

When Are Deliveries Profitable?

Considering Order Value and Transport Capacity in Demand Fulfillment for Last-Mile Deliveries in Metropolitan Areas

The authors suggest a way of making attended last-mile deliveries in metropolitan areas more profitable by considering a given transport capacity and the expected value of orders.

DOI 10.1007/s12599-014-0321-9

The Authors

Prof. Dr. Catherine Cleophas

Research Group Advanced Analytics
RWTH Aachen University
Kackerstraße 7
52072 Aachen
Germany
catherine.cleophas@rwth-aachen.de

Prof. Dr. Jan Fabian Ehmke (✉)

Department Information Systems
Freie Universität Berlin
Garystraße 21
14195 Berlin
Germany
janfabian.ehmke@fu-berlin.de

Received: 2013-07-13

Accepted: 2014-02-03

Accepted after three revisions by the editors of the special focus.

Published online: 2014-04-16

This article is also available in German in print and via <http://www.wirtschaftsinformatik.de>: Cleophas C, Ehmke JF (2014) Wann sind Lieferaufträge profitabel? Berücksichtigung des Auftragswertes und der Transportkapazität in der Tourenplanung für die „letzte Meile“ in Ballungsräumen. WIRTSCHAFTSINFORMATIK. doi: 10.1007/s11576-014-0412-8.

© Springer Fachmedien Wiesbaden 2014

1 Introduction

The ongoing boom of online retailing has led to an ever-growing number of last-mile deliveries. In the United States, for example, online retailing has been

the fastest growing retail sector with an overall growth of about 15 % in recent years (US Census Bureau 2012). While last-mile deliveries are one of the most important sections of the supply chain, they are also regarded as very expensive and inefficient (Gevaers et al. 2010). Attended home deliveries pose a particular challenge, as consumers need to be at home to physically and legally receive their delivery. For attended home deliveries, e-commerce businesses and consumers have to mutually agree on delivery time windows to avoid costly failures of deliveries.

Since profit margins for e-commerce businesses are often low and consumer expectations regarding service quality and reliability of deliveries are high (Allen et al. 2007), the delivery of small shipments requires careful planning. The economic challenges these businesses face have been underlined by the failures of the online grocery stores *Webvan* (bankruptcy in 2001) and *Publix Direct* (shut down in 2003). Nevertheless, *Amazon*, *eBay*, and *Walmart* have recently begun to offer same day delivery for a flat fee starting as low as \$5 (Wirthman 2013). Also groceries can be ordered online: *Peapod*, for example, serves major cities in the U.S., while grocery stores like *REWE* in Germany or *Tesco* in the United Kingdom are experimenting with online ordering services and varying pricing of deliveries. *Peapod* offers attended home deliveries for more than 10,000 products, including fresh groceries, such as cheeses and milk; delivery fees range from \$6.95 to \$9.95.

Intuitively, the short-term profitability of e-commerce businesses delivering small shipments can be improved in two ways. On the one hand, efficient vehicle routing and optimal usage of transport capacities, combined with adequate

shipping charges, can minimize delivery costs. However, there are clear limits to consumer acceptance of shipping charges. Businesses often need to subsidize deliveries, since tight time windows increase costs of vehicle routing. Furthermore, costs of vehicle routing are usually only roughly anticipated in the moment of order acceptance. On the other hand, order acceptance techniques that prefer profitable delivery requests and take expected costs of delivery into account can maximize revenue from order value. We refer to the combination of cost minimal routing with value-based order acceptance techniques as value-based demand fulfillment.

While classical demand fulfillment aims to maximize the number of orders or the expected order value, vehicle routing is about minimizing delivery costs. As shown in the following literature review, existing approaches to demand fulfillment often neglect the expected demand, the order value, or the expected costs of delivery. In this contribution, our research objective is to provide a more integrated approach, interweaving the concepts of demand fulfillment and vehicle routing. When transport capacities are limited, value-based demand fulfillment may help maximize revenues by offering desirable time windows only to consumers with highly profitable orders. Note that in future work, the definition of “highly profitable” may be extended to include customer long-term value in the sense of customer relationship management. Here, we concentrate on the short-term value of individual orders, ignoring the potential impact of a rejected order on customer value.

In this paper, we consider the following two research questions: How could an iterative solution approach, enabling value-based demand fulfillment, be designed? And: What would be its poten-

tial effect on costs and revenues? The objective of the iterative solution approach is to maximize the expected value of orders, while considering given transport capacities that limit costs of transportation. To this end, based on historical order data, we derive a set of expected delivery requests that we would be able to accommodate with a given fleet. As a result, we have the number and the location of time slots that would allow for efficient delivery tours. Using revenue management techniques, order acceptance controls then assign or reject actual requests to these “reserved” time slots, maximizing the overall value of orders. When order acceptance is finished, actual delivery tours are planned for the set of accepted requests, minimizing total costs of transportation.

Our solution approach supports e-commerce businesses (1) in reserving transport capacities for specific delivery areas and time windows with a high expected order value (demand forecast), (2) in the booking of delivery time slots, i.e., deciding to accept or reject delivery requests posed by consumers (fulfillment controls), and (3) in the computation of cost-minimal delivery tours (vehicle routing). To this end, we combine concepts of revenue management – maximizing the expected overall revenue earned from orders – with techniques of time-dependent vehicle routing – ensuring the feasibility of delivery tours and minimizing costs of delivery in metropolitan areas.

Order forecasting is based on concepts from demand management as summarized by Talluri and Van Ryzin (2004). Demand forecasts are aligned with transport capacities, taking expected travel times into account. Order acceptance is employed first for capacity planning based on acceptance mechanisms known from vehicle routing and home delivery (Ehmke and Campbell 2014), considering varying transport capacity in metropolitan areas due to rush-hour phenomena. This is complemented by the application of the Estimated Marginal Seat Revenue heuristic (EMSR; Belobaba 1987) to select for valuable orders. The iterative solution approach is demonstrated based on order data for a fictitious delivery service in the metropolitan area of Stuttgart, Germany. A computational study investigates the potential of value-based demand fulfillment using a simulation model of consumer demand.

The following section provides a brief literature review on concepts of demand fulfillment and vehicle routing. Section 3 outlines our value-based demand fulfillment approach. Subsequently, we present a simulation study, illustrating the potential and limitations of our approach (Sect. 4). The paper concludes with an outlook pointing out further possibilities of interweaving revenue management and vehicle routing.

2 Literature Review

The fulfillment process for consumer direct businesses can be structured into the phases (1) order capture and promise, (2) order sourcing and assembly, and (3) order delivery (Campbell and Savelsbergh 2005). The corresponding planning problems have different objectives and are often handled distinctly and sequentially. For Phase 1, the number of orders is maximized, and for Phase 3, costs of delivery are minimized given a fixed set of orders and transport capacities. A more integrated approach may help reducing inefficiencies arising from insufficient alignment between the distinct planning problems by considering limited transport capacities as early as in Phase 1. For instance, from an integrated perspective, in order to efficiently employ fixed transport capacities, we would accept a low-value order only if the likelihood of a high-value order’s arrival is rather low.

