Chapter 12
Social and Environmental Impact of Advances in Economically Driven Transport Optimization: Case Study in Automobile Distribution

Thomas Wensing

Abstract Contemporary optimization methods have shown to save costs and increase revenues in many operative fields of application. While human planners tend to focus on parts of the overall problem, these methods use the computational power of modern computers to deeply explore the solution space and thus enable decision-making on a superior level.

The methods itself are well explored by the operations research community, where much less is known about their effect on problem aspects that are not directly focused. This study examines the impact of improvements in optimization methods on the economic, social, and environmental dimension within the context of a realistic case in automobile distribution.

Two planning methods are compared. The first adapts a step-by-step planning technique typically used by human planners; the second addresses the problem from an overall perspective. The comparison is based on two scenarios. One assumes that a fixed amount of transport orders has to be fulfilled, while the other considers a freight market from which transport opportunities can be freely selected for fulfillment.

When the workload is fixed, advancements appear to be beneficial in the economic, social, and environmental dimension at the same time. In contrast the economic dimension is improved disproportionately in the freight market scenario. It can be shown that the objectives of the economic dimension are in conflict to a certain extent with those of the other two dimensions.

Keywords Transport operations • Finished vehicle logistics • Capacitated vehicle routing problem

---

T. Wensing (✉)
INFORM GmbH, Aachen, Germany
e-mail: thomas.wensing@inform-software.com

M. Brandenburg et al. (eds.), *Social and Environmental Dimensions of Organizations and Supply Chains*, Greening of Industry Networks Studies 5,
DOI 10.1007/978-3-319-59587-0_12

12.1 Introduction

It is a widespread tendency to consider advances in operative transportation optimization that enable more efficient transport order fulfillment as beneficial for all parties involved. The transport company profits from lower costs and may also create higher revenues, e.g., when truck utilization is improved. At the same time, the drivers' productivity is increased which can justify better payments and/or less working time. Customers potentially benefit from faster fulfillment and more reliable arrival estimates. Finally, well-utilized distance-minimal trips also reduce emissions, noise, and road utilization which creates a positive effect on the environment.

It seems like a clear issue that any initiative to improve operative transport efficiency is beneficial for multiple parties involved. The aim of this study is to examine this assumption more carefully in terms of testing the following hypothesis:

Hypothesis 12.1 *Optimization advances in transport operations planning – in terms of increasing the potential to create efficient trips – induce improvements in an economic, social, and environmental sense at the same time.*

Against the backdrop of distributing cars via road from a terminal to a network of dealerships, it will be shown that the hypothesis can be supported if the workload may not be changed. When orders emerge from a freight market and may thus also be rejected, the three dimensions do not share a common objective anymore. The observed phenomenon relates to the well-studied rebound effect in the energy sector; see Greening et al. (2000).

The paper is organized as follows. Section 2 defines the underlying problem of planning automobile transports on a daily basis. The relevant literature is reviewed in Sect. 3, and this study's contribution to it is clarified. Section 4 establishes the study case by describing the data generation process, performance indicators, and scenarios that will be explored. Section 5 presents two planning methods, a basic method that creates one trip at a time and an advanced method that creates a whole transport plan of multiple trips simultaneously from a holistic perspective. In order to test Hypothesis 12.1, the solution quality of both methods is experimentally compared in Sect. 6. Finally, Sect. 7 summarizes the results and implications.

12.2 Problem Definition

Road transports of cars require special trucks and trailers called car transporters that cannot be used for other types of freight. As a result, the empty mile factor is much higher than in the general freight sector. This circumstance in connection with the inherent problem complexity of feasibly matching cars with suitable transporters creates interesting opportunities for decision support methods.

The practical planning problem considered in this study emerges from the short-term dispatch operations of a car transporting company. It comprises the day-to-day planning of car transports from a single terminal to a number of contracted dealerships. The task can be subdivided into four levels as depicted in Fig. 12.1. On level A, cars are selected for shipment either from the available stock of contracted cars or from a market place. Each selected car has then to be assigned to a transporter (level B). Knowing the set of cars to deliver, a stop sequence can be created for each transporter (level C). Finally, the packing problem is solved on level D, i.e., it is decided, where each car should be placed on the particular transporter.

