A simulation framework to evaluate urban logistics schemes

W.J.A. van Heeswijk[a,b], M.R.K. Mes[a], J.M.J. Schutten[a], and W.H.M. Zijm[a]

[a]Department of Industrial Engineering and Business Information Systems, University of Twente
[b]Department of Intelligent and Autonomous Systems, Centre for Mathematics and Computer Science

June 21, 2018

Abstract

The domain of urban freight logistics is becoming increasingly complex. Many urban supply chains are composed of small and independent actors; for such actors it is particularly difficult to efficiently organize the highly fragmented supply chains that they must deal with. As a consequence, a high number of transport movements is required to satisfy the high service level demands of receivers, which has negative effects on the environment as well as on the quality of life in urban areas. Both companies and local administrators try to improve transport efficiency and reduce external costs, but the effects of such interventions are difficult to predict, especially when applied in combination with each other (an urban logistics scheme). This paper proposes an agent-based simulation framework that quantifies the effects of urban logistics schemes on multiple actors. We provide a detailed mathematical representation of the framework. The working of the framework is demonstrated with numerical experiments that evaluate a variety of urban logistics schemes. We show that most schemes yield significant environmental improvements; achieving financial sustainability, however, is challenging. Also, interventions such as subsidies and access restrictions do not always have the intended effect. Our experiments underline the applicability of the framework to quantify the effects of urban logistics schemes.

1 Introduction

Throughout the world, ongoing urbanization results in an ever increasing consumption of goods in urban areas, resulting in large flows of goods towards cities (MDS Transmodal Limited 2012). Due to competition with e-commerce, the roles of retailers have changed over time. Many retailers have adopted Just-in-Time ordering principles to cut their on-site inventories, thus requiring more frequent and accurately timed deliveries. Another development is the resurrected interest in small, specialized shops, seemingly reversing the trend of ever-expanding chains of mega-stores. The resulting fragmentation of freight flows makes it challenging for carriers to efficiently conduct their transport (Dablanc 2011).

Inefficiently organized freight flows have hazardous environmental and societal effects, e.g., more traffic jams, higher emissions, more noise pollution, and reduced traffic safety (Ploos Van Amstel 2015). Both governments and companies try to find solutions in an effort to reduce the negative impacts of urban freight logistics. For example, local administrators may impose access restrictions on heavy trucks or set zone access fees. Companies may adopt new methods and resources to improve the efficiency of their operations. Smaller companies are often limited by the scale of their operations; often collaboration is necessary to significantly reduce costs. The solutions of companies and governments may affect each other, e.g., a minimum load factor imposed by the local administrator may direct carriers towards horizontal collaboration efforts. It is therefore imperative that interventions are analyzed in conjunction with each other. In the context of this paper, we define an urban logistics scheme as a set of company-driven initiatives and administrative policies, with the aim to improve efficiency and/or reduce environmental and societal costs. The focus of this paper is on small actors (e.g., shippers, carriers, receivers); such actors are often unable to efficiently organize their logistics activities internally due to a lack of volume, and need to cooperate to achieve economies of scale.

Despite the eminent need for a better organization of urban logistics, the vast majority of schemes implemented in practice fails after a short life-span (Browne et al. 2005). A major reason for this inefficiency is that the involved actors typically have divergent objectives (Bektaş et al. 2015). Furthermore, administrators often provide subsidies to support initiatives, but actors typically revert back to their former decisions once subsidies are halted (Ploos Van Amstel 2015). A last shortcoming is that analysis before implementation often focuses solely on the processes within the city boundaries, ignoring the considerable impact of upstream decisions. The framework presented in this paper explicitly addresses these aspects.

A key success factor for urban logistics schemes is the right combination of company involvement and governmental policies (Browne et al. 2005). Actors must be willing to permanently change their behavior, without requiring perpetual external cash flows. Traditional optimization techniques may be used to find viable system-wide solutions, yet these are not guaranteed to be stable when their success depends on multiple decision makers. Agent-based simulation (ABS) is suitable to evaluate such schemes, as it is capable of monitoring and altering the behavior of autonomous actors under conditions that may be flexibly adjusted (Taniguchi et al. 2014). We present an ABS framework to evaluate the effectiveness of urban logistics schemes that include both governmental policies and company-driven initiatives.

The remainder of this paper is structured as follows. In Section 2, we contrast our work to both other agent-based studies and alternative quantitative solution methods. Section 3 describes our ABS framework. Section 4 presents numerical experiments that both demonstrate the framework and provide insights into urban logistics schemes. We end with the main conclusions in Section 5.

2 Literature on evaluating urban logistics initiatives

As mentioned in Section 1, both companies and administrators actively intervene in urban freight flows to improve the cost efficiency and environmental impact of urban freight logistics. Quak (2011) distinguishes four classes of initiatives in urban logistics: (i) physical infrastructure initiatives (e.g., consolidation centers, special lanes for electric vehicles), (ii) transport-reorganizing initiatives (e.g., using alternative transport modes), (iii) governmental policies (e.g., vehicle access restrictions, zone access fees, subsidization), and (iv) company-driven initiatives (e.g., improved routing algorithms, better use of real-time data, collaboration structures). To quantify the impact of such initiatives, a variety of operations research methods might be applied. For literature reviews on modeling efforts in urban logistics, we refer to Anand et al. (2012) and Bektaş et al. (2015). Relatively few papers in the context of urban logistics adopt an operations research perspective (Crainic et al. 2009). This section gives a non-exhaustive overview.

The Vehicle Routing Problem (VRP) is at the base of many transport problems studied in literature. Also in the context of urban logistics, a number of VRP variants exist. Cattaruzza et al. (2015) and Kim et al. (2015) provide overviews of VRP solution methods in urban logistics; we highlight a few applications in various problem classes. A first variant of applications embeds time-dependent travel times, since urban areas are often subject to congestion during peak hours. Examples of works on this variant are Van Woensel et al. (2008) and Kok et al. (2010). A second class is the VRP with access time constraints that allow trucks to drive within the city only within designated time intervals, see Muñuzuri et al. (2012) for an example. The final variant that we mention is the emission routing problem, which aims to minimize emissions within certain constraints (Bektaş et al. 2016). Recent works on this topic are Suzuki (2011), Ehmke et al. (2016), and Cheng et al. (2016).

Urban consolidation centers (UCCs) enable to decouple and bundle inbound freight flows. They have an important role in many urban logistics initiatives. Many cities have at most a single UCC, yet systems with more facilities – and sometimes multiple layers of facilities – exist as well. The multi-echelon distribution system is a general description for such structures; Cuda et al. (2015) provide a recent literature review on planning methods in such distribution systems. Some notable works on logistics planning in (multi-)echelon distribution systems are Crainic et al. (2009), Hemmelmayr et al. (2012), Baldacci et al. (2013), and Dellaert et al. (2016).

The solution methods mentioned so far are designed from the perspective of isolated decision makers (typically carriers); as such they are not able to capture system-wide effects and autonomous decision-making by multiple actors. To address this issue, Macharis (2007) describes the use of multi-criteria, multi-actor evaluation in the field of logistics. This method defines an objective function for every actor involved – often comprising both financial and non-financial components – and computes the impact of schemes for each actor involved. Exemplary applications of such evaluations may be found in Macharis et al. (2012), Stathopoulos et al. (2012), and Verlinde et al. (2014). Although multi-actor evaluation considers the system-wide impact of schemes and illuminates discrepancies, it does not specify the decision policies of the involved actors and possible interactions between them. Therefore, it does not allow to analyze evolutionary systems in which actors can adjust their behavior based on a changing state of the system.

A solution method that does take into account the strategic interactions between various actors is game theory. In particular, cooperative game theory is relevant to study collaboration initiatives in urban logistics. Examples of game-theoretical applications in logistics are Reyes (2005), Yang and Odani (2007), Dahlberg (2015), and Hafezalkotob and Makui (2015). Although game theory provides a framework to model and evaluate the interaction between multiple actors, it also has several drawbacks. Most pronounced is the large computational effort required to solve instances with many actors and the inflexibility to model different schemes are problematic when large numbers of schemes need to be evaluated.

Taniguchi et al. (2014) and Anand et al. (2016) state that ABS is the most applicable method to study the behavior of and interaction between the various actors in urban logistics. We mention some notable ABS studies in the field of urban logistics. Taniguchi et al. (2007) present an agent-based model in which carriers use a routing algorithm that dynamically adjusts to the prevailing travel times. By adopting this more advanced routing algorithm, both carriers and shippers are able to benefit from the resulting cost reductions, while emissions are also reduced. A prerequisite for success is that a gain-sharing mechanism is installed, such that both carriers and shippers profit from the efficiency gains. Tamagawa et al. (2010) perform an ABS study in which they heuristically solve a VRP and iteratively update the actions of five agent types by means of learning mechanisms. The numerical experiments revolve around the effects of road pricing and truck bans. Van Duin et al. (2012) investigate the financial model and the environmental impact of UCCs, taking into account UCC service fees, road pricing, and subsidizing. Based on numerical experiments, they conclude that for a UCC to be financially viable, additional revenues from permanent subsidies or value-adding services are required. Wangapisit et al. (2014) evaluate the use of consolidation centers by imposing parking constraints and providing subsidies to carriers. They state that these interventions significantly contribute to the reduction of emissions. Teo et al. (2014) study the effects of road pricing and a load factor control system, using customer complaints on late deliveries as a driver to alter the selection of carriers. The authors claim that the studied interventions have a positive environmental effect. A recent line of research (see, e.g., Anand et al. (2016) and Marcucci et al. (2017)) emphasizes the importance of validating the agent descriptions with real-life counterparts, possibly even directly deriving utility functions from their input. We are unaware of any realistic-sized models of this type; although improving the credibility of the model, such a validation step is highly time-consuming.

The contribution of this paper is twofold. First, in contrast to the ABS studies mentioned in this literature overview, we take into account the transport process outside the city, better reflecting decisions made by shippers and carriers. Second, we explicitly include various forms of cooperation between companies, while existing studies tend to have a strong focus on testing governmental policies. As practice shows that successful schemes require both administrative policies and company-driven initiatives, a framework that explicitly includes both aspects is essential for proper evaluation of these schemes. Van Heeswijk et al. (2016) sketch a preliminary version of this framework. In this paper, we present a rigorous mathematical formulation of the model that allows others to also implement our framework. Furthermore, we perform a numerical study with the framework and provide the key insights for the tested urban logistics schemes.

3 Framework design

In this section, we outline the design of our ABS framework. The mathematical formulation of the framework provides concrete guidelines for its implementation. To not overly complicate the mathematical presentation, we introduce the following simplifications, which may easily be omitted in actual implementations of the framework: all routes are completed within a single time interval in the model, pickup routes outside the city are not specifically described, and the network contains a single UCC location in a single city.

Figure 1 shows a high-level overview of the framework. Decisions are divided into three levels: strategic (long-term), tactical (medium-term), and operational (short-term). Strategic decisions about for example administrative policies are fixed for the complete simulation run, which typically represents multiple years. At the tactical level (e.g., decisions updated every month), costs, prices and subsidies are updated; subsequently, decisions such as outsourcing to the UCC are made. At the operational level (e.g., daily), orders are placed and agent make decisions like dispatching and routing. The supply chain is structured as follows: (i) receivers communicate their order placements to the shippers, (ii) shippers dispatch a subset of the orders requested by receivers, (iii) carriers transport the orders that are dispatched by shippers, (iv) the UCC receives a subset of the orders transported by the carriers, and (v) the UCC dispatches a subset of orders in inventory for delivery.

The outline of this section is as follows. We describe the roles of the agent types in Section 3.1. In Section 3.2, we formally define the state of the system. Section 3.3 introduces the operational decisions per agent type. The transition function that describes the state change over time is given in Section 3.4. We define the cost functions for each agent type and the environmental performance indicators in Section 3.5. In Section 3.6, we discuss various solution methods that may be applied within the framework on the levels of strategic, tactical, and operational decision-making.

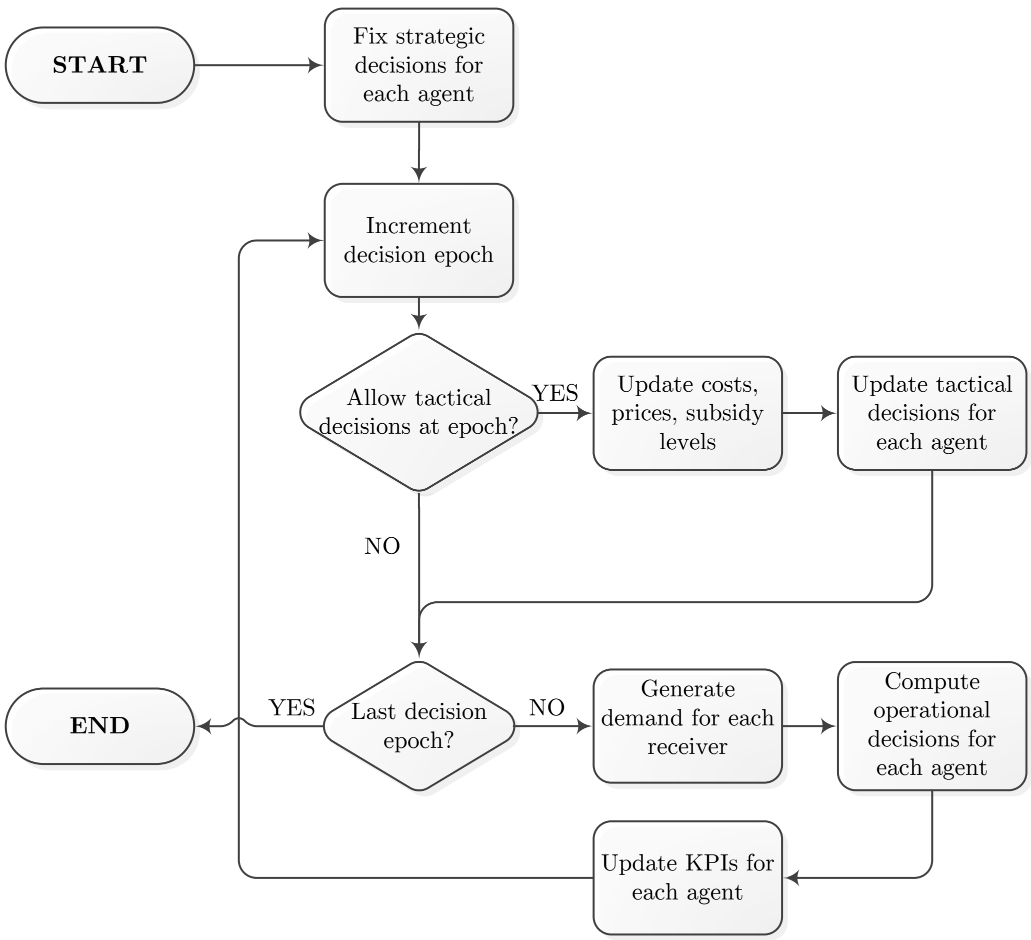

Figure 1: Conceptual flowchart of the simulation framework.

