pronounced for the sophisticated incentive schemes (from about 8% with $x = 0.1$ to about 13% for $x = 0.2$).

Next, we examine what happens when customers are more restrictive in terms of which delivery windows are acceptable to them. So far, we have assumed that there are eight delivery windows with positive probability of being selected by a customer. In the results reported in Table 5, the customers only have four delivery windows with positive probability of being selected. As this setup may be somewhat more realistic in certain applications, we are presenting the expanded set of results, that is, also varying the number of time slots considered for incentives.

Table 4 Base Results for Varying Incentive Impact Levels

UTILITY	*PROBPAT*	*NOINC*	*CHPFLT*	*CHPLP*	*BSTLP*	*BST*
0.1	1	1,224.72	1,288.38	1,335.63	1,362.08	1,349.29
	2	1,235.63	1,350.65	1,309.03	1,331.62	1,349.29
	3	1,228.93	1,323.40	1,329.03	1,288.34	1,349.29
Average		1,229.76	1,320.81	1,324.56	1,327.35	1,349.29
% improvement		0.00	7.40	7.71	7.94	9.72
0.2	1	1,224.72	1,304.28	1,405.99	1,379.40	1,349.29
	2	1,235.63	1,365.10	1,391.44	1,398.97	1,349.29
	3	1,228.93	1,336.17	1,372.35	1,407.83	1,349.29
Average		1,229.76	1,335.18	1,389.93	1,395.40	1,349.29
% improvement		0.00	8.57	13.02	13.47	9.72

Table 5 Base Results with Four Acceptable Delivery Windows

PROBPAT	$\|U\|$	*NOINC*	*CHPFLT*	*CHPLP*	*BSTLP*	*BST*
1	1	1,172.35	1,220.49	1,243.77	1,249.69	1,262.66
	2	1,172.35	1,183.79	1,251.40	1,229.05	1,262.66
	3	1,172.35	1,116.55	1,229.27	1,205.68	1,262.66
	4	1,172.35	1,148.06	1,277.97	1,201.59	1,262.66
Average		1,172.35	1,167.22	1,250.60	1,221.50	1,262.66
% improvement		0.00	−0.44	6.67	4.19	7.70
2	1	1,172.31	1,223.86	1,255.17	1,288.69	1,262.66
	2	1,172.31	1,185.25	1,268.94	1,225.02	1,262.66
	3	1,172.31	1,147.39	1,279.52	1,256.90	1,262.66
	4	1,172.31	1,076.21	1,268.73	1,223.43	1,262.66
Average		1,172.31	1,158.18	1,268.09	1,248.51	1,262.66
% improvement		0.00	−1.21	8.17	6.50	7.71
3	1	1,178.54	1,210.13	1,230.66	1,287.55	1,262.66
	2	1,178.54	1,176.17	1,259.50	1,226.81	1,262.66
	3	1,178.54	1,166.94	1,245.79	1,230.95	1,262.66
	4	1,178.54	1,149.36	1,258.99	1,247.86	1,262.66
Average		1,178.54	1,175.65	1,248.74	1,248.29	1,262.66
% improvement		0.00	−0.25	5.96	5.92	7.14

Several interesting observations can be made when examining these results. First, on average, the simple incentive scheme leads to a decrease in profits for all the probability patterns, whereas the sophisticated incentive schemes continue to result in increased profits (between 6% and 8%). A more careful examination of the results shows that the simple incentive scheme performs poorly when the number of time slots receiving an incentive gets larger but that it does reasonably well if the number of time slots receiving an incentive is small (one or two). This demonstrates that with fewer acceptable delivery windows, and thus fewer delivery options, a judicious choice of incentives is a must.

