Master's Thesis

Allocating CO_2 Emissions to Customers on a Distribution Route

S.K. Naber, 303215

April 18, 2012

Supervisors:
Dr. W. van den Heuvel
R. Spliet MSc.
Erasmus Universiteit Rotterdam

Dr. A.W. Veenstra
D.A. de Ree MSc.
TNO innovation for life

Abstract

Currently, there is no generally accepted method to allocate CO_2 emissions to customers on a distribution route. When such a method would be available, the Logistics Service Providers (LSPs) could compete with each other in terms of the amount of CO_2 they emit. Moreover, their customers would be able to choose an environmentally friendly way of transportation. This thesis provides a comparison of one relatively simple and four more advanced allocation methods. The methods based on the theory of cooperative games have been used to share the costs of joint delivery, but this problem differs from the sharing of emissions in several ways. Moreover, earlier comparative studies do not take into account the fact that customers may want to be served within specific time windows. Four ways of penalizing customers for this are suggested as an extension to the existing allocation methods. The methods are applied to both a business case and some hypothetical ones. Their results are compared in terms of fairness, robustness and computational effort. The cooperative game theoretic methods generate more fair allocations, but the simple method is more robust and requires less computation time. Which allocation method is preferred will be dependent on the preferences of the user.

Keywords: CO_2 emissions, distribution route, cooperative game theory, time window penalties

Contents

Chapter 1

Introduction

Over the last decades, issues concerning the environment have received an increasing amount of attention. People from all over the world have become aware of the harm that human activities bring to our planet. These environmentally minded people may be called pessimists by the ones who are not convinced (yet), but the evidence is overwhelming. According to Solomon et al. [10] the climate changes irreversibly due to carbon dioxide (CO_2) emissions. They show that the effects of CO_2 on global warming are largely irriversible for 1,000 years after emissions stop.

The warming of the earth causes both precipitation changes and global sea level rise. Changes in the amount of rainfall can limit human water supplies, worsen the harvest, increase the fire frequency, change ecosystems and cause desertification. In northern Africa and southern Europe the precipitation will decrease with 20% when the global temperature rises with 2°C (see Solomon et al.). Most developed countries have enough funds at their disposal to protect themselves against floods. On the contrary, the consequences to poor countries can be catastrophic. Dasgupta et al. [9] investigated the impact of continued sea level rise for 84 developing countries. Although only a relatively small number of countries (e.g. Vietnam, Egypt, the Bahamas) will face severe impacts, the overall magnitudes for developing countries are sobering. Within this century, hundreds of millions of people are likely to be forced to move due to sea level rise.

A significant part of the globally emitted CO_2 is caused by the transport sector. Although its exact share is hard to determine, the sector is probably responsible for 20 to 30% of the total amount. As other sectors become more efficient in terms of emissions (e.g. energy is saved through better isolation of buildings), the transport sector is likely to get a relatively larger impact on global warming. The fact that several transport modes become more environmentally friendly does not outweigh the increase in freight traffic due to continuous globalization. Figure 1.1Yearly amount of CO_2 per modality (left axis) and transport sector emissions as a % of total man-made emissions (right axis)figure.1.1, which has been retrieved from Fuglestvedt et al. [8], shows the amount of emissions that can be owed to the transport sector as a percentage of the total man-made emissions on its right axis. The steady increase over the 20^{th} century is likely to continue. On the left axis, the graph shows the yearly amount of CO_2 in teragrams (1 teragram = 10^9 kilograms) for rail, road, aviation and shipping. From 1950 to 2000 (and probably also afterwards), there is an enormous increase in the emissions that can directly be related to road transport (mostly trucks). With about 75% of the emissions caused by the whole transport sector in the year 2000, road transport has a large impact on global warming.

As trucks are gaining popularity, the importance of determining the CO_2 they emit increases too. Although reliable methods to calculate emissions of truck routes do exist, a generally accepted one is still missing. Therefore, some environmentally minded LSPs use their own CO_2 calculation method for this purpose. Even though their approaches may be getting more advanced, the correctness remains questionable as most LSPs do not fully explain the methodology they use. This implies that comparing LSPs on the amount of CO_2 they emit seems to be meaningless. The fact that one LSP claims to emit less CO_2 than another one, does not mean that it actually operates more environmentally friendly. Therefore, the degree of 'greenness' of an LSP is not

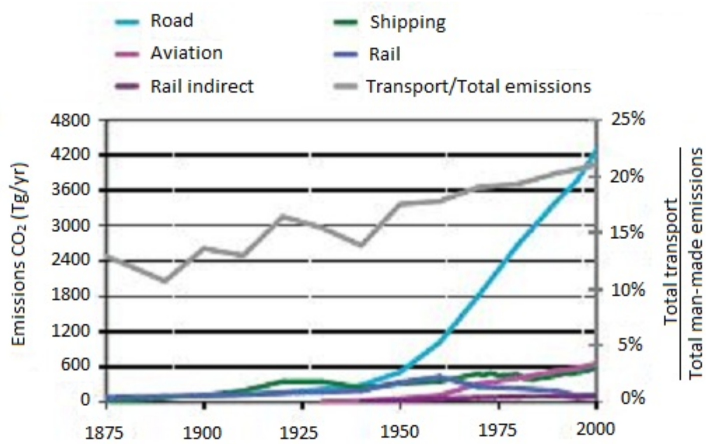

Figure 1.1: Yearly amount of CO_2 per modality (left axis) and transport sector emissions as a % of total man-made emissions (right axis)

easily determined.

Besides the total amount of CO_2 emitted by the trucks of an LSP, an increasing number of people wants to be informed about the amount of CO_2 that is emitted due to their individual actions. They are interested in the so-called 'carbon footprint' of their activities. For most products, the CO_2 involved in their transportation determines an important share of their carbon footprint. To determine the amount of CO_2 that should be assigned to an order because of its road transportation, one should know whether the order occupies the whole load of a vehicle. If it does, the amount of CO_2 that should be allocated to the order is just the total amount of CO_2 emitted in the route. If the order does not occupy the whole truck load, it can be carried in a distribution route in which multiple customers are served. Then, the total amount of CO_2 emitted should somehow be distributed among them. The LSPs that already allocate emissions to orders that are transported jointly often use simple approaches. For example, the amount of CO_2 allocated to a customer is proportional to its order size or distance from the distribution center. Because a generally accepted method for sharing emissions is unavailable, the LSPs cannot be held responsible for using these simple approaches. Due to the small amount of scientific research performed within this field of study, the situation is even more complicated than for the determination of the total emissions of a route.

In this thesis, the allocation of emissions to customers in a distribution route is investigated using the following structure. First of all, the problem definition and the research questions to be answered are presented in the next chapter. The relevant literature for this problem is largely retrieved from research in which the theory of cooperative games is used and it is reviewed in Chapter 3Literature Reviewchapter.3. Chapter 4Core and Pseudo-core of a Gamechapter.4 is completely dedicated to one of the most important elements of cooperative game theory named the core. Afterwards, the allocation methods and some ways to penalize time windows are explained in Chapter 5Methodologychapter.5 and their implementation is discussed in Chapter 6Implementationchapter.6. To evaluate the quality of these methods, they are examined in the test cases defined in Chapter 7Test Caseschapter.7. The results of these tests and the corresponding conclusions can be found in Chapters 8Resultschapter.8 and 9Conclusionschapter.9 respectively. Finally, some suggestions for further research are presented in the last chapter.

Chapter 2

Problem Definition

As explained in Chapter 1Introductionchapter.1, it is both difficult and important to determine the carbon footprint of a product or activity. It was stated that the transport sector, and especially road transport, is responsible for a large share of the CO_2 emissions worldwide. In this research, the focus is on the pick-up and delivery of goods by trucks visiting multiple customers in a route. The total amount of CO_2 emitted in a truck route is calculated by applying a function proposed by Ligterink et al. [19]. If a truck would carry a single order from one location to the other and then return to its starting-point, such a function would be sufficient. However, when multiple customers are served within one route, the total emissions of the route should somehow be distributed among them. As there is no generally accepted method available for allocating emissions to customers on a distribution route, this thesis will provide a comparison of several solution methods for this emission allocation problem.

This chapter presents an overview of some important aspects of the problem and is organized in the following way. First of all, the need for a general method to allocate CO_2 is verified by providing some insight into the current practice of allocating CO_2. Then, the motivation for the research is explained in further detail in the second section. The differences with allocating costs are highlighted in a separate section. Section 2.4Influence of Time Windowssection.2.4 is about the extension of the problem with time windows and Section 2.5A Good Allocation Rulesection.2.5 is dedicated to defining the properties that constitute a good allocation rule. The assumptions made in this research are defined in Section 2.6Assumptionssection.2.6 and the chapter ends with a formulation of the research questions to be answered in this thesis.

2.1 Current Practice

Even though it is not obligatory (yet), some LSPs already calculate and allocate their own emissions as an extra service to their customers. This can either be an indication of the emission before transport (e.g. calculated by a web-tool) or an actual amount of CO_2 kilograms on the waybill afterwards. The airline company KLM provides their customers with a web-tool that calculates the amount of CO_2 that is caused by their freight. The method takes into account the weight of a customer's cargo, the route and flown distances for this route, the historic fuel consumption, the aircraft type and the ratio of passengers to cargo. KLM has received assurance of the calculation method from KPMG Sustainability. Scandinavian Airlines also provides an internet tool to calculate your share of emissions when you travel or send cargo with them. Their methodology has been reviewed by Deloitte & Touche. This company concluded that there are no indications that the method does *not* provide a reasonable output. A company that presents the amount of CO_2 of an order on the waybill is Mars Netherlands. On their website, they claim to be the first food manufacturer in Europe to measure CO_2 emissions at order placement level and to state them on waybills. They rely on calculations made by Kuehne and Nagel.

Note that KLM, Scandinavian Airlines and Mars Netherlands are just examples of companies that calculate and allocate emissions in their own way. Because the actual calculations are not shown, the trustworthiness of the methods is very questionable. Although the companies show that their method has received some kind of approval from another party, this does not guarantee that it functions properly. First of all, it should be mentioned that the parties that gave their approval were hired by the companies themselves. Second, the example below illustrates that the results between these tools can differ quite a lot.

Suppose one wants to send 2,000 kilograms of freight from Amsterdam to Helsinki. Both Scandinavian Airlines and KLM have direct flights between these cities. The tool of Scandinavian Airlines states that your pollution can range from 3.17 to 5.20 tonnes of CO_2 (dependent on the type of airplane). When you fill in the exact same order in the tool of KLM, the pollution is only 2.32 tonnes of CO_2. As neither of the companies explains their methodology properly on their website, the reason for the different CO_2 value cannot be discovered. It is also somewhat strange that both web-tools do not ask you to fill in the volume of your cargo. When your freight fills up e.g. the whole airplane, you should get all the emissions of the flight, whereas if it fills only half of the airplane, the airline company may transport some other freight as well. Moreover, the total amount of CO_2 emitted during a flight should be shared by all freight carried. This implies that it is very questionable to allocate an amount of CO_2 to a package without considering the other cargo carried during the flight. Consequently, all web-tools calculating such individual emissions in advance are only able to give a rough estimation of the actual amount. For more information on these two tools, one can visit the websites. [1]

2.2 Motivation

The above example clearly shows the problem with companies calculating and allocating their emissions themselves. Currently, the correctness and fairness of the calculation of the total amount of emissions, let alone an allocation of these emissions, cannot be guaranteed. Additionally, it is very questionable whether calculation results of two different LSPs are comparable or not. Even though the methods of the LSPs may seem to be reasonable, there is no reliable measure to check them. Therefore, there is a need for a generally acknowledged approach, which is preferably developed by an independent party such as TNO. When the same method is used by several LSPs, customers can actually compare these providers not only based on their tariffs and quality of service but also in terms of emissions. In the future, it may even be obligatory to inform a customer about the amount of emissions caused by a product or whole company. Governments might also introduce taxes on such emissions, in which case a generally accepted method is also needed.

Of course, one can argue that an LSP does not need to inform its customers about the exact amount of CO_2. One could also state that consumers should be content with the CO_2 indication they currently receive from some of the LSPs. However, the competition in the field of logistics can be ruthless and consumers are increasingly demanding. As taking social responsibility by limiting your contribution to climate change is ‘hot’ nowadays, many consumers are willing to pay more for environmentally friendly ways of transportation. For LSPs, this creates the opportunity to distinguish themselves from others in being ‘greener’ instead of cheaper or faster. Although some LSPs may actually care about the environment, the most important thing for such companies is to make a profit or at least to keep up with the competition. When they can attract more customers by being ‘green’, they will find it beneficial to do so.