3.1 Agent types

We design the ABS framework in a manner that allows to simultaneously evaluate company-driven initiatives and administrative policies. Five types of agents are distinguished: receivers, shippers, line-haul carriers, the UCC operator, and the administrator. We proceed to briefly describe their roles; Figure 2 shows the actions, monetary flows, and information flows between the agents, whereas Table 1 summarizes the roles and main actions of each agent type. The actions of agents and information exchanges between agents are defined in detail in Section 3.3 and Section 3.4. The cost functions corresponding to the monetary flows are defined in Section 3.5. We note that the real-life counterparts of the agents are not necessarily rational decision makers, particularly when solutions require drastic changes in behavior. However, the simulation results yield insights into the effects on real-life agents if they would change their behavior, as such providing directions for change.

The **receivers** (denoted by the set $\mathcal{R}$) generate order demands subject to some (stochastic) process. Receivers may order from multiple suppliers at a single decision epoch. For each receiver we may define fixed epochs at which orders may be placed (e.g., two epochs per week); these ordering patterns implicitly reflect factors such as internal consolidation, storage costs, and stockouts. This implies that no interventions can be taken to improve internal consolidation at the receiver level. Although arguably being a strong assumption, modeling order optimization at the receiver level would require many assumptions with respect to their operations as well. Both to prevent indefinite postponement of order shipping and to accurately model practice, receivers specify a delivery window for each order. In the standard situation, orders are directly delivered by the carrier to the address of the receiver, but a receiver can also opt to select the UCC as its delivery address. In that case, the UCC is responsible for the last-mile distribution. Besides the costs associated with physically receiving goods, receivers also dedicate time to value-adding services, e.g., putting clothes on hangers or unpacking bulk shipments to have shelf-ready products. These services may be outsourced to the UCC as well.

The **shippers** in the framework are denoted by the set $\mathcal{S}$. They act on incoming orders from the receivers and hire carriers to transport orders. As carriers charge relatively lower prices for higher volumes (i.e., a lower price per volume unit), shippers have an incentive to bundle multiple orders before shipping. However, the shippers should dispatch the orders in time, such that the carrier is able to meet the delivery windows. For the sake of generality, we have separated the roles of the carrier and the shipper. In practice, the shipper and carrier might be the same actor. The shipper then has a similar incentive, as consolidation allows a more efficient utilization of its transport resources.

Line-haul **carriers**, denoted by the set $\mathcal{C}$, pick up goods at the shippers, and transport them either directly to the receivers or to the UCC. The line-haul arc is represented by an entrance vertex (near the pickup area) and an exit vertex (near the city). They may decide to outsource their last-mile distribution to the UCC completely when this yields a financial benefit or is enforced by regulation. In some cases, carriers may also deliver individual orders to the UCC. When the shipper dispatches an order before the earliest delivery time, the carrier must outsource this order to the UCC, as the carriers do not offer intermediate storage capacity. If the receiver sets the UCC as its delivery address, the carrier must use the UCC as well for all orders destined to that receiver, but in that case it is the receiver who pays for the last-mile delivery service. The carrier uses a price function based on volume and line-haul distance (i.e., the distance between cities, ignoring distance variations due to routing) that reflects economies of scale. In the typical setting that we study, a dispatched truck will visit multiple cities during a single tour, yet we focus on a single city only. Consequently, the load destined for the city is often considerably less than the truck's capacity.

The **UCC** receives goods from the line-haul and is responsible for their last-mile distribution. We use a generic set notation $\mathcal{H}$ for UCCs, but for simplicity only consider a single UCC for the outline of the framework. Goods are transported via the UCC if either the carrier outsources its last-mile delivery or the receiver selects the UCC as its delivery address. At the UCC, orders stemming from multiple carriers are bundled and may be temporarily held to account for future consolidation opportunities. The UCC uses its own vehicles to perform the last-mile distribution. For the UCC to be competitive, the efficiency gains on the last-mile ideally compensate for the necessary reload costs. Additional revenues might be obtained from performing value-adding services or redistributing external costs (e.g., by taxes and subsidies).

Finally, the **administrator** implements governmental policies to influence the behavior of agents. Since such policies are typically implemented for a longer time, we model them as strategic decisions. The financial gains stemming from these policies might be redistributed to subsidize agents. As mentioned, many schemes fail due to being overly dependent on subsidies. We therefore look for schemes that are financially sustainable without requiring permanent external cash flows. One way to achieve this involves temporary subsidies, paid until a new equilibrium state of the system is reached. Another sustainable solution can be achieved when the income from specific policies suffices to cover the subsidy expenses. In the latter case, the administrator essentially monetizes and redistributes external costs.

Table 1: Summary of agent types.

Agent type	Description	Main actions
Receiver	Located in urban area, places orders with shippers.	- Place orders with shippers (based on stochastic process) - Decide whether to select UCC for last-mile delivery - Decide whether to outsource value-adding services to UCC
Shipper	Facility outside urban area, ships based on receiver orders.	- Select carrier for transport - Time dispatch of shipments
Carrier	Conducts transport on request of shipper, starts routes from depot outside urban area.	- Perform transport (pickup, line-haul, and delivery) - Decide whether to (partially) outsource transport
UCC	Located at edge of urban area, conducts last-mile distribution.	- Time dispatch of orders - Perform last-mile distribution - Perform value-adding services for receivers
Administrator	Reduce emissions and freight traffic in urban areas by intervening.	- Impose vehicle access restrictions - Collect fees stemming from policies - Allocate subsidies

3.2 State of the system

In this section, we provide the notation required to define the state of the system and the actions of the agents. We use discrete-event simulation to model the behavior of the system over a sequence of equidistant time intervals that typically represent several hours up to a day. Let $\mathcal{T} = \{0, 1, \ldots, T\}$ be the set of discrete decision epochs. As stated before, we distinguish between three levels of decision making. Strategic decisions are fixed at $t = 0$ and remain in effect until T. Tactical decisions are made at designated tactical decision epochs $t \in \mathcal{T}^{tac} \subset \mathcal{T}$. Operational decisions are made at every $t \in \mathcal{T}$.

We define an abstract network consisting of a pickup area with shipper locations and the home

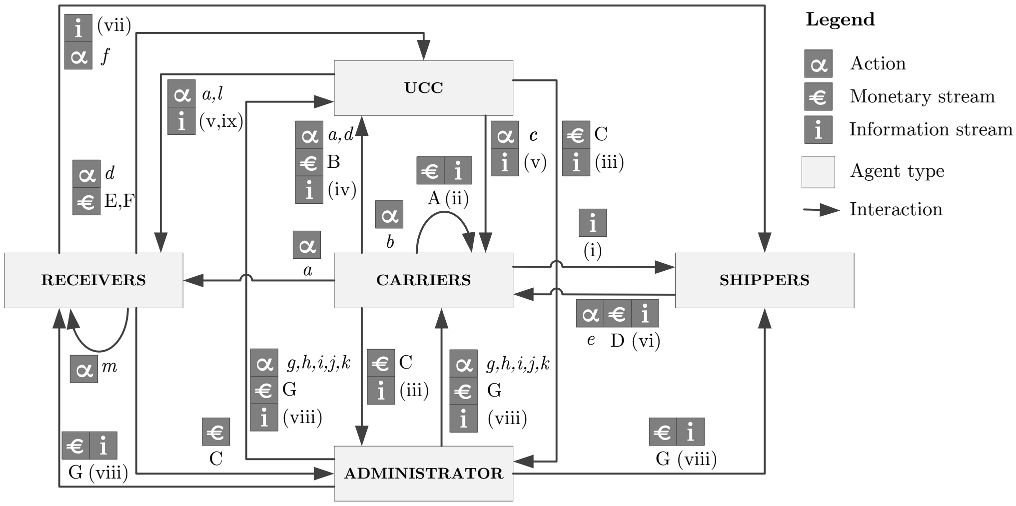

ID	ACTION
a	Transport orders
b	Combine transport
c	Enforce advance notice
d	Select UCC as delivery address
e	Release orders
f	Place orders
g	Set costs road pricing
h	Set tariff for urban zone
i	Set minimum load size
j	Set access time restrictions
k	Set vehicle access restrictions
l	Perform value-adding services
m	Synchronize delivery windows

ID	MONETARY FLOW
A	Gain-sharing joint transport
B	Costs delivery service UCC (for carrier)
C	Costs policies
D	Transport costs
E	Costs delivery service UCC (for receiver)
F	Costs value-adding services
G	Subsidies

ID	INFORMATION FLOW
(i)	Shipping price
(ii)	Joint transport costs
(iii)	External costs
(iv)	Arrival time order
(v)	Last-mile distribution price
(vi)	Transport request
(vii)	Order properties
(viii)	Policy settings
(ix)	Price value-adding services

Figure 2: Overview of the (inter)actions, monetary flows, and information flows for each agent type.

depots of carriers, a line-haul distance, and the urban area with receiver locations and a single UCC. Figure 3 shows a typical network setting for the framework. The network is comprised by a vertex set $\mathcal{V}$ and an arc set $\mathcal{A}$. Vertex set $\mathcal{V}$ is defined as $\mathcal{V} = \mathcal{V}^R \cup \mathcal{V}^S \cup \mathcal{V}^C \cup \mathcal{V}^H$, i.e., the set consists of subsets of locations of receivers $\mathcal{V}^R$, shippers $\mathcal{V}^S$, carriers $\mathcal{V}^C$, and the UCC $\mathcal{V}^H$. As we consider a single UCC location for the sake of illustration, we define $\mathcal{V}^H = \{v^{ucc}\}$, with v^{ucc} denoting the UCC location in the network. Every arc $a \in \mathcal{A}$ connects a vertex pair (v_a, v'_a), with $v_a, v'_a \in \mathcal{V}$ and $v_a \neq v'_a$.

We proceed with the order properties. Let $\mathcal{L} = \{\frac{1}{y}, \frac{2}{y}, \ldots, 1\}$ (with integer $y > 1$) be the set of possible order volumes, with a volume of 1 being equal to the volume of the smallest vehicle in the model (e.g., a delivery van). Let $\mathcal{T}^e$ be the set of possible earliest delivery times (relative to t) and $\mathcal{T}^w$ be the set of possible widths of the delivery window. We define an order as a request to ship a certain load, with an *order type* being a unique combination of the delivery window $[t + t^e, t + t^l]$ (with $t^e \in \mathcal{T}^e$ and $t^l - t^e \in \mathcal{T}^w$), the current position of the order $v \in \mathcal{V}$ (which indicates the agent responsible for handling the order at t), the receiver $r \in \mathcal{R}$, the carrier $c \in \mathcal{C}$, the shipper $s \in \mathcal{S}$, the order volume $l \in \mathcal{L}$, and an indicator $\gamma \in \{0, 1\}$ that specifies whether delivery takes place via the

UCC ($\gamma = 1$ sets the UCC as the delivery address for the carrier). The indicator γ can be specified either by the receiver or by the carrier. To denote which agent is responsible for the order, we expand the vertex notation with a superscript (rec, car, shp, ucc) and a subscript $r \in \mathcal{R}, c \in \mathcal{C}, s \in \mathcal{S}$, e.g., v_s^{shp} when shipper S holds an order. The order type of an order changes during the simulation. As the delivery window is defined relative to the decision epoch t, it must be updated after each time step. Furthermore, decisions regarding the use of the UCC may alter the order type. The transition functions that describe how orders change type are defined in Section 3.4. The number of orders of a specific type in the system at time t is denoted by $I_{t,t^e,t^l,v,r,c,s,l,\gamma} \in \mathbb{N}$. The vector I_t describes the number of every order type in the system at time t and is defined as

$$I_t = [I_{t,t^e,t^l,v,r,c,s,l,\gamma}]_{\forall t^e,t^l,v,r,c,s,l,\gamma \in \mathcal{T}^e \times \mathcal{T}^w \times \mathcal{V} \times \mathcal{R} \times \mathcal{C} \times \mathcal{S} \times \mathcal{L} \times \{0,1\}} \ . \tag{1}$$

For notational convenience, we omit the set notations in subscripts from here on. The notation I_t describes all orders in the system at t. As explained before, we indicate the orders relevant for a specific agent (for which we use the term *order position*) with the appropriate vertex, e.g., a subscript v_c^{car} implies that carrier c is currently responsible for the order. Thus, carrier c at time t has an order position $I_{t,t^e,t^l,v_c^{car},r,c,s,l,\gamma}, \forall t^e, t^l, r, s, l, \gamma$. As order types are defined by a sequential decision-making process that involves multiple agent types, several indices of the order type are left undefined during the process. For instance, a receiver will not specify which carrier delivers the order; this index is left blank in the specification of the order type or the action – which we denote by a dot in place of the unspecified index – until it is specified by the shipper. The decision-making process that defines the order types is described in Section 3.3.