In the following literature review, we briefly elaborate on existing solution approaches for Phase 1 (Sect. 2.1) and Phase 3 (Sect. 2.2). To this end, we have focused on established review papers, books, and well-cited standard papers for the particular areas, as well as on most related recent work in the areas of revenue management and vehicle routing. Order sourcing (Phase 2) is not within the scope of this paper and therefore not discussed further.

2.1 Order Capture and Promise

Sophisticated order capture and promise requires appropriate forecasts of demand and techniques that distinguish valuable orders from less valuable orders. Stadtler (2005) differentiates demand planning and demand fulfillment. Demand planning is defined as the process of forecasting demand; demand fulfillment is described as order promising and due date

setting. In the context of this paper, we implement demand fulfillment by controlling the availability of delivery time slots per delivery area depending on order value. We base this decision on demand planning by forecasting delivery requests and order values according to particular delivery areas and time slots.

Demand forecasting is a task of supply chain management in general as well as revenue management in particular. For attended last-mile deliveries, the length, the number, and the potential overlap of time windows are determined based on demand forecasts (“time window design”). Agatz et al. (2011) present an approach that anticipates transportation costs in time window design. Given service requirements and average weekly demands for each zip code area of a delivery region, time windows for each zip code area are determined, minimizing the expected costs of delivery. Other papers investigate the impact of time window design on resulting costs of delivery and customer satisfaction. Punakivi and Saranen (2001) report that completely flexible, unattended delivery services may reduce costs by up to 1/3 relative to attended deliveries with two-hour time windows. Campbell and Savelsbergh (2005) find that the expansion of one-hour time windows to two hours increases profits by more than 6%. Ehmke and Campbell (2014) investigate different designs of time windows in a metropolitan setting, highlighting that shorter time windows are much more difficult to handle in terms of costs and reliability.

Demand forecasts and/or time window design are usually assumed as given when it comes to order promise. In this paper, however, we assume that only the length and the structure of time windows has already been fixed. In contrast to the above approaches, we forecast the expected number and value of orders for each delivery time window and area by static demand models using exponential smoothing as explained, for instance, by Talluri and Van Ryzin (2004). This allows for a more detailed alignment with the order promise and order delivery phases.

Building on demand forecasts, demand fulfillment is about deciding in advance which delivery requests to fulfill, given already accepted and still expected delivery requests. Existing literature differs in the extent to which demand forecasts are taken into account. Work in the area of vehicle routing refers

to “acceptance mechanisms”, which decide on when to open and when to close delivery time slots while considering limited transport capacities. Campbell and Savelsbergh (2005) introduce several acceptance mechanisms that determine whether a delivery request can be feasibly accommodated in any of the predefined delivery time slots. The SLOT mechanism accepts requests in a first-come-first-serve manner, as long as a given number of deliveries for a particular area and time slot is not exceeded. The more sophisticated DYN mechanism actually schedules delivery tours when checking the feasibility of a new delivery request with respect to already accepted requests. Ehmke and Campbell (2014) extend the SLOT and DYN mechanisms for congested metropolitan areas by including detailed information on expected congestion in rush-hour time windows.

While the sources cited above consider expected transportation costs in detail, they often ignore the value of orders and the expected value of future orders and their role in maximizing overall revenues. This is the main objective of revenue management, which describes the idea of maximizing revenue by selling units of the same product at different prices to different consumers. The aim is to both fully exploit consumers’ willingness-to-pay and to sell leftover capacity at a discount. Talluri and Van Ryzin (2004) offer a thorough introduction to mathematical methods in this field. As conditions for the applicability of revenue management in addition to distinguishable demand segments, inflexible capacity and a fixed expiry date are frequently cited. Assuming that the number of delivery vehicles is limited, the condition of inflexible capacity also applies to attended deliveries. Also, the delivery slot at a particular time of day loses its value when this time has passed.

From a revenue management perspective, there are a vast number of techniques that support decisions on order acceptance to support “profitable-to-promise demand fulfillment” (Hahn and Kuhn 2012). The probably most well-known heuristic for revenue optimization is EMSR introduced by Belobaba (1987), which extends Littlewood’s Rule (Littlewood 1972). Based on demand forecasts, EMSR assigns a number of units to sell to specific “fare classes”. Each class indicates a different fare at which units reserved for this class can be

sold. If orders are expected for units offered at a high fare, capacity is reserved by assigning it to the corresponding fare class. The idea of reserving capacity for late, valuable delivery requests follows an assumption frequently phrased in revenue management literature, namely that requests can be segmented by request time and that high-value requests arrive after low-value requests. While much more complex approaches have been developed since, EMSR is still widely used in the industry due to its simplicity and robust results (Talluri and Van Ryzin 2004). We use EMSR to decide which delivery time slots to offer to which consumers. Instead of using fare classes, we sort customers according to their order value.

A few approaches implement the idea of integrated demand fulfillment and revenue management. The most similar to ours is the work by Yang et al. (2012), who formulate a combination of demand fulfillment and revenue management with regard to e-commerce deliveries. The authors propose to realize demand planning through a multinomial logit model of choice and to control fulfillment via dynamic pricing of transport charges. However, while that study focuses on appropriate pricing of delivery, we concentrate on the decision of which requests to accept based on order value. Another more general example is the work by Quante et al. (2009), who summarize and compare models and planning systems applicable for integrated demand fulfillment and revenue management. In contrast, our contribution demonstrates the transfer of standard revenue management techniques to the context of a specific application, namely attended delivery services, and the investigation of its potential.

Considering the works citing Quante et al. (2009) up to this point, a focus on made-to-order manufacturing and other application areas of classical revenue management becomes evident. So far, little consideration is given to the combination of order acceptance, maximizing revenue from order values, and cost-efficient vehicle routing.

2.2 Order Delivery

The last phase of the fulfillment process of consumer direct businesses is to schedule and deliver orders. The cost-efficient allocation of orders to delivery vehicles, as well as the optimal sequencing of orders, is achieved through

vehicle routing and scheduling procedures. The corresponding optimization problem is known as the vehicle routing problem with time windows (VRPTW). The VRPTW aims at determining the distance(cost)-minimum tour plan, where (1) every tour departs and terminates at a central depot, (2) every consumer is visited exactly once by one vehicle, and (3) every consumer is serviced within the delivery time window. In this paper, we assume that delivery time windows are “hard”, i.e., we do not allow the violation of delivery time windows. In other words, order promises are scheduled straightforwardly. If a vehicle arrives early at a consumer, delivery needs to wait until the time window opens.

Due to its wide applicability in transportation and logistics and its computational complexity, the VRPTW has been extensively studied in the last 25 years. Baldacci et al. (2012) have recently reviewed exact optimization approaches and model formulations. Especially heuristic and metaheuristic approaches have been investigated intensively, as the VRPTW is NP-hard and exact solution approaches may take very much computational time for instances of realistic problem size. Heuristics usually aim at minimizing the number of vehicles as a primary objective and minimizing distances or travel times as a secondary objective. Bräysy and Gendreau (2005a) review route construction and local search algorithms. Bräysy and Gendreau (2005b) summarize metaheuristic solution approaches.