It is important to point out that LSPs will not be likely to invest in ‘greener’ ways of transportation when they cannot convince their (potential) customers of the fact that they emit less

[1]

Emission calculation tools:
http://sasems.port.se/emissioncalc.cfm?sid=cargo&lang=1&utbryt=0 (Scandinavian Airlines)
http://csr.afklcargo.com/index.cfm/co2/what-you-can-do/co2%20calculator/index.cfm (KLM)

CO_2 than their competitors. A fair comparison can only be made when the same method is applied by multiple LSPs (preferably all of them). In that case, LSPs can actually show that they are operating more environmentally friendly than their competitors. Moreover, customers can consciously choose for a 'green' mode of transportation.

For every party involved, the advantage of a widely accepted allocation method is listed below.

1. **LSPs** can gain a competitive advantage through which they can attract new customers.
2. **Customers** can choose an environmentally friendly way of transportation.
3. **Governments** are able to charge taxes on the amount of CO_2 caused by a product or company as a whole.
4. In the end, **the planet** will benefit from less global warming.

2.3 Differences with Allocating Costs

Some of the cooperative game theoretic methods that will be used in this research have already been used to share the costs of joint delivery. Moreover, they have been compared with each other in earlier studies. If the cost allocation problem would be identical to the emission allocation problem, the need to test and compare these allocation methods extensively with each other would be small. However, there are some important differences between allocating costs and allocating emissions which will be discussed in this section.

2.3.1 Strong Dependence on Route Length

The main expenses of an LSP include truck, driver and fuel costs. Mostly, the truck is owned by the LSP and depreciation costs need to be considered. When the fleet size of the LSP is not sufficient, the costs of hiring a truck should be taken into account. These depreciation or hiring costs can be dependent on the duration and/or distance of the route. Personnel costs will mostly be time dependent and fuel costs are closely related to the route length. The duration of a route is dependent on its distance, but from the differences in shortest and fastest routes suggested by nagivation systems it may be clear that the dependence cannot be characterized as a one-to-one relationship. As the time of both the trucks and its drivers is often more valuable than the amount of money saved by taking the shortest route, LSPs tend to minimize the duration of their trips. For the amount of emissions, the distance (to be) travelled is much more important. Moreover, the amount of CO_2 emitted by a truck is often expressed as a value per kilometer (see e.g. Albrecht or Sullivan et al.). It increases virtually linearly with the length of the route. This implies that the route choice of the LSP has a larger influence on the amount of emissions than on the costs of joint delivery.

2.3.2 Emission Allocation Problem is not always Subadditive

Suppose there are two different distribution routes of which the orders would also fit in a single truck load. For an LSP it is very likely that the costs of the separate routes are higher than the costs of a single route visiting the same set of customers. This is due to the fact that one possibility of merging all customers into a single route, is to drive exactly the same path as in the smaller routes (including at least one intermediate visit to the depot). Due to the triangle inequality, excluding the unnecessary visit(s) to the depot cannot lead to an increase in the amount of kilometers and time required to perform the route. Because costs are generally time and/or kilometer dependent, driving the same path as in the separate routes is a upper bound on the costs made in the merged route. In mathematical terms, such a cost allocation problem is said to be subadditive (the exact definition is given in Section 3.2Theory of Cooperative Gamessection.3.2).

An emission allocation problem is not necessarily subadditive. The amount of CO_2 emitted within a route is not only dependent on the amount of kilometers, but also on the weight of the truck. Even though the effect of the payload may be relatively small, the dependence makes that the triangle inequality described above does not hold for the amount of emissions. Suppose again that two routes are merged into one by driving exactly the same path as in the separate routes. Because the orders of the customers in the second route are carried along the whole path required to visit the customers of the first one, the amount of CO_2 emitted in the single route might be higher than the amount of CO_2 emitted in the two separate routes.

2.3.3 Allocated CO_2 = Total CO_2

In order to make a profit, an LSP will often request more money from its customers than the actual costs that are made to serve them. Therefore, the customers will not only share the costs of delivery, but also the profits of the LSP. On the contrary, the sum of all allocated emissions should exactly equal the total emission of the route. The fact that the total cost does not need to equal the amount demanded from the customers, makes the sharing of costs a lot easier than the sharing of emissions. As an example, consider the medium-sized transportation company Alwex Transportation AB described in Sichwardt [30]. The company distinguishes eight order weight classes and six tariff zones based on the amount of kilometers from the terminal. For every distance and weight of an order, the price can simply be read from a table.

One should notice that when there is no LSP involved in the cooperation, for example in a logistics collaboration between two or more companies, the total costs should equal the sum of the costs allocated to the different parties involved. In the literature, cooperative game theoretic methods have been applied to such kinds of collaborations.

2.3.4 VRPTW Objective does not match with the Quantity to be shared

Another important difference between allocating costs and allocating emissions, is that collaboration between companies is typically based on reducing costs and not necessarily on reducing emissions. Because costs are largely time dependent, most LSPs determine their routes by minimizing the time they require instead of their distance. This implies that the objective of the Vehicle Routing Problem with Time Windows (VRPTW) is quite similar to the quantity to be shared. When allocating emissions, this is generally not true. To observe the effect of the mismatch in objective and the quantity that is shared, consider a simple network consisting of three customers, such as the one displayed in Figure 2.1Network with three customersfigure.2.1. The numbers beside the arrows connecting the customers and the depot indicate the driving times in hours. To simplify matters, assume that the amount of kilometers and emissions increase linearly in these driving times.

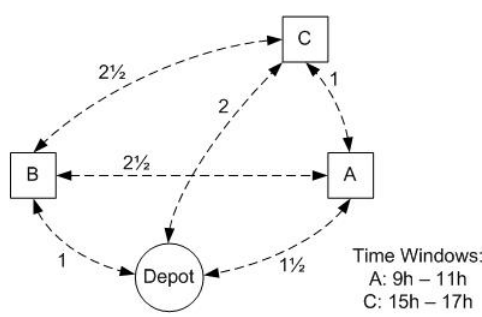

Figure 2.1: Network with three customers

As the time window of customer A ends before the time window of customer C begins, the route should be such that customer A is visited before C. As customer B can be served at any time, it can be visited before A, after C or in between the visits to customers A and C. The three feasible routes that are constructed in this way are displayed in Figure 2.2Possible routes to visit customers A, B and Cfigure.2.2.

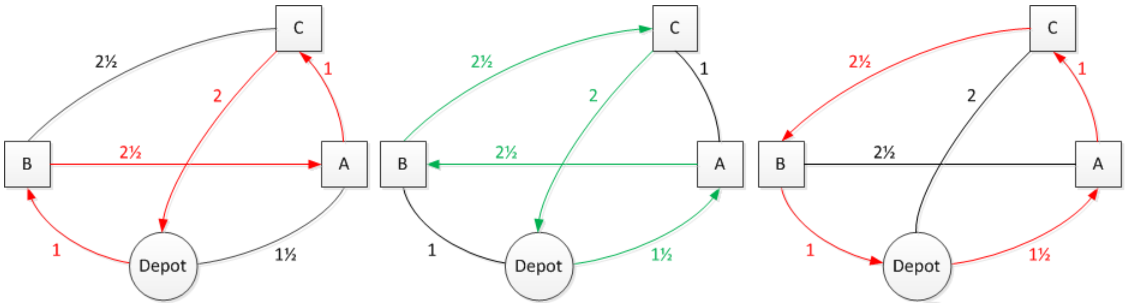

Figure 2.2: Possible routes to visit customers A, B and C

Although the route displayed on the left and right side include fewer kilometers to be driven, these will not be preferred by LSPs that minimize the duration of their routes. Even though the driving times in the red routes are a lot lower (6 and 6.5 hours instead of 8.5), the waiting time of 3 hours between A and C causes the green route to take less time. The additional CO_2 emitted in the green route should be allocated to the customers, even though the LSP would have been able to take one of the red routes. When these differences are large, it may be hard to convince customers of their emission share.

2.4 Influence of Time Windows

In the previous example, customers A and C wanted to be served within specific time windows. These time windows forced the LSP to take another route than it would have preferred without considering time windows. In practice, the LSPs need to cope with such time windows as well. Then, they may be set for convenience of customers only, but may also be the result of legislation. For example, in some cities trucks are not allowed to drive in the main center when shops are open. There will often be some margins on the boundaries of customers' time windows. Because these margins are not clearly defined and may also be customer and situation dependent, the program used to solve the VRPTW cannot take them into account. Instead, the program ensures a solution that completely satisfies all time window requirements.

Whenever time windows lead to detours or waiting time at some point in the route, another solution approach can be desirable. In this section, a detour is defined and its existence is discussed. Particular attention is devoted to extremely large detours.

2.4.1 Definition of a Detour

The LSP always has a preference for a specific route. In this thesis, it is assumed that this is the route that takes the least amount of time. In the preferred trip, the customers are visited in a specific sequence. Time windows or any other extra requirements may be such that the LSP should deviate from this sequence. Then, the LSP is said to be forced to take a detour.

2.4.2 Existence of Detours

Whether or not the LSP has to take a detour is dependent on the objective of the VRPTW and on the combination of time windows of the customers on the route.

VRPTW Objective

To what extent time windows lead to detours is partly dependent on the objective of the VRPTW. When routes are constructed by a program that minimizes their distance, the driver should always take the shortest route. Time windows will often lead to waiting time at some point in the route, because the driver is only supposed to take a detour if this is inevitable. Therefore, time windows are expected to have a minor effect on the amount of kilometers driven. When the objective is to minimize the duration of the route though, a detour will be beneficial if the extra driving time is smaller than the waiting time. The longer the waiting time and the smaller the extra driving time, the more likely it is that a specific detour will be taken. Other objectives might be to minimize the costs or emissions of routes, but for LSPs the most popular objective is to minimize the duration of a route. As costs depend more heavily on the duration of a trip than on its distance, this objective is quite similar to minimizing total costs. However, it is much easier to drive the fastest route (especially when using a navigation system) than to drive the route that minimizes the total costs of the LSP. An additional advantage of using time as the VRPTW objective is that drivers may sometimes have some working hours left to perform other tasks. Moreover, the trucks are earlier available for their next routes.

Combination of Time Windows

Whether a time window causes a detour or not also depends on the situation concerned. To be more precise, the combination of the customers' time windows is important. If only one customer in a route has a time window, the driver can just start driving at that specific moment for which he will be on time at that customer. However, if the customers who are normally supplied just before and just after this customer have a completely different time window, the need for a detour will be hard to prevent. The customers on a route may also have very narrow time windows, but their combination can still be such that the preferred trip is possible.

2.4.3 Extremely Large Detours

It is important to note that poorly chosen time windows may lead to a situation in which the emission allocated to a customer is larger than the emission of driving directly from the depot to the customer and back (i.e. its stand-alone emission). Although this may be hard to explain to customers, one should keep in mind that cooperation is particularly based on saving costs and not necessarily on reducing emissions. Therefore, such an undesirable situation may happen once in a while.

As explained in Section 2.3Differences with Allocating Costssection.2.3, the subadditivity of the emission allocation problem cannot be guaranteed. Especially when time windows lead to extremely large detours, the problem is often not subadditive. As an example think of a truck that has to drive to a specific city, then to a completely different area and finally back to the city which was visited first. This solution may involve more CO_2 emissions than serving the second customer individually. An example of such a situation will be provided in Chapter 4Core and Pseudo-core of a Gamechapter.4.

2.5 A Good Allocation Rule

To be able to decide whether one allocation method performs better than another one, the main desired properties considered in this thesis are fairness and robustness. If customers want to be served within specific time windows, the method preferably also stimulates customers to broaden such time windows. For practical use, it is important that an allocation method has a low computational effort and that it can be understood by customers. All of these features are described in more detail below.

2.5.1 Fairness

The allocation of emissions should be such that every customer has a 'fair' advantage of the joint delivery. Of course, fair can be interpreted in several ways and what seems to be fair for one customer may not be fair for another one. Although it will remain a matter of opinion, an attempt to characterize a fair solution will be presented in Chapter 3Literature Reviewchapter.3.

As the fairness of an approach might also depend on the situation concerned, there will possibly be no single best method in this respect.

2.5.2 Robustness

Besides its fairness, the robustness of an approach is important. This term can be interpreted in multiple ways. In general, robustness can be defined as 'the ability of a method to perform well under changing conditions'. For a method to allocate emissions, it is preferred that slightly different situations do not result in completely different allocations.