Finally, we consider the situation in which we assume that customers walk away with probability p_i^t if delivery window t is not presented to them as a delivery option, that is, the probability p_i^t is not redistributed among the other feasible slots. Realize, though, that the walkaway probability p_i^t can still be reduced as a result of incentives to other time slots. The incentive optimization problem can be adapted easily to accommodate this variant. Let V represent the set of time slots with positive probability p_i^t for which a feasible insertion exists, and let F represent the set of time slots with positive probability p_i^t for which no feasible insertion exists. The objective of the incentive optimization problem becomes

$$\max \sum_{t\in U}(r_i - C^t - I^t)(p_i^t + xI^t) + \sum_{t\in V}(r_i - C^t)(p_i^t - z) \qquad (20)$$

subject to

Table 6 Base Results with Walkaway Allowed

PROBPAT	$\|U\|$	*NOINC*	*CHPFLT*	*CHPLP*	*BSTLP*	*BST*
1	1	1,027.60	1,320.30	1,396.57	1,370.63	1,349.29
	2	1,027.60	1,263.03	1,374.65	1,401.23	1,349.29
	3	1,027.60	1,250.99	1,451.21	1,377.79	1,349.29
	4	1,027.60	1,157.40	1,380.27	1,404.54	1,349.29
Average		1,027.60	1,247.93	1,400.68	1,388.55	1,349.29
% improvement		0.00	21.44	36.31	35.13	31.30
2	1	1,081.44	1,318.23	1,386.96	1,282.36	1,349.29
	2	1,081.44	1,230.43	1,371.81	1,382.63	1,349.29
	3	1,081.44	1,191.74	1,365.79	1,408.27	1,349.29
	4	1,081.44	1,213.28	1,397.78	1,427.16	1,349.29
Average		1,081.44	1,238.42	1,380.59	1,375.11	1,349.29
% improvement		0.00	14.52	27.66	27.15	24.77
3	1	1,086.12	1,240.96	1,351.72	1,353.17	1,349.29
	2	1,086.12	1,275.58	1,349.34	1,395.38	1,349.29
	3	1,086.12	1,255.49	1,311.73	1,337.72	1,349.29
	4	1,086.12	1,202.58	1,360.29	1,356.42	1,349.29
Average		1,086.12	1,243.65	1,343.27	1,360.67	1,349.29
% improvement		0.00	14.50	23.68	25.28	24.23

$$z \le p_i^t \quad \forall\, t \in V, F \tag{21}$$

$$\sum_{t\in U} xI^t = z(|V| + |F|) \tag{22}$$

$$0 \le I^t \le B \quad \forall\, t \in U. \tag{23}$$

The time slots in the set F do not appear in the objective function since we assume that there is no profit when a customer walks away. However, the time slot with the smallest probability in F may be binding, unless a more restrictive probability is found in V (Equation (21)). Again, the quadratic term can be approximated by a piecewise linear function as before to create a linear program.

It is amazing to see, in Table 6, the negative effect of walkaways on profit and, at the same time, the power of incentives to recapture that loss of profit. For probability pattern 1, the profit drops more than 16% compared to the base results without walkaways, but the sophisticated incentive schemes are able to completely recapture the loss and substantially increase the profit (for an improvement over no use of incentives of over 35%)! This shows that using incentives to dissuade customers from walking away may be even more important in increasing profit than simply reducing delivery costs.

3. The Home Delivery Problem with Wider Slot Incentives

Delivery costs are impacted significantly by the stringent one-hour delivery windows. In Campbell and Savelsbergh (2005), we demonstrated that expanding a one-hour delivery window to two hours can increase profits by more than 6% and can be increased an additional 5% if further expanded to three hours. Consequently, instead of using incentives to encourage the use of one-hour delivery windows preferred from a scheduling perspective, we may alternatively consider using incentives to encourage accepting wider delivery windows, which increases our scheduling flexibility and therefore may increase profits.

3.1. Problem Definition and Assumptions

To study the benefits of using incentives to encourage the consumer to accept wider time windows, we define the HDPWI.