The model keeps track of the agents that utilize the UCC. The variable $\gamma_{t,r}^{rec} \in \{0,1\}$ indicates whether receiver r sets the UCC as its delivery address at decision epoch t; the variable $\gamma_{t,c}^{car} \in \{0,1\}$ has the same purpose for carriers. These variables might be fixed for longer periods of time; we discuss strategic and tactical decision making in Section 3.6. In total there are $|\mathcal{R}|$ receivers and $|\mathcal{C}|$ line-haul carriers; the vectors $\gamma_t^{rec} = [\gamma_{t,1}^{rec}, \ldots, \gamma_{t,|\mathcal{R}|}^{rec}]$ and $\gamma_t^{car} = [\gamma_{t,1}^{car}, \ldots, \gamma_{t,|\mathcal{C}|}^{car}]$ describe which receiver and carrier agents are committed to the UCC at time t. The state of the system at t is defined by $S_t = [I_t, \gamma_t^{rec}, \gamma_t^{car}]$; this description provides all information required for the agents to make decisions and to compute the performance indicators.

3.3 Definition of dispatching decisions

In some form, every agent – except the administrator – is faced with a dispatching decision based on its order position. For a given agent (again indicated by the appropriate vertex in the subscripts), the dispatching decision may be generally defined by $x_{t,t^e,t^l,v,r,c,s,l,\gamma} \leq I_{t,t^e,t^l,v,r,c,s,l,\gamma} \forall t^e, t^l, r, c, s, l, \gamma$. Depending on the context, a dispatching decision is either a physical transfer of orders or an information transfer of order properties. The dispatching decision of one agent transfers the responsibility for the dispatched orders to other agents; the dispatching decisions are made sequentially by the various agent types. As shown earlier, the current decision maker is indicated by the vertex in the subscript of an action or order position. Physically, orders move downstream from shipper to receiver via the carrier (and possibly via the UCC). In the basic version of the framework, information transfers of order properties only take place between receivers and shippers, i.e., through

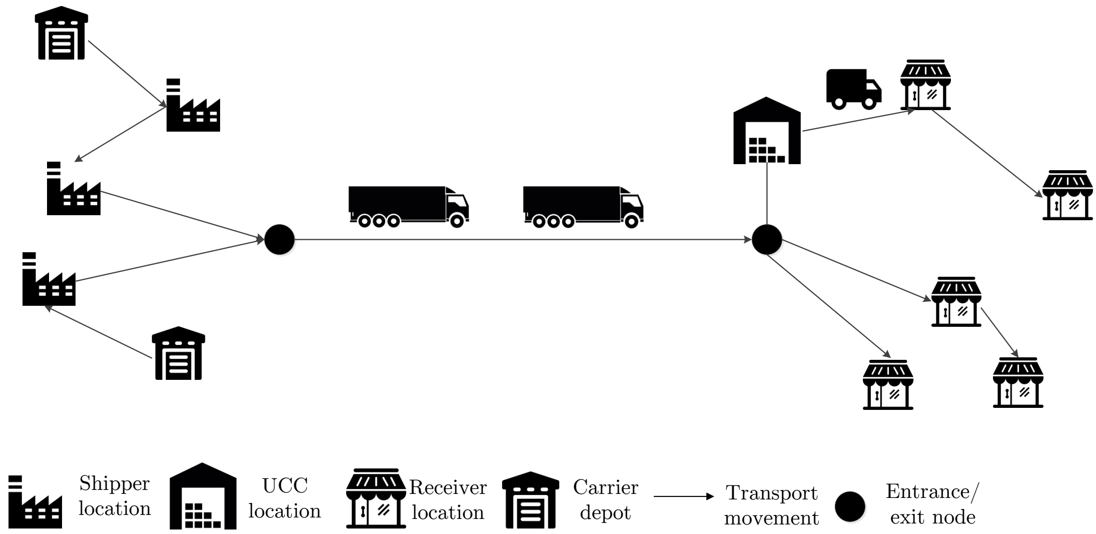

Figure 3: Example of a network with a UCC enabling carriers to outsource last-mile distribution.

the placement of orders. Both types of actions have in common that they update the (physical or virtual) order positions of other agents.

The dispatching decision for the receiver $r \in \mathcal{R}$ is the information transfer of its demand arising between time $t-1$ and time t, i.e., it places an order at the shippers that is equal to its demand. The carrier is undefined at this point. When placing the order, the receiver specifies a delivery window $[t+t^e, t+t^l]$. A value $t^e > 0$ implies that the order must be held either at the shipper location or at the UCC before the earliest delivery time. The dispatching decision by receiver $r \in \mathcal{R}$ is defined as:

$$x_{t,t^e,t^l,v_r^{rec},r,\cdot,s,l,\gamma} := I_{t,t^e,t^l,v_r^{rec},r,\cdot,s,l,\gamma} \qquad \forall t^e, t^l, s, l, \gamma \ . \tag{2}$$

The dispatching decision for the shipper is a physical dispatching decision. Although Figure 1 also describes the possibility to pre-announce order arrivals to the UCC (i.e., an information transfer), we ignore this action to simplify the mathematical formulation of actions and transitions. By accumulating orders, the shipper may achieve lower average transport costs, yet the delivery windows imposed by the receivers must be respected. At every decision epoch t, the shipper s must select a carrier $c \in \mathcal{C}$ to perform the transport for the dispatched orders (i.e., one carrier for the complete shipment at t). We introduce the binary variable $z_{t,c,s} \in \{0,1\}$ to describe this choice. The dispatching decision of shipper $s \in \mathcal{S}$ is given by

$$[x_{t,t^e,t^l,v_s^{shp},r,c,s,l,\gamma}]_{\forall t^e,t^l,r,c,l,\gamma} \ , \tag{3}$$

s.t.

$$\sum_{c\in\mathcal{C}} z_{t,c,s} \cdot x_{t,t^e,t^l,v_s^{shp},r,c,s,l,\gamma} \leq I_{t,t^e,t^l,v_s^{shp},r,c,s,l,\gamma} \qquad t^l > 1, \forall t^e, r, l, \gamma \;, \tag{4}$$

$$\sum_{c\in\mathcal{C}} z_{t,c,s} \cdot x_{t,0,1,v_s^{shp},r,c,s,l,\gamma} = I_{t,0,1,v_s^{shp},r,c,s,l,\gamma} \qquad \forall r, l, \gamma \;, \tag{5}$$

$$x_{t,t^e,t^l,v_s^{shp},r,c,s,l,0} = 0 \qquad t^e > 0, \forall t^l, r, c, l \;, \tag{6}$$

$$\sum_{c\in\mathcal{C}} z_{t,c,s} = 1 \;, \tag{7}$$

$$x_{t,t^e,t^l,v_s^{shp},r,c,s,l,\gamma} \in \mathbb{N} \qquad \forall t^e, t^l, r, c, l, \gamma \;, \tag{8}$$

$$z_{t,c,s} \in \{0,1\} \qquad \forall c \;. \tag{9}$$

Constraint (4) states that no more orders of any order type are shipped than that are placed by the receiver. Constraint (5) ensures that orders are shipped in time, such that the carrier or UCC can make the delivery to the receiver before the last allowed delivery time (given our assumption that all routes, including the line-haul component, are completed within a single time interval). Constraint (6) states that – if the order does not have the UCC as its delivery address – the shipper cannot ship orders before the earliest delivery date; this is because the carrier does not offer intermediate storage capacity. With Constraint (7), we enforce that exactly 1 carrier is selected to conduct the transport of the shipment, i.e., all orders that are dispatched by the shipper at time t. Constraint (8) and Constraint (9) are domain constraints.

The dispatching decision of the carrier is equivalent to transporting all orders that are physically released by the shippers to carrier c. To simplify the mathematical presentation, we impose that orders that are released by the shipper at t are picked up and delivered by the carrier between t and $t+1$. Recall that orders with $t^e > 0$ can only be delivered to the UCC; unlike the UCC and the shipper, the carrier does not offer a temporary storage service. We define the dispatching decision of a carrier $c \in \mathcal{C}$ as follows:

$$x_{t,t^e,t^l,v_c^{car},r,c,s,l,\gamma} := I_{t,t^e,t^l,v_c^{car},r,c,s,l,\gamma} \qquad \forall t^e, t^l, r, s, l, \gamma \;. \tag{10}$$

Finally, the definition of the dispatching decision for the UCC is similar to that of the shipper, although without the constraints relating to the carrier selection:

$$[x_{t,t^e,t^l,v^{ucc},r,c,s,l,1}]_{\forall t^e,t^l,r,c,s,l,1} \;, \tag{11}$$

s.t.

$$x_{t,0,t^l,v^{ucc},r,c,s,l,1} \leq I_{t,0,t^l,v^{ucc},r,c,s,l,1} \qquad \forall t^l, r, c, s, l \ , \tag{12}$$

$$x_{t,0,1,v^{ucc},r,c,s,l,1} = I_{t,0,1,v^{ucc},r,c,s,l,1} \qquad \forall r, c, s, l \ , \tag{13}$$

$$x_{t,t^e,t^l,v^{ucc},r,c,s,l,1} = 0 \qquad t^e > 0, \forall t^l, r, c, s, l \ , \tag{14}$$

$$x_{t,t^e,t^l,v^{ucc},r,c,s,l,1} \in \mathbb{N} \qquad \forall t^e, t^l, r, c, s, l \ . \tag{15}$$

3.4 Transition function

Having defined the state of the system and the corresponding actions, we proceed to describe the transition function. We first introduce the stochastic order arrival process. Let $\tilde{I}_{t,t^e,t^l,v_r^{rec},r,\cdot,s,l,\cdot} \in \mathbb{N}$ denote the number of arrivals of a given order type during the time interval $(t-1, t]$. We denote all arrivals during this interval by $\omega_t = [\tilde{I}_{t,t^e,t^l,v_r^{rec},r,\cdot,s,l,\cdot}]_{\forall t^e,t^l,r,s,l}$ and the set of all possible arrivals by Ω_t. Let W_t be a random variable with $\omega_t \in \Omega_t$ being a realization of W_t.

As we consider a system with interacting agents, the actions of one agent affect the order positions of other agents. Dispatching decisions may also transform order types, e.g., when the carrier decides to deliver via the UCC. We describe the transition function as a sequential updating procedure. The influence that the actions of one agent have on the order position of another agent is represented by the vertex subscripts. For example, in the dispatching decision of a receiver r (communicating an order placement to the shipper), the subscript v_r^{rec} is specified. Processing this decision updates the order position of the shipper, updating the subscript v_s^{shp} to implies that the decision of the receiver affects the order position of the shipper.

We split the transition function into two components. First, we address the transitions occurring at a given decision epoch t as a result of the dispatching actions of the agents. Second, we describe the transition that occurs when making the discrete time step from t to $t+1$, comprising the order arrival process and the update of time indices. At various points, we use the post-decision order position, which is the position after applying a dispatching decision on order position for a given agent type, but before the arrival of new orders (i.e., before the discrete time step). Notations that relate to the post-decision order position are defined with a superscript x.

We start with an initial order position I_t, in which positive numbers of order types may exist for the receivers (newly generated orders according to the realization of W_t), the shippers (existing order positions), and the UCC (existing order position). From this initial position, we proceed to move to the post-decision position I_t^x via a sequential updating procedure that we outline below.

The demand of receiver $r \in \mathcal{R}$ for a given order type is defined by the variable $I_{t,t^e,t^l,v_r^{rec},r,\cdot,s,l,\cdot} \in \mathbb{N}$; the carrier and the use of the UCC are initially undefined, hence the dots $\cdot$ in the notation. If $\gamma_{t,r}^{rec} = 1$ (which may be fixed at the tactical or strategic level), this means that the receiver will use the UCC as its delivery address, thus we must set $\gamma = 1$ for all its orders. If $\gamma_{t,r}^{rec} = 0$, the receiver is indifferent as to whether it will be delivered directly or via the UCC, meaning that this choice remains open to the carrier. The index for γ is left blank in that case. Thus, *iff* $\gamma_{t,r}^{rec} = 1$, we perform the following update for receiver r

$$I_{t,t^e,t^l,v_r^{rec},r,\cdot,s,l,1} := I_{t,t^e,t^l,v_r^{rec},r,\cdot,s,l,\cdot} \qquad \forall t^e, t^l, s, l \ .$$

Subsequently, we reset all order positions of receivers without a specified value for γ (i.e., $\gamma^{rec}_{t,r} = \cdot$) by performing the update $I_{t,t^e,t^l,v^{rec}_r,r,\cdot,s,l,\cdot} := 0 \qquad \forall t^e, t^l, s, l$.

The orders placed by the receiver are converted into an action according to Equation (2). After obtaining the action of the receiver, it is communicated to the shipper, which prompts an update of the order position of the shipper. At this stage, the carrier c is still undefined. To obtain the order position of a shipper $s \in \mathcal{S}$ at t, the order requests by the receivers are added to the already existing order positions of shippers, which follows from the post-decision position at $t-1$. To consistently describe the update at t, we therefore first show how we compute the post-decision order position at $t-1$. The shipper selects an action as defined in Equation (3). By dispatching the orders, we can compute the post-decision order position for shipper s:

$$I^x_{t-1,t^e,t^l,v^{shp}_s,r,\cdot,s,l,\gamma} = I_{t-1,t^e,t^l,v^{shp}_s,r,\cdot,s,l,\gamma} - \sum_{c\in\mathcal{C}} x_{t-1,t^e,t^l,v^{shp}_s,r,c,s,l,\gamma} \qquad \forall t^e, t^l, r, l, \gamma \ .$$

The order position of the shipper at t follows from the post-decision order position at $t-1$ and the orders posed by the receivers at t. To take into account the time update of delivery windows relative to the current decision epoch, we must update the order types of existing orders in the order position. Most time indices t^e and t^l are simply reduced by 1, as they are relative to t. However, orders with an earliest delivery time $t^e = 0$ or $t^e = 1$ relative to $t-1$ both correspond to $t^e = 0$ at the current decision epoch t. We incorporate this update with the help variable $\tilde{t}^e$. The update of the order position at t is defined as follows:

$$\begin{aligned}
I_{t,0,t^l,v^{shp}_s,r,\cdot,s,l,\gamma} &= \sum_{\tilde{t}^e=0}^{1} (I^x_{t-1,\tilde{t}^e,t^l+1,v^{shp}_s,r,\cdot,s,l,\gamma}) + x_{t,0,t^l,v^{rec}_r,r,\cdot,s,l,\gamma} \qquad \forall t^l, r, l, \gamma \ , \\
I_{t,t^e,t^l,v^{shp}_s,r,\cdot,s,l,\gamma} &= I^x_{t-1,t^e+1,t^l+1,v^{shp}_s,r,\cdot,s,l,\gamma} + x_{t,t^e,t^l,v^{rec}_r,r,\cdot,s,l,\gamma} \qquad t^e > 0, \forall t^l, r, l, \gamma \ .
\end{aligned}$$

We proceed with the update at the carrier level based on its decision whether to use the UCC. Based on the decision variables $\gamma^{car}_{t,c}$ and $\gamma^{rec}_{t,r}$, the value for γ may be updated. If $\gamma = 1$ due to the preceding decision of the receiver, the carrier cannot set $\gamma = 0$. However, if $\gamma^{rec}_{t,r} = 0$, the index γ has not been defined and will be determined by the carrier. The update of the order types for carrier c is defined by

$$\begin{aligned}
I_{t,t^e,t^l,v^{car}_c,r,c,s,l,0} &:= I_{t,t^e,t^l,v^{car}_c,r,c,s,l,\cdot} \qquad \forall t^e, t^l, r, s, l \ \text{ if } \ \gamma^{car}_{t,c} = 0 \ , \\
I_{t,t^e,t^l,v^{car}_c,r,c,s,l,1} &:= I_{t,t^e,t^l,v^{car}_c,r,c,s,l,\cdot} \qquad \forall t^e, t^l, r, s, l \ \text{ if } \ \gamma^{car}_{t,c} = 1 \ .
\end{aligned}$$

After completing the update for γ, the number of orders for types characterized by $\gamma = \cdot$ are set to 0: $I_{t,t^e,t^l,v^{car}_c,r,c,s,l,\cdot} := 0 \qquad \forall t^e, t^l, r, s, l$.