Several types of restrictions limit the possible combination of cars that may be loaded on a trailer. For simplicity, this study only considers the capacity aspect of the transporter, where the following capacity model is used. Each car is mapped to a proper size category so that a feasible load of a transporter can be expressed by a combination of these categories, called a load pattern. For example, such a pattern may state that a transporter is capable of holding five medium- and four large-sized cars. The overall capacity is represented by the set of all feasible load patterns. To allow substitution, each car may also belong to multiple size categories of which only one needs to be covered by a pattern.

Examples of further restrictions that have been addressed in recent INFORM projects are the following. It may not be allowed to assign two vehicles to the same transporter, e.g., if different brands must not be mixed. A transporter may be unsuitable to approach a certain dealership, e.g., because it is too large to maneuver on the site. Road segments like tree-lined avenues or narrow bridges may not be passed by a transporter with its top deck loaded. Dealerships may exhibit opening hours beyond which a transporter cannot be received. Finally, the driver must take frequent breaks from driving and working, required by law in many countries and induced by common sense everywhere.

For the purpose of this study an optimization model is considered that covers levels A, B, and C. The objective is to maximize overall profits defined by total revenue for successful deliveries minus total direct costs that are incurred per trip, per dealer stop, and per kilometer traveled. There is a homogeneous fleet of trans-

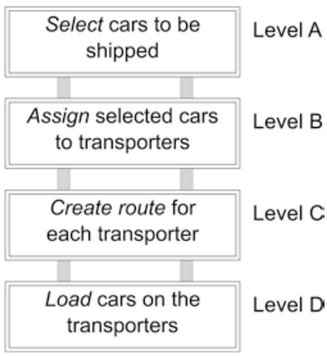

Fig. 12.1 Levels of day-to-day planning based on (Agbegha et al. 1998)

porters which is assumed to be sufficient to cover any amount of trips per day. The capacity of the transporters is modeled via load patterns as described above. No further constraints are considered.

12.3 Literature

There is a substantial literature on providing decision support to the operative planning of finished vehicle transports.

The earliest works are due to Agbegha et al. (1998) and Agbegha (1992). They focus on the packing aspect of the problem, i.e., level D in Fig. 12.1. They assume that a transporter's capacity can be modeled as a set of slots that are suitable to hold certain vehicles, where a slot may block another slot, in the sense that the former must be cleared before a vehicle assigned to the latter can be unloaded. The model and solution approach aim to minimize reload operations for a given set of vehicles to be delivered in a fixed sequence. The authors develop an exact branch-and-bound algorithm to determine blocking-minimal assignments. The same problem is later revisited by Lin (2010).

Tadei et al. (2002) study the problem of building delivery loads against the backdrop of a real-world case in Italy. They address levels A, B, and C of Fig. 12.1 as follows. The selection (level A) is controlled via delivery revenues and an urgency factor for each car reflecting the costs of postponing the order fulfillment to the next day. The capacity aspect (level B) is linearized by introducing single-dimension *length equivalents* for both the transporters and typical vehicle models. The routing aspect (level C) is eliminated from the problem by creating tight regional clusters, where a trip may only visit locations that belong to the same cluster. The minimization of the number of different dealerships to be visited is part of the objective, where the actual routing costs within a region are neglected. The problem is heuristically solved by combining a matheuristic, using a MIP solver for the loading problem, with a neighborhood search method.

Dell'Amico et al. (2014) address levels B, C, and D, where they extend the focus of Tadei et al. (2002) in two ways: They represent the capacity requirements (levels B and D) in a more detailed fashion and consider routing costs (level C) explicitly without defining regional clusters. Their capacity model combines the approaches described above in the sense that they divide the available space on the transporter into platforms. Each platform is modeled as in Tadei et al. (2002), and a platform blocks another one analogously to the slots in Agbegha et al. (1998) and Agbegha (1992), where blocking is completely prohibited here. The problem is solved by an iterative local search heuristic using a branch-and-bound algorithm for the loading problem.

Cordeau et al. (2015) study the problem considered by Dell'Amico et al. (2014) in a multi-period context with uncertainty. They use a framework similar to that of Dell'Amico et al. (2014) to solve the daily planning problems with preselected vehicles. In addition they also address level A by a rule-based selection routine that distributes the available cars over the remaining days and particularly selects the volumes to be considered in the current run.