The basic setup is the same as for the HDPTI. Thus, we assume that we know the probability p_i^t that a customer i will choose a delivery in time slot t. When a request for service arrives, the vendor may now offer incentives for the selection of a two-hour delivery window up to an amount of B. The two-hour delivery windows are also assumed to cover the entire day, but to overlap, for example, 8:00–10:00, 9:00–11:00, ..., 18:00–20:00. The probability of selecting a two-hour delivery window is initially zero but increases by an amount equal to the incentive offer multiplied by rate y. An increase in probability of a two-hour time slot is compensated for by a decrease in probability of *all* one-hour time slots with p_i^t values greater than zero. As in the HDPTI, the probability of one-hour time slots is decreased by equal amounts. The customer then selects the time slot based on these modified probabilities. The objective is to maximize the total profit resulting from executing the set of delivery routes, that is, total revenue minus incentives costs and delivery costs.

Not all two-hour delivery windows are viable candidates to receive an incentive. It seems reasonable to assume that a customer would not be interested in a wider time slot if either of the component one-hour time slots have a zero probability of being selected. (A zero probability indicates a delivery in the window is impossible or highly undesirable, and agreeing to a larger delivery window would signal that this is no longer the case.) Similarly, it is also reasonable to eliminate a wider time slot from consideration when it is infeasible for the vendor to make the delivery to the customer in one of the component time slots.

3.2. Modeling with Wider Time Slots

In the HDPTI, an increase in the probability of a time slot due to an incentive is compensated for by a decrease in probability of the time slots in V (by equal amounts). In the HDPWI, an increase in the probability of a wider time slot is compensated for by a decrease in the probability of other time slots (again by equal amounts), but now the set V of other time slots consists of *all* one-hour time slots with positive probability. In this way, the two incentive models are fairly similar in terms of how money is traded for probability. The HDPWI has a significant advantage over the HDPTI, though, from a computational perspective. Since the set of time slots for which the probability decreases are the same regardless of the set of (wider) time slots that are being considered for incentives, the selection of the time slots that may receive incentives becomes much easier.

In the incentive optimization problem presented below, we will refer to the set of two-hour windows under consideration for an incentive as W (as opposed to U), while V continues to represent the one-hour slots not receiving an incentive for which the probability may be reduced (now containing all one-hour time slots with positive probability). Therefore, let

- $W =$ set of two-hour time slots under consideration for an incentive
- $V =$ set of one-hour time slots with positive probability
- $C^{(t,\,t+1)} =$ cost of inserting a delivery in the two-hour time slot spanning one-hour time slots t and $t+1$, that is, $C^{(t,\,t+1)} = \min(C^t, C^{t+1})$.

We want to find

- $I^{(t,\,t+1)} =$ the incentive for the two-hour time slot spanning one-hour time slots t and $t+1$
- $z =$ the probability removed from one-hour time slots in V

so as to maximize expected profitability.

The incentive decision for customer i can be represented by the following optimization problem:

$$\max \sum_{(t,\,t+1)\in W} (r_i - C^{(t,\,t+1)} - I^{(t,\,t+1)})yI^{(t,\,t+1)} + \sum_{t\in V} (r_i - C^t)(p_i^t - z) \qquad (24)$$

subject to

$$z \le p_i^t \quad \forall\, t \in V \qquad (25)$$

$$\sum_{(t,\,t+1)\in W} yI^{(t,\,t+1)} = z|V| \qquad (26)$$

$$0 \le I^{(t,\,t+1)} \le B \quad \forall\, (t, t+1) \in W. \qquad (27)$$

In the objective, Equation (24), the first portion represents the expected profit from wider time slots receiving incentives, and the second portion represents the modified expected profit from the one-hour slots. Because all two-hour time slots in W initially have zero probability of being selected, the probability of each wider slot is no longer dependent on initial p_i^t values, and wider slots do not change the constraints concerning the amount of incentives that can be awarded (Equation (25)). In the model for the HDPTI, the decision of which time slots should receive incentives or even be included in U is based on a combination of cost and original probability values. Here, as all wider time slots originally have an identical probability of zero, it is always optimal to include the two-hour time slots with the smallest $C^{(t,\,t+1)}$ values in W. These time slots have the highest $r_i - C^{(t,\,t+1)}$ values and thus can contribute the most to the profit.

Even though it is possible to consider all viable two-hour windows in the incentive optimization problem, it is easy to show that it will never be the case that all two-hour windows should receive an incentive in an optimal solution.