For routing and scheduling in metropolitan areas, the standard VRPTW needs to include time-varying information on typical congestion in terms of time-of-day dependent travel times. Kok et al. (2012) investigate the impact of time-dependent travel times on the reliability of tour plans. Since more accurate travel time information has become available, there has been an increasing number of papers that include time-dependent travel times. Fleischmann et al. (2004), Hashimoto et al. (2008), and Figliozzi (2009) implement adapted local search heuristics to solve the time-dependent VRPTW. Haghani and Jung (2005) present a genetic algorithm, Donati et al. (2008) refer to ant colonies, and Ichoua et al. (2003) as well as Maden et al. (2010) develop solution frameworks based on tabu search metaheuristics. In

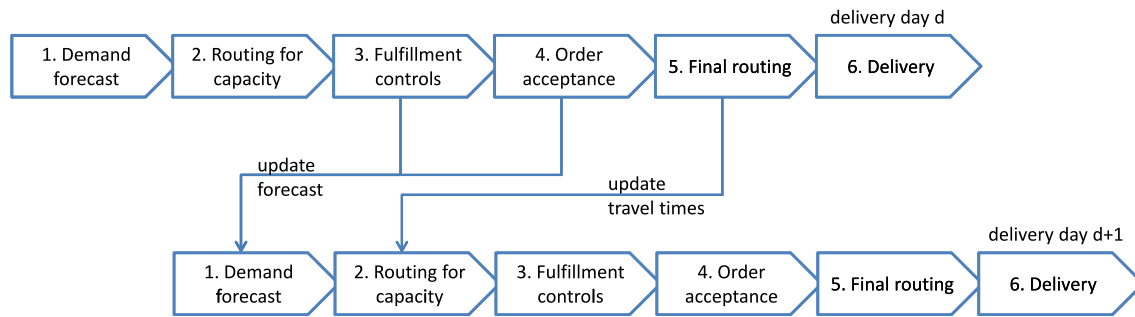


Fig. 1 Process combining routing and fulfillment controls over two delivery days

this paper, we refer to the latter in the extended version by Ehmke et al. (2012a), who deploy the framework with large amounts of taxi speed data for routing in metropolitan areas.

3 An Iterative Solution Approach for Value-Based Demand Fulfillment

The idea of valued-based demand fulfillment is to maximize revenues while considering limited transport capacities by improving the phase of order capture and promise. To this end, we iteratively align forecasts of order locations and order values with expected transportation costs by reserving transport capacities where valuable and efficient. Within the scope of order promise, order acceptance considers the value of actual orders as well as reserved transport capacities.

In the following subsection, we describe the proposed process of combining demand forecasting and vehicle routing for value-based demand fulfillment. Subsequently, the three major planning components demand forecasting, vehicle routing, and fulfillment control planning are presented in detail.

3.1 Overview of the Iterative Solution Approach

Our demand fulfillment model assumes that a given metropolitan area is organized into distinct delivery areas. We offer multiple, fixed time windows for delivery and flexibly control the availability of delivery time slots depending on order value. Historical order data is available for each combination of time slot and delivery area. Furthermore, we operate with a fixed transport capacity determined by a given fleet size.

Based on these assumptions, we substantiate value-based demand fulfillment

as illustrated in Fig. 1. Our solution approach follows the standard fulfillment process for consumer direct processes, as discussed in Sect. 2. We regard the use of a standard model as beneficial for future research and discussion. The solution approach iteratively interweaves planning activities of order forecasting and management (steps 1, 3, 4) with transport planning (steps 2 and 5):

- Order capture: (1) Using historical order data, we forecast the number of delivery requests for each combination of value, time slot, and delivery area. (2) These expected delivery requests serve as input for an initial routing, which accepts all the requests that can be accommodated with the given transport capacity. As a result, we have identified the subset of fictitious requests that we would be able to maintain in terms of the number of requests for each combination of time slot and delivery area (“reserved time windows”).
- Order promise: (3) Using EMSR, we compute fulfillment controls that assign a certain number of possible deliveries resulting from step (2) to expected requests of a certain value, maximizing the overall expected value of accepted requests. (4) Based on the capacities calculated in step (2) and fulfillment controls resulting from step (3), actual delivery requests are accepted or rejected, considering the availability of reserved time slots as well as the occurrence and value of still expected requests. As a result, we have the actual, final set of accepted requests (and their time windows).
- Order delivery: (5) Given the final set of accepted requests, we finalize routing. Deliveries are scheduled and sequenced in a cost-minimal way, considering the given time windows. (6) The final step is the physical delivery of orders.

Assuming independent planning for separate delivery days, this process can be repeated for each day. New information collected in the process is used to update the demand forecast and the travel times for future planning. Note that only step (4) needs to be executed online when the consumer is posing a request, while the other steps can be conducted offline in advance and thus allow for more complex solution approaches.

3.2 Demand Forecast (Step 1)

Step 1 is about forecasting demand based on historical order data. Only given an estimate of future demand, fulfillment controls can be computed in a way that maximizes expected revenue. For this purpose, the number of delivery requests expected for differentiable value segments has to be forecasted. In this section, we first present the consumer model used for the analysis documented in this paper. We then explain the proposed segmentation of demand based on expected order value according to discrete value buckets. Finally, we devote some remarks to the problem of censored demand observations. To merely demonstrate the potential of our approach, the concept of forecasting is demonstrated here using a simple generic method. The development of more sophisticated models and statistical methods for this particular purpose is a possible direction of future research.

The complexity of demand forecasting largely depends on the complexity of the assumed demand model. For instance, the multinomial logit model proposed by Yang et al. (2012) allows for the consideration of customers choosing dynamically between given alternatives. Such a model calls for the estimation of more parameters than one that assumes consumers to behave statically, requesting one option and aborting the request if the desired option is not available. To merely

demonstrate the potential of value-based demand fulfillment, we take a static view of consumers in this paper.

In our model, delivery requests $r_{b,d,s}$ arrive for an order in a value bucket b to a delivery area d and for a particular time slot s . Request acceptance depends on a control variable $a_{b,d,s}$. If the desired combination is not available ($a_{b,d,s} = 0$), the request is not accepted and the consumer will abort the order. When the request has been accepted, the order and its value are fixed – there are no cancellations. The number of accepted orders for one combination is denoted as $o_{b,d,s}$. This model suffices to illustrate the potential of value-based demand fulfillment; it may be improved and extended by further research as pointed out in the final section.

As a straightforward approach to demand forecasting, we propose to generate a set of discrete value buckets b , based on an analysis of the value of past orders. The number of defined buckets directly affects the level of detail with which demand fulfillment can be controlled. However, if the range of value is partitioned too finely, the problem of small numbers arises, as only few observations are available per bucket. The availability of time slots as decided in step (3) will depend on the bucket an order value falls into.

Given a parameter B describing the number of desired value buckets, the lower and upper bound of each value bucket b , L_b and U_b , can be defined as follows given a set of observed order values $v(o_i)$:

$$L_b = (\max(v(o_i))/B) \times (b - 1),$$

$$U_b = (\max(v(o_i))/B) \times b.$$

The result is a set of value buckets ranked by increasing value from $b = 1$ to B . An order o_i is assigned to value bucket b if $L_b < v(o_i) \leq U_b$.