A real-world case in China addressing levels B and C is studied by Hu et al. (2015). Analogously to this paper, they model the transporter's capacity via loading pattern sets. All available cars must be delivered. Splitting loads for the same dealer is not allowed, and the objective is to minimize the total travel distance of all transporters used. The problem is solved heuristically by an evolutionary algorithm.

The contribution of this study to the existing literature is twofold. Firstly, the daily problem scope addressed in Hu et al. (2015) is extended to the selection of vehicles (level A) and studied moreover within a multiple-period context. Secondly, the sociological and environmental impact of economically driven optimization advances is studied by the example of a realistic case in vehicle routing.

12.4 Case

The overall case is based on 269 cities in the German state of North Rhine-Westphalia that are delivered with cars by a transport company. Every working day, around 250 cars newly arrive and are ready for shipment, which means a rate of 1.5 cars per 100,000 inhabitants per day. Each city's demand is proportional to its population. Details on data generation are given in Sect. 4.1. According to the German national agency for road transport, the Kraftfahrt-Bundesamt, 633,643 cars were licensed in North Rhine-Westphalia in 2015, which is around 2000 cars per working day including Saturdays, i.e., the study considers a market-share of 12.5%.

Hypothesis 12.1 is tested within the context of two scenarios; see Sect. 4.2. One scenario comprises a fixed set of transport orders that must be fulfilled, whereas the other scenario also allows the rejection of orders. The impact is monitored by three performance indicators that represent the transport company's profitability (economic dimension), driver productivity (social dimension), and environmental load (environmental dimension); see Sect. 4.3. For a clearer focus, the customer perspective is not directly represented.

One may argue that an increase of productivity will always have the negative social effect of reducing the required workforce. In contrast, this study follows the premise that there is no reason to fulfill a task with more effort than is necessary. Increasing an employee's performance capability is thus preferred over generating a long-term stable amount of work from equivalent sets of tasks.

12.4.1 Data Generation

Daily transport orders are sampled as follows. There are 292 dealerships in total, where there is at least one dealer in each of the 269 cities plus one or more additional dealers per 200,000 inhabitants in the larger cities. A dealership is chosen with 30% probability to exhibit demand on the particular day, so that on every day there are new arrivals for approximately 100 cities. The amount of cars per selected dealership is drawn from a normal distribution with mean $\mu = 3.5 \cdot pp$ and covariance $Cov = 0.5$, where pp is the associated population in 100,000 inhabitants. Samples are rounded up or down with equal probability, where numbers lower than or equal to zero cause that the according dealership is skipped. Two model sizes (medium and large) are distinguished that arrive in equal fractions. The transport company uses transporters of a single type that may carry nine medium-sized cars, seven large-sized cars, or a mixed load of five medium- and three large-sized cars, where a medium-sized car may substitute a large-sized car. The payment per delivery consists of a fixed sum, randomly sampled from €[25,40], and a kilometer-dependent fraction, where the price per kilometer is randomly sampled from €[0.7, 0.9].

For simplicity, the distance (d) between two dealerships is approximated by the direct distance (dd) with correction factors: $d = 1.2 \cdot dd + 5$ [km]. Each trip creates fixed costs of €50, plus €20 per dealer stop plus €2 per kilometer driven from terminal via the assigned dealers back to the terminal.

12.4.2 Scenarios

Two scenarios are considered, both comprising 10 days:

- In *Scenario 1* the transport company has to immediately ship a car on the day of its arrival on the terminal. Such a requirement may directly be prescribed by transport contracts, but it can also result from tightly limited terminal space combined with a high transshipping rate, where cars have to be sent out quickly to create space for new arrivals.
- *Scenario 2* considers the same transport orders as Scenario 1, but the transport company may freely reject or accept to transport a newly arriving car and thus realize or lose its revenue. Rejected cars will be taken care of by competitors, so they do not reappear on the next day.

12.4.3 Performance Indicators

The following figures are used as performance indicators for the economic, social, and environmental dimension of the considered application.

- The *economic* impact is measured by the *overall profit* that is created for the transport company. Overhead expenses are excluded for the purpose of this study, so that the profit is defined as revenue minus direct costs, i.e., fixed costs per trip plus costs per stop and per kilometer driven.