Observation 4. A two-hour time slot will not receive an incentive if it has a profit that is less than the average profit of the one-hour slots, that is, $I^{(t,\,t+1)} = 0$ when

$$r_i - C^{(t,\,t+1)} - \frac{\sum_{t\in V}(r_i - C^t)}{|V|} \le 0. \qquad (28)$$

This is not only true when a two-hour time slot is considered by itself, but also when it is considered in conjunction with other two-hour time slots.

As before, we can approximate the quadratic terms with a piecewise linear function and transform the incentive optimization problem into a linear program. As a result, incentives for wider time slots can also be computed within a few seconds.

3.3. Computational Results

The setup used for the computational experiments is the same as before, with only minor changes. There is no longer any reason to include *BSTLP* because, as we have discussed above, we know the best set W of size l is the one with the l least cost insertions. There is also a small change in the implementation of *BST*. In *BST* tests, all customers receive a delivery in a wider time slot if one is viable, and if none are viable, the customer receives a delivery during the cheapest one-hour window. We have conducted the same set of experiments as for HDPTI and have observed fairly similar behavior. Therefore, we only present a few key results tables.

The results for the 25 instances in the basic data set are presented in Table 7. Before discussing the performance of the incentive schemes, we observe that the potential value of enticing customers to select two-hour time windows is huge, because the upper bound values (*BST*) are significantly higher than the values without incentives (*NOINC*), with increases in profit of more than 20% in all cases. Although the incentive schemes perform well, they are not yet capable of fully capitalizing on these opportunities. The simple incentive scheme increases profit by about 8%–10%, where the more sophisticated incentive scheme increases profit by about 12%–14%. It may be the case that with wider time slots it is key to convince customers to accept a wider delivery window but that it is less important which wider delivery window they select.

Table 7 Base Results

PROBPAT	$\|W\|$	*NOINC*	*CHPFLT*	*CHPLP*	*BST*
1	1	1,224.72	1,393.22	1,349.92	1,508.32
	2	1,224.72	1,330.39	1,412.91	1,508.32
	3	1,224.72	1,325.87	1,388.19	1,508.32
	4	1,224.72	1,328.10	1,445.42	1,508.32
Average		1,224.72	1,344.40	1,399.11	1,508.32
% improvement		0.00	9.77	14.24	23.16
2	1	1,235.63	1,350.33	1,335.97	1,508.32
	2	1,235.63	1,342.77	1,387.17	1,508.32
	3	1,235.63	1,325.38	1,353.92	1,508.32
	4	1,235.63	1,346.81	1,423.03	1,508.32
Average		1,235.63	1,341.32	1,375.02	1,508.32
% improvement		0.00	8.55	11.28	22.07
3	1	1,228.93	1,296.80	1,377.95	1,508.32
	2	1,228.93	1,387.41	1,358.76	1,508.32
	3	1,228.93	1,331.12	1,428.42	1,508.32
	4	1,228.93	1,322.13	1,386.06	1,508.32
Average		1,228.93	1,334.37	1,387.80	1,508.32
% improvement		0.00	8.58	12.93	22.73

Table 8 presents the results for the situation in which fewer requests for delivery arrive. The results reveal that, compared to the base case results, the potential for incentives is much smaller. In the base case, the upper bound values (*BST*) were about 22%–23% higher than the values without incentives (*NOINC*), whereas in this case, the upper bounds are only 13%–14% higher. On the other hand, the incentive schemes do a slightly better job of bridging the gap between *NOINC* and *BST*. (The average *CHPLP* value deviates from the average *BST* value by 5% rather than 8%.)

Finally, the results in Table 9 demonstrate that incentives help increase revenues by dissuading customers from walking away, as with the HDPTI. Furthermore, the results reinforce that the more time slots are considered for incentives, the more important it appears to be to use more sophisticated incentives schemes. For all probability patterns, the performance of *CHPFLT* degrades substantially when $|W|$ is equal to 3 or 4.