The major advantage of methods that are robust in this sense is that customers tend to trust them more. Consider a customer with a constant weekly demand. Such a customer probably expects to receive a relatively stable amount of emissions. When the amount would fluctuate a lot instead, customers may ask the LSP to explain such changes. Of course, a valid explanation is that the amount of CO_2 emitted depends on the route and on the other orders that should be delivered or picked up. Indeed, it is defendable to allocate less emissions to an exact same order when it has more synergy with the other orders in the route. It would even be strange if modifications in the route would not lead to any fluctuation in the amount of CO_2 allocated at all. However, this does not imply that customers will accept any fluctuation in the amount of CO_2 they receive. An LSP may even forfeit its credibility if the amount of emissions allocated to customers varies a lot.

Trustworthiness is not the only advantage of robustness. The property also enhances the possibilities to give reliable forecasts of the amount of CO_2 that will be allocated to a specific order. If such forecasts are available at all LSPs, then customers can select the LSP that will (probably) emit the least amount of CO_2 due to their order. If customers receive their amount of CO_2 after the route has been planned, they can only find the most environmentally friendly LSP by trial and error. As it stimulates the LSPs to operate in a more environmentally friendly way and requires less effort from the customers, it would be best to forecast the CO_2 emissions before the route is planned. To be able to do this, the most appropriate method to allocate the emissions should be determined first. The comparison of allocation methods provided in this research can contribute to a well considered choice. Forecasting the amount of emissions that will be allocated to a specific customer is considered to be beyond the scope of this thesis.

2.5.3 Stimulance to broaden Time Windows

As mentioned before, customers might want to be served within specific time windows. Such restrictions limit the LSP in its possibilities to serve its customers and the fastest trips to do so may no longer be possible. Especially when time windows force the LSP to wait at some point in the route or take a detour, one would like to stimulate the customers to broaden their time windows. Of course, it would be even better if they do not set time windows at all. If a method allocates more emissions to customers with smaller time windows and they are notified about this, then environmentally minded customers are likely to broaden their time windows. The extent to which customers would actually increase the length of their time windows in return for emission reductions can hardly be forecasted and will depend on the type of customers.

2.5.4 Low Computational Effort

The method generating the most fair and robust CO_2 allocations may require a large computational effort. As the purpose of the research is to find a method which can be applied in practice, this is undesirable. Probably, there will be a trade-off between the fairness and robustness of a solution method on one hand and its computation time on the other. Because the computation times of most methods increase in the size of a problem, simple methods having a small computation time will probably gain popularity as the problem is extended.

2.5.5 Understandability

Finally, it is preferable that an allocation method is such that customers can be able to understand the method or at least the reasoning behind it. Most parties would like to have some insight into the way the allocation is done. Especially when customers do not agree on the size of their share, they might ask questions on how it is determined. In such a case, a simplification of the (idea behind the) method can be sufficient. When customers know what factors are (most) relevant for the amount of emissions, they can also contribute to a reduction in the total amount of CO_2 emissions. In this respect, one may think of increasing the length of the time window within which they want to be visited.

2.6 Assumptions

In this research, the following assumptions have been made:

1. The CO_2-function is able to calculate the emissions of a truck route properly.
2. The LSP itself is not responsible for any of the CO_2 emitted in its routes (i.e. all emissions should be allocated to its customers).
3. LSPs prefer to minimize the duration of their routes.
4. When customers dictate time windows, they should be served within these time windows.

2.7 Research Questions

Based on the problem description above, the three main research questions in this thesis are:

1. What constitutes a 'fair' solution to the problem of allocating emissions to customers in a distribution route?
2. Which emission allocation method is preferred based on the characteristics defined in Section 2.5A Good Allocation Rulesection.2.5? Does this depend on the situation concerned?
3. Which way of penalizing customers for their time windows is preferred?

Note that these research questions are answered by the combination of a literature review and the actual implementation and comparison of the different allocation methods.

Chapter 3

Literature Review

In order to get insight into the findings of other researchers, performing an extensive literature study is required. As mentioned in Chapter 2Problem Definitionchapter.2, an important feature of a good solution approach is the fairness of the allocations it generates. Therefore, the way to judge the fairness of an allocation is discussed first. As the sharing of CO_2 emissions is still in its infancy, the literature on this topic in particular is very scarce. However, scientific work about solving gain and cost sharing problems is available. In this field of research, it is common to use the theory of cooperative games. Because of the resemblance of such problems with the sharing of emissions, it seems appropriate to investigate this theory. Section 3.2Theory of Cooperative Gamessection.3.2 is dedicated to the theory of cooperative games and several allocation methods based on this theory are introduced. Because it would be too time consuming to explore all of them extensively, the most promising ones are selected.

3.1 How to judge Fairness?

As stated in the Problem Definition, one of the most important properties of an allocation method is its fairness. At the same time, it is far from straightforward what kind of method should be labeled as fair. To determine the fairness of an allocation, one may use relatively simple statistics such as the variance, coefficient of variation or min-max ratio. Jain et al. [18] come up with a more complicated measure to judge the fairness of an allocation. Their Fairness Index Measure (FIM) is population size independent, metric and scale independent, bounded and continuous and they show that none of the statistics mentioned earlier possesses all these desired properties. The FIM is able to rate the fairness of any allocation by an index number ranging from 0 to 1. A value of 0 implies that the allocation is completely unfair while a value of 1 implies the opposite. In general, fairness is interpreted as allocating an equal share to every participant. If another allocation is regarded as fair instead, the FIM is easily modified. As long as the user knows what is desired in terms of fairness, the FIM is able to generate a number that indicates to what extent a particular allocation respects this fairness.

The problem with allocating emissions to customers in a distribution route is exactly the fact that there is no general agreement on what type of allocation is fair. Some will argue that the allocation should be based on order size, while others would prefer an allocation based on the distance from the depot. When a weighted function of such properties could be used to generate a fair allocation, the FIM would be applicable. However, from the fact that CO_2 emissions are often expressed as grammes per kilometer (see e.g. Albrecht [1], Sullivan et al. [32], among others), one may conclude that the amount of CO_2 largely depends on the length of the route. On the contrary, the relationship between emissions and customer characteristics such as order size/volume and distance from the depot is rather weak. Therefore, it seems inappropriate to define a fair allocation by using a weighted function of such properties.

One may conclude that even the most advanced fairness measure is not sufficient to the problem.

This result stresses the need for a completely different approach. As similar allocation problems have often been tackled by using cooperative game theory, this field of study is investigated in the next section.

3.2 Theory of Cooperative Games

In 1947, cooperative games were introduced by Von Neumann and Morgenstern [24]. The theory of cooperative games can offer guidelines to what kind of allocation methods should be regarded as fair. In this section, the cooperative game that is used in this research is defined first. Then, the desired properties of a solution approach are discussed and finally some allocation methods and their characteristics are described.

3.2.1 Cooperative Game in the Emission Allocation Problem

Let N be the set of all n customers to be served within a route. This total set of customers (or players) is referred to as the grand-coalition and any subset S is called a coalition. Such a coalition includes at least one player and at most all players except for one (i.e. $1 \leq |S| < |N|$). The emission of the route that would be chosen by the LSP to serve the customers in S is denoted by $e(S)$ and is also referred to as the stand-alone emission of S. The same notation is used for a route in which only a single customer i is served. The emission allocation problem can be denoted by (N, e). According to Hougaard [17], such a problem is called *essential* if the amount of CO_2 emitted by serving the players in the grand-coalition in a single route $e(N)$ is smaller than the total amount of emissions incurred with serving them individually. If for all coalitions in N the emissions do not decrease when other players are added, the game is said to be *monotone*. In the following subsections, three important properties of the cooperative game in the emission allocation problem are explained.

Core

Already in 1959, Gillies [13] introduced one of the most crucial concepts in cooperative game theory: the core. For an emission allocation problem, the core is a set of allocations for which no subset S has an incentive to quit the collaboration based on the amount of emissions the players in S receive. In other words, no subset S can allocate its stand-alone emission $e(S)$ in such a way that every player in S gets a lower amount of emissions than it would have in the grand-coalition.

Let x_i be the amount of emissions allocated to customer i. Then, the core is defined as the set of allocations for which all x_i satisfy the following restrictions:

$$\begin{array}{rcll} \sum_{i \in S} x_i & \leq & e(S), & S \subset N \\ \sum_{i \in N} x_i & = & e(N) & \\ x_i & \geq & 0, & i \in N. \end{array}$$

The first constraints describe both the individual and the group rationality conditions. A solution is individually rational when none of the customers is assigned a higher emission than its stand-alone emission. Group rationality implies that every subset of two or more customers is unable to obtain a lower amount of CO_2 by separating itself from the large set of customers. As a single customer is also a subset of the large set, both types of rationality are captured in those constraints.

The second requirement states that the sum of emissions allocated to the customers in N should equal the total CO_2 emitted in the route. Both allocating more and allocating less CO_2 than the actual amount is undesired. When the actual and allocated emissions are exactly the same, the solution is said to be *efficient*.

The final inequalities make sure that every customer gets a nonnegative share of the total amount of emission. Although this may be intuitive, the constraints should not be left out.

Concavity

There is an important distinction between *concave* and *non-concave* problems. A problem (N, e) is said to be concave if for all coalitions S, $S' \subset N$, the following holds:

$$e(S \cup S') + e(S \cap S') \leq e(S) + e(S'). \tag{3.1}$$

If this condition holds for all subsets that do not overlap (i.e. $S \cap S' = \emptyset$), the problem is said to be *subadditive*. This property implies that cooperation between any pair of non-overlapping coalitions yields an emission that is smaller (or at least not larger) than the sum of their stand-alone emissions. As explained in Section 2.3Differences with Allocating Costssection.2.3, subadditivity does not always hold for the emission allocation problem. Because concavity is a stronger requirement than subadditivity, the problem is not guaranteed to satisfy this property either.

Balancedness

Bondareva [3] and Shapley [27] showed that the core of a cooperative game is non-empty if and only if the game is *balanced*. A collection of non-empty subsets $\beta = \{S_1, ..., S_m\}$ of N is said to be balanced if there exists a positive number δ_j for every subset S_j such that for every player in N the following holds:

$$\sum_{j: i \in S_j} \delta_j = 1, \qquad \forall i \in N. \tag{3.2}$$

The collection δ is referred to as a system of balancing weights. The game (N, e) as a whole is said to be balanced if for all such systems the following holds:

$$\sum_{S \subseteq \beta} \delta_S e(S) \geq e(N). \tag{3.3}$$

A few years later, Shapley [28] also showed that a concave allocation problem always has a non-empty core. As the concave allocation problems are a special class of balanced games, this should not be surprising.

3.2.2 Properties of Allocation Methods

Some of the desired properties of an allocation method already defined in the theory of cooperative games are mentioned below. As these properties will be used to determine which of the methods should be regarded as fair, they will also be referred to as the fairness properties.

Stability

A solution is stable if and only if it is in the core. This implies that no subgroup of players can be better off by separating themselves from the total group of players. If the core is non-empty, stability implies efficiency and individual rationality [34].

Efficiency

A method is efficient if for any route, the sum of allocated emissions to the customers is exactly equal to the total amount of CO_2 emitted. As this is one of the constraints defining the core, every stable solution is efficient.

Dummy Player

A dummy player refers to a customer for which the marginal emission of adding it to a route is equal to the emission of handling it separately (i.e. its stand-alone emission). Such a customer is not beneficial for the others and should therefore receive its stand-alone emission. Then, both this customer and the ones already in the route are not harmed and do not benefit either. When the customer would have an advantage of joining the route, the others would be worse off and vice versa. If the customer would be able to reduce its emission in this way, this would be an example of 'free-rider' behaviour which is undesired.

Monotonicity

Another desirable property when considering fairness is the monotonicity property. It may refer to multiple variants.

1. **Coalitional monotonicity**
   When an allocation method is S-monotonic for a subset S, the players in S do not get more (less) emissions when the stand-alone emission of the subset decreases (increases). If this property is satisfied for all subsets, the allocation method is said to be coalitionally monotonic. Young [36] showed that a stable allocation method cannot be coalitionally monotonic for games with more than 4 players. In an example of 5 customers, the author shows that increasing both the value of the grand-coalition and the stand-alone value of one of its subsets can lead to a decrease of the value allocated to two players that belong to the subset.

2. **N-monotonicity**
   An allocation method is N-monotonic if none of the players in N get more (less) emissions when the total emissions of the route decrease (increase). If one or more players in N would get more emissions due to an emission reduction, they will try to block such an emission reduction.

3. **Cross-monotonicity** or **Population monotonicity**
   When a method possesses the cross-monotonicity or population property, the addition of a customer to a route does not lead to an increase in the amount of emission allocated to customers already in that route. This implies that the advantage of customers in a route is 'guaranteed' as long as no one quits the collaboration.