When past orders have been recorded according to value bucket, delivery area, and time slot, future orders can be forecasted using any time series approach. As our model does not consider seasonal patterns or events of any kind, a simple exponential smoothing technique can be applied. The use of such techniques for revenue management demand forecasts is extensively explained in Cleophas et al. (2009). It entails calculating an updated forecast $f'_{b,d,s}$ by combining the old forecast $f_{b,d,s}$ with estimated requests $r'_{b,d,s}$

using a smoothing parameter α with $0 \leq \alpha \leq 1$:

$$f'_{b,d,s} = (1 - \alpha) \times f_{b,d,s} + \alpha \times r'_{b,d,s}.$$

However, note that as in classical revenue management, the problem of censored demand arises: Some requests may be aborted if the offered time slots do not fit the consumers' requirements. If the desired time slot had been available, these aborted requests would have been realized as orders in this time slot. Given that it was not available, the demand could not be observed, i.e., historical observations are censored.

To compensate for this phenomenon, a so-called unconstraining routine is needed, preparing the historical observations before they can be used to forecast expected demand. The static model of demand described here makes unconstraining comparatively easy: If theoretical acceptance exceeded the accepted orders, we assume that the demand corresponded to orders.

$$a_{b,d,s} > o_{b,d,s} \rightarrow r'_{b,d,s} := o_{b,d,s}.$$

If the acceptance limit was fully utilized, a compensation for censored demand can be calculated using the mean imputation method as described by Zeni (2001). This method refers to the mean of historically observed orders $\hat{o}_{b,d,s}$ as a lower-bound estimate of uncensored demand:

$$a_{b,d,s} = o_{b,d,s} \\ \rightarrow r'_{b,d,s} := \max(\hat{o}_{b,d,s}, o_{b,d,s}).$$

Note that a case in which accepted orders exceeded acceptance limits is not feasible.

3.3 Routing for Capacity Control and Final Routing (Step 2 & 5)

Vehicle routing procedures are used within steps 2 and 5. For step 2, they are utilized to decide whether to accept or reject fictitious requests in delivery tours for a given fleet. This is a simple way to ensure that transportation capacity is "reserved" where and when valuable requests are expected. For step 5, the final routing of actual requests is conducted, and vehicle routing procedures provide cost-minimum delivery tours.

To project when and where to reserve transport capacities in step 2, we use the DYN algorithm as proposed by Ehmke and Campbell (2014). As mentioned above, DYN is an "acceptance algorithm" aiming to maximize the number of orders that a logistics service

provider can accommodate while considering a given number of delivery vehicles. DYN works as follows. The input is a temporally ranked or value-ranked set of forecasted requests, which contain information on the expected delivery time slot s and delivery area d . Forecasted requests are processed in a sequential manner according to their ranking. Based on the set of acceptable orders A and the current delivery request $r_{b,d,s}$, a solution of the resulting time-dependent VRPTW is computed and checked for feasibility with respect to the given time windows and number of vehicles. If there is a feasible solution, i.e., if all time windows and the number of vehicles are satisfied, $r_{b,d,s}$ is considered as acceptable and added to A . As a result, A contains the set of acceptable (forecasted) requests that can be accommodated with given transport capacities. For further processing, we count the number of acceptable requests per time slot and delivery area and interpret them as reserved transport capacities $c_{d,s}$.

The DYN mechanism can be implemented with any (time-dependent) VRPTW heuristic. To reflect the impact of varying traffic conditions in metropolitan areas, we use the time-dependent routing framework proposed by Maden et al. (2010) in a variant with real travel time data by Ehmke et al. (2012a). Whenever a new request $r_{b,d,s}$ is posed, requests of set A and the new request are assigned to vehicles minimizing (1) the number of required vehicles and (2) overall travel times, considering time window restrictions. In particular, the parallel insertion heuristic by Potvin and Rousseau (1993) produces an initial solution. Standard metaheuristic operators as proposed by Maden et al. (2010) improve the initial solution until a given run time limit is reached. If the resulting solution satisfies the given number of vehicles, $r_{b,d,s}$ is accepted and added to A . This is repeated until all delivery requests have been processed. Note that DYN is usually an online problem that needs to be solved quickly. However, demand forecasts and transport capacities can be aligned offline in our setting.

We also use the above VRPTW framework to conduct the final routing in step (5), which schedules the actual set of orders that have been accepted by the EMSR heuristic in step (3). The objectives are the same as above. However, we ignore the given number of vehicles in the final routing, since we eventually need to fulfill all delivery requests that have

been accepted in step (3), because we are committed to them in any case.

3.4 Computing Fulfillment Controls and Order Acceptance (Step 3 & 4)

Using the demand forecast $f_{b,d,s}$ (step 1) and transport capacities $c_{d,s}$ (step 2) as input, acceptance controls $a_{b,d,s}$ can be calculated for each combination of value bucket b , delivery area d , and time slot s using a deterministic version of the EMSR heuristic described by Belobaba (1987). In applying EMSR to the context of this paper, we substitute the concept of fare classes for one of the value buckets. Following this concept, the acceptance of a delivery request depends on the acceptance limit assigned to the value bucket the order falls into.

For example, let there exist two value buckets including orders from 0 to 99 € vs. 100 € and above. Only for the more valuable bucket requests are acceptable according to plan. This means that a consumer with an order value of 100 € or above is offered the desired time slot, while a consumer with an order value of only 75 € cannot have his or her order delivered at the desired time slot.

The maximum number of orders to be accepted is calculated starting with the lowest value bucket $b = 1$. The demand expected for each bucket, $r'_{b,d,s}$, is compared to the demand expected for all more valuable buckets combined, $\sum_{i=b+1}^B (r'_{i,d,s})$. As transport capacity is limited, delivery requests from this bucket are only accepted if there is capacity left after the expected demand for more valuable buckets has been subtracted. The expected leftover capacity accordingly indicates the maximum number of orders to accept; the resulting limits are the acceptance controls:

$$a_{b,d,s} := c_{d,s} - \sum_{i=b+1}^B (r'_{i,d,s}).$$

Order acceptance is realized using serial nesting as described in Vinod (2006). This means that the condition for accepting an additional delivery request is that the new number of orders from this value bucket and all lower buckets does not exceed the computed limit. This condition can be formalized as follows:

$$\sum_{i=1}^b (o_{i,d,s}) + 1 \leq a_{b,d,s}.$$

4 Computational Study

The computational study documented in this section illustrates the potential of our integrated solution approach. After detailing the simulation model, we describe the results for three sets of experiments. First, we compute benchmarks for the success of vehicle routing and demand fulfillment, demonstrating that even for simple models and methods, order revenue can be significantly improved. As these first experiments rely on an unattainably accurate forecast, we go on to show the sensitivity of performance with regard to forecast accuracy. Finally, the third set of experiments illustrates the impact of demand composition on the potential of value-based demand fulfillment.

4.1 Simulation Model

The implemented simulation follows the discrete, event-based, stochastic paradigm. Each simulation run includes the following steps: After the offline calculation of capacity and availability, delivery requests are processed as individual events and are turned into orders upon acceptance. After all request events have been processed, the final vehicle routing is conducted. In the next simulation run, the same steps are iterated given a stochastically varied set of delivery requests.