- The *social* dimension is represented by the driver's productivity, which is defined as the ratio of *revenue per working time* required for order fulfillment in this context. It is assumed that transporters travel at 60 km/h on average including breaks. Each trip creates an organizational overhead of 30 min, and 20 min is required to access and leave a dealership. Finally, handling per car, i.e., loading and unloading, takes 10 min overall. The productivity is obviously directly correlated with the maximum payment that the transport company may afford which makes it an important social factor.
- Emissions, noise, and resource consumption are summarized by *travel distance per order distance* here, to characterize the *environmental* impact in relation to the basic effort of a transport order. Travel distance (of a trip) refers to the total distance covered by the transporter when going from the terminal via the dealerships to visit back to the terminal. Order distance is the distance from the terminal to a car's destination if it gets delivered without detour.

12.5 Solution Methods

Two solution methods are compared, a greedy algorithm (Sect. 5.1) that uses a trip-for-trip strategy to create a solution, and a more advanced method (Sect. 5.2) that creates solutions from a holistic perspective.

12.5.1 Greedy Algorithm

Experienced dispatchers know what quality they may expect of trips that go to a certain dealer region. Once they have enough cars at hand to form a *good* trip, they just create it more or less independently from other opportunities and remaining truck capacities.

The greedy algorithm used in this study adapts this widespread manual planning technique. Algorithm 12.1 outlines the heuristic principle. Candidate trips are built by first selecting the dealership to which the most profitable direct trip can be created with the remaining cars. To these seeded trips, the nearest (primary criterion) and most profitable (secondary criterion) cars are added that still find place on the transporter. The resulting trip is locally searched, i.e., the drops are arranged according to the distance-minimal sequence, and non-beneficial stops are removed. In Scenario 1, the trip is then assigned to the transporter and thus added to the solution in any case. In scenario 2, it is only added if it positively contributes to the overall profits. Otherwise the trip is dissolved, and the dealership is not further considered as seed for a trip.

In case that the best possible solution is dominated by one- or two-stop trips, the greedy algorithm is already suitable to create solutions very close to optimality.

```
Create set seeds containing all dealerships;
Create new empty solution sol;
for each transporter t do
    Create new empty trip t_best;
    Create variable d_best to remember a seed;
    for each dealership d in seeds do
        Create new empty trip t_cand;
        for each unassigned car uc going to d sorted by revenue do
            if uc may be assigned to t_cand then
                assign uc to t_cand;
            end
        end
        if benefit(t_cand) > benefit(t_best) then
            t_best = t_cand;
            d_best = d;
        end
    end
    for each unassigned car uc sorted by (1) proximity to d_best and (2)
      revenue do
        if uc may be assigned to t_best then
            assign uc to t_best;
        end
    end
    search(t_best);
    if benefit(t_best) > 0) then
        sol.add(t_best);
    else
        remove d_best from seeds;
        if seeds is not empty then
            reconsider t;
        end
    end
end
```

Algorithm 12.1 Greedy algorithm

12.5.2 Holistic Optimizer

The advanced method is a specific configuration of the vehicle routing solver for automobile distribution used within INFORM. It consists of a set of construction, local search, perturbation, and recombination methods. Algorithm 12.2 depicts the overall procedure. Each iteration may comprise the independent construction of new solutions as well as the deduction of solutions from the pool by perturbation or recombination. All methods are strengthened by a local search method that tries to improve a given solution by relocating and swapping cars either within a trip or between two trips.

In the following, the three subroutines construction, perturbation, and recombination are described on a principle level. Algorithm 12.3 shows the construction subroutine. New solutions are created via one of several implemented construction methods. One of them is the greedy method outlined above; others use different variants of best insertion procedures. All methods can be randomized so that multiple calls may lead to different solutions. Each new solution is locally searched until no more improvements can be found. It is added to the pool of solutions if its objective meets the current acceptance threshold and deleted otherwise.

The perturbation subroutine (Algorithm 12.4) comprises several methods to remove assignments from a given solution, e.g., the random removal of dealership visits from a trip. After perturbation, the new solution is locally searched. It replaces the original solution in the pool if it exhibits a lower objective value.