Uniqueness

It is preferable to have a method that guarantees a unique allocation. When applying a method results in a set of solutions instead of just one, it should somehow be defined which solution in the set is preferred over the others. For example, finding the core cannot be regarded as a solution method. Solutions within the core may be preferred over solutions outside the core, but there should be a clear way of picking a solution when the core is empty or includes more than one solution. When a method is not clear in this respect, the exact same situation may lead to a different allocation a second time. Obviously, this is undesired and it reduces the credibility of a method.

Anonymity or Symmetry

When two different customers with the same characteristics (order size, location etc.) are added to the same route, they should receive the same amount of CO_2 emissions.

Additivity

The additivity property concerns two cooperative games with the same set of customers N in a different area. If merging the games (N, v) and (N, w) results in an allocation in which every customer receives the same amount of emissions as before the merge, the allocation method is said to be additive. In that case, it does not matter whether the games are evaluated separately or jointly. Let $x_i(N, v)$ be the emission allocated to customer i in game (N, v). Then, an allocation method is said to be additive if for every pair of games (N, v) and (N, w), the following holds:

$$x_i(N, v) + x_i(N, w) = x_i(N, v + w), \qquad \forall i \in N. \tag{3.4}$$

3.2.3 Allocation Methods

In this section, the allocation methods are introduced very briefly. For a more detailed and mathematical description of the methods used in this research, see Chapter 5Methodologychapter.5. The first two methods are very popular and well known in the field of cooperative games and the other two are relatively new. In the last subsection, some other popular allocation methods are mentioned along with a motivation for not exploring them further within this thesis.

Shapley Value

Already in the early 50's, Shapley [26] introduced the Shapley value. To find the Shapley value of a player, one must compute its marginal emission for all possible permutations of players and then take the average. Shapley showed that the method fulfills the efficiency, symmetry, additivity and dummy player property listed above. It can even be proven that this is the only method that possesses these four properties. Furthermore, the Shapley value is unique and even though the method does not guarantee a core solution, its allocation often belongs to the core (see e.g. Frisk et al. [11]). Concavity of the game even ensures that the allocation belongs to the core [17].

Nucleolus

The Nucleolus was introduced by Schmeidler [25] in 1969. It is a frequently used allocation method that searches for the 'mid-point' of the core. The Nucleolus possesses the symmetry and dummy player property, but fails to satisfy the additivity property. Furthermore, Schmeidler proved that its solution is unique.

Lorenz Allocation

The Lorenz allocation has been defined in 2004 by Arin et al. [2] as the method that searches for that core allocation for which the absolute amount of emissions is as equal as possible for all players. In 1989, Dutta and Ray [5] already introduced the 'egalitarian' allocation as a solution to reach social equality while respecting individual differences. This method is similar to the Lorenz allocation, except for the fact that it is based on the 'Lorenz' core instead of the usual core. In this thesis, the method suggested by Arin et al. is used. The disadvantage of this approach and the one to be described next is that relevant literature is scarce.

Equal Profit Method

In 2006, the Equal Profit Method (EPM) was introduced by Frisk et al. [11]. The method tries to find that core allocation for which the percentual profits of the players w.r.t. their stand-alone values are as equal as possible. Although some later authors refer to and make use of the method, it has not been explored in much detail.

Other Allocation Methods

Due to time limitations, it is not possible to investigate every allocation method defined in the literature. Below, some methods that have not been explored are briefly explained.

First of all, a popular allocation method that has been proposed by Tijs [33] is the Tau-value. According to Hougaard [17], this method is not directly comparable with the other allocation methods because apart from the efficiency and unique solution axiom, it is characterized by different properties. It possesses for example the minimal right property (which is weaker than the additivity property) and the restricted proportionality property. Hougaard stresses the important trade-off between the monotonicity requirements and the stability requirements. A solution approach can be considered as fair if it is either monotonic or guarantees a core solution. As the Tau-value does not satisfy any of the monotonicity properties and is not stable either, it is not explored any further in this thesis.

In some allocation methods, a distinction is made between separable and non-separable costs. Examples of such methods are the Equal Charge Method (ECM), Alternative Cost Avoided Method (ACAM) and Cost Gap Method (CGM) introduced by Tijs and Driessen [34]. Even though they all satisfy the symmetry and efficiency property, these methods are not useful. This is due to the fact that they are based on the idea that some types of costs can directly be assigned to players, while other types should be shared. When sharing emissions though, there is no part that is directly caused by anyone.

Several modifications of the Nucleolus have been suggested as well. Grotte [14] proposed the normalized Nucleolus and some disruption Nucleoli have been studied by Gately [12], Littlechild and Vaidya [20], among others. As the differences between these versions are relatively small, it is defendable to investigate the original version only.

3.2.4 Comparison of the Allocation Methods

To summarize this section, a comparison of the allocation methods using the properties defined in this section is given in the first paragraph. Afterwards, some other comparison studies and their shortcomings are discussed.

Comparing the Nucleolus and Shapley Value using the Fairness Properties

A comparison of the Shapley value and the Nucleolus in terms of the desired properties that are mentioned above is shown in Table 3.1Fairness properties of the Nucleolus and Shapley valuetable.3.1. After a ‘Yes’ indicating that the method does satisfy the property or a ‘No’ indicating that it does not, the source of this specific information is given. Due to their recent introduction, the EPM and Lorenz allocation have not been characterized by these properties yet. Frisk et al. do state that the EPM solution is only defined for balanced problems and for these problems its solution is guaranteed to lie in the core. From its definition, it is clear that the Lorenz allocation can be characterized in the same way.

Table 3.1: Fairness properties of the Nucleolus and Shapley value

	Nucleolus	**Shapley value**
Stability	Yes [25]	If game is concave [17]
Efficiency	Yes [34]	Yes [26]
Dummy player	Yes [21]	Yes [26]
Coalitional monotonicity	No [36]	Yes [29]
N-monotonicity	No [23]	Yes [22]
Cross-monotonicity	No [15]	If game is concave [31]
Uniqueness	Yes [25]	Yes [26]
Symmetry	Yes [21]	Yes [26]
Additivity	No [34]	Yes [26]

Summarizing, the Nucleolus fails to satisfy all the monotonicity properties and it is not additive. Its major advantage is that it is stable, which is in general not true for the Shapley value. However, when the game is concave the Shapley value satisfies all the properties.

Other Comparative Studies

Even though there is a large amount of literature on allocation problems and the different approaches that can be used to tackle them, only a few comparative studies have been published. All authors are very reluctant to draw any definite conclusions though.

Recently, Borderé [4] made a comparison of the Equal Profit Method, Tau-value, Shapley Value and Nucleolus using some practical test cases. He decided that the choice of an allocation method should be dependent on the situation concerned, but does not explicitly mention what method should be preferred in which case. The author does conclude that when there is little synergy between the customers it is important to find a stable solution (i.e., one that is in the core). If the level of synergy in a coalition is high however, stability is almost always guaranteed and the focus can be more on other desired properties. Another conclusion is that the allocation problem is relatively simple for games with a small number of players.

Engevall et al. [7] compare the Nucleolus with some of its modified versions. Because they use only one test instance, they are also reluctant to draw any definite conclusions.

Chapter 4

Core and Pseudo-core of a Game

This chapter is dedicated to the core of a game, which is a central concept in cooperative game theory. It is important to note that in this research, a pseudo-core is used instead of the actual core. This simplification has been made for computational reasons.

4.1 The Core

In the Literature Review, the core has been explained briefly. Its mathematical definition including the rationality and efficiency constraints are presented there.

The reason why solutions in the core are preferred over others is mainly because they are both individually and group rational. When an allocation is rational for every subset of customers, no one has the incentive to quit the collaboration. When sharing costs, a subset may directly quit the cooperation if its allocated costs are higher than its stand-alone costs. By separating themselves from the other customer(s), all players in the subset can be better off. When sharing emissions, customers will probably not stop cooperating when their allocated emission is not individually or group rational. However, rationality should be preferred and could be an important indicator of the fairness of a solution.

4.1.1 Existence of Core Solutions

One should note that it is possible that a solution in the core does not exist. This is more likely to happen when time windows are considered, but can also be true for problems without time windows.

Problems without Time Windows

Generally, the amount of emissions allocated to a customer reduces if one or more customers are added to its route. However, as mentioned in Section 2.3.2Emission Allocation Problem is not always Subadditivesubsection.2.3.2 the emission allocation problem may not be subadditive. This implies that the amount of CO_2 emitted in two separate routes can be lower than the emissions of a single route visiting the same set of customers. It has been explained in Section 2.3.2Emission Allocation Problem is not always Subadditivesubsection.2.3.2 that such a situation can be due to the fact that the orders of the customers visited last are carried over a longer distance. Because a heavier truck load implies a larger amount of emissions per kilometer, more CO_2 may be emitted in the combined route. Because the sum of the stand-alone emissions of the subsets is smaller than the total amount of emissions, it is not possible to allocate at most their stand-alone emission to both of these subsets. This implies that the core is empty.

Problems with Time Windows

While the existence of an empty core without considering time windows seems to be a rare exception, an empty core can easily be caused by poorly chosen time windows. This was already mentioned briefly in Section 2.4.3Extremely Large Detourssubsection.2.4.3 and will now be illustrated by an example.

Suppose that an LSP needs to deliver an order to three customers labeled as A, B and C. It has one depot and one truck at its disposal. Assume that the situation can be visualized by the symmetric graph in Figure 4.1Network with three customersfigure.4.1. The distances are given in hours, so it takes for example half an hour to drive from A to B.

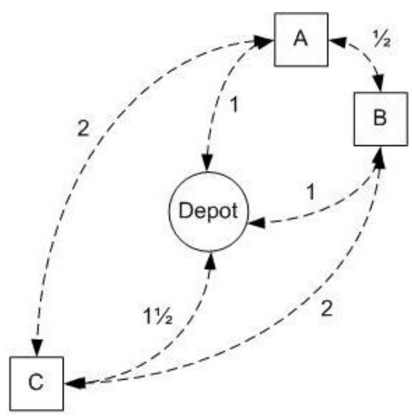

Figure 4.1: Network with three customers

When time windows are not considered, there are four routes that minimize the time travelled. Those are indicated with green arrow in Figure 4.2Fastest routes without time windowsfigure.4.2. All of them have a driving time of 5 hours.

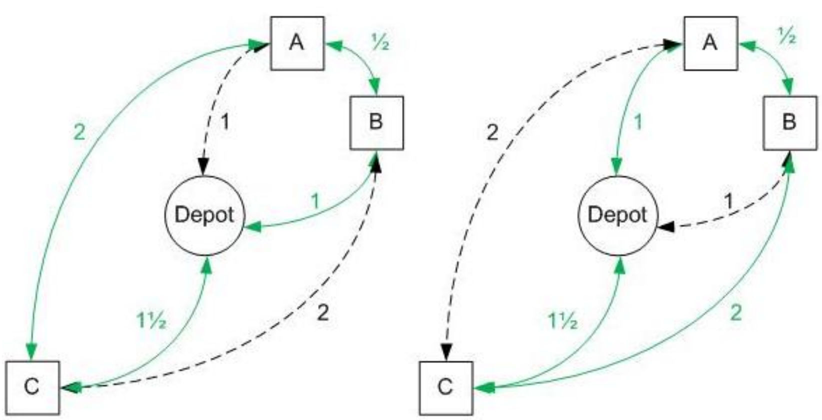

Figure 4.2: Fastest routes without time windows

A solution is in the core if every subset has an allocated emission which is smaller than or equal to its stand-alone emission. To simplify the calculation of emissions in this example, suppose that the amount of emissions woulde be twice the travel time. Denote the emission allocated to customers A, B and C as x_A, x_B and x_C, respectively.

Then, the core is defined by the following restrictions:

$$\begin{aligned} x_A &\leq 4 \\ x_B &\leq 4 \\ x_C &\leq 6 \\ x_A + x_B &\leq 5 \\ x_A + x_C &\leq 9 \\ x_B + x_C &\leq 9 \\ x_A + x_B + x_C &= 10 \\ x_A, x_B, x_C &\geq 0. \end{aligned}$$

As it takes 2 hours to drive up and down to customer A, an upper bound for x_A is 4. The route including only A and B takes 2.5 hours, so the allocation to those customers should be at most 5. As the total route takes 5 hours, 10 units of emission should be divided among all of them. An example of an allocation in the core is $x_A = 2.5$, $x_B = 2.5$ and $x_C = 5$.

Now, suppose that all customers want to be served within a specific time window, e.g. within the ones shown in Figure 4.3Route with time windowsfigure.4.3. If one wants to respect these customers' desires, it is no longer possible to drive one of the routes displayed in green. Instead, the LSP is forced to drive the longer route indicated with the blue arrows.