Our simulation model refers to a metropolitan area that is divided up into nine delivery areas ($d := 1..9$). The number of delivery areas included is somewhat arbitrary as in this study, demand data is largely fictitious. However, as described below, we rely on income distributions, which were available for areas defined according to a segmentation of the area that follows the directions of the compass as shown in Fig. 2. In further applications of the proposed solution approach, the division of delivery areas would be subject to data analysis: The aim is to find a perfect balance between demand segmentation (with internally homogeneous areas) and data aggregation (with areas that are as large as possible). Figure 2 shows the spatial segmentation into geographically different delivery areas. The letters represent the center of gravities, which serve as fictitious locations for forecasted orders in the corresponding delivery area. Furthermore, eight time slots ($s := 1..8$) are

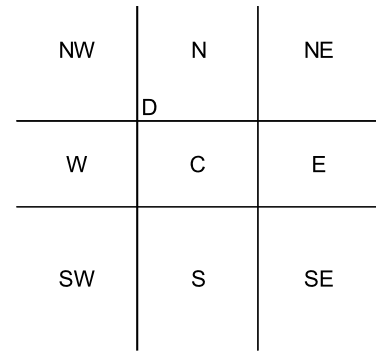


Fig. 2 Delivery areas with center of gravity (NW, N, NE, W, C, E, SW, S, SE) and depot location (D) for the fictitious application in the metropolitan area of Stuttgart. Delivery areas differ in size due to varying population size. The area C refers to the inner city of Stuttgart. The gravity center of each delivery area serves as a fictitious consumer address

available for delivery that is conducted by a fleet of four delivery vehicles.

Demand is generated using a Monte Carlo approach over 100 simulation runs. As demand is normally distributed, we provide confidence intervals based on the student distribution and a 5 %-error probability (see Kelton and Law 2000). For each run, the number of requests r is drawn from the normal distribution with a mean of 130 and 1 % variation. While variation was fixed arbitrarily to provide a reasonable degree of fluctuation, the mean number of requests to be generated was calibrated to an expected capacity-to-demand ratio of 2.5 given an expected capacity c of 52. This ensures competition for time slots, a condition for acceptance controls to be meaningful. Expected capacity was estimated from preliminary vehicle routing experiments.

To provide demand segmentation, delivery requests are created for one of three population segments P , distinguished by the average order value V_P . For the base demand scenario as used in Sects. 4.2 and 4.3, the three segments place orders that have an average value of 30, 50, and 70, respectively, with 10 % variation. Once more, the data used here is artificially designed to display both a useful degree of differentiation and variation. In real world applications, the distribution, mean, and variation of order values would be derived from an analysis of historical order data.

The desired time slot s and the delivery area d are linked to the population segment P . These characteristics are

drawn based on discrete probability distributions $p_s(P)$ and $p_d(P)$. This distribution was designed for each segment to display differences in location preference and daily schedule. **Tables 1** and **2** show the probability distributions of requests from each population segment occurring for a particular time slot s and delivery area d .

While the distributions used for the simulation study are artificial, their motivation is intuitive: Different urban areas house different proportions of socio-demographic population segments. From those, students can be expected to follow a daily schedule differing from that of professionals and pensioners, for example.

Consequentially, some combinations of time slots and area can expect more valuable demand than others: A relatively high share of demand can be expected for three of the eight delivery time slots, whereas a relatively low share of demand can be expected for another three. High value demand concentrates especially on one particular time slot, approximating a “common favorite hour of delivery”. The distribution of demand across delivery areas loosely follows data on income distributions for Stuttgart, as given in Landeshauptstadt Stuttgart (2012). This results in a distribution that is diverse enough to approximate the consequences of controlling demand fulfillment for a set of areas with varying income levels.

The result of these parameters is a set of conditional probabilities. Whenever a consumer is generated, the demand segment, time slot, and delivery area are determined based on a random number. When the delivery area has been set, the precise address is drawn from a uniform distribution across all addresses in this area. Then, the order value is drawn from another normal distribution described per segment by an expected value and deviation.

To include a basic test for sensitivity, we designed three demand scenarios, differing in the average order value of the included demand segments. The least valuable demand segment accounts for 70 %, the medium segment for 20 %, and the most valuable segment for 10 % of demand. We kept these values conservative as the success of value-based approaches can be expected to increase with the share of valuable demand. Two further demand scenarios have been designed for the purpose of the corresponding experiment presented in Sect. 4.4 as detailed there.

Table 1 Probability distributions for time slot choice per population segment $p_s(P)$

	s							
	1	2	3	4	5	6	7	8
$P = 1$	10 %	10 %	20 %	20 %	20 %	10 %	5 %	5 %
$P = 2$	5 %	5 %	5 %	15 %	20 %	20 %	20 %	10 %
$P = 3$	5 %	5 %	5 %	5 %	10 %	20 %	30 %	20 %

Table 2 Probability distributions for time slot choice per population segment $p_d(P)$

	d								
	1	2	3	4	5	6	7	8	9
$P = 1$	14 %	8 %	11 %	15 %	10 %	9 %	14 %	9 %	10 %
$P = 2$	13 %	13 %	11 %	8 %	10 %	9 %	16 %	6 %	14 %
$P = 3$	13 %	13 %	11 %	9 %	6 %	9 %	13 %	13 %	13 %

For vehicle routing, we consider time-dependent travel times from computational experiments that have been conducted for the area of Stuttgart, Germany, by Ehmke et al. (2012b) and are based on a large historical travel time database. The nodes of a digital roadmap represent actual consumer locations. For travel time determination, consumer locations are distinguished by two different zones, namely the downtown (area “C” in Fig. 2) and the suburban areas of Stuttgart. We thus capture varying proximity of consumers, which is much closer downtown than in suburban areas. Typical traffic patterns such as morning and afternoon peak hours with traffic jams on trunk roads are derived as follows. Straight-line distances between nodes are multiplied with a correction factor of 1.5, which is a rough approximation of a typical shortest path distance on the underlying detailed road network. These distances are then converted into travel times based on typical speeds within and between zones. Multipliers for the adjustment of downtown paths are provided by “cluster 6”, while multipliers for the adjustment of suburban paths are provided by “cluster 5” (Ehmke et al. 2012b). The clusters mainly comprise inner city and trunk roads, respectively, and thus serve as a simplified representation of a typical metropolitan road network. Multipliers reflect the typically varying extent of speed variation at different times of the day.

In **Table 3**, the variables of the computational study are summarized.