Finally, the recombination subroutine (Algorithm 12.5) comprises a set of methods to interchange or combine promising structures of two or more (parent) solu-

Algorithm 12.2 General outline of the holistic optimizer

```
Create new empty set of solutions pool;
for 1 to numIterations do
    construction();
    perturbation(pool);
    recombination(pool);
    parameterAdjustment();
end
```

Algorithm 12.3 Construction subroutine

```
for 1 to numConst(iteration) do
    select construction method cm;
    newSolution = cm.createSolution();
    search(newSolution);
    if obj(newSolution) < threshold(pool) then
        pool.add(newSolution);
    end
end
```

```
for 1 to numPert(iteration) do
    select perturbation method pm;
    origSolution = selectSingle(pool);
    newSolution = pm.perturbSolution(origSolution);
    search(newSolution);
    if obj(newSolution) < obj(origSolution) then
        pool.replace(origSolution, newSolution);
    end
end
```

Algorithm 12.4 Perturbation subroutine

```
for 1 to numRec(iteration) do
    select recombination method rm;
    setOfSolutions = selectMultiple(pool);
    newSolution = rm.recombineSolution(setOfSolutions);
    search(newSolution);
    if obj(newSolution) < obj(origSolution) then
        pool.replace(setOfSolutions, newSolution);
    end
end
```

Algorithm 12.5 Recombination subroutine

tions from the pool. Here, for example, new child solutions are generated by mixing the trips of two or more parents. The new solution replaces the whole set of parent solutions if it is better than the best of its parents.

12.5.3 Numerical Example

Figures 12.2 and 12.3 illustrate the solution quality of the methods described in this section by the example of 25 cars that should be transported from Aachen to 11 destinations in North Rhine-Westphalia. One drop is depicted as a small circle, two drops as a larger circle, and three drops as two nested circles. For each delivery, a relatively high revenue of €100 is realized so that it is beneficial in any case to fulfill all transport orders.

The greedy algorithm (Fig. 12.2) first considers the southernmost trip with three stops and a total of nine drops that create revenues of €900 at costs of 50€ for organizing the trip plus 60€ for visiting three dealers plus 312€ for driving 156 km. The trip thus generates profits of €478. Second, it creates the northernmost trip with four stops and also nine drops inducing profits of €278. Finally, it finds a last trip cover-