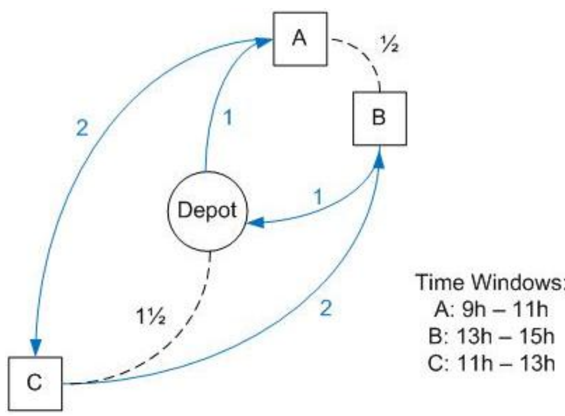

Figure 4.3: Route with time windows

As the routes including only one or two of the three customers can still be driven without taking a detour, the individual and group rationality constraints of the core remain the same. One should notice though that serving A and B together implies a waiting time of at least 1.5 hours. The truck has to be at A before 11h, it will take half an hour to drive from A to B, but the truck will not be (un)loaded at B before 13h. The value of the route with three customers does increase with 2 units of emission due to the time windows. The constraint considering the efficiency of the game is modified into:

$$x_A + x_B + x_C = 12. \tag{4.1}$$

As this efficiency constraint makes it impossible to find an allocation that fulfills all the above constraints, the core is empty. The allocation to customer C should be at most 6, the allocation to A and B together should be at most 5 and the total emission that needs to be allocated is 12. As this is a contradiction, there is no solution. Another way to see this is to sum the values of the constraints regarding the group rationality ($5+9+9=23$). If and only if this value is larger than or equal to twice the value of the efficiency constraint ($2 \text{ x } 12 = 24$), then the core is non-empty. In this case, one or more customers should accept an emission which is higher than they could obtain in a route with fewer customers.

4.2 The Pseudo-core

The 'pseudo-core' is a simplified version of the core for which the set of mathematical restrictions defining its boundaries is almost identical to those defining the core. The only simplification made in the pseudo-core is that the right-hand sides of the group rationality constraints contain an upper bound on the stand-alone emissions of every subset.

When determining these stand-alone emissions it is assumed that the sequence in which customers are served is the same for all subsets. Mostly, the LSP would indeed choose to visit the customers in the same sequence in case some of them are left out. Sometimes, it would be more efficient to rearrange the sequence. Then, the approximation will result in an overestimation of the actual amount of emissions. Because for routes including only two customers, rearranging their sequence results in the same route but then reversed, the VRPTW program would be indifferent between these two options. Therefore, the fact that such a route is not re-optimized when using the pseudo-core cannot affect the stand-alone emission for subsets of two customers. An overestimation can only occur for subsets including at least three customers. As an example, consider the situation depicted on the left side of Figure 4.4Network with four customers (left) and one of the fastest ways to visit them (right)figure.4.4 including four customers to be served within one route.

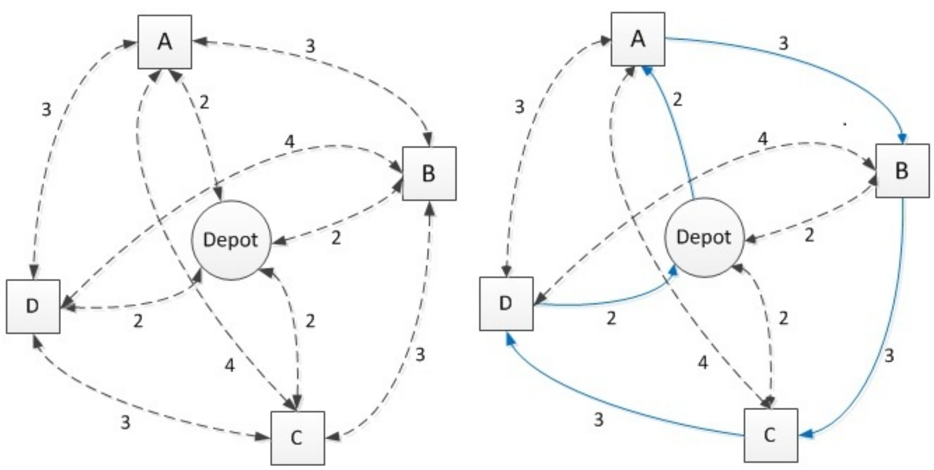

Figure 4.4: Network with four customers (left) and one of the fastest ways to visit them (right)

When assuming that the numbers next to the arrows indicate driving times in hours, one of the fastest ways to serve the customers is A - B - C - D. However, as every customer is equally far away from the depot and from each other, it does not matter which customer is served first. Moreover, as time windows are not taken into account, the direction (clockwise or counter-clockwise) does not make any difference for the total duration of the trip either. Suppose that the computer program used to solve the Vehicle Routing Problem (VRP) proposes to take the sequence shown on the right side of Figure 4.4Network with four customers (left) and one of the fastest ways to visit them (right)figure.4.4.

As in the example above explaining the phenomenon of the empty core, suppose furthermore that the emissions are twice the driving times. By denoting the emissions allocated to customers A, B, C and D by x_A, x_B, x_C and x_D respectively, the core can be defined by the following set of constraints:

$$
\begin{aligned}
x_A, x_B, x_C, x_D &\leq 8 \\
x_A + x_B &\leq 14 \\
x_A + x_C &\leq 16 \\
x_A + x_D &\leq 14 \\
x_B + x_C &\leq 14 \\
x_B + x_D &\leq 16 \\
x_C + x_D &\leq 14 \\
x_A + x_B + x_C &\leq 20 \\
x_A + x_B + x_D &\leq 20 \\
x_A + x_C + x_D &\leq 20 \\
x_B + x_C + x_D &\leq 20 \\
x_A + x_B + x_C + x_D &= 26 \\
x_A, x_B, x_C, x_D &\geq 0.
\end{aligned}
$$

With the exception of the constraints related to subsets {A,B,D} and {A,C,D}, the formulation of the pseudo-core is identical to this one. For these two subsets, assuming the same sequence of customers does not result in the fastest trip. Instead, it would be more efficient to visit customer A in second place in subset {A,B,D} and customer D in second place in subset {A,C,D}. For subset {A,B,D} this is illustrated in Figure 4.5The stand-alone emission of subset {A,B,D} is 20 (left), but its estimation is 22 (right)figure.4.5.

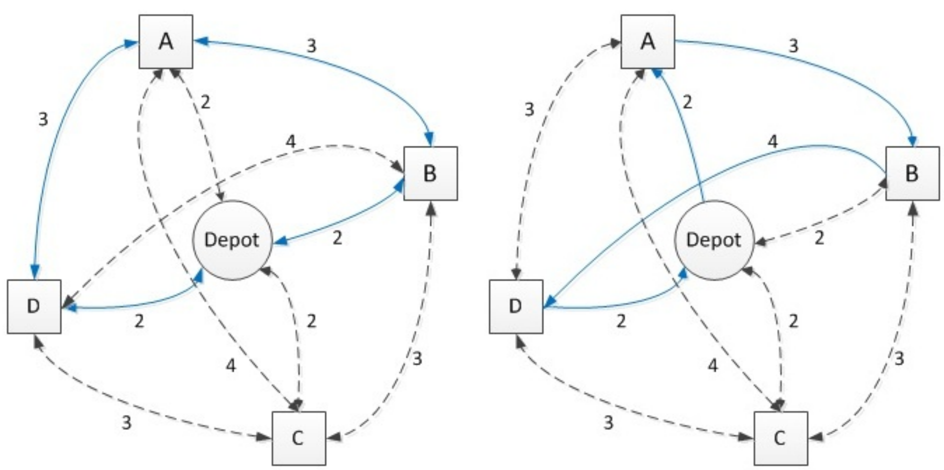

Figure 4.5: The stand-alone emission of subset {A,B,D} is 20 (left), but its estimation is 22 (right)

As the players are identical in terms of distance from the depot and from each other, an allocation in which the total amount of emissions is divided by four seems to be most fair. Indeed, if the actual stand-alone values are used to define the core, the allocation methods that will be described in the next chapter would all come up with this solution. If the pseudo-core is used instead, this equivalence among the players is no longer guaranteed. One may argue that because customers B and C have better alternatives, it is defendable to give them a little less emissions than the other two. Note however that the only reason that B and C have better alternatives is because the VRP program came up with a route that turned out to be profitable for them. If it would have selected the sequence C - D - A - B as the optimal route instead, then it would be the other way around and customers A and D would have deserved a larger benefit according to the pseudo-core.

Even though the VRPTW solution is not unique in the above example, an overestimation of the stand-alone emissions of one or more subsets can also happen in case the solution is unique. If the distance between customers A and B would be increased slightly (e.g. from 3 to 3.01), then the optimal sequences would be B - C - D - A and its reverse. In case time windows are defined in such a way that one of these is preferred over the other, there is only one optimal route. The stand-alone emission of the subsets of 3 customers including both A and B would still be overestimated in the pseudo-core.

From a theoretical point of view, it would of course be better to use the actual stand-alone emissions. However, the computational effort of solving a VRP for every subset in order to find the route with the smallest possible duration is too large. From the previous example, one can observe that a difference between the core and pseudo-core may exist, but its impact is expected to be small. Because one of the features of a good solution approach is a reasonable computation time, it is justified to simplify the calculation of the stand-alone emissions by allowing a small deviation.

Chapter 5

Methodology

There may be several ways to divide an amount of emissions amongst different customers in a route. This chapter discusses the allocation methods which are compared in this research. It also defines some penalty methods to penalize customers in case they want to be served within specific time frames. As already argued in Chapter 2Problem Definitionchapter.2, problems with time windows differ from those without time windows. Because the suggested ways to deal with time windows are an extension to the allocation methods, they are discussed last.

5.1 Allocation Methods

In this section, a relatively simple allocation method is explained first. Afterwards, an extensive description of the four promising cooperative game theoretic methods which were already briefly touched upon in Chapter 3Literature Reviewchapter.3 is given.

5.1.1 Star-method

According to the Star-methodology, the total amount of CO_2 should be divided among customers in such a way that every customer gets a share that is proportional to e.g. its direct distance from the depot. The method derives its name from the shape that appears if one would draw direct lines between every customer and the distribution center. As an example, such a 'star' is shown in Figure 5.1Star-methodfigure.5.1, where the customers are indicated by the letters A to F.

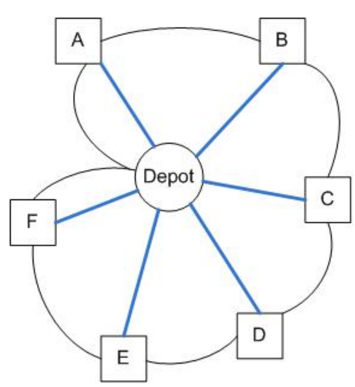

Figure 5.1: Star-method

The property used to base the allocation on can typically be anything and a combination of multiple customer characteristics is also possible. The direct distance between the customer and the depot is just an example of such a characteristic. If one would allocate emissions based on this value only, then a customer's share is calculated by dividing its own direct distance to the depot by the sum of direct distances of all customers in the route to the depot. To derive the CO_2 that should be allocated to the customer, this share is multiplied by the total amount of CO_2 emitted in the route.

Every customer may also get an emission share proportional to the volume and/or weight of its order. As the amount of CO_2 partly depends on the truck load, the weight of a package influences the amount of emissions directly. To what extent the loading space of a truck is filled in terms of volume does not influence the amount of CO_2 emitted directly. When smaller packages are ordered though, there is more space left for other orders. An identical reasoning applies to weight. Just like a maximum volume, every truck also has a maximum allowable weight. Which of these two is most restrictive depends on the type of products transported.

From the above, one may conclude that determining an emission allocation based on a single property like direct distance, size or weight of the order may be too simple. To take into account multiple of these characteristics requires a better insight into their influence for the amount of emissions. A more natural thing to do is to base the emission share on the stand-alone emission of the customer. Instead of using e.g. the distance to the depot, one considers the amount of CO_2 that would have been emitted if the customer would have been served individually. For every customer, such a share is calculated by dividing its stand-alone emission by the sum of all stand-alone emissions. Then, every customer gets the amount of CO_2 for which all customers have the same emission reduction percentage compared to their stand-alone emissions. By using the symbols defined in Chapter 3Literature Reviewchapter.3, the emission that is allocated to customer i is mathematically defined as:

$$x_i = \frac{e(i)}{\sum_{i \in N} e(i)} e(N), \qquad \forall i \in N. \tag{5.1}$$

Despite its simplicity, the Star-method can serve as a good benchmark for more sophisticated approaches. The advantages are the low computational effort and the fact that it is relatively easy to understand for the parties involved. A disadvantage is that some customers may not regard this approach as fair, because the solution can be outside the core. This is due to the fact that the method does not take into account the synergy that customers have. In other words, it is individually rational but not always group rational.