We proceed with the inventory update of the UCC. The inventory at time t depends on the post-decision order position at $t-1$ and the dispatching decisions by the carriers taken at t, as defined in Equation (10). We first describe how we compute the post-decision order position of the

UCC. Again, we first define the post-decision order position at $t-1$. For a given (pre-decision) order position, the UCC makes its dispatching decision as defined in Equation (11), which is used to compute its post-decision order position:

$$I^x_{t-1,t^e,t^l,v^{ucc},r,c,s,l,1} = I_{t-1,t^e,t^l,v^{ucc},r,c,s,l,1} - x_{t-1,t^e,t^l,v^{ucc},r,c,s,l,1} \qquad \forall t^e, t^l, r, c, s, l \ .$$

The inventory position of the UCC at t depends on the existing inventory from $t-1$ and the dispatching decisions by the carriers taken at t, as defined in Equation (10). Based on this information, we update the inventory position of the UCC as follows:

$$I_{t,0,t^l,v^{ucc},r,c,s,l,1} = \sum_{\tilde{t}^e=0}^{1} (I^x_{t-1,\tilde{t}^e,t^l+1,v^{ucc},r,c,s,l,1}) + x_{t,0,t^l,v^{car}_c,r,c,s,l,1} \qquad \forall t^l, r, c, s, l \ ,$$
$$I_{t,t^e,t^l,v^{ucc},r,c,s,l,1} = I^x_{t-1,t^e+1,t^l+1,v^{ucc},r,c,s,l,1} + x_{t,t^e,t^l,v^{car}_c,r,c,s,l,1} \qquad t^e > 0, \forall t^l, r, c, s, l \ .$$

Having defined the sequential procedure to obtain the post-decision order position of the system, we proceed to describe the transition from post-decision order position I^x_{t-1} to the next pre-decision order position I_t, comprising the processing of new order arrivals and the update of time indices. First, the receiver orders are generated based on the realization of the random variable W_{t+1}. At this point, the carrier c and the indicator γ are still undefined. For receiver r, its order position at $t+1$ is given by

$$I_{t+1,t^e,t^l,v^{rec}_r,r,\cdot,s,l,\cdot} = \tilde{I}_{t+1,t^e,t^l,v^{rec}_r,r,\cdot,s,l,\cdot} \qquad \forall t^e, t^l, s, l \ .$$

For each shipper s, we need to update the time indices of its order position, i.e., the post-decision position at decision epoch t. This time update goes as follows:

$$I_{t+1,0,t^l,v^{shp}_s,r,\cdot,s,l,\gamma} = \sum_{\tilde{t}^e=0}^{1} I^x_{t,\tilde{t}^e,t^l+1,v^{shp}_s,r,\cdot,s,l,\gamma} \qquad \forall t^l, r, l \ ,$$
$$I_{t+1,t^e,t^l,v^{shp}_s,r,\cdot,s,l,\gamma} = I^x_{t,t^e+1,t^l+1,v^{shp}_s,r,\cdot,s,l,\gamma} \qquad t^e > 0, \forall t^l, r, l \ .$$

The time update for the UCC is comparable:

$$I_{t+1,0,t^l,v^{ucc},r,c,s,l,1} = \sum_{\tilde{t}^e=0}^{1} I^x_{t,\tilde{t}^e,t^l+1,v^{ucc},r,c,s,l,1} \qquad \forall t^l, r, c, s, l \ ,$$
$$I_{t+1,t^e,t^l,v^{ucc},r,c,s,l,1} = I^x_{t,t^e+1,t^l+1,v^{ucc},r,c,s,l,1} \qquad t^e > 0 \forall t^l, r, c, s, l \ .$$

3.5 Cost functions and Key Performance Indicators

In this section, we provide the cost (profit) functions and Key Performance Indicators (KPIs) of the agents. We start by introducing some notation required to describe the price functions P (describing income) and cost functions C (describing expenses). For a variety of parameters and variables, we use the superscript hd to refer to costs for handling operations (e.g., (un)loading by the driver), rc for receiving (e.g., working hours spent on receiving, quality control of goods, and allocation to shelves), sp for shipping (e.g., loading operations), and sb for income from subsidies. In the description, we restrict ourselves to subsidies for using the UCC; schemes that subsidize other initiatives may be modeled in a comparable manner. The used order volumes depend on the context of the function. For example, the volume transported by carrier c at time t to the UCC is given by

$$l_{t,\cdot,\cdot,v_c^{car},\cdot,c,\cdot,\cdot,1} = \sum_{t^e,t^l,r,s,l} x_{t,t^e,t^l,v_c^{car},r,c,s,l,1} \cdot l \enspace .$$

We proceed with the notation required to describe the routes. Let $\mathcal{Q}_c$ denote the set of vehicles operated by carrier (or UCC) $c \in \mathcal{C} \cup \mathcal{H}$. A vehicle $q \in \mathcal{Q}_c$ has a vehicle capacity ≥ 1. The smallest vehicle defined in the simulation model has a capacity of 1, order sizes and the capacities of other vehicles are defined relative to this volume. To ease the notation, we assume that all routes starting at t are completed at $t+1$. For the same reason, we do not explicitly include the collection of orders at shipper locations. We denote a route started by vehicle q of carrier c at time t by $\delta_{t,c,q}^{car} = \{\delta_{t,c,q}^{car,lh}, \delta_{t,c,q}^{car,lm}\}$, with the components referring to line-haul transport and last-mile distribution respectively. This distinction is used to assign distinct properties (e.g., fuel usage, road pricing, driver wage) to the associated travel distances η^{lh} and η^{lm}. We let $\Delta_{t,c}^{car}$ denote the set of routes for carrier c at decision epoch t and use Δ_t^{ucc} to describe the set of routes for the UCC. The UCC only has to deal with last-mile distribution, such that $\delta_{t,q}^{ucc} = \{\delta_{t,q}^{ucc,lm}\}$. We use $\Delta_t = [[\Delta_{t,c}^{car}]_{\forall c \in \mathcal{C}}, \Delta_t^{ucc}]$ to denote all routes starting at time t.

For the carriers and the UCC, handling costs depend on the subsets of locations visited. The carrier incurs handling costs at each shipper of the subset $\mathcal{S}' \subseteq \mathcal{S}$ that it visits to collect shipments at a given decision epoch. Both carriers and receivers incur handling costs at each location in the subset of receivers visited $\mathcal{R}' \subseteq \mathcal{R} \cup \mathcal{H}$ (note that the UCC is a receiver from the perspective of the carrier). For each agent that conducts transport, the subsets $\mathcal{S}'$ and $\mathcal{R}'$, as well as the distances η^{lh} and η^{lm} are derived from the route sets $\Delta_{t,c}^{car}$ and Δ_t^{ucc}. To clarify which aspect of a route is responsible for a particular cost or price component (i.e., the costs per km driven or the loading costs per shipper visited), in our definitions of cost and price functions we use the subsets $\mathcal{S}'$ and $\mathcal{R}'$ and the distances η^{lh} and η^{lm} rather than the generic route notation.

We now introduce the cost functions for the agents. The outcomes of these cost functions over the full time horizon $\mathcal{T}$ serve as KPIs for the agents. To be consistent in our notation, we express the performance of each agent in terms of costs; price components are included as negative costs. Although agents aim to minimize their costs over the full planning horizon, they make periodic decisions based on incomplete information.

The objective of the shipper is to minimize the sum of transport costs and shipping costs. Shippers can influence these costs by selecting the set of orders to ship at every decision epoch and select the cheapest carrier to conduct the transport of this order set, see Equation (3). The costs for shipper s at time t are given by

$$C^{shp}_{t,s}(l_{t,\cdot,\cdot,v^{shp}_s,\cdot,c,s,\cdot,\cdot},\eta^{lh}) = C^{shp,tr}_{t,s}(l_{t,\cdot,\cdot,v^{shp}_s,\cdot,c,s,\cdot,\cdot},\eta^{lh}) + C^{shp,sp}_{t,s}(l_{t,\cdot,\cdot,v^{shp}_s,\cdot,c,s,\cdot,\cdot}) \ .$$

Carriers attempt to maximize profit (determined by the transport price, subsidy income, transport costs, and outsourcing costs) by selecting the route set $\Delta^{car,x}_{t,c}$ that minimizes costs at every decision epoch. In addition, carriers can choose whether they perform the full transport themselves or outsource the last-mile transport to the UCC. In case of outsourcing, the carrier only pays for outsourcing volume that is not already outsourced by the receiver, i.e., only the volume characterized by $\gamma = \cdot$. The carrier also receives subsidies for this volume only. The cost function of a carrier c at time t is given by

$$C^{car}_{t,c}(l_{t,\cdot,\cdot,v^{car}_c,\cdot,c,\cdot,\cdot,\cdot}, l_{t,\cdot,\cdot,v^{car}_c,\cdot,c,\cdot,\cdot,1},\eta^{lh},\eta^{lm},\gamma^{car}_{t,c},\mathcal{S}',\mathcal{R}') =$$
$$\begin{cases} C^{car,tr}_{t,c}(\eta^{lh},\eta^{lm},\mathcal{S}',\mathcal{R}') - P^{car,tr}_{t,c}(l_{t,\cdot,\cdot,v^{car}_c,\cdot,c,\cdot,\cdot,\cdot},\eta^{lh}) & \text{if } \gamma^{car}_{t,c} = 0 \ , \\ C^{car,tr}_{t,c}(\eta^{lh},0,\mathcal{S}',h) + C^{car,lm}_{t,c}(l_{t,\cdot,\cdot,v^{car}_c,\cdot,c,\cdot,\cdot,\cdot},\eta^{lh},\mathcal{S}') & \\ \quad -P^{car,tr}_{t,c}(l_{t,\cdot,\cdot,v^{car}_c,\cdot,c,\cdot,\cdot,\cdot},\eta^{lh}) - P^{car,sb}_{t,c}(l_{t,\cdot,\cdot,v^{car}_c,\cdot,c,\cdot,\cdot,\cdot}) & \text{if } \gamma^{car}_{t,c} = 1 \ . \end{cases}$$

The objective of the UCC is to maximize profit, which is determined by the prices it charges, subsidy incomes, receiving costs, costs for performing value-adding services, and transport costs. If the receiver decides to select the UCC as its delivery address, it is the receiver that pays the UCC for the last-mile distribution costs. If the carrier decides to outsource its last-mile distribution, it only pays the UCC for the volume that has not already been outsourced by the receiver. To influence its profit, at every decision epoch it selects a subset of orders to dispatch as defined in Equation (11), which yields a corresponding route set Δ^{ucc}_t. The cost function for the UCC at time t is defined by

$$\begin{aligned}
&C^{ucc}_t(l_{t,\cdot,\cdot,v^{ucc},\cdot,\cdot,\cdot,\cdot,1},\eta^{lm},\mathcal{R}') = \\
&C^{ucc,rc}_t(l_{t,\cdot,\cdot,v^{ucc},\cdot,\cdot,\cdot,\cdot,1}) + C^{ucc,tr}_t(l_{t,\cdot,\cdot,v^{ucc},\cdot,\cdot,\cdot,\cdot,1},\eta^{lm},\mathcal{R}') + \\
&\sum_{r\in\mathcal{R}'} C^{ucc,rec,val}_t(l_{t,\cdot,\cdot,v^{rec}_r,r,\cdot,\cdot,\cdot,1}) - \\
&P^{ucc,sb}_t(l_{t,\cdot,\cdot,v^{ucc},\cdot,\cdot,\cdot,\cdot,1}) - \sum_{r\in R'} \gamma^{rec}_{t,r} \cdot P^{ucc,rec,tr}_t(l_{t,\cdot,\cdot,v^{ucc},r,\cdot,\cdot,\cdot,1}) - \\
&\sum_{c\in C'} \gamma^{car}_{t,c} P^{ucc,car,tr}_t(l_{t,\cdot,\cdot,v^{car}_c,\cdot,c,\cdot,\cdot,1}) - \sum_{r\in\mathcal{R}'} P^{ucc,val}_t(r) \ .
\end{aligned}$$

The performance of the receiver is measured as the sum of receiving costs. These depend on whether delivery takes place via the UCC. If the receiver does not mandate delivery via the UCC, it incurs costs for every carrier (possibly including the UCC) that delivers goods at the receiver. Furthermore, the receiver must perform its value-adding services in-house and incurs a cost for that. If the receiver mandates deliveries via the UCC, the receiver pays the UCC for last-mile delivery, but may incur lower receiving costs due to receiving bundled orders from only one carrier (i.e., the UCC). The receiver also incurs costs for outsourcing its value-adding services. The costs for a receiver r at time t are given by:

$$
\begin{aligned}
&C_{t,r}^{rec}(l_{t,\cdot,\cdot,v_r^{rec},r,\cdot,\cdot,\cdot,\cdot}, l_{t,\cdot,\cdot,v_r^{rec},r,\cdot,\cdot,\cdot,1}, \gamma_{t,r}^{rec}) = \\
&\begin{cases}
C_{t,r}^{rec,rc}(l_{t,\cdot,\cdot,v_r^{rec},r,\cdot,\cdot,\cdot,\cdot}) + C_{t,r}^{rec,val} & \text{if } \gamma_{t,r}^{rec} = 0 \ , \\
C_{t,r}^{rec,rc}(l_{t,\cdot,\cdot,v_r^{rec},r,\cdot,\cdot,\cdot,1}) + C_{t,r}^{rec,lm}(l_{t,\cdot,\cdot,v_r^{rec},r,\cdot,\cdot,\cdot,1}) + & \text{if } \gamma_{t,r}^{rec} = 1 \\
C_{t,r}^{rec,val,ucc} - P_{t,r}^{rec,sb}(l_{t,\cdot,\cdot,v_r^{rec},r,\cdot,\cdot,\cdot,1}) \ . &
\end{cases}
\end{aligned}
$$

The performance of the administrator is measured with the following KPIs: (i) the number of vehicles per type that enter the urban area, (ii) the total distance covered within the urban area per vehicle type, (iii) the income from policies minus the provided subsidies, and (iv) the emission levels for CO_2, SO_2, NO_x, and particulate matter (PM). The first two KPIs indirectly capture effects such as noise hindrance and the contribution to road congestion. The third KPI is an indicator for the financial feasibility of a scheme.