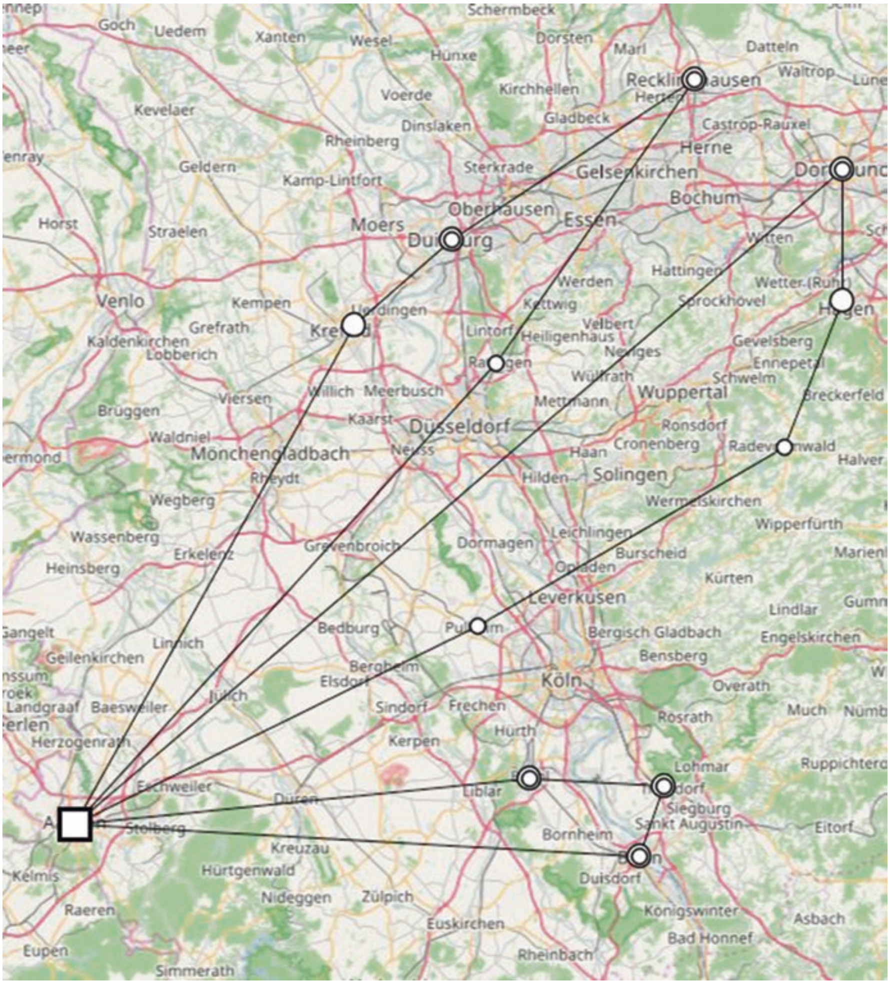

Fig. 12.2 Example solution of the greedy algorithm

ing the remaining seven cars that are delivered to four destinations at a profit of €40. In total, the greedy algorithm arrives at profits of €796 for delivering all 25 cars.

The holistic optimizer (Fig. 12.3) also considers the southernmost trip in its solution, where the two other trips are rearranged. Here, the nine-drop trip going to the northeasternmost dealers creates profits of €210, which is €68 less than the second trip created by the greedy algorithm. However, the remaining seven cars are covered by a trip that creates profits of €182 (vs. €40) leading to an improved overall result of €870.

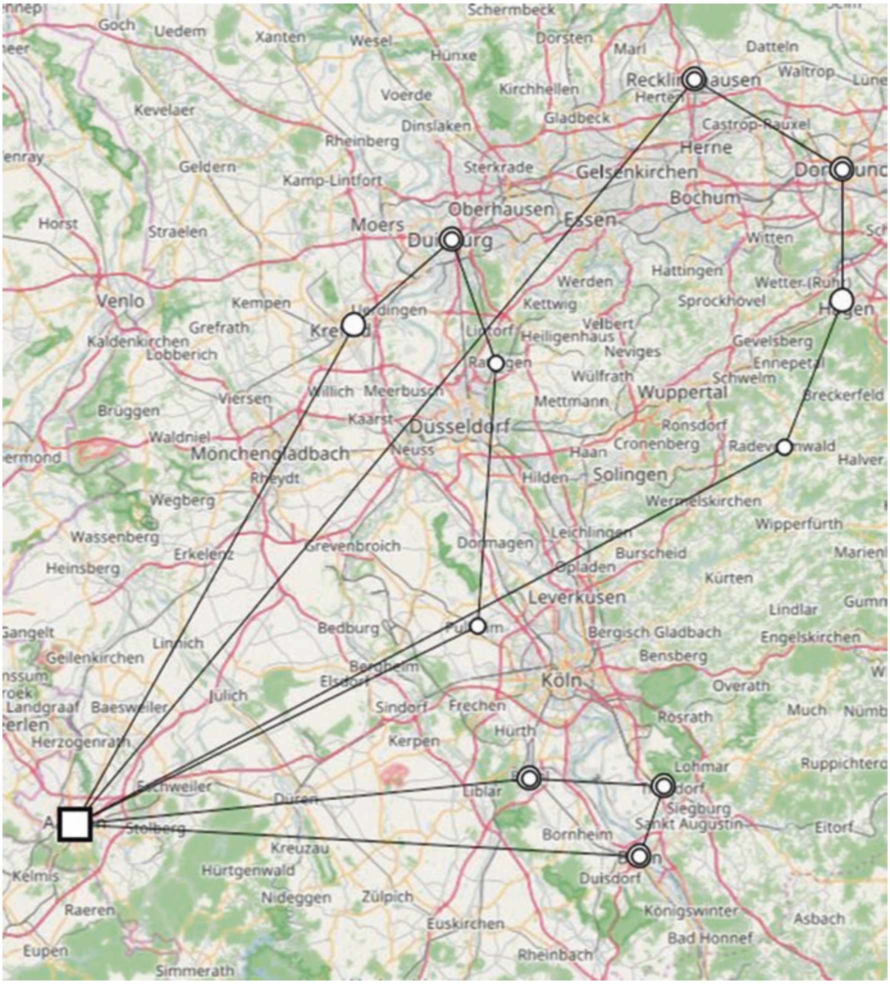

Fig. 12.3 Example solution of the holistic optimizer

12.6 Experiments

On basis of Sects. 4 and 5, Hypothesis 12.1 is replaced by the more focused Hypothesis 12.2 for the experimental part of this study.