5.1.2 Equal Profit Method

A more advanced way to share the benefits of cooperation equally, is the Equal Profit Method (EPM) introduced by Frisk et al. [11]. Because some customers may be more valuable for the grand-coalition than others, they argue that it is not always fair to give every customer the same emission reduction percentage (as is done in the Star-method). To incorporate the fact that some customers may deserve a larger percentual benefit than others, the method always generates a solution in the core (if it exists). In case the core is empty, the method is unable to come up with a solution. This is due to the fact that the boundaries of the core are restrictions in the Linear Programming (LP) problem that needs to be solved. Mathematically, it is defined as follows:

$$\begin{array}{llcll}
\text{Minimize } f & & & & \\
\text{s.t.} & f & \geq & \frac{x_i}{e(i)} - \frac{x_j}{e(j)}, & \forall (i,j) \in N \\
 & x(S) & \leq & e(S), & S \subset N \\
 & x(N) & = & e(N) & \\
 & x_i & \geq & 0, & \forall i \in N.
\end{array}$$

The relative CO_2 reduction of a single customer i joining a coalition can be defined as $1 - \frac{x_i}{e(i)}$. This value is zero if the stand-alone emission of the customer equals the amount of CO_2 allocated to it in the coalition considered. In this case, the customer does not benefit (in terms of emissions) from joining the coalition. For every pair of customers i and j, the difference in relative savings can be expressed as $\frac{x_i}{e(i)} - \frac{x_j}{e(j)}$. The goal of the EPM is to minimize the difference in relative savings for each pair of customers. Therefore, the objective of the LP problem is to minimize the maximum difference in relative savings (f). When the value of f equals zero, all customers benefit an equal percentage and the solution is identical to the one generated by the Star-method. If and only if these allocations are identical, the solution of the Star-method is in the core of the game. This implies that when the solution of the EPM is known, one can easily check whether the solution of the Star-method is in- or outside the core.

A disadvantage of the EPM is that the number of constraints increases rapidly in the number of customers per route. If the computational effort gets too large, one may consider techniques such as constraint generation to reduce the number of constraints.

5.1.3 Lorenz Allocation

Instead of minimizing the difference in relative savings (which is done by the EPM), one might also be interested in minimizing the absolute difference in the emissions allocated to customers. For customers i and j, this difference is simply expressed as $x_i - x_j$. The Lorenz allocation can be found by solving almost the same LP problem as provided in the section about the EPM. The only difference is that the relative savings are replaced by the differences in absolute allocated emissions.

$$
\begin{array}{llcll}
\text{Minimize } g & & & & \\
\text{s.t.} & g & \geq & x_i - x_j, & \forall (i,j) \in N \\
 & x(S) & \leq & e(S), & S \subset N \\
 & x(N) & = & e(N) & \\
 & x_i & \geq & 0, & \forall i \in N
\end{array}
$$

As mentioned in Chapter 2Problem Definitionchapter.2, fairness can be interpreted in many ways. One may argue e.g. that a core solution is always fair because every customer cannot be better off in a smaller subset of the same customers. If one strives for distributional equality between customers, then the core-solution which divides the emissions most equally among them may be preferred.

The solution for which every customer gets an equal amount of emissions is called the 'equal split' allocation. When such a solution is in the core, this is the Lorenz allocation. As there is no difference in emissions between any pair of customers in this case, g is equal to zero. If the equal split allocation does not belong to the core, the Lorenz allocation will be that core-allocation which is closest to it. In such a situation, the allocation will always be on the boundary of the core.

Just like the EPM, the allocation method will only be able to generate a solution if the core is non-empty and the allocation found will always be in the core. As the properties of the EPM and the Lorenz allcoation are more or less the same, choosing one over the other will be a matter of taste.

5.1.4 Nucleolus

The Nucleolus is an allocation method in which the minimum benefit of a subset is maximized over all subsets. The benefit of a subset is the difference between its stand-alone emission and the amount of emissions allocated to it. By maximizing the minimum difference over all subsets, the Nucleolus tries to satisfy the subset gaining the smallest benefit as much as possible. If the core is non-empty, the resulting allocation is its mid-point.

The benefit of a subset $b(S, x)$ (also referred to as excess) can be quantified by the gap between the emission allocated to the subset $x(S)$ and its stand-alone emission $e(S)$ (i.e. $b(S, x) = e(S) - x(S)$). The imputation set $I(e)$ of the game (N, e) consists of all efficient allocations which are individually rational. Let x and y be two imputations in $I(e)$. According to the definition of the

Nucleolus, the imputation x is better than y if its minimum benefit is larger. This is equivalent to:

$$\min(b(S,x), S \subset N) > \min(b(S,y), S \subset N). \tag{5.2}$$

If the imputation with the largest minimum benefit is unique, then this is the Nucleolus. When two or more imputations have the largest minimum benefit, a more difficult approach is required in order to choose one of them. Let $\theta(x)$ be a vector consisting of the benefits of all subsets for imputation x in a non-decreasing order. This implies that if $i < j$ then $\theta_i(x) \leq \theta_j(x)$ for all $1 \leq i < j \leq n$. The vector $\theta(x)$ is said to be lexicographically larger than $\theta(y)$ if $\theta_1(x) > \theta_1(y)$ or if there is a positive integer q such that $\theta_i(x) = \theta_i(y)$ whenever $i < q$ and $\theta_i(x) > \theta_i(y)$ for $i = q$. The Nucleolus is the imputation with the lexicographically greatest vector of benefits.

When the stand-alone emissions of all subsets are known, the Nucleolus can be found by solving one or more successive LP problems. The first problem to be solved is the following one:

$$\begin{array}{llcll}
\text{Maximize } \delta_1 & & & & \\
\text{s.t.} & x_i & \leq & e(i), & i \in N \\
& x(S) + \delta_1 & \leq & e(S), & S \subset N \\
& x(N) & = & e(N) & \\
& x_i & \geq & 0, & \forall i \in N.
\end{array}$$

When there is a unique solution to this LP problem, this solution is the Nucleolus of the game. As long as the solution is not unique, the problem is reduced by freezing the binding and active constraints including δ_1. According to Engevall et al. [6], these constraints can be found by using the dual formulation of the problem in the following way.

Let $\Pi^1(S)$ be the dual variable belonging to the constraint $x(S) + \delta_1 \leq e(S)$. By maximizing δ_1, at least one of these constraints is turned into an equality. The p_1 binding constraints that have a strictly positive dual variable ($\Pi^{1*}(S) > 0$) are not only binding, but also active. These constraints define the p_1 smallest elements of the lexicographically largest vector of benefits $\theta(x)$. If the solution is not unique, then at least $p_1 + 1$ elements of the lexicgraphically greatest vector should be determined. This is done by solving another LP problem in which the binding and active constraints are fixed. The value of δ_1 is also fixed and a new variable δ_2 to be maximized is introduced. This process is repeated until a unique allocation has been found. As an example, the second LP problem is given below.

$$\begin{array}{llcll}
\text{Maximize } \delta_2 & & & & \\
\text{s.t.} & x_i & \leq & e(i), & i \in N \\
& x(S) + \delta_2 & \leq & e(S), & S \in \{S \subset N \mid \Pi^{1*}(S) = 0\} \\
& x(S) + \delta_1^* & = & e(S), & S \in \{S \subset N \mid \Pi^{1*}(S) > 0\} \\
& x(N) & = & e(N) & \\
& x_i & \geq & 0, & \forall i \in N
\end{array}$$

Because the group rationality constraints of the core are not included in the LP problem(s) that should be solved to find the Nucleolus, this method is able to generate a solution in case the core does not exist. The allocation should belong to an imputation set including those allocations that are efficient and individually rational. If the core is empty, the allocated emission to at least one subset is larger than its stand-alone emission. In that case, the maximum value of δ_1 and possibly also of subsequent δ's will be negative.

Just like in the EPM, the number of constraints increases rapidly in the number of customers per route. Again, a technique like constraint generation might be needed for large instances of the problem. Because multiple LP problems may need to be solved, the computational effort of the Nucleolus will probably be larger than for the EPM and Lorenz allocation. If solving sequential LP problems takes too much computation time, a method that approaches the Nucleolus (e.g. one proposed by Engevall et al. [7]) can also be used. To illustrate the explanation of the Nucleolus, a numerical example is given below.

Numerical Example

Consider the following LP problem:

$$\begin{array}{lrcl}
\text{Maximize } \delta_1 & & & \\
\text{s.t.} & x_A, x_B, x_C & \leq & 5 \\
& x_A + \delta_1 & \leq & 5 \\
& x_B + \delta_1 & \leq & 5 \\
& x_C + \delta_1 & \leq & 5 \\
& x_A + x_B + \delta_1 & \leq & 9 \\
& x_A + x_C + \delta_1 & \leq & 8 \\
& x_B + x_C + \delta_1 & \leq & 9 \\
& x_A + x_B + x_C & = & 11 \\
& x_A, x_B, x_C & \geq & 0.
\end{array}$$

Solving this problem does not result in a unique solution. One can derive that the optimal value of δ_1 equals 1 and the binding constraints are the ones including only x_B and the one belonging to the subset of x_A and x_C. This implies that x_B should be equal to 4 and customers A and C should get 7 together. However, it does not matter whether x_A is 3 and x_C is 4 or the other way around. All allocations in between are also feasible, as long as they sum up to 7. To find the optimal division of emissions between these customers, a second LP problem needs to be solved. In this problem δ_1 is fixed at its maximum value of 1 and δ_2 is the new variable to be maximized.

$$\begin{array}{lrcl}
\text{Maximize } \delta_2 & & & \\
\text{s.t.} & x_A & \leq & 5 \\
& x_B & = & 4 \\
& x_C & \leq & 5 \\
& x_A + \delta_2 & \leq & 5 \\
& x_B + \delta_1^* & = & 5 \\
& x_C + \delta_2 & \leq & 5 \\
& x_A + x_B + \delta_2 & \leq & 9 \\
& x_A + x_C + \delta_1^* & = & 8 \\
& x_B + x_C + \delta_2 & \leq & 9 \\
& x_A + x_B + x_C & = & 11 \\
& x_A, x_B, x_C & \geq & 0
\end{array}$$

By filling in the value of x_B and removing all redundant constraints, the problem can be simplified to:

$$\begin{array}{lrcl}
\text{Maximize } \delta_2 & & & \\
\text{s.t.} & x_A, x_C & \leq & 5 \\
& x_A + \delta_2 & \leq & 5 \\
& x_C + \delta_2 & \leq & 5 \\
& x_A + x_C & = & 7 \\
& x_A, x_C & \geq & 0.
\end{array}$$

The maximum value of δ_2 turns out to be 1.5 and both x_A and x_C should be equal to 3.5. As the solution is unique, the allocation [3.5, 4, 3.5] is the Nucleolus.

'Irrelevant Core Alternatives'

An important difference between the Nucleolus on the one hand and the EPM and Lorenz allocation on the other hand is related to the concept of 'irrelevant core alternatives' (see e.g. Hougaard [17]). As already mentioned, the allocation generated by the Nucleolus is the mid-point of the core (if it exists). Both the EPM and the Lorenz allocation try to find a core-solution which is, according to their reasoning, (closest to) the most desirable outcome. Whenever this 'most desirable' outcome is not in the core, the EPM and Lorenz allocation will be on the boundary of the core. Then, it does not matter how far the core stretches out to the other 'irrelevant' side. As the Nucleolus always tries to find its mid-point, an extension of the core to any side influences its outcome.

Uniqueness

Whereas the solution generated by the EPM and Lorenz allocation are not necessarily unique, Schmeidler proved that the Nucleolus always results in a single value. By solving the LP problem to find the Lorenz allocation, the largest gap between the absolute emissions of two players is minimized. For the other customers, the only requirements are that their allocated emission is within the core and that it is not smaller than the minimum or larger than the maximum allocated amount. As there can be several solutions that fulfill these requirements, the allocation can be rather arbitrary. In case the gap between the minimum and maximum allocated amount is zero however, the solution will always be unique because it can only be the equal split allocation. The same reasoning applies to the EPM as it has been defined by Frisk et al. [11].