Table 8 Base Results with Less Demand

PROBPAT	$\|W\|$	*NOINC*	*CHPFLT*	*CHPLP*	*BST*
1	1	1,022.69	1,102.00	1,098.40	1,167.39
	2	1,022.69	1,104.00	1,112.52	1,167.39
	3	1,022.69	1,089.02	1,109.55	1,167.39
	4	1,022.69	1,092.69	1,128.73	1,167.39
Average		1,022.69	1,096.93	1,112.30	1,167.39
% improvement		0.00	7.26	8.76	14.15
2	1	1,034.43	1,099.46	1,087.61	1,167.39
	2	1,034.43	1,100.17	1,113.22	1,167.39
	3	1,034.43	1,087.90	1,111.62	1,167.39
	4	1,034.43	1,091.90	1,111.80	1,167.39
Average		1,034.43	1,094.86	1,106.06	1,167.39
% improvement		0.00	5.84	6.92	12.85
3	1	1,033.82	1,093.83	1,108.98	1,167.39
	2	1,033.82	1,107.24	1,098.99	1,167.39
	3	1,033.82	1,084.76	1,133.39	1,167.39
	4	1,033.82	1,082.74	1,092.75	1,167.39
Average		1,033.82	1,092.14	1,108.53	1,167.39
% improvement		0.00	5.64	7.23	12.92

Table 9 Base Results with Walkaway Allowed

PROBPAT	$\|W\|$	*NOINC*	*CHPFLT*	*CHPLP*	*BST*
1	1	1,027.60	1,359.43	1,369.87	1,517.42
	2	1,027.60	1,317.84	1,390.13	1,517.42
	3	1,027.60	1,250.31	1,474.26	1,517.42
	4	1,027.60	1,245.76	1,365.73	1,517.42
Average		1,027.60	1,293.34	1,400.00	1,517.42
% improvement		0.00	25.86	36.24	47.67
2	1	1,081.44	1,305.60	1,323.24	1,517.42
	2	1,081.44	1,362.00	1,350.44	1,517.42
	3	1,081.44	1,295.58	1,380.43	1,517.42
	4	1,081.44	1,258.51	1,387.59	1,517.42
Average		1,081.44	1,305.42	1,360.43	1,517.42
% improvement		0.00	20.71	25.80	40.31
3	1	1,086.12	1,236.27	1,347.68	1,517.42
	2	1,086.12	1,312.00	1,329.42	1,517.42
	3	1,086.12	1,267.88	1,308.24	1,517.42
	4	1,086.12	1,214.50	1,345.79	1,517.42
Average		1,086.12	1,257.66	1,332.78	1,517.42
% improvement		0.00	15.79	22.71	39.71

4. The Value of Optimization

In recent studies of online, dynamic routing problems (e.g., Bayraksan 2000 for the traveling salesman problem), questions concerning the value of optimization are raised and addressed. For example, it is not clear that when simple insertion techniques are used to add stops to a tour (as this results in extremely fast response times) there is value in starting from an optimized routing solution versus starting from a heuristic routing solution. The study by Bayraksan shows that only after 40% of the orders have materialized does it become valuable to start from an optimized traveling salesman tour.

In the HDPTI (or the HDPWI), there is a similar question that can be asked. Because it is not clear that the insertion costs based only on the first few arrivals adequately reflect the final costs of servicing the order, it is not clear that these insertion costs can be used to compute useful incentives. This issue was raised earlier in the discussion of the computational results for Table 1. Thus, we have performed an experiment to specifically address and study this issue. In the results presented in Table 10, the first 25% of customers that place an order do not receive an incentive. Only after these first 25% of customers have been scheduled, did we start using the insertion costs to compute incentives for the remaining 75% of requests where feasible.