Hypothesis 12.2 *Changing the planning principle from the greedy algorithm to the holistic optimizer positively influences profits, driver productivity, and environmental load even if the only explicit objective is to maximize profits.*

To test the hypothesis, each 10-day scenario is once planned by both methods outlined above, where an additional third experiment is conducted on Scenario 2,

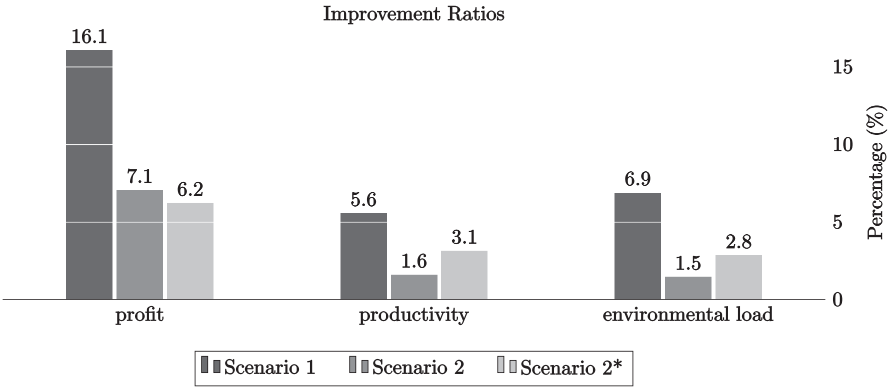

Fig. 12.4 Comparison of performance indicators

referred to as Scenario 2* in the following. The greedy algorithm terminates in a few seconds, while the holistic optimizer requires a run time of approximately 10 min per day of a scenario on a single core of an Intel i7-6820HQ CPU with 2.70 GHz. The improvements within the three considered dimensions are compared in Fig. 12.4. Table 12.1 summarizes the results of the three experiments in detail.

In Scenario 1, every order must be fulfilled on the same day that it becomes available. The only way to improve is therefore to reduce costs. The holistic optimizer creates significant cost savings of 6.9% which increase profits by 16.1%. The improvement potential is partly due to a structural deficit of the greedy algorithm that the holistic optimizer overcomes. By myopically seeking the best dealership to seed a trip in every step, the greedy algorithm tends to create an increasingly heterogeneous set of remaining cars. Since every car must be shipped in this scenario, the according orders finally get covered by very inefficient trips at the cost of the overall solution quality. Besides an increased profitability, there are also substantial improvements in the social and environmental dimensions. Driver productivity is raised by 5.6%, and environmental load is dropped by 6.9%, i.e., at levels comparable to the cost reduction. Since the increase of profits is leveled by the fixed revenue, it is not comparable with that in productivity, emissions, and costs. The results of Scenario 1 therefore support Hypothesis.

Two experiments on Scenario 2 reveal a different picture. Profits are increased by 7.1% when the holistic optimizer is applied to the plain scenario, i.e., at levels comparable to the cost reduction in Scenario 1. However, improvements in the social (1.6%) and environmental dimensions (1.5%) significantly lag behind. The holistic optimizer accepts 3.3% more transport orders than the greedy algorithm in its solution. Therefore, one may speculate that these additional 81 orders only marginally improve profitability at the cost of worsening the possible social and environmental benefits.