5.1.5 Shapley Value

For the calculation of the Shapley value, all possible permutations of a set of customers are needed. For each permutation, the marginal emission of every customer is calculated. This implies that the first customer gets the amount of emissions allocated that is needed to serve the customer on its own; the second customer gets the additional CO_2 that is emitted when it is added to the route and so on. After all emissions are known for every permutation, an average over all the permutations is calculated for each customer. The resulting number is the Shapley value. If $m_i(S)$ denotes the marginal emission of adding customer i to subset S, the emission that should be allocated to this customer is calculated by applying the following formula:

$$x_i = \sum_{S \subseteq N \setminus \{i\}} \frac{s!(n-s-1)!}{n!} m_i(S), \qquad \forall i \in N. \tag{5.3}$$

Although it is an easy method to understand and implement, the computational effort may get too large when a lot of customers are supplied in one route. This is because the number of coalitions that should be considered increases rapidly in the number of customers. Shapley [26] showed that the method generates an allocation which satisfies the efficiency, dummy player and symmetry property explained in Chapter 3Literature Reviewchapter.3. There is a unique solution for every problem instance, but it does not necessarily lie in the core. Note that this implies that a solution might not be rational for every participating customer in a coalition. The test instances will be needed to investigate how often and in what cases the Shapley value is outside the core.

5.1.6 Characterization of the Allocation Methods

In the Literature Review, several properties that contribute to the fairness of an allocation method were described. Earlier research could be used to find out which of these properties are satisfied by the Nucleolus and Shapley value. Such a characterization could not be found for the Star-method, EPM and Lorenz allocation. Using the definitions of these allocation methods, an attempt to characterize them in terms of the same properties is presented below. The results of this discussion are summarized in Table 5.1Fairness properties of the allocation methodstable.5.1. In the last paragraph of this section, the methods are also compared in terms of their understandability.

Table 5.1: Fairness properties of the allocation methods

	EPM	Lorenz	Nucleolus	Shapley	Star
Stability	Yes	Yes	Yes	If game is concave	No
Efficiency	Yes	Yes	Yes	Yes	Yes
Dummy player	Yes	Yes	Yes	Yes	No
Coalitional monotonicity	n.a.	n.a.	No	Yes	Yes
N-monotonicity	n.a.	n.a.	No	Yes	Yes
Cross-monotonicity	n.a.	n.a.	No	If game is concave	No
Uniqueness	No	No	Yes	Yes	Yes
Symmetry	n.a.	n.a.	Yes	Yes	Yes
Additivity	No	If 'equal split allocation'	No	Yes	Yes

Comparing the Allocation Methods using the Fairness Properties

Because the EPM and Lorenz allocation always generate an allocation within the core if it is non-empty, they are said to be *stable* and *efficient*. The Star-method does not guarantee a core allocation and therefore it is not stable. The total emission allocated does equal the total CO_2 emitted on a route, so the method does satisfy the efficiency property. As mentioned in Section 5.1.4Uniquenesssection*.25, the solution of the Lorenz allocation and EPM is not guaranteed to be *unique*. The Star-method does possess the uniqueness property. Furthermore, because every customer always receives an equal percentual profit w.r.t. to its stand-alone emission, the *symmetry* axiom is satisfied by the Star-method, but the *dummy player* property is not. Because the EPM and Lorenz allocation are stable methods, they will allocate the stand-alone emission to a customer if its addition to the grand-coalition does not involve an emission reduction. Therefore, free-rider behaviour of a dummy player is not possible.

Due to the inability of the EPM and Lorenz allocation to generate a unique solution, the monotonicity and symmetry properties are not defined. The formulation of the monotonicity properties are such that the emissions allocated to a customer may not increase or decrease when the game changes. Because their allocations are not guaranteed to be unique, an allocated emission may increase or decrease even if the game remains the same. Because the methods may allocate the emissions to customers that do not receive the minimum or maximum amount of emissions or percentual benefit rather arbitrary, identical customers may receive a different amount of emissions. This implies that the symmetry property cannot be satisfied either. Therefore, the cells of these properties in Table 5.1Fairness properties of the allocation methodstable.5.1 are filled with n.a. (which refers to 'not applicable') for the EPM and Lorenz allocation.

Because the Star-method only considers the stand-alone emission of a single customer, the *coalitional monotonicity* property implies that if the stand-alone emission of a single customer increases (decreases) then its allocated emission should not decrease (increase). Because the total amount of emissions to be shared does not change and every customer receives an equal percentage of its stand-alone emission, a customer will never get a higher amount of emissions when its stand-alone value decreases and vice versa. A method satisfies the *N-monotonicity* property if the emissions allocated to a customer cannot increase (decrease) if the total amount of emissions decreases (increases). In the Star-method, every customer will benefit from a reduction of the total emissions because the profit percentage will increase. The *cross-monotonicity* property implies that the allocated emissions to a customer may only decrease if players are added to the route. This property is not satisfied by the Star-method because the addition of a player may involve a larger increase in the total amount of emissions than is allocated to the added customer. In that case, all customers that were already in the route get a higher amount of emissions than before.

For the Star-method and EPM, the allocation of two separate problems may in total be different from the allocation generated for the single problem that combines them. For the Lorenz allocation, one may conclude that it does satisfy the *additivity* property if its solution is the equal split

Table 5.2: Additivity example: two allocation problems to be merged

Total emission to be divided = 70	Customer A	Customer B
Stand-alone emission	80	60
Star & EPM	40	30
Lorenz	35	35

Total emission to be divided = 90	Customer A	Customer B
Stand-alone emission	70	140
Star & EPM	30	60
Lorenz	45	45

Table 5.3: Additivity example: merged allocation problem

Total emission to be divided = 160	Customer A	Customer B
Stand-alone emission	150	200
Star & EPM	68.6	91.4
Lorenz	80	80

allocation for both of the problems. For other problems, it cannot be guaranteed because the allocations may not be unique. To show that the Star-method and EPM are not additive, consider a simple example in which there are only 2 customers: A and B. The two problems to be merged are displayed in Table 5.2Additivity example: two allocation problems to be mergedtable.5.2.

If summing the allocated emissions will always result in the same allocation, the method is said to be additive. As can be observed from the merged allocation problem in Table 5.3Additivity example: merged allocation problemtable.5.3, the Star-method and EPM do not even satisfy the property for a problem with only two customers. However, they would have resulted in the same allocation if the same percentual profit would have been given in the two separate problems.

Understandability of the Allocation Methods

The extent to which customers are likely to understand the allocation methods can be deduced from their definitions. Due to its simplicity, it may be clear that the Star-method will be most easily understood by customers on a distribution route. As is done by Borderé [4], the Shapley value can be explained relatively easy by using a small example. Even though the methodology of the EPM and Lorenz allocation are rather complicated, the idea behind the methods will be understood by many. The reasoning and the computation of the Nucleolus is much more difficult. In case it is important that customers understand how their amount of emissions has been computed, this method is not preferred.

5.2 Penalizing Time Windows

In the Problem Definition, the effect of time windows was already explained briefly. It was argued that a problem with time windows probably requires a different solution approach as there may be detours and the problem is not necessarily subadditive. From an extensive literature survey, no scientific research in which any of the allocation methods described above is applied to a VRPTW could be retrieved. Therefore, finding an answer to the third research question will be one of the most interesting parts of this thesis.

When customers want their order to be delivered (or picked-up) in specific time windows, the fairness of the above methods needs to be revised. Especially when only some of the customers in a route have a time window, the allocation may no longer be fair. If such time windows imply that a different route with more emissions needs to be driven, these customers should be 'punished' for this. Then, the difference in emissions between the route with and without time windows should somehow be distributed among them. The same is true for situations in which all customers have a time window.

As said, there is no literature available about penalizing customers for the additional costs or emissions caused by their time window(s). It is suggested to apply one of the allocation methods described above to the route that would have been driven if there were no time windows involved. If this route is the same as the one that does take into account the time windows, there is no need to change the allocation. Then, the time windows clearly do not have any impact and the amount of extra emissions due to time windows is zero. Whenever there is a difference in the amount of emissions, the following four ways of dividing the extra emissions among the customers can be applied. The first two methods are relatively simple and intuitive while the second two require some more explanation. All approaches can be used for any number of time windows.

An important remark to be made is that because the LSP minimizes the duration of its trips and not the amount of CO_2 that is emitted, the time windows may in some cases cause a reduction in the total amount of emissions. For such cases, penalty methods 1, 3 and 4 assign an equal negative penalty (i.e. benefit) to all customers. However, the total amount of emission allocated to a customer should never be negative. If the negative penalty happens to be larger than the emission that was allocated to a customer in the route without time windows, the customer should get an allocated emission of zero. In such a situation, the remaining benefit is spread out equally over the other customers in the route.

5.2.1 Simple Penalty Methods

The penalty methods described in this section are rather simple and may not be fair in all cases. However, the computational effort of these methods is expected to be lower than of the more advanced methods described in the next section. Just like for an allocation method, the practical use of a penalty method is dependent on its computation time.

1. Equal Penalty for every Customer with Time Window(s)

Because all customers with a time window extend the problem with a restriction, one could argue that the extra emissions should be divided equally among all customers having at least one time window.

2. Penalty based on Length of Time Window(s)

Instead of punishing all customers with one or more time windows equally, one may want to allocate the extra emissions based on the length of the customers' time windows. The reasoning behind such a penalty is that a narrow time frame can be very restrictive on the sequence in which the customers are visited, while a wider one can be such that is has no effect at all on the chosen route. A customer having a time window of e.g. one hour may receive twice as much of the extra emissions to be divided as a customer with a time window of two hours. When customers are informed about such a method, they will generally be more motivated to extend the length of their time windows. In case customers have multiple time windows, the total length of these time windows is used. As stated above, when time windows lead to a reduction in the total amount of emissions, this benefit is spread out equally over the customers by the other three methods. For this particular penalty method, a smaller time window will imply that a customer will get a smaller benefit.

5.2.2 Advanced Penalty Methods

The previous methods do not take into account whether a customers' time window influences the amount of CO_2 emitted in the route. The following two methods attempt to penalize only these customers whose time windows have an actual 'impact' on the amount of emissions. These methods are also referred to as the impact methods.

3. Equal Penalty for every Customer with Time Window(s) with Impact

One may argue that it is not fair to give customers a penalty in case their time windows does not influence the specific route that is driven. As the LSP is assumed to minimize the duration of its trips, a customers' time window is said to have impact if it is possible to take a route that requires less time in case it would have been ignored. If this is possible, the customer prevents the LSP from taking another (faster) trip. Therefore, such customers should be penalized for their behaviour. In this approach, all these customers get an equal share of the additional emissions caused by all time windows.

4. Penalty based on Impact

An even more advanced way of sharing the extra emissions is to base a customer's penalty on the size of the impact that its time window has. The more time it takes to perform a trip as a consequence of a customer's time window, the higher its penalty should be. For every customer i, a factor is derived from the difference in time between the route with all time windows and the route without considering the time window of customer i. To calculate the percentage of the additional emissions that should be allocated to the customer, one may divide this factor by the sum of the factors of all customers in the route. To compute the penalty, this value only needs to be multiplied by the total amount of additional emissions.

Practical Issues concerning the Impact Methods
A disadvantage of the last two methods is that for every customer with a time window the route has to be re-optimized to check whether ignoring its time window leads to a time reduction. This may require a large computational effort. To overcome this problem, it is suggested to only try to relocate customer i within its route. Consider a route in which all customers A to D have a time window. Originally, they are visited in the following sequence: Depot - A - B - C - D - Depot. To observe the impact of the time window of customer B, one should temporarily ignore it. Then, it is interesting to check whether it is possible to visit B before A (i.e. Depot - B - A - C - D - Depot), after C and after D. Whenever such an option does not violate a time window of any other customer in the route and requires less time than the original trip, one may conclude that the time window of B has a negative influence on the trip.

Another difficulty of the third and fourth method is that one should be able to determine whether a specific sequence of customers can be driven within less time than the original trip. In other words, one should be able to determine the minimum duration of serving a specific sequence of customers. When time windows are not considered, this duration is just the sum of the travel and service times. In case time windows need to be taken into consideration, waiting time may also be involved. Then, the departure time at the depot can influence the duration of a trip. Because determining the departure time that minimizes the duration of a trip is a lot easier when customers have only one time window, this situation is discussed first. Afterwards, the procedure to find the optimal departure time if customers have multiple time windows is described.

Single Time Window per Customer
When customers specify a single time window within which they want to be served, determining the optimal departure time is relatively easy. Consider the situation as depicted in Figure 5.2Departing at moment B minimizes the trip durationfigure.5.2.