3.6 Solution methods for decision-making

In this section, we describe various solutions methods that may be applied in the framework. We distinguish between solution methods on the strategic level (Section 3.6.1), the tactical level (Section 3.6.2), and the operational level (Section 3.6.3). The descriptions of the solution methods are deliberately kept at an abstract level, as the applicability of methods depends on factors such as the required level of detail and the computational budget.

3.6.1 Solution methods for strategic decisions

Strategic decisions are embedded in the framework as pre-defined input; they are fixed at $t = 0$ and agents are committed to them for the full planning horizon. Typically, strategic decisions may be incorporated into the network by defining arc properties. We discuss implementation guidelines for two classes of strategic decisions, namely administrative policies and company-driven initiatives.

Administrative policies are typically described by forms of access restrictions or access pricing. Such policies are incorporated in the framework by (i) defining restrictions and pricing functions as arc properties, (ii) defining the cost allocation structure, and (iii) specifying rules for subsidies. Although we model the decision to award subsidies as a strategic one, allocation rules and time thresholds may be defined in such a way that a dynamic distribution of subsidies remains possible.

Company-driven initiatives that may permanently affect agents are, e.g., buying vehicles of a certain type or establishing a communication structure between supply chain partners. To reflect such strategic decisions, the characteristics of the agents may be adjusted in the model. Initiatives that reorganize either the infrastructure or the transport process can be incorporated by modifying the network.

We conclude with a brief outline on setting up coalition structures in the framework, e.g., synchronizing delivery windows or pooling orders and resources. First, the agents that may enter the coalition should be specified. Second, the action space of the coalition needs to be defined, including rules for the allocation of tasks to agents. Third, a single objective function must be specified for the coalition. Fourth – as coalitions require rational agents that are willing to cooperate – an appropriate gain-sharing mechanism (e.g., from the field of cooperative game theory) should be incorporated.

3.6.2 Solution methods for tactical decisions

Tactical decisions in the simulation framework are made only at a pre-defined subset of decision epochs $\mathcal{T}^{tac} \subset \mathcal{T}$. These decisions commit agents for at least the time period between two adjacent decision epochs. Examples of tactical decisions are committing to outsource last-mile distribution to the UCC or adjusting the price levels of services. In this subsection, we discuss a sampling procedure to make decisions on the tactical level.

Although the range of tactical decisions might be diverse, the common characteristic of tactical decisions is that they affect the performance of agents over a time period of medium length. Therefore, each agent should evaluate how its future performance is affected by its tactical decisions. One approach to evaluate the impact of tactical decisions is to measure performance over some past time interval and compare it to the performance that might have been achieved under other tactical decisions. However, if the state of the system changes over time (e.g., altered prices or subsidies), lookahead estimates typically yield more accurate expected values.

We exemplify the lookahead procedure as follows. Consider a receiver that must decide whether to select the UCC as its delivery address. The receiver aims to minimize its expected receiving costs until the next tactical decision moment. For this purpose, N sample paths of order arrivals of length τ^{sample} are generated, with $n \in \{1, \dots, N\}$ representing the index for the sample path and $t^n \in \{0, \dots, \tau^{sample}\}$ representing the time index for the sample states of path n. The obtained sets of sample arrivals are denoted by $\{\tilde{\omega}_0^n, \dots, \tilde{\omega}_{\tau^{sample}}^n\}, \forall n \in \{1, \dots, N\}$. From these samples we derive order volumes for the receiver that are denoted by $l_{t^n,\cdot,\cdot,v_r^{rec},r,\cdot,\cdot,\cdot,\cdot}$ (the total volume) and $l_{t^n,\cdot,\cdot,v_r^{rec},r,\cdot,\cdot,\cdot,1}$ (the volume delivered via the UCC). Let $\tilde{\gamma}_r^{rec} \in \{0,1\}$ be a binary variable indicating whether receiver r uses the UCC in the lookahead procedure; to compare both options we compute the costs for $\tilde{\gamma}_r^{rec} = 0$ as well as for $\tilde{\gamma}_r^{rec} = 1$. In mathematical form, the following minimization problem is solved to update the tactical decision:

$$\gamma_{t,r}^{rec} = \underset{\tilde{\gamma}_r^{rec} \in \{0,1\}}{\arg\min} \ \tilde{C}^{rec}(\tilde{\gamma}_r^{rec}, l_{t',\cdot,\cdot,v_r^{rec},r,\cdot,\cdot,\cdot,\cdot}, l_{t',\cdot,\cdot,v_r^{rec},r,\cdot,\cdot,\cdot,1}) \ .$$

The sample costs $\tilde{C}^{rec}(\tilde{\gamma}_r^{rec}, l_{t^n,\cdot,\cdot,v_r^{rec},r,\cdot,\cdot,\cdot,\cdot})$ are computed as follows:

$$\tilde{C}^{rec}(\tilde{\gamma}_r^{rec}, l_{t^n,\cdot,\cdot,v_r^{rec},r,\cdot,\cdot,\cdot,\cdot}) =$$
$$\begin{cases} \frac{1}{N}\sum_{n=1}^{N}\sum_{t^n=0}^{\tau^{sample}} \left(C_{t^n,r}^{rec,rc}(l_{t^n,\cdot,\cdot,v_r^{rec},r,\cdot,\cdot,\cdot,\cdot}) + C_{t^n,r}^{rec,val} \right) & \text{if } \tilde{\gamma}_r^{rec} = 0 \ , \\ \frac{1}{N}\sum_{n=1}^{N}\sum_{t^n=0}^{\tau^{sample}} \Big(C_{t^n,r}^{rec,rc}(l_{t^n,\cdot,\cdot,v_r^{rec},r,\cdot,\cdot,\cdot,\cdot}) + C_{t^n,r}^{rec,lm}(l_{t^n,\cdot,\cdot,v_r^{rec},r,\cdot,\cdot,\cdot,\cdot}) & \\ +C_{t^n,r}^{rec,val,ucc} - P_{t^n,r}^{rec,sb}(l_{t^n,\cdot,\cdot,v_r^{rec},r,\cdot,\cdot,\cdot,\cdot}) \Big) & \text{if } \tilde{\gamma}_r^{rec} = 1 \ . \end{cases}$$

The sampling procedure assumes that agents respond instantly to changes in the system such as new price levels. These levels are adjusted before updating the tactical decisions. In practice, decision makers might display a more gradual approach in altering their tactical decisions. Such behavior may be reflected by incorporating learning mechanisms such as Q-learning, which place weight on both past observations and future expectations of state-action values. This results in a smoother transition of the state of the system. Although learning mechanisms may reflect the behavior of real-life actors more accurately, the downside is that more computational effort is likely required to identify the steady state (if such a state exists) of the system.

3.6.3 Solution methods for operational decisions

All agents aim to optimize their own objective functions. As we model a dynamic environment, operational decision problems are subject to incomplete information. Exact solution methods for stochastic problems usually require a large computational effort, which is why we typically resort to heuristic solutions in ABS. Various policies may be used to tackle the operational decision problems. The specific policies to be used depend on the instance, the learning objectives of the simulation study, and the computational budget.

Both the shippers and the UCC face a variant of the Delivery Dispatching Problem with time windows (DDP-TW) that is described in Van Heeswijk et al. (2017). In the DDP-TW, the agent must decide at each decision epoch which subset of orders to dispatch. The goal is to dispatch the orders in batches that allow for efficient delivery tours. Shippers face a simple variant of the problem as we model only one line-haul arc, UCCs face a more complicated variant as their routing costs depend on the subset of locations visited. Possible solution methods are learning a dispatching policy with ADP, a lookahead policy based on sampling, or a heuristic policy; Van Heeswijk et al. (2017) show all three variants.

Carriers, as well as the UCC, are required to solve a vehicle routing problem (VRP). Orders arrive via the line-haul arc, implying a single entry point to the city that functions as the starting point of the route. For orders that are delivered via the UCC, the UCC location is the starting point of the route. In Section 2 we gave an overview of VRP methods specifically designed to deal with the context of urban logistics. As we split the decision problem into a periodic dispatch decision and a vehicle routing problem, solution methods for the static VRP may be applied.

4 Numerical experiments

In this section, we demonstrate the working of the simulation framework using numerical experiments. The purpose of the experiments is threefold. First, the experiments provide a tangible application of the framework. Second, we provide metrics with respect to the computational time required to complete the experiments. Third, we give insights into effective urban logistics schemes.

The section is constructed as follows. Section 4.1 describes the properties of the instance and the used parameter settings. Section 4.2 discusses the setup and implementation details of the experiments. Section 4.3 outlines the solution methods used by the agents. Finally, Section 4.4 discusses the results of the numerical experiments.

4.1 Instance properties

We describe the properties of the test instance and the parameter settings, focusing on the general outline and motivating the important design choices.

We represent the city by a square area of 1km^2 that contains 100 uniformly distributed receiver locations. Furthermore, we assume a 10km^2 square area that is located 100km away from the urban area (i.e., the line-haul distance); this area contains 30 shipper locations and 30 carrier depots, which are uniformly spread.

We proceed to describe the receiver properties. We assume that each receiver uses a decentralized supply system, i.e., it places orders with multiple shippers. According to Cherrett et al. (2012), such receivers approximately receive about 12 distinct deliveries per week; we randomly generate 6-20 orders per receiver per week (3-5 order days and 2-4 orders per order day), with the delivery window for each order being a random day of the week. The volume per order is drawn from a set of 10 volumes (expressed in m^3) in the range $[0.24, 6.44] \cap \mathbb{Q}$ (Van Heeswijk 2017), assuming a triangular distribution with mean 1.67m^3. The unloading time at each receiver (from the perspective of the carrier) is uniformly drawn from the range of 7 to 34 minutes (Schoemaker et al. 2006, Allen et al. 2008). However, the receiver itself is not necessarily involved in the entire unloading process, and might therefore spend less time on receiving than the carrier does on unloading. We estimate that the receiving time lies between 2 minutes and the unloading time for the carrier, again drawing from a uniform distribution. We set the staff costs for a retailer at €15.3/hour (Van Duin et al. 2010); in conjunction with the time spent on unloading, this determines the receiving costs. The monthly cost for value-adding services – when performed in-house – is uniformly drawn from the range €30-100.

Next, we discuss the transport-related properties. Carriers use medium-sized trucks with a load capacity of 28m^3 and the UCC uses light trucks with a load capacity of 18m^3. We assume that the light trucks are exempted from any access restrictions imposed by the administrator. The financial and environmental properties of these truck types are based on Quak and de Koster (2009), Boer et al. (2011), and Roca-Riu et al. (2016). Table 2 summarizes these properties.

For the UCC, we also set parameters for the costs of their non-transport operations and prices for their services. Following Van Duin et al. (2010), we set a cost range of 11-19 €/m^3 handled at the UCC. The costs are initialized at €19 (assuming an initial state in which the UCC has no customers) and reduces to €11 when all inbound volume is handled by the UCC. Between these bounds, the costs are computed linearly with respect to the ratio between volume handled by the UCC and total inbound volume. For the value-adding services, we assume that the costs are initially 80% of the in-house costs, but linearly decrease with the volume ratio to 50% of the costs. The price for value-adding services as charged to the receiver is determined based on the cost price for the UCC; we use a default profit margin of 30%. The receiver price for the basis service of bundled deliveries is set at €70 per month and may linearly decrease to €60 per month (Van Heeswijk 2017). For the carriers, each outsourced stop costs €12 to €18.

Table 2: Vehicle properties for light trucks (UCC) and medium-sized truck (carriers).

Vehicle type	Light truck	Medium-sized truck
Load capacity (m^3)	18	28
Costs urban transport (€/km)	1.56	1.70
Costs line-haul transport (€/km)	–	1.24
Urban driving speed (km/hour)	25	25
Line-haul driving speed (km/hour)	–	50
CO_2 (g/km)	504	943
$PM_{2.5}$ (mg/km)	36	56

4.2 Experimental setup

We describe the setup of our experiments. We select a set of six variables that each have two levels (default and intervention); our 2^k factorial design therefore yields $2^6 = 64$ experiments. We perform 10 replications per experiment, the reported results are averages. The settings are summarized in Table 3. Variable A sets the access time window for medium-sized trucks; under the intervention level, these trucks are only allowed in the urban area during a two-hour interval each day. Variable B sets the subsidy level for carriers that outsource their last-mile distribution to the UCC; in the intervention setting they get 30% of the price paid to the UCC reimbursed by the local administrator. Variable C is the zone-access fee, which we set at €10 per medium-sized truck that enters the city. Variable D sets the profit margin of the UCC that determines its price levels, with the intervention setting reflecting a non-profit UCC. Variable E indicates whether or not the carriers form a coalition in which all transport jobs are bundled. Finally, variable F relates to the dispatch threshold that shippers maintain before shipping accumulated sets of orders; the intervention setting is 18m^3.