4.2 The Potential of Value-Based Demand Fulfillment

To demonstrate the potential of value-based demand fulfillment, we consider the following four simulation scenarios. All scenarios make use of the basic demand scenario as described in Sect. 4.1. They differ by the extent to which value-based demand fulfillment is realized, though:

- *No VB*: This scenario employs *no* value-based (VB) techniques. Capacity control is realized by vehicle routing based order acceptance mechanisms as detailed in Sect. 3.3, accepting or rejecting requests for a particular delivery area and time slot without regarding the order value. There is no demand forecast and accordingly, demand fulfillment is not controlled further.
- *VB with Perfect DF*: This scenario employs value-based (VB) capacity planning and fulfillment control techniques based on a perfect, deterministic forecast. Capacity planning is realized by considering forecasted requests in the sequence of their value, starting with the most valuable. Based on the resulting transport capacity, actual delivery requests are accepted in the sequence of their value, assuming flawless demand fulfillment controls.
- *VB Routing with DF*: This scenario employs value-based (VB) capacity planning and demand fulfillment (DF) based on a perfect, aggregated forecast: Based on our knowledge of the simulation set-up, we set $r'_{b,d,s} := r_{b,d,s}$ in each simulation run. As described in

Table 3 Summarizing the parameterization of demand and traffic in the simulation

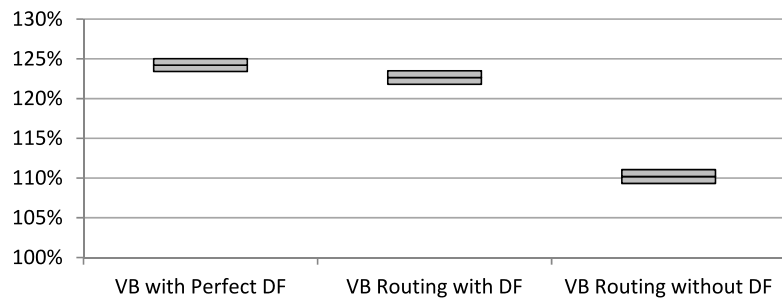
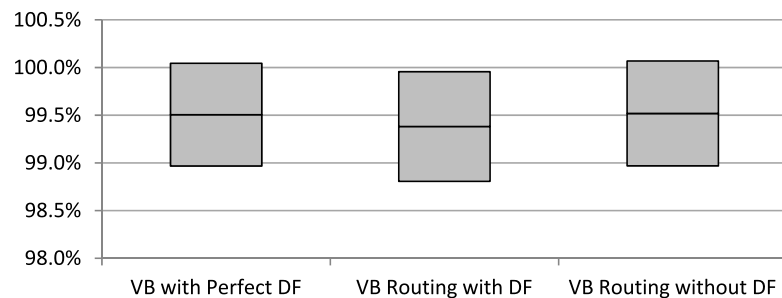
Variable	Distribution	Operationalization
Overall request volume r	Normal distribution with mean = 130, variation = 1 %	Set in proportion to capacity is 2.5:1 to ensure spill
Population segments P	3 population segments differentiated by average order value and probability of delivery area and time slot	Request volume ratio 7:2:1, delivery area and time slot probability distributions given in Tables 1 and 2
Order value	Normal distribution with mean varying per population segment P	Mean varies to account for three demand scenarios for sensitivity analysis, variation is 10 %
Requested time slot s	Empirical discrete distribution across time slot per demand segment	See Table 1
Delivery area d	9 areas following points of compass	See Table 2
Travel times	Empirical distribution	Following data analysis by Ehmke et al. (2012b)

Sect. 3.2, demand is forecasted for five value buckets segmenting the range of possible order values. Capacity control is realized by considering forecasted requests per bucket, starting with the most valuable bucket. Based on the resulting capacity, EMSR is used to calculate demand fulfillment, which is implemented as requests arrive sequentially.

- *VB Routing without DF*: This scenario employs only value-based (VB) capacity planning without demand fulfillment (DF). Once more, demand is forecasted for value buckets and capacity planning is realized by treating forecasted requests per bucket. The idea is to demonstrate how, even without further efforts to control demand fulfillment, value-based forecasts of demand can improve the results of capacity planning in terms of the resulting order value.

Figure 3 compares the overall value of fulfilled orders expressed in percent of the results derived from the *No VB* scenario. The results for the *VB with Perfect DF* scenario illustrate that given the underlying demand structure, potentially 24 % more value may be realized in terms of orders using value-based approaches. The results for the *VB Routing with DF* scenario are barely significantly inferior, however, note that while the forecast is aggregated to five value buckets, it is still unrealistically accurate. Finally, while the results for the scenario *VB Routing without DF* do not reach such levels, they still show that an additional 10 % of order value can be attained by merely planning transport capacity according to a forecast of demand value.

Since demand fulfillment controls limit the number of accepted delivery requests, the resulting transportation costs are reduced as well. As shown in **Fig. 4**, the working time generated by the vehicle

**Fig. 3** Potential of VB capacity planning and demand fulfillment: Value fulfilled in % of No VB**Fig. 4** Potential of VB capacity planning and demand fulfillment: Working time in % of No VB

routes tends to be less than what can be expected for the *No VB* scenario. As working time varies more than fulfilled value does, the differences between the scenarios are not significant. However, there seems to be a tendency for slightly lower working times when demand fulfillment controls are in place – this appears logical, as these controls limit the number of accepted orders most severely.

4.3 The Impact of Forecast Accuracy

To illustrate the impact of forecast accuracy, we consider some scenarios where the forecast has been deliberately skewed. We benchmark the following scenarios by

comparing them to the results of *VB Capacity Planning with DF*. Their set-up is the same with regard to capacity planning and demand fulfillment; it deviates only with regard to the forecast as described below:

- *High Volume FC*: Initially perfect forecast increased by 10 % across all value buckets.
- *Low Volume FC*: Initially perfect forecast reduced by 10 % across all value buckets.
- *High Value FC*: The value of forecasted orders increased by 10 % before aggregation into value buckets.
- *Low Value FC*: The value of forecasted orders reduced by 10 % before aggregation into value buckets.

Figure 5 illustrates that any skewing of the forecast decreases the overall value of fulfilled orders. The comparison is drawn by expressing value in terms of what might have been achieved using an accurate, aggregate forecast as done in *VB Capacity Planning with DF*.

Note that the scenarios *Low Volume FC* and *Low Value FC* do not result in as much of a loss of order value as do *High Volume FC* and *High Value FC*. The lowest value that can be expected from a too low forecast is that achieved using first-come-first-serve controls as seen in the *No VB* scenario. A too high forecast in terms of value or volume, however, leads to the rejection of orders that could actually have been accommodated. Thereby, an overestimation of demand can have an extremely negative impact on performance.

The above argument is underlined by a comparison of the number of accepted orders resulting from the skewed forecast scenarios. **Figure 6** shows these orders in percent of the *No VB* scenario. While an underestimation of demand volume will lead to nearly the same number of orders as the *No VB* scenario, the overestimation of value decreases the number of accepted orders by more than 35 %. This lack of order volume cannot be compensated by actual order value – performance falls short of expectations.

4.4 The Impact of Demand Composition

To show the impact of different demand set-ups, we consider two alternative demand scenarios as described in Sect. 4.1: Heterogeneous and homogeneous demand. For each demand scenario, we compute orders as given by the set-up of *VB Capacity Planning with DF*.