To explore the effect in more depth, an additional experiment (Scenario 2*) is conducted for which the rejection rate of the holistic optimizer is increased by inter-

Table 12.1 Summary of the experimental results

	Scenario 1			Scenario 2			Scenario 2*		
	Greedy	Holopt	Ratio	Greedy	Holopt	Ratio	Greedy	Holopt	Ratio
Deliveries	2,563	2,563	±0%	2,405	2,486	3.3%	2,405	2,366	−1.6%
Revenue	370,296	370,296	±0%	349,898	360,450	2.9%	349,898	344,484	−1.6%
Trips	361	352	−2.6%	333	340	2.1%	333	322	−3.4%
Stops	815	733	−11.2%	698	686	−1.7%	698	618	−12.9%
Distance	115,403	107,959	−6.9%	102,215	103,492	1.2%	102,215	97,569	−4.8%
Costs	276,916	259,028	−6.9%	245,350	247,964	1.1%	245,350	232,998	−5.3%
Working time	168,163	158,809	−5.9%	150,215	152,272	1.4%	150,215	143,249	−4.9%
Profit	93,380	111,268	16.1%	104,548	112,486	7.1%	104,548	111,486	6.2%
Productivity	2.202	2.332	5.6%	2.329	2.367	1.6%	2.329	2.405	3.1%
Env. load	0.322	0.301	−6.9%	0.301	0.297	−1.5%	0.301	0.293	−2.8%

nally reducing all revenues to 90% of the original value. Each order thus has to contribute more substantially to the overall result to be accepted for fulfillment, so that the ones that only marginally improved the result in Scenario 2 are most likely rejected now. The solution is evaluated with the original revenues.

As expected, the adjustment leads to a decrease in overall revenues; in fact they even drop by 1.6% compared to the greedy algorithm's solution. Improvements in profitability are only slightly reduced to a ratio of 6.2% which shows that the now rejected cars would indeed only marginally increase overall profits. As a further consequence, improvement ratios of productivity and environmental load are almost doubled to 3.1% and 2.8%.

12.7 Conclusion

The aim of this study was to explore the side effects of economically driven advances in operative transport optimization. Experiments were conducted against the backdrop of a realistic case in automobile distribution that mainly revealed two insights.

When there is a fixed volume of orders to fulfill, advances turn out to be beneficial throughout the economic, social, and environmental dimensions in the considered field of application. Comprehensive benefits are realized even when the objective is to only maximize profits. This is mainly due to the fact that improvements are only possible by reducing costs in this case, which are highly correlated with individual trip efficiency and the reduction of environmental load.

The situation is more ambivalent if revenues are also subject to the optimization. The higher level of efficiency that the optimization advances establish is at least partly used to extend the overall workload. Transport orders whose relatively low margins caused them to be rejected by the greedy algorithm can now be fulfilled profitably. While these additional orders slightly increase overall profits, they worsen presumably the individual trip efficiency. Therefore, the unadjusted holistic optimizer creates significantly less improvement in the social and environmental dimensions when applied to the market-based scenario.

Even though all three dimensions are better served by the holistic method, the additional experiment on Scenario 2 shows that they are in conflict to some extent. By forcing the optimizer to be stricter with the acceptance of transport orders, productivity and environmental load are substantially improved – at the cost of losing profit.

Current developments in the transportation industry suggest that methods like the holistic optimizer presented in this paper will increasingly be used in the near future to support day-to-day planning. To prosper in a competitive industry, companies will have to use the according advances to stabilize or even increase their profits. By showing how economic, social, and environmental aspects may be in conflict, this study stresses the importance of explicitly regarding the effects of advances in optimization methods in all relevant dimensions.

References

Agbegha G (1992) An optimization approach to the auto-carrier problem. PhD thesis, Case Western Reserve University

Agbegha G, Ballou R, Mathur K (1998) Optimizing auto-carrier loading. Transp Sci 32(2):174–188

Cordeau J-F, Dell'Amico M, Falavigna S, Iori M (2015) A rolling horizon algorithm for auto-carrier transportation. Transp Res B Methodol 76:68–80

Dell'Amico M, Falavigna S, Iori M (2014) Optimization of a real-world auto-carrier transportation problem. Transp Sci 49(2):402–419

Greening LA, Greene DL, Difiglio C (2000) Energy efficiency and consumption – the rebound effect – a survey. Energ Policy 28(6):389–401

Hu Z-H, Zhao Y, Tao S, Sheng Z-H (2015) Finished-vehicle transporter routing problem solved by loading pattern discovery. Ann Oper Res 234(1):37–56

Lin C (2010) An exact solving approach to the auto-carrier loading problem. J Soc Trans Traffic Stud 1:93–106

Tadei R, Perboli G, Della Croce FD (2002) A heuristic algorithm for the auto-carrier transportation problem. Transp Sci 36(1):55–62