Table 3: Variable levels (default and intervention).

Variable	**Description**	**Default setting**	**Intervention setting**
A	Access time window	24h/day	2h/day
B	Carrier subsidies	0%	30%
C	Zone access fee	€0	€10
D	UCC margin value-adding services	30%	0%
E	Carrier coalition	No	Yes
F	Shipper dispatch threshold	0m^3	18m^3

We proceed to describe the implementation details of the framework. The framework is coded as a discrete-event simulation model in Delphi XE6 and runs on a computer with 8GB RAM and a 2.90GHz Intel Core i7 processor. The planning horizon contains 1250 decision epochs and corresponds to a five-year period (assuming 250 working days per year); the set of tactical decisions contains 50 epochs that are separated by equidistant time intervals. Subsidies are allocated during the first two years and then reset to 0. All reported results are taken from the last two years of the period, implying a one-year transition period. The uncertainty in the system is represented by randomly generating the information variables W_t.

4.3 Choice of solution methods

We describe the solution methods that are used by the agents on the three levels of decision-making, starting with the strategic level. These decisions correspond to the variables mentioned in Table 3. The access time window, zone-access fee, and carrier subsidies are incorporated as network properties; recall that the subsidies are set to 0 after two years. The volume thresholds for shippers and the profit margin of the UCC are fixed input. For the carrier coalition, we assume that all carriers bundle their transport jobs and that one randomly selected carrier executes the transport at a given decision epoch. To compute the KPIs, we simply measure the coalition-wide costs and prices divided by the number of carriers. Thus, we implicitly assume that the coalition partners are able to agree on a fair gain-sharing mechanism.

The tactical decisions in the numerical experiments are adjusting the price levels of the UCC and the decision whether to outsource last-mile distribution to the UCC (both by carriers and by receivers). The lookahead procedure corresponding to these decisions has been described in Section 3.6.2.

We proceed with the operational decisions. For shippers to determine the dispatch timing of accumulated orders, they use the fixed volume threshold (which is set at the strategic level). For example, a shipper may dispatch all its orders when their total volume exceeds 18m^3. An exception to this rule occurs when the set of accumulated orders contains an urgent order for which further postponement would violate its delivery window. If the set contains at least one urgent order, the shipper dispatches all orders that can be delivered within their delivery windows.

Based on the transport jobs communicated by the shippers, the carriers have to solve routing problems. Each pickup tour starts at the carrier depot and culminates at the entrance vertex of the line-haul segment. Delivery tours start at the exit vertex of the line-haul segment and – after delivering at the receivers – end at the same vertex. Both routing problems are independently solved with the Clarke-Wright savings algorithm and improved with a 2-opt heuristic.

The UCC makes two operational decisions, namely the dispatching decision and the routing decision. To solve its dispatching decisions, the UCC uses a sampling procedure as described in Section 3.6.2. To limit the computational effort, we apply a k-means clustering algorithm that assigns receivers to 1 of 10 clusters that may be visited, rather than enumerating every possible combination of receivers to visit. Like the carriers, the UCC also uses the Clarke-Wright algorithm in combination with 2-opt to solve their routing problems, with each tour both starting and ending at the UCC location.

4.4 Results of the numerical experiments

The full experimental results are shown in Appendix A; this section describes the key insights. We start with discussing the isolated effects, i.e., the effect of only a single intervention compared to the default scheme. Table 4 summarizes the isolated effects of the individual interventions on both the financial performance of the agents and three environmental performance indicators. The only interventions that have a substantial environmental impact are subsidizing carriers (leading to increased use of the UCC), forming carrier coalitions (improving the transport efficiency), and to a lesser extent the zone access fee (leading to increased use of the UCC). Finally, we point out that positive financial effects for one agent type typically imply negative effects for other agent types, illustrating the challenge of finding schemes that are beneficial for everyone. This is particularly visible when subsidizing carriers; the financial performance of the UCC improves considerably, but this ultimately is at the expense of the carriers.

Table 4: *Isolated effects per intervention. Positive percentages indicate a positive change, i.e., higher profits/lower costs for financial indicators, and lower emissions/less vehicles for environmental indicators.*

Variable	**Description**	**Financial**				**Environmental**		
		Carrier	Receiver	Shipper	UCC	PM	# vehicles	Vehicle dist.
A	Access time window	-8%	-2%	0%	0%	0%	1%	1%
B	Carrier subsidies	-85%	4%	2%	25%	30%	18%	16%
C	Zone access fee	22%	1%	-2%	1%	7%	1%	1%
D	UCC margin value-adding service	0%	-2%	-1%	0%	-1%	0%	1%
E	Carrier coalition	766%	0%	0%	-2%	46%	55%	55%
F	Shipper dispatch threshold	10%	-1%	-2%	0%	0%	-2%	-2%

We continue by describing the main effects (the average over all schemes with and without the intervention) and the two-way interaction effects (the average over all schemes with and without two given interventions). Figure 4 and Figure 5 display all statistically significant main financial effects and -environmental effects respectively, along with their 95% confidence bounds. The full overview of all effects (including non-significant ones) can be found in Appendix B.

First, we discuss the most important main effects. Again, we see that the carrier coalition (Var. E) has the greatest impact by far, strongly improving the performance for carriers. The efficiency gain reduces usage of the UCC, leading to significantly larger losses for the UCC. Subsidizing carriers (Var. B) strongly improves the financial performance of the UCC; also the receivers benefit significantly. However, the carriers themselves are ultimately worse off. When carriers commit to the UCC, receivers are more inclined to also use the UCC for value-adding services. This lock-in effect makes it non-beneficial for carriers to unilaterally opt out, yet they end up having higher costs. Five main effects are statistically significant for the shippers, yet this is more due to the very small variance of the simulation results for this group of agents, rather than the magnitude of the effects.

In terms of environmental impact, only three interventions have significant effects. Subsidizing carriers has a positive effect on emissions, but negatively affects the number of vehicles and the total vehicle distance. This can be explained by the increased use of the UCC; the vehicles have lower emissions, but more of them are needed to transport the same volume. This is particularly the case in schemes including carrier coalitions (comprising half of the schemes that determine the main effect); in those cases inbound trucks are usually loaded efficiently. A similar effect can be observed for the zone access fee, but in this case the effect on emissions is not significant. The carrier coalition has positive effects on all environmental indicators, due to the increased transport efficiency.

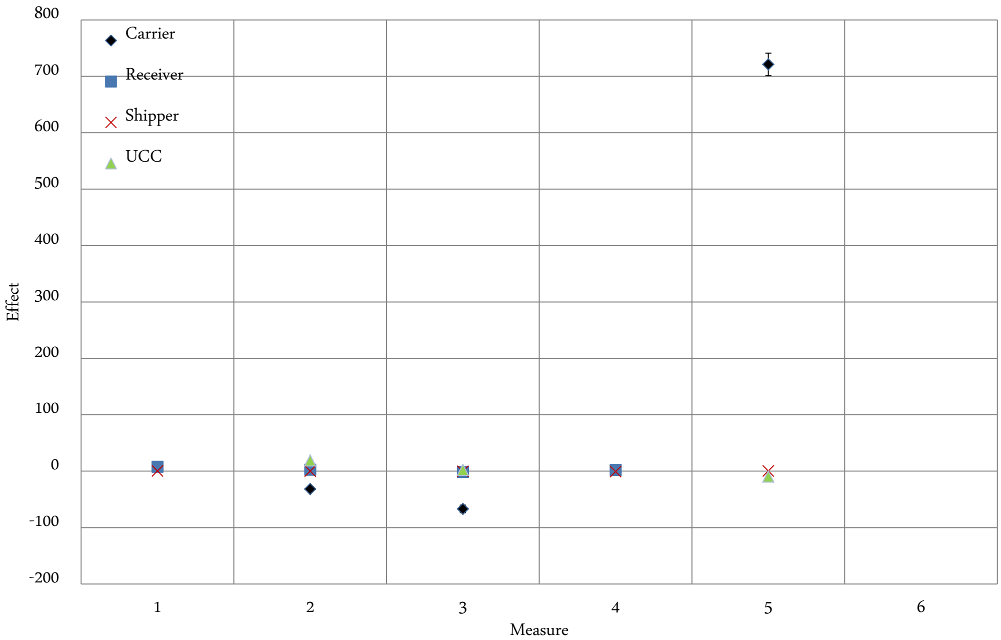

Figure 4: Financial main effects per agent type.

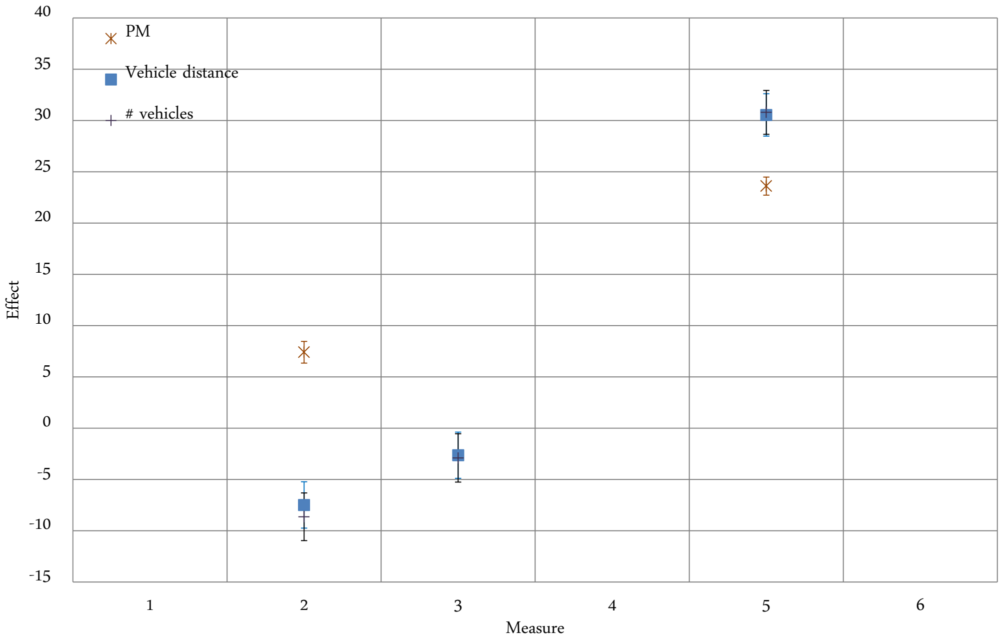

Figure 5: Environmental main effects per agent type.

Next, we discuss the two-way interaction effects, shown in Figure 6 (financial) and Figure 7 (environmental). Most notably, 6 out of 15 interaction effects have a positive and significant effect on the carriers. The strongest positive effect combines subsidizing with the carrier coalition; this is likely due to the higher commitment of receivers. For the UCC itself this interaction has a negative impact. The strongest negative effect combines the zone access fee with the carrier coalition. In this case, the fee is insufficient to alter the behavior of the coalition (as direct delivery almost always remains cheaper than outsourcing), so the coalition simply incurs the fees. Again, effects on the shipper are significant but marginal.

Interestingly, for the environmental performance there are only negative interactions that are significant. The strongest effects are combining subsidization or a zone access fee with the carrier coalition. This implies that more use is made of the UCC, while actually being less efficient. Also two combinations including the access time window yield negative effects.

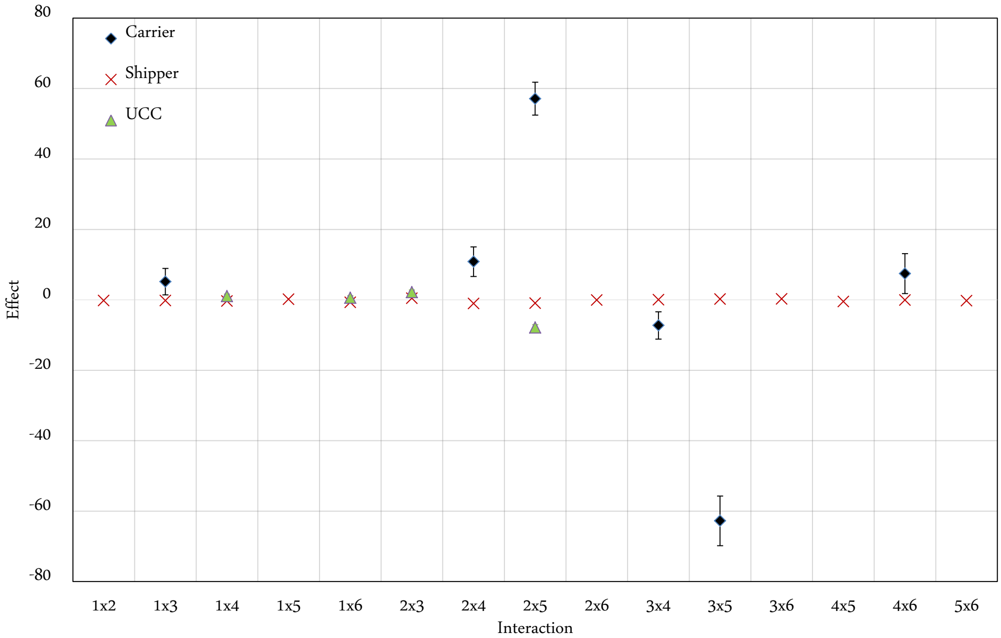

Figure 6: Financial interaction effects per agent type.