Table 10 Results from Varying the Start of Incentive Computation

STARTPCT	*PROBPAT*	*NOINC*	*CHPFLT*	*CHPLP*	*BSTLP*	*BST*
0%	1	1,224.72	1,304.28	1,405.99	1,379.40	1,349.29
	2	1,235.63	1,365.10	1,391.44	1,398.97	1,349.29
	3	1,228.93	1,336.17	1,372.35	1,407.83	1,349.29
Average		1,229.76	1,335.18	1,389.93	1,395.40	1,349.29
% improvement		0.00	8.57	13.02	13.47	9.72
25%	1	1,224.72	1,275.20	1,301.16	1,301.49	1,324.39
	2	1,235.63	1,299.16	1,361.18	1,363.24	1,386.43
	3	1,228.93	1,298.51	1,324.95	1,324.38	1,346.68
Average		1,229.76	1,290.96	1,329.10	1,329.70	1,352.50
% improvement		0.00	4.98	8.08	8.13	9.98

The idea is that only after 25% of orders have been processed will the insertion costs adequately reflect the true costs of servicing an order and can therefore be used to compute meaningful incentives. For ease of comparison, Table 10 includes the results from the case in which all orders are candidates for incentives (right from the start) as well as the case in which the first 25% are not considered for incentives. *STARTPCT* reflects the percent of expected arrivals that occurs before incentives are used. At most two incentives are offered to each potential delivery.

Surprisingly, we see that not offering incentives to the first 25% of the customers hurts the profit for almost all of the incentive methods and for almost all of the probability patterns. This demonstrates that offering incentives can be critical even in the early stages of building the schedule. It should be noted that the same experiment with the HDPWI had very similar results.

5. Insights

The primary objective of this study was to determine if it is possible to increase the profitability of home delivery operations using incentive schemes. Our computational experiments have demonstrated that even relatively simple incentive schemes have the potential to do this. A summary of the insights obtained is given below:

- The use of incentive schemes can substantially reduce delivery costs and thus enhance profits.
- Incorporating intelligence into incentive schemes enhances their performance.
- Incentive schemes may substantially reduce the number of walkaways.
- It is sufficient to provide incentives to only a few delivery windows (≤ 3).
- The more time slots are considered for incentives, the more important it becomes to use more sophisticated incentive schemes.
- It is easier to develop incentive schemes that encourage customers to accept wider delivery windows, rather than those that encourage customers to select specific time slots.
- The use of incentives can be critical even in the early stages of building a delivery schedule.

6. Discussion

Consumer direct service models give rise to fascinating and challenging optimization problems. We focused on the use of time slot incentives to reduce delivery costs, which in turn should result in increased profits. We have shown that optimization models to compute incentives can be designed and implemented given assumptions on customer behavior and the impact of incentives. Our simulation studies indicate that these optimization models have the potential to substantially increase profits.

We realize that the results obtained in our simulations are affected by our assumptions regarding customer behavior and incentive impact and our choice of parameters. The results may be different for different assumptions regarding customer behavior and incentive impact and different parameter choices. This shows that we have only scratched the surface and that more research is needed to better understand and fully exploit the potential of incentives.

Because the use of incentives to reduce delivery costs in attended home delivery environments is a new phenomenon, little information is available about the impact of incentives on customers' behavior. Therefore, we have had to make various assumptions when modeling customer behavior. In future work, we intend to consider additional models for capturing the relationship between incentives and their impact on time slot selection.

To illustrate, consider our assumption that the total increase in probability created by incentives is divided up and removed *equally* from all the time slots with positive probability not receiving an incentive. Mathematically, this is expressed as

$$z \leq p_t \quad \forall\, t \in V$$
$$\sum_{t \in U} x I^t = z|V|,$$

where V is the set of time slots with positive probability not receiving an incentive and z represents the reduction in probability. As probabilities are nonnegative, our assumption restricts the total amount of incentives to $(\min_{t \in V} p_t \times |V|)/x$. A simple extension reduces the probabilities for all time slots in V at an equal rate until they reach 0. That is, when the probability of one of the time slots in V reaches 0, we continue to reduce the probability of the time slots that still have a positive probability. Mathematically, this can be expressed as follows:

$$z \leq p_t + v_t \quad \forall\, t \in V$$

$$\sum_{t\in U} xI^t = z|V| - \sum_{t\in V} v_t,$$
$$v_t \geq 0 \quad \forall\, t \in V,$$

where v_t is an auxiliary variable capturing the "negative" portion of z. Alternatively, we can assume that the total increase in probability created by incentives is divided up and removed *proportionally* from all the time slots with positive probability not receiving an incentive. Mathematically, this is expressed as

$$\sum_{t\in U} xI_t = \sum_{t\in V} yp_t$$
$$0 \leq y \leq 1,$$

where y represents the proportion of probability that is reduced for each time slot in V. These are just three of the many possible models for the change in probability as a result of incentives. It is not clear yet which model best captures customer behavior. Therefore, we selected one as a starting point for our initial investigations.