- *Homogeneous Demand*: For this scenario, the average expected value of orders for the three demand segments was adjusted to render the demand segments more similar. The value of individual orders is now drawn from three normal distributions displaying average values of 40, 50, and 60, respectively. As a result, the variation coefficient of the demand for this scenario is 0.18.
- *Heterogeneous Demand*: For this scenario, the average expected value of orders for the three demand segments was adjusted to render the demand segments less similar. The value of individual orders is now drawn from three normal distributions displaying

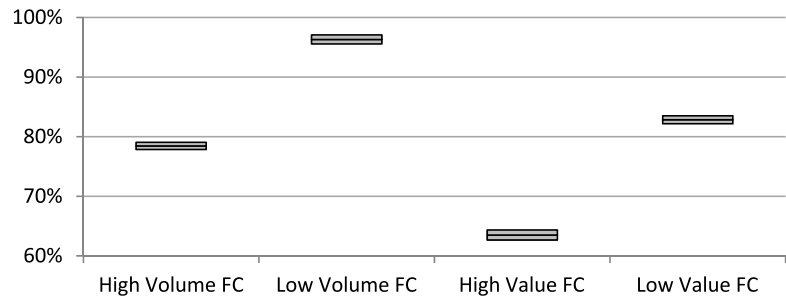


Fig. 5 Impact of forecast accuracy: Value fulfilled in % of VB Capacity Planning with DF

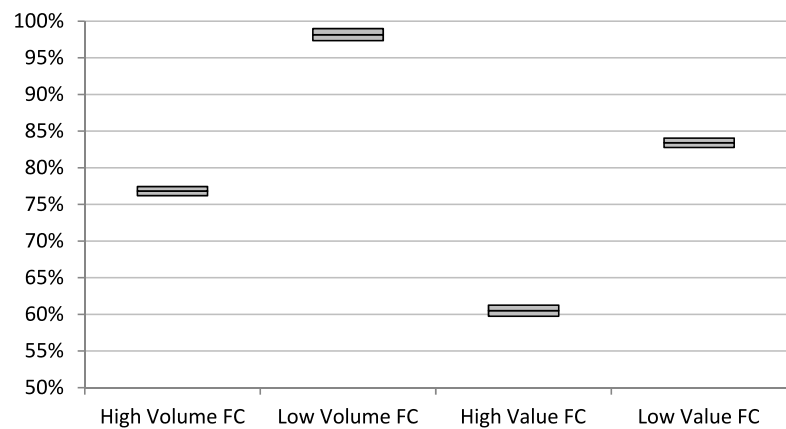


Fig. 6 Impact of forecast accuracy: Orders in % of VB Capacity Planning with DF

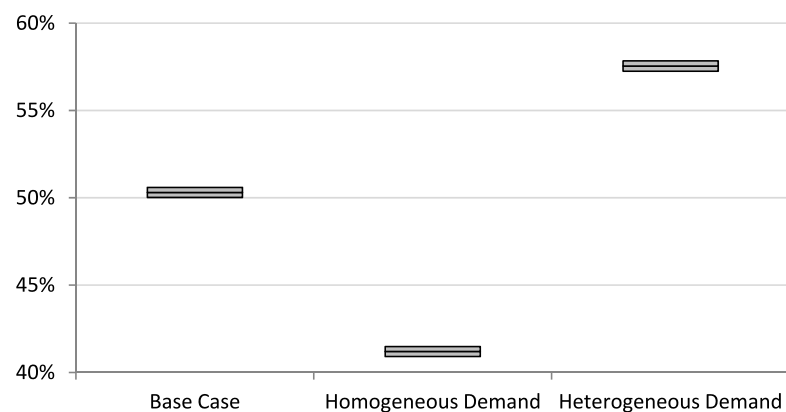


Fig. 7 Impact of demand composition: Value fulfilled in % of potential

average values of 30, 50 and 70, respectively. As a result, the variation coefficient of the demand for this scenario is 0.63.

As the resulting order value is also determined by the potential given by the demand scenario, a straightforward comparison of results can be misleading in this case. Accordingly, **Fig. 7** shows the impact of demand composition in terms of percent of potential value. For each scenario, the overall potential value has

been calculated by summing up the value of demand generated for the simulation. This was used to scale the value of accepted orders.

Figure 7 supports the following intuitive assumption: The more heterogeneous the demand is, the more potential exists for the application of value-based demand fulfillment. Even for a highly accurate forecast, less than 45 % of potential value can be turned into orders when demand is homogeneous. Yet when de-

Abstract

Catherine Cleophas, Jan Fabian Ehmke

When Are Deliveries Profitable?

Considering Order Value and Transport Capacity in Demand Fulfillment for Last-Mile Deliveries in Metropolitan Areas

The paper aims to optimize the final part of a firm's value chain with regard to attended last-mile deliveries. It is assumed that to be profitable, e-commerce businesses need to maximize the overall value of fulfilled orders (rather than their number), while also limiting costs of delivery. To do so, it is essential to decide which delivery requests to accept and which time windows to offer to which consumers. This is especially relevant for attended deliveries, as delivery fees usually cannot fully compensate costs of delivery given tight delivery time windows. The literature review shows that existing order acceptance techniques often ignore either the order value or the expected costs of delivery. The paper presents an iterative solution approach: after calculating an approximate transport capacity based on forecasted expected delivery requests and a cost-minimizing routing, actual delivery requests are accepted or rejected aiming to maximize the overall value of orders given the computed transport capacity. With the final set of accepted requests, the routing solution is updated to minimize costs of delivery. The presented solution approach combines well-known methods from revenue management and time-dependent vehicle routing. In a computational study for a German metropolitan area, the potential and the limits of value-based demand fulfillment as well as its sensitivity regarding forecast accuracy and demand composition are investigated.

Keywords: Demand fulfillment, Vehicle routing, Revenue management, Simulation, City logistics

mand is heterogeneous, more than 55 % of potential value can be turned into actual orders. Note that this observation holds in spite of a constant ratio of demand to capacity across all scenarios.

5 Conclusions and Outlook

For the near future, we expect that order acceptance systems integrating traffic information, demand forecasting, and vehicle routing functionality will become more and more relevant due to the ongoing, increasing competition in e-commerce. Especially transportation costs in crowded metropolitan areas cannot be neglected anymore. Thus, in this paper, we have combined well-known techniques of revenue management and time-dependent vehicle routing in order to improve the profitability of attended home delivery applications in metropolitan areas. To this end, we have demonstrated how vehicle routing and demand fulfillment controls can be integrated in order to reserve transport capacities where and when they are expected to be profitable and cost-efficient.

Computational experiments underline challenges and risks of combining revenue management and vehicle routing techniques for value-based demand fulfillment. An accurate forecast is mandatory to improve revenues. Given an accurate forecast, the potential benefits with regard to overall delivered order value are significant, while transportation costs do not get out of control. However, this finding is subject to two caveats: On the one hand, the quality of the forecast must be carefully controlled so as not to fall prey to an overestimation of expected demand. As shown, if the expected value or volume of demand is too high, restrictive fulfillment controls can considerably hurt performance. On the other hand, the success of value-based fulfillment depends on an existent segmentation of demand. If demand is homogeneous and cannot be systematically differentiated by value, the potential of the approach proposed in this paper is limited. Also note that the considered fleet size should be aligned with the expected demand in order to alleviate rejection of too many requests and to ensure that transportation resources can be used efficiently.

Future research directions appear as follows. Models of consumer behavior included in the solution approach need to be extended to advanced e-commerce

business models. For the sake of simplicity, we did not model a flexible choice of time slots according to consumer-individual priorities of time windows. In practice, however, e-commerce businesses offer a choice of differently priced delivery time slots. Furthermore, when deciding on the acceptance of an order, one may not only want to investigate its short-term value, but also consider the history of sales and potential future requests of a consumer, i.e., the long-term consumer value.