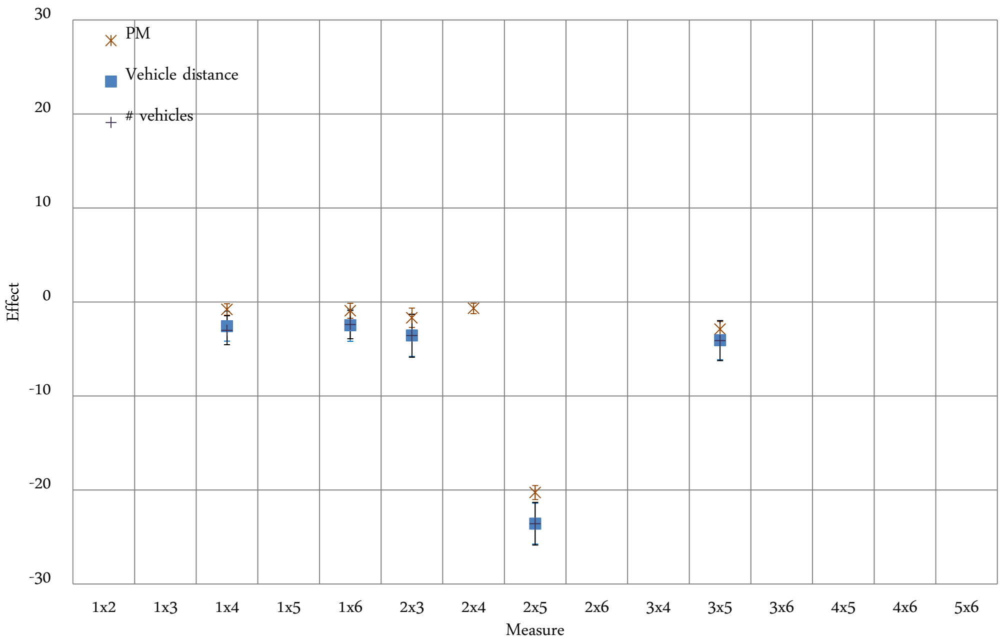

Figure 7: Environmental interaction effects per agent type.

We reflect on the obtained insights from a practical point of view. First, the carrier coalition is very effective, as the consolidation of goods yields significant savings both on the line haul and the last mile. Although theoretical savings are possible, we emphasize that implementation concerns – such as gain-sharing – are not addressed in these experiments. Also, we stress that this study only addressed a single city; in reality, trucks may carry higher loads that are destined for multiple city, implying a higher transport efficiency for individual carriers. Second, the order in which carriers and receivers commit to the UCC has a large impact on the eventual steady state of the system, which is particularly visible with the subsidy intervention. Clearly, in reality a system will not be that rigid; carriers might unilaterally opt out or renegotiate terms with the UCC. Still, the order of attracting users is very important for real-life UCCs and administrators, as (i) carriers potentially generate the bulk of the revenue and (ii) it should be verified whether the UCC remains a financially acceptable alternative for carriers after subsidies are halted. Third, in line with many observations in practice, we note that none of the schemes yields a profitable scheme for the UCC. However, we emphasize that UCCs typically handle smaller volumes than the generic volume range for these experiments; the prices charged in our simulation are therefore relatively low compared to the order volumes. Fourth, the interventions generally affect carriers and the UCC the strongest; the impact on receivers and shippers is generally limited. We summarize the key managerial insights from our simulation study as follows:

- Increased collaboration between carriers may effectively address inefficiency in urban transport and achieve substantial reductions in both emissions and the number of trucks in urban areas.

- UCCs may yield considerable environmental benefits, but even with supportive administrative measures are unlikely to achieve financial sustainability.

- Carriers are affected the most by the interventions, but also have the most direct influence on the efficiency of the transport system. Their support for any scheme should therefore be ensured.

We conclude with some remarks on the computational time. On average, a simulation run with 1250 decision epochs – of which 50 are tactical epochs – took 68.4 minutes to complete. A tactical decision epoch takes 50 times longer to solve than a operational epoch due to the lookahead procedure. Heuristic approaches to estimate the lookahead costs could therefore considerably reduce the reported computational time. The computational time largely depends on the number of order types; the current experiments distinguished 240,000 order types.

5 Conclusions

This paper presents an ABS framework that enables the evaluation of a vast array of urban logistics schemes, taking into account the autonomous decision making processes of five agent types and their corresponding KPIs (both financial and environmental), actions, monetary flows, and information flows. To enable other researchers to apply our framework, we mathematically define the agent actions, transition functions, and cost functions. Compared to other ABS studies in urban freight transport, our framework distinguishes itself as follows. First, we explicitly take into account line-haul transport and upstream decisions; other studies typically ignore decisions made outside the city boundaries. Second, our framework explicitly focuses on evaluating the synergy company-driven initiatives and administrative policies, as in practice such combinations often yield the most successful schemes. In particular, we emphasize cooperative freight bundling, as many shipments in urban freight transport are simply too small to efficiently utilize transport resources. Third, we define three levels of decision-making to reflect the various impacts of decisions over time, and suggest applicable solution methods for each level.

Numerical experiments performed with the framework indicate that initiatives that bundling through cooperation – such as a UCC or forming a coalition of carriers – yield considerable environmental benefits. However, achieving financial sustainability for all agents involved remains challenging. In particular, we did not identify any schemes in which the UCC is profitable in the long run. Subsidizing carriers to use the UCC and implementing a zone-access fee appear to be effective supportive interventions for the UCC, but are still insufficient from a financial perspective. Finally, carriers are both influenced the most by interventions and have the largest direct effect on transport efficiency.

The experiments also demonstrate that the framework enables to evaluate large numbers of schemes within reasonable computational time. For long-term success, it is important that schemes are supported by the involved actors; the framework helps to identify such schemes out of many possible combinations of interventions and parameter settings. If a scheme yields both good environmental results and financial performance, it may be explored in further depth. As such, the framework is primarily useful as a tool to quantify the effects of schemes and filter out the most promising ones.

Acknowledgments

This study is partially funded by Dinalog, the Dutch Institute for Advanced Logistics. The icons to create Figure 3 are obtained from http://www.flaticon.com and are freely available under the Flaticon Basic License.

References

Allen, J., Browne, M., Cherrett, T., and McLeod, F. (2008). Review of UK urban freight studies. Technical report, Green Logistics project, Universities of Westminster and Southampton, London Southampton, UK.

Anand, N., Quak, H. J., Van Duin, J. H. R., and Tavasszy, L. A. (2012). City logistics modeling efforts: Trends and gaps - a review. *Procedia-Social and Behavioral Sciences*, 39:101–115.

Anand, N., Van Duin, J. H. R., and Tavasszy, L. A. (2016). Framework for modelling multi-stakeholder city logistics domain using the agent based modelling approach. *Transportation Research Procedia*, 16:4–15.

Baldacci, R., Mingozzi, A., Roberti, R., and Calvo, R. W. (2013). An exact algorithm for the two-echelon capacitated vehicle routing problem. *Operations Research*, 61(2):298–314.

Bektaş, T., Crainic, T. G., and Van Woensel, T. (2015). From managing urban freight to smart city logistics networks. Technical report, CIRRELT 2015-17, CIRRELT, Montréal, Canada.

Bektaş, T., Demir, E., and Laporte, G. (2016). Green vehicle routing. In *Green Transportation Logistics*, pages 243–265. Springer.

Boer, E., Otten, M. B. J., and Essen, H. (2011). *Comparison of various transport modes on a EU scale with the STREAM database*. CE Delft, Delft, The Netherlands.

Browne, M., Sweet, M., Woodburn, A., and Allen, J. (2005). Urban freight consolidation centres. Technical report, University of Westminster, London, UK.

Cattaruzza, D., Absi, N., Feillet, D., and González-Feliu, J. (2015). Vehicle routing problems for city logistics. *EURO Journal on Transportation and Logistics*, 6(1):1–29.

Cheng, C., Qi, M., Wang, X., and Zhang, Y. (2016). Multi-period inventory routing problem under carbon emission regulations. *International Journal of Production Economics*, 182:263–275.

Cherrett, T., Allen, J., McLeod, F., Maynard, S., Hickford, A., and Browne, M. (2012). Understanding urban freight activity: Key issues for freight planning. *Journal of Transport Geography*, 24:22–32.

Crainic, T. G., Ricciardi, N., and Storchi, G. (2009). Models for evaluating and planning city logistics systems. *Transportation Science*, 43(4):432–454.

Cuda, R., Guastaroba, G., and Speranza, M. G. (2015). A survey on two-echelon routing problems. *Computers & Operations Research*, 55:185–199.

Dablanc, L. (2011). City distribution, a key element of the urban economy: guidelines for practitioners. In Macharis, C. and Melo, S., editors, *City distribution and urban freight transport: Multiple perspectives*, pages 37–56. Edward Elger, Cheltenham, UK.

Dahlberg, J. (2015). Cooperative transportation planning and cost allocation. Licentiate thesis, SE-601 74, Linköping University, Norrköping, Sweden.

Dellaert, N., Saridarq, F. D., Van Woensel, T., and Crainic, T. G. (2016). Branch & price based algorithms for the two-echelon vehicle routing problem with time windows. Technical report, CIRRELT-2016-45, CIRRELT, Montréal, Canada.

Ehmke, J. F., Campbell, A. M., and Thomas, B. W. (2016). Vehicle routing to minimize time-dependent emissions in urban areas. *European Journal of Operational Research*, 251(2):478–494.

Hafezalkotob, A. and Makui, A. (2015). Cooperative maximum-flow problem under uncertainty in logistic networks. *Applied Mathematics and Computation*, 250:593–604.

Hemmelmayr, V. C., Cordeau, J.-F., and Crainic, T. G. (2012). An adaptive large neighborhood search heuristic for two-echelon vehicle routing problems arising in city logistics. *Computers & Operations Research*, 39(12):3215–3228.

Kim, G., Ong, Y.-S., Heng, C. K., Tan, P. S., and Zhang, N. A. (2015). City vehicle routing problem (city VRP): A review. *Intelligent Transportation Systems, IEEE Transactions*, 16(4):1654–1666.

Kok, A. L., Hans, E. W., Schutten, J. M. J., and Zijm, W. H. M. (2010). A dynamic programming heuristic for vehicle routing with time-dependent travel times and required breaks. *Flexible Services and Manufacturing Journal*, 22(1-2):83–108.

Macharis, C. (2007). Multi-criteria analysis as a tool to include stakeholders in project evaluation: The MAMCA method. In Haezendonck, E., editor, *Transport Project Evaluation: Extending the Social Cost–Benefit Approach*, pages 115–131.

Macharis, C., Turcksin, L., and Lebeau, K. (2012). Multi actor multi criteria analysis (MAMCA) as a tool to support sustainable decisions: State of use. *Decision Support Systems*, 54(1):610–620.

Marcucci, E., Le Pira, M., Gatta, V., Inturri, G., Ignaccolo, M., and Pluchino, A. (2017). Simulating participatory urban freight transport policy-making: Accounting for heterogeneous stakeholders preferences and interaction effects. *Transportation Research Part E: Logistics and Transportation Review*, 103:69–86.

MDS Transmodal Limited (2012). DG MOVE European Commission: Study on urban freight transport. Technical report, MDS Transmodal Limited, Chester, UK.

Muñuzuri, J., Cortés, P., Guadix, J., and Onieva, L. (2012). City logistics in Spain: Why it might never work. *Cities*, 29(2):133–141.

Ploos Van Amstel, W. (2015). City Logistics: Working on livable cities through sustainable city logistics. Technical report, Amsterdam University of Applied Sciences, Amsterdam, The Netherlands.

Quak, H. J. (2011). Urban freight transport: The challenge of sustainability. In Macharis, C. and Melo, S., editors, *City distribution and urban freight transport: Multiple perspectives*, pages 37–56. Edward Elger, Cheltenham, UK.

Quak, H. J. and de Koster, M. B. M. (2009). Delivering goods in urban areas: How to deal with urban policy restrictions and the environment. *Transportation Science*, 43(2):211–227.

Reyes, P. M. (2005). Logistics networks: A game theory application for solving the transshipment problem. *Applied Mathematics and Computation*, 168(2):1419–1431.

Roca-Riu, M., Estrada, M., and Fernández, E. (2016). An evaluation of urban consolidation centers through continuous analysis with non-equal market share companies. *Transportation Research Procedia*, 12:370–382.

Schoemaker, J., Allen, J., Huschebek, M., and Monigl, J. (2006). Quantification of urban freight transport effects I. Technical report, Published online, BESTUFS.net.

Stathopoulos, A., Valeri, E., and Marcucci, E. (2012). Stakeholder reactions to urban freight policy innovation. *Journal of Transport Geography*, 22:34–45.

Suzuki, Y. (2011). A new truck-routing approach for reducing fuel consumption and pollutants emission. *Transportation Research Part D: Transport and Environment*, 16(1):73–77.

Tamagawa, D., Taniguchi, E., and Yamada, T. (2010). Evaluating city logistics measures using a multi-agent model. *Procedia-Social and Behavioral Sciences*, 2(3):6002–6012.

Taniguchi, E., Thompson, R. G., and Yamada, T. (2014). Concepts and visions for urban transport and logistics relating to human security. In Taniguchi, E., Fwa, T. F., and Thompson, R. G., editors, *Urban transportation and logistics: Health, safety, and security concerns*, pages 1–30. CRC Press, Boca Raton, USA.

Taniguchi, E., Yamada, T., and Okamoto, M. (2007). Multi-agent modelling for evaluating dynamic vehicle routing and scheduling systems. *Journal of the Eastern Asia Society for Transportation Studies*, 7:933–948.

Teo, J. S. E., Taniguchi, E., and Qureshi, A. G. (2014). Evaluation of load factor control and urban freight road pricing joint schemes with multi-agent systems learning models. *Procedia-Social and Behavioral Sciences*, 125:62–74.

Van Duin, J. H. R., Quak, H. J., and Muñuzuri, J. (2010). New challenges for urban consolidation centres: A case study in The Hague. *Procedia-Social and Behavioral Sciences*, 2(3):6177–6188.

Van Duin, R. J. H., Van Kolck, A., Anand, N., and Taniguchi, E. (2012). Towards an agent-based modelling approach for the evaluation of dynamic usage of urban distribution centres. *Procedia-Social and Behavioral Sciences*, 39:333–348.