Similarly, we have chosen to represent the impact of an incentive as a linear function, that is, $p_t + xI^t$. Other functional forms may better capture customer behavior, but our choice provided a reasonable starting point for our investigations. In our simulations, we have chosen the coefficient x to be either 0.1, 0.2, or 0.3. That is, we assume that offering an incentive of one dollar increases the probability of selecting a certain time slot by 0.1, 0.2, and 0.3, respectively. Even though these values may seem large, they were selected based on feedback from home grocery delivery providers. Their experience indicates that many customers are fairly indifferent about a set of time slots and that very small incentives often have a significant impact on time slot selection.

As mentioned above, we have only scratched the surface. More research is needed to better understand and fully exploit the potential of incentives.

Acknowledgments

This work was partially supported by the National Science Foundation through Grant 0237726 (Campbell).

References

Bayraksan, G. 2000. An introductory analysis of online traveling salesman. Unpublished master's thesis, Dept. of Operations Research and Industrial Engineering, University of Texas, Austin, TX.

Bent, R., P. Van Hentenryck. 2004. Scenario-based planning for partially dynamic vehicle routing with stochastic customers. *Oper. Res.* **52**(6) 977–987.

Campbell, A., M. Savelsbergh. 2004. Efficient insertion heuristics for vehicle routing and scheduling problems. *Transportation Sci.* **38** 369–378.

Campbell, A., M. Savelsbergh. 2005. Decision support for consumer direct grocery initiatives. *Transportation Sci.* **39**(3) 313–327.

Cox, B. 2001. Internet grocer seeks an angel. *InternetNews* (April 16) http:// www.internetnews.com.

Farmer, M., G. Sandoval. 2001. Webvan delivers its last word: Bankruptcy. *News.com* (July 9) http://news.com.com.

Feo, T., M. Resende. 1995. Greedy randomized adaptive search procedures. *J. Global Optimization* **6** 109–133.

Green, H. 2003. Fresh direct. *Business Week* (November 24) 104.

Johnson, C. A., K. Delhagen, E. H. Yuen. 2003. US eCommerce Overview: 2003 to 2008. Report, Forrester Research. http:// www.forrester.com.

Larsen, A., O. Madsen, M. Solomon. 2002. Partially dynamic vehicle routing—Models and algorithms. *J. Oper. Res. Soc.* **53** 637–646.

Lin, I., H. Mahmassani. 2002. Can online grocers deliver? Some logistics considerations. *Transportation Res. Record* **1817** 17–24.

Madsen, O. B. G., K. Tosti, J. Vælds. 1995. A heuristic method for dispatching repair men. *Ann. Oper. Res.* **61** 213–226.

Nemhauser, G. L., L. A. Wolsey. 1988. *Integer and Combinatorial Optimization*. John Wiley and Sons, New York.

Punakivi, M. 2000. The value added of route optimisation in home delivery. Working paper, The ECOMLOG Research Program, Department of Industrial Engineering and Management, Helsinki University of Technology, Helsinki, Finland.

Saranen, J., J. Småros. 2001. An analytical model for home delivery in the new economy. Working paper, The ECOMLOG Research Program, Department of Industrial Engineering and Management, Helsinki University of Technology, Helsinki, Finland.

Savelsbergh, M. W. P. 1986. Local search for routing problems with time windows. *Ann. Oper. Res.* **4** 285–305.

Yrjölä, H. 2001. Physical distribution considerations for electronic grocery shopping. *Internat. J. Physical Distribution Logistics Management* **31** 746–761.