With respect to revenue management techniques, more sophisticated approaches than EMSR, such as the implementation of bid price controls and/or dynamic programming, would allow for a more finely grained control of value-based consumer acceptance. However, this would require a somewhat more complex forecast, e.g., a distribution on value and number of expected consumers. Furthermore, more advanced techniques of vehicle routing would allow for a detailed consideration of expected demand and consumer locations, e.g., by including information on the likelihood of an individual consumer posing a request. While we concentrated on the impact of the forecast's accuracy on solution quality, we did not investigate the impact of the routing technique on the quality of value-based demand fulfillment. Instead, we utilized a state-of-the-art time-dependent routing framework. Future work could explore to what extent the success of demand fulfillment depends on the particular quality of the vehicle routing technique and the quality of the underlying travel time data.

Finally, our contribution merely intends to demonstrate the potential of even simple first steps at combining the two areas of vehicle routing and value-based order acceptance. Apart from applying more advanced techniques from these areas, a parallel rather than iteratively integrated solution of the described problem is conceivable.

References

- Agatz N, Campbell AM, Fleischmann M, Savelsbergh M (2011) Time slot management in attended home delivery. *Transportation Science* 45(3):435–449
- Allen J, Thorne G, Browne M (2007) BESTUFS good practice guide of urban freight transport. Manual. http://www.bestufs.net/download/BESTUFS_II/good_practice/English_BESTUFS_Guide.pdf. Accessed 2013-06-30

- Baldacci R, Mingozzi A, Roberti R (2012) Recent exact algorithms for solving the vehicle routing problem under capacity and time window constraints. *European Journal of Operational Research* 218(1):1–6
- Belobaba PP (1987) Air travel demand and airline seat inventory management. Dissertation, Flight Transportation Laboratory, Massachusetts Institute of Technology, Cambridge
- Bräysy O, Gendreau M (2005a) Vehicle routing problem with time windows. Part I. Route construction and local search algorithms. *Transportation Science* 39(1):104–118
- Bräysy O, Gendreau M (2005b) Vehicle routing problem with time windows. Part II. Metaheuristics. *Transportation Science* 39(1):119–139
- Campbell AM, Savelsbergh M (2005) Decision support for consumer direct grocery initiatives. *Transportation Science* 39(3):313–327
- Cleophas C, Frank M, Kliewer N (2009) Recent developments in demand forecasting for airline revenue management. *International Journal of Revenue Management* 6(3):252–269
- Donati AV, Montemanni R, Casagrande N, Rizoli AE, Gambardella LM (2008) Time dependent vehicle routing problem with a multi ant colony system. *European Journal of Operational Research* 185(3):1174–1191
- Ehmke JF, Steinert A, Mattfeld DC (2012a) Advanced routing for city logistics service providers based on time-dependent travel times. *International Journal of Computational Science* 3(4):193–205
- Ehmke JF, Meisel S, Mattfeld DC (2012b) Floating car based travel times for city logistics. *Transportation Research Part C: Emerging Technologies* 21(1):338–352
- Ehmke JF, Campbell AM (2014) Customer acceptance mechanisms for attended home deliveries in metropolitan areas. *European Journal of Operational Research* 233(1):193–207
- Fleischmann B, Gietz M, Gnutzmann S (2004) Time-varying travel times in vehicle routing and scheduling problem with time windows. *European Journal of Operational Research* 66(3):331–340
- Figliozzi MA (2009) A route improvement algorithm for the vehicle routing problem with time dependent travel times. In: Proceedings of the 88th transportation research board annual meeting, Washington, DC
- Gevaers R, van de Voorde E, Vanelslander T (2010) Characteristics and typology of last-mile logistics from an innovation perspective in an urban context. In: Proc of WCTR 2010, Lisbon, Portugal
- Haghani A, Jung S (2005) A dynamic vehicle routing problem with time-dependent travel times. *Computers & Operations Research* 32(11):2959–2986
- Hahn GJ, Kuhn H (2012) Designing decision support systems for value-based management: a survey and an architecture. *Decision Support Systems* 53:591–598
- Hashimoto H, Yagiura M, Ibaraki T (2008) An iterated local search algorithm for the time-dependent vehicle routing problem with time windows. *Discrete Optimization* 5(2):434–456
- Ichoua S, Gendreau M, Potvin J-Y (2003) Vehicle dispatching with time-dependent travel times. *European Journal of Operational Research* 144(2):379–396
- Kelton WD, Law A (2000) Simulation modeling and analysis. McGraw Hill, Boston
- Kok AL, Hans EW, Schutten JMJ (2012) Vehicle routing under time-dependent travel times: the impact of congestion avoidance. *Computers & Operations Research* 39(5):910–918
- Landeshauptstadt Stuttgart (2012) Kleinräumige Einkommensverteilung in Stuttgart 2009. <http://www.stuttgart.de/item/show/305805/1/publ/20908>. Accessed 2013-06-30
- Littlewood K (1972) Forecasting and control of passenger booking. In: AGIFORS symposium proceedings
- Maden W, Eglese R, Black D (2010) Vehicle routing and scheduling with time-varying data: a case study. *Journal of Operations Research Society* 61:515–522
- Potvin JY, Rousseau JM (1993) A parallel route building algorithm for the vehicle routing and scheduling problem with time windows. *European Journal of Operational Research* 66(3):331–340
- Punakivi M, Saranen J (2001) Identifying the success factors in e-grocery home delivery. *Journal of Retail and Distribution Management* 29(4):156–163
- Quante R, Meyr H, Fleischmann M (2009) Revenue management and demand fulfillment: matching applications, models, and software. *OR Spectrum* 31:31–62
- Stadtler H (2005) Supply chain management and advanced planning – basics, overview and challenges. *European Journal of Operational Research* 163(3):575–588
- Talluri KT, Van Ryzin GJ (2004) The theory and practice of revenue management, vol 68. Springer, Heidelberg
- US Census Bureau (2012). Quarterly retail e-commerce sales 2nd quarter 2012. News release. http://www.census.gov/retail/mrts/www/data/pdf/ec_current.pdf. Accessed 2012-09-30
- Wirthman L (2013) Amazon, Ebay, Walmart same-day deliver, but should you? *Forbes BrandVoice*. <http://www.forbes.com/sites/ups/2013/04/04/amazon-ebay-walmart-same-day-deliver-but-should-you/>. Accessed 2013-06-30
- Yang X, Strauss AK, Currie C, Eglese R (2012) Choice-based demand management and vehicle routing in e-fulfilment. Working paper, University of Warwick. <http://www2.warwick.ac.uk/fac/soc/wbs/subjects/orms/about/people/strauss/publications/jointdemandmanagementfinalplainr1.pdf>. Accessed 2013-10-15
- Vinod B (2006) Advances in inventory control. *Journal of Revenue and Pricing Management* 4(4):367–381
- Zeni R (2001) Improved forecast accuracy in airline revenue management by unconstraining demand estimates from censored data. Dissertation. State University of New Jersey