Van Heeswijk, W. J. A. (2017). *Consolidation and coordination in urban freight logistics*. PhD thesis, University of Twente, Enschede, The Netherlands.

Van Heeswijk, W. J. A., Mes, M. R. K., and Schutten, J. M. J. (2016). An agent-based simulation framework to evaluate urban logistics schemes. *Lecture Notes in Computer Science*, 9855:369–383.

Van Heeswijk, W. J. A., Mes, M. R. K., and Schutten, J. M. J. (2017). The delivery dispatching problem with time windows for urban consolidation centers. *Transportation Science, Articles in Advance*, pages 1–19.

Van Woensel, T., Kerbache, L., Peremans, H., and Vandaele, N. (2008). Vehicle routing with dynamic travel times: A queueing approach. *European Journal of Operational Research*, 186(3):990–1007.

Verlinde, S., Macharis, C., Milan, L., and Kin, B. (2014). Does a mobile depot make urban deliveries faster, more sustainable and more economically viable: Results of a pilot test in Brussels. *Transportation Research Procedia*, 4:361–373.

Wangapisit, O., Taniguchi, E., Teo, J. S. E., and Qureshi, A. G. (2014). Multi-agent systems modelling for evaluating joint delivery systems. *Procedia-Social and Behavioral Sciences*, 125:472–483.

Yang, D. and Odani, M. (2007). An analysis on city freight cooperative transportation system using game theory. *Journal of the Eastern Asia Society for Transportation Studies*, 7:989–1001.

Appendix A

Table 5: *Overview of results per scheme, relative to default scheme. (averages over 10 replications)*

Scheme	**Variable**						**Financial performance**				**Environmental performance**		
	A	B	C	D	E	F	Rec	Car	Shp	UCC	PM	Distance	# trucks
0							0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00
1						x	-1.04%	-1.57%	9.96%	-0.06%	-0.03%	-1.74%	-2.04%
2					x		1.64%	-0.29%	765.65%	-1.59%	45.86%	55.16%	55.23%
3					x	x	2.32%	1.22%	752.38%	-1.55%	46.31%	54.92%	55.18%
4				x			-1.68%	-0.57%	-0.08%	-0.09%	-0.58%	0.17%	0.52%
5				x		x	-2.24%	2.17%	-2.82%	0.12%	2.89%	1.85%	1.81%
6				x	x		1.87%	2.06%	704.72%	-1.62%	47.93%	57.50%	57.69%
7				x	x	x	0.45%	0.62%	744.93%	-1.43%	45.77%	54.16%	54.40%
8			x				1.22%	-1.97%	21.64%	0.53%	6.89%	1.29%	0.60%
9			x			x	-1.90%	0.08%	-9.46%	0.59%	3.49%	1.81%	1.59%
10			x		x		1.61%	0.33%	569.42%	-1.57%	46.42%	55.83%	55.95%
11			x		x	x	4.75%	-0.10%	604.71%	-1.73%	46.53%	56.34%	56.40%
12			x	x			-1.05%	0.99%	4.48%	0.99%	8.70%	4.56%	3.56%
13			x	x		x	1.48%	-0.15%	-38.30%	0.00%	1.82%	5.32%	6.46%
14			x	x	x		-4.48%	-0.97%	589.83%	-1.40%	45.22%	54.34%	54.34%
15			x	x	x	x	4.07%	0.07%	596.23%	-1.42%	45.94%	54.97%	55.12%
16		x					4.39%	1.91%	-85.05%	25.08%	29.56%	17.53%	16.27%
17		x				x	3.18%	4.69%	-92.30%	27.03%	33.49%	21.25%	20.16%
18		x			x		15.70%	0.30%	708.64%	11.09%	31.10%	17.59%	16.39%
19		x			x	x	9.52%	-1.48%	698.46%	6.11%	36.03%	31.77%	30.94%
20		x		x			2.99%	-2.28%	-71.37%	22.76%	24.99%	12.56%	11.20%
21		x		x		x	6.23%	-2.41%	-81.06%	24.82%	24.98%	12.54%	11.43%
22		x		x	x		3.74%	-2.55%	770.98%	3.54%	37.61%	37.31%	36.95%
23		x		x	x	x	-1.16%	1.85%	815.10%	-1.16%	46.21%	53.99%	54.18%
24		x	x				8.77%	-0.15%	-92.86%	28.33%	30.86%	17.30%	16.02%
25		x	x			x	6.29%	1.31%	-94.11%	28.60%	32.12%	18.80%	17.54%
26		x	x		x		12.57%	0.33%	632.35%	12.51%	31.14%	17.63%	16.47%
27		x	x		x	x	8.32%	1.54%	627.12%	12.93%	32.42%	19.17%	17.97%
28		x	x	x			4.10%	1.32%	-93.85%	28.54%	32.27%	18.99%	17.64%
29		x	x	x		x	5.76%	-1.55%	-91.20%	28.58%	30.00%	16.27%	14.64%
30		x	x	x	x		14.10%	-2.04%	637.26%	12.22%	29.31%	15.44%	14.11%
31		x	x	x	x	x	10.90%	1.15%	621.22%	12.86%	31.97%	18.63%	17.47%
32	x						-2.32%	0.49%	-8.44%	0.15%	-0.05%	0.61%	1.10%
33	x					x	-5.02%	-1.23%	2.64%	0.04%	-0.40%	-0.48%	-0.57%
34	x				x		-4.13%	3.14%	705.40%	-1.64%	48.44%	57.58%	57.71%
35	x				x	x	0.71%	2.17%	726.84%	-1.74%	47.39%	56.07%	56.15%
36	x			x			-3.25%	1.51%	-4.70%	0.06%	1.75%	1.88%	2.01%
37	x			x		x	-1.40%	0.75%	3.42%	0.12%	0.67%	0.27%	0.38%
38	x			x	x		-6.50%	2.33%	787.84%	-0.87%	45.28%	51.76%	51.86%
39	x			x	x	x	-1.93%	-0.75%	690.45%	-1.66%	45.21%	54.33%	54.59%
40	x		x				-2.05%	-1.30%	11.64%	0.75%	5.78%	1.34%	1.14%
41	x		x			x	-1.04%	-0.64%	-7.96%	0.24%	1.22%	0.01%	0.06%
42	x		x		x		-0.44%	2.01%	648.31%	-1.18%	46.04%	53.40%	53.31%
43	x		x		x	x	4.11%	1.51%	581.97%	-1.52%	47.13%	56.34%	56.61%
44	x		x	x			0.13%	-3.10%	2.64%	0.09%	-0.96%	-2.40%	-2.71%
45	x		x	x		x	-2.12%	0.53%	11.33%	0.71%	9.21%	4.02%	2.97%
46	x		x	x	x		-0.98%	0.13%	547.49%	-1.86%	46.79%	56.69%	56.76%
47	x		x	x	x	x	2.64%	0.79%	605.08%	-1.57%	46.30%	55.00%	55.18%
48	x	x					5.37%	1.73%	-83.05%	24.37%	28.89%	17.15%	15.81%
49	x	x				x	5.01%	0.23%	-87.42%	26.11%	28.82%	16.31%	15.17%
50	x	x			x		6.01%	1.20%	804.32%	-1.19%	46.05%	54.29%	54.28%
51	x	x			x	x	9.93%	1.63%	692.66%	11.78%	32.46%	19.22%	17.94%
52	x	x		x			4.47%	0.76%	-90.03%	27.40%	30.52%	17.46%	16.20%
53	x	x		x		x	2.59%	1.15%	-91.88%	28.04%	31.65%	18.47%	17.12%
54	x	x		x	x		7.14%	-0.36%	727.60%	10.12%	32.08%	20.70%	19.70%
55	x	x		x	x	x	6.39%	-4.25%	765.07%	10.65%	27.37%	13.13%	11.73%
56	x	x	x				6.24%	1.14%	-93.80%	28.56%	31.88%	18.52%	17.48%
57	x	x	x			x	7.70%	0.74%	-93.27%	28.32%	31.67%	18.27%	16.95%
58	x	x	x		x		8.71%	1.15%	623.98%	12.94%	32.04%	18.72%	17.41%
59	x	x	x		x	x	10.48%	1.06%	626.34%	12.85%	31.84%	18.47%	17.22%
60	x	x	x	x			13.78%	-0.09%	-93.40%	28.18%	30.89%	17.34%	16.03%
61	x	x	x	x		x	9.42%	-0.09%	-92.46%	28.39%	30.70%	17.20%	15.99%
62	x	x	x	x	x		10.21%	0.19%	631.10%	12.80%	31.11%	17.60%	16.34%
63	x	x	x	x	x	x	6.62%	-3.93%	643.92%	12.42%	27.75%	13.58%	12.00%

Appendix B

Table 6: *Overview of main effects and two-way interaction effects.*

Var	**Receiver**			**Shipper**			**Carrier**			**UCC**			**PM**			**Vehicle distance**			**# vehicles**		
	Low	Avg	Hgh	Low	Avg	Hgh	Low	Avg	Hgh	Low	Avg	Hgh	Low	Avg	Hgh	Low	Avg	Hgh	Low	Avg	Hgh
A	-1.22	-0.81	-0.40	0.20	0.21	0.22	-6.13	-0.91	4.32	-0.02	0.57	1.16	-1.15	-0.36	0.42	-2.97	-1.26	0.45	-3.03	-1.29	0.46
B	7.37	7.88	8.39	-0.20	-0.19	-0.18	-36.47	-31.82	-27.17	18.22	18.93	19.64	6.34	7.40	8.46	-9.75	-7.49	-5.23	-10.96	-8.64	-6.32
C	1.95	2.53	3.11	-0.44	-0.43	-0.43	-73.49	-66.92	-60.35	1.73	2.37	3.01	-0.76	0.19	1.15	-4.90	-2.64	-0.38	-5.26	-2.91	-0.55
D	-1.44	-1.07	-0.71	-0.88	-0.87	-0.86	-4.58	2.59	9.76	-0.71	-0.15	0.41	-1.42	-0.63	0.17	-1.98	-0.46	1.05	-1.50	-0.23	1.04
E	1.65	2.21	2.78	0.16	0.18	0.19	701.27	721.20	741.13	-10.53	-9.66	-8.80	22.72	23.60	24.49	28.47	30.54	32.60	28.65	30.78	32.92
F	-0.24	0.10	0.43	-0.01	-0.01	0.00	-9.32	-4.41	0.50	-0.22	0.30	0.82	-0.92	-0.15	0.62	-1.70	-0.17	1.36	-1.67	-0.43	0.82
AxB	0.62	1.10	1.58	-0.20	-0.19	-0.18	-4.94	-0.32	4.30	-0.07	0.55	1.17	-0.86	-0.15	0.55	-2.29	-0.63	1.03	-2.40	-0.68	1.05
AxC	0.15	0.62	1.08	-0.23	-0.22	-0.20	1.42	5.18	8.94	-1.23	-0.60	0.03	-0.75	0.01	0.76	-1.24	0.48	2.19	-1.28	0.50	2.27
AxD	0.48	0.94	1.41	-0.35	-0.34	-0.34	-6.60	-3.01	0.58	0.45	1.04	1.63	-1.41	-0.80	-0.20	-4.17	-2.59	-1.00	-4.55	-3.01	-1.48
AxE	-1.30	-0.87	-0.44	0.15	0.16	0.17	-6.81	-1.01	4.79	-0.34	0.21	0.77	-1.07	-0.41	0.24	-2.75	-1.10	0.54	-2.86	-1.16	0.54
AxF	0.19	0.63	1.08	-0.71	-0.70	-0.69	-9.03	-4.35	0.32	0.03	0.61	1.19	-1.74	-0.94	-0.15	-4.18	-2.48	-0.78	-3.92	-2.42	-0.91
BxC	0.20	0.75	1.30	0.42	0.43	0.44	-1.59	4.28	10.14	1.49	2.16	2.83	-2.72	-1.69	-0.65	-5.80	-3.57	-1.34	-5.87	-3.58	-1.29
BxD	-0.75	-0.24	0.28	-1.04	-1.03	-1.02	6.65	10.85	15.05	-0.78	-0.18	0.42	-1.24	-0.68	-0.13	-2.38	-0.84	0.71	-1.87	-0.61	0.65
BxE	-0.22	0.47	1.15	-0.98	-0.97	-0.95	52.45	57.12	61.80	-8.72	-7.88	-7.03	-21.03	-20.28	-19.53	-25.76	-23.58	-21.40	-25.87	-23.60	-21.32
BxF	-1.95	-1.42	-0.88	-0.06	-0.05	-0.04	-4.56	0.25	5.05	-0.20	0.40	1.00	-0.50	0.10	0.70	-1.99	-0.39	1.21	-1.94	-0.73	0.48
CxD	0.61	1.03	1.44	0.00	0.01	0.02	-11.13	-7.25	-3.38	-0.55	0.05	0.64	-0.55	-0.03	0.50	-1.33	0.11	1.55	-1.31	-0.23	0.86
CxE	-0.46	0.06	0.58	0.19	0.21	0.22	-69.79	-62.75	-55.71	-0.11	0.78	1.67	-3.70	-2.89	-2.09	-6.14	-4.07	-1.99	-6.24	-4.11	-1.99
CxF	-0.26	0.22	0.69	0.27	0.28	0.29	-2.38	0.97	4.31	-0.91	-0.31	0.29	-0.79	-0.11	0.57	-0.95	0.65	2.25	-0.25	1.03	2.30
DxE	-1.86	-1.35	-0.84	-0.48	-0.46	-0.45	-2.49	4.30	11.10	-0.68	-0.15	0.37	-0.89	-0.33	0.23	-1.80	-0.37	1.06	-1.40	-0.09	1.22
DxF	-0.49	0.10	0.68	-0.09	-0.08	-0.07	1.77	7.44	13.11	-1.17	-0.38	0.41	-0.35	0.50	1.35	-1.23	0.91	3.05	-1.24	1.01	3.25
ExF	0.09	0.61	1.13	-0.24	-0.23	-0.23	-3.32	0.51	4.33	-0.60	-0.08	0.45	-0.72	-0.21	0.30	-2.01	-0.54	0.93	-2.04	-0.81	0.42