The four customers in this trip all have one time window indicated by a green block. For simplicity, the examples in this section do not include any service times at customers. However, the same reasoning applies to situations in which they are taken into account. The Distribution Center (DC) is open during the whole time period. The blue lines indicate the driving time between the stops and the red lines are used for waiting time. The trip can be read from the top to the bottom of the figure and starts at either the earliest possible departure time (A), the latest possible departure time (B) or any moment in between. In this example, the trip will always be finished at the same moment in time. Therefore, departing at the latest possible moment will

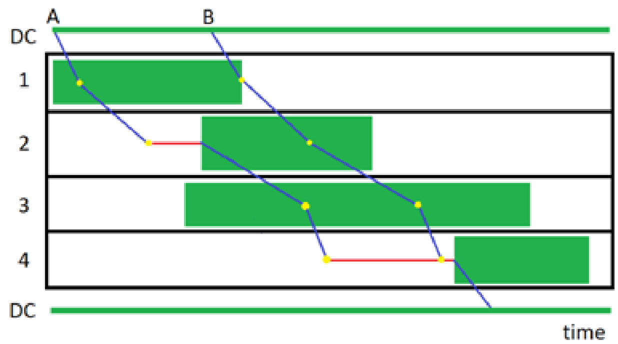

Figure 5.2: Departing at moment B minimizes the trip duration

minimize the total time.

Of course, different departure times do not always lead to identical arrival times at the depot. Consider the same example, but with a small modification in the time window of customer 4. The situation is depicted in Figure 5.3Departing at moment B, C or any moment in between minimizes the trip durationfigure.5.3. Because the last customer can be served earlier, the waiting time can be reduced. When the truck departs from the depot at C, B or any moment in between, the waiting time has been removed completely. Therefore, these moments all minimize the time required to perform the trip. Again, leaving the depot at the latest possible departure time is an optimal solution. It turns out that even though there may be multiple moments of departure for which the duration is minimized, the latest feasible departure time will always be one of them if customers have only one time window. This departure time can be found by the approach described in the section below.

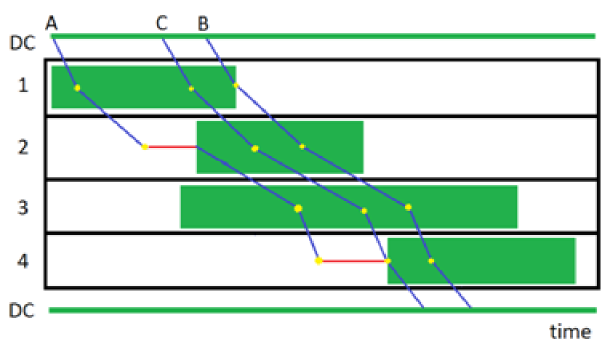

Figure 5.3: Departing at moment B, C or any moment in between minimizes the trip duration

Multiple Time Windows per Customer

When customers are allowed to have multiple time windows, the best departure time is determined less easily. This is due to the fact that one does not know in which time window every customer should be served in order to have the optimal solution. From the example with one time window per customer we do know that it is possible to reduce waiting time by delaying the departure time at the depot. The algorithm used to find the optimal departure time with multiple time windows per customer is derived from this idea. It can be described by the following steps:

1. Take the earliest possible departure time and calculate the total waiting time that would occur when the truck would depart at that time.

   (a) If there is no waiting time at all, the solution is optimal and no other departure times need to be considered.

(b) If there is waiting time at some point in the route, store the solution and its total waiting time (it is not yet known whether the departure time will correspond to the minimum trip duration). Then, go to step 2.

2. Find the first customer in the route where the truck needs to wait and try to reduce this waiting time as much as possible by delaying the departure time of the truck at the depot.
   (a) If it is possible to remove the waiting time completely by delaying the departure time of the truck at the depot, delay the trip by this amount and go to step 3.
   (b) Otherwise, determine the maximum possible delay of the trip according to the time windows of all predecessors. When the visit to these customers is delayed by this amount of time, one of them will be served at the very end of its time window (otherwise, the trip could have been delayed further). Store the solution and go to step 4.
3. Check whether there is still waiting time at some point in the route. If there is, go back to step 2. Otherwise, the optimal solution has been found and no other departure times need to be considered.
4. Despite the possible waiting time reduction that has been achieved, the current departure time may not correspond to the minimum trip duration. The departure time will be such that it minimizes the duration for the specific set of time windows used, but there might be another combination of time windows resulting in a faster trip. In order to investigate whether such a combination exists, one goes through the trip found in the last step until one comes across a customer who is served at the very end of its time window (this particular customer made it impossible to delay the trip any further).
   (a) If this customer has at least one later time window for which the trip is still feasible w.r.t. all the other customers' time windows, the trip will be delayed such that this customer is visited at the beginning of its next time window. Determine the corresponding departure time and calculate the waiting time that would occur when the truck would start at that time.
      i. If there is no waiting time at all, then the solution is optimal and no other departure times need to be considered.
      ii. If there is waiting time at some point in the route, go back to step 2.
   (b) In case the customer does not have any later time window for which the trip is still feasible, no other departure times have to be considered and the best option found up to now is the optimal one.

For a better understanding of this procedure, Figure 5.4Departing at moment D minimizes the trip durationfigure.5.4 shows an example in which 4 customers all have 2 time windows. When departing at the earliest possible departure time (point A), waiting time is included at customers 2 and 4. First, try to reduce the waiting time at customer 2 by delaying the departure time as much as possible. Departing at point C would completely remove the waiting time at customer 2, but customer 1 cannot be served at the time the truck would be visiting it. To make sure that customer 1 can be served, the truck should depart at point B. This implies that the largest possible delay is the difference between point A and B. Because the departure time corresponds to the minimum duration that has been found up to now, it should be stored. Now, other combinations of time windows will be investigated. Because customer 1 was the one that made it impossible to delay the trip any further, one should check whether its moment of delivery can be shifted to the beginning of its next time window. It turns out that when the truck departs at the corresponding departure time (point D), there is no waiting time in the trip. Therefore, the departure time that corresponds with the minimum duration has been found.

Point E is included in the figure to show the difference between a single and multiple time windows per customer. If at least one of the customers has more than one time window, the minimum duration is not guaranteed to be found at the latest possible departure time at the depot. In this example, it would be better to visit customer 2 in its first time window instead of in its last. However, if the optimal departure time would have implied that every customer should be visited in its last time window, then departing at point E would have minimized the trip duration.

Generally, the following can be conluded. If the optimal combination of time windows is known, then the latest possible departure time for which the customers are served within these time windows belongs to the optimal departure times. Because such a departure time is only the ‘latest possible’ if it cannot be delayed any further, one of the customers must be visited at the very end of its time window. Therefore, finding the optimal departure time could also be done by comparing the durations of the trips in which customers are served at the very end of their time windows only. In the example in Figure 5.4Departing at moment D minimizes the trip durationfigure.5.4 one would have to consider at most the eight departure times that correspond to visiting the four customers at the very end of their two time windows. Whenever such a departure time would result in a trip without waiting time, the minimum duration would have been found and the other possibilities would not need to be considered.

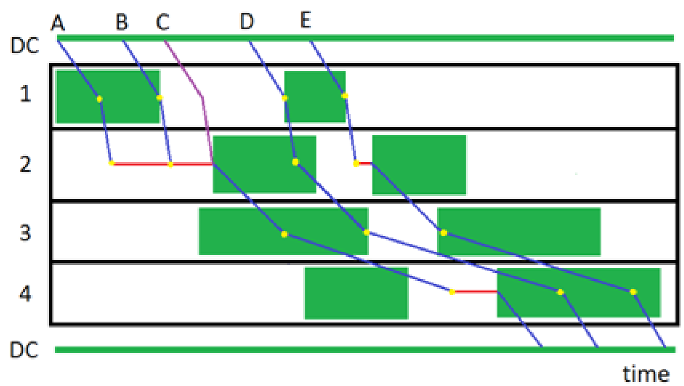

Figure 5.4: Departing at moment D minimizes the trip duration

Tricoire et al. [35] suggest a similar algorithm to find the departure time that minimizes the duration of a trip in which customers have multiple time windows. They also provide the reader with a mathematical proof. In Chapter 6Implementationchapter.6, the implementation of the algorithm described above is explained in a flow chart.

Chapter 6

Implementation

In this chapter, the implementation of the allocation methods is discussed. Because all methods are extensions of the program RESPONSE, which is written in the Visual Basic programming language, a brief description of this program is given first. Afterwards, the implementation of the allocation methods is discussed. The final section is about the approaches to penalize customers for setting time windows.

6.1 RESPONSE

The computer program RESPONSE has been developed by TNO to design, evaluate and optimize logistics networks. It is able to solve a VRP with multiple time windows per depot and per customer. The routes are constructed by using an evolutionary strategy heuristic developed by Homberger and Gehring [16]. They propose two evolutionary metaheuristics for the VRPTW. Their primary objective is to minimize the amount of vehicles used and their secondary objective is to minimize total travel distance. In RESPONSE, the objective of the shortest path algorithm used to solve the VRP is to minimize the duration of trips instead of their distance. For a more detailed description of the evolution strategies, we refer to the article of Homberger and Gehring.

6.2 Allocation Methods

In this section, the implementation of the five allocation methods that can be applied to problems without time windows is explained. The approach to include the Star-method, which was already implemented by TNO, is discussed first. Then, the implementation of the Shapley value is explained briefly. As the EPM, Lorenz Allocation and Nucleolus all derive an allocation by solving one or more LP problems, their implementation is discussed jointly.

6.2.1 Star-methodology

Until now, RESPONSE was only able to allocate emissions according to the Star-methodology. After the routes are constructed, the program can allocate the CO_2 emitted in every route to the customers served in it. In the Graphical User Interface (GUI), the user can select whether the allocation should be based on the star-distance, -time, -costs, -emission, -load, -volume or any combination of these by checking the accompanying checkboxes (see Figure 6.1Star-methodology in RESPONSEfigure.6.1). Then, the chosen properties are used to calculate the share of the total route emission that should be allocated to every customer. As an example, the star-time is the time needed to drive from the depot to a specific customer and back to the depot. This is often just twice the time needed to drive from the depot to the customer (or back). When the allocation is based on time only, the share of a customer equals its own star-time divided by the sum of the star-times of all customers in its route. The emission allocated to the customer is then equal to this share multiplied by the total emission of the route. When multiple boxes are checked, e.g. both time and distance, then the share of a customer is its star-time multiplied by its star-distance divided by the sum of these values for all customers in the route. As has been explained in Chapter 5Methodologychapter.5, in this research the Star-method will be based on the star-emission of customers only.

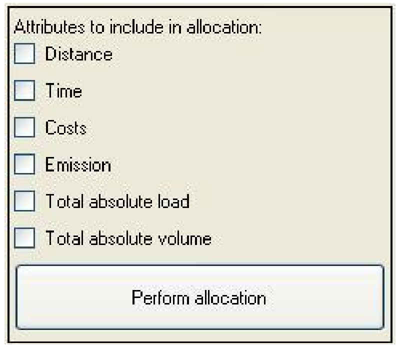

Figure 6.1: Star-methodology in RESPONSE

6.2.2 Shapley Value

As explained in Chapter 5Methodologychapter.5, the Shapley value is the average marginal emission of a customer over all possible permutations of customers. In order to apply this method, the marginal emission m_i of adding customer i to all possible subsets not already including this customer should be known first. Then, calculating the Shapley value for customer i is just a matter of filling in the formula given in Section 5.1.5Shapley Valuesubsection.5.1.5.

6.2.3 Nucleolus, EPM and Lorenz Allocation

The Nucleolus, EPM and Lorenz allocation generate an allocation within the core (if it exists) by solving one or more LP problems. The Nucleolus is the mid-point of the core, the EPM computes the core-allocation with the most equal percentual benefit for every customer and the Lorenz allocation is the core-allocation for which the emissions allocated to every customer are as equal as possible.

Before being able to solve any of the LP problems, some computations need to be made. According to the definition of the core, any subset of players should not get more CO_2 allocated than it would emit on its own. As all three methods guarantee a solution in the core (if it exists), the emission of serving every subset of customers needs to be known. Because it takes too much computational effort to optimize the route for every subset of customers, it is assumed that every subset of customers would be served in the same sequence as in the large coalition. This implies that the stand-alone emission of a subset of three or more customers can be larger than its actual value. Implementing the methods in this way guarantees a solution in the 'pseudo-core' which has been explained in more detail in Chapter 4Core and Pseudo-core of a Gamechapter.4.

When all variables needed are available in Visual Basic, the program calls another computer program called AIMMS. This program is specialized in solving LP problems and it is integrated with Visual Basic very easily. For the EPM and Lorenz allocation, the solution is retrieved from AIMMS, the computation time is determined and the allocation is stored in an Excel file. The calculation of the Nucleolus might require some additional steps. As explained in Chapter 5Methodologychapter.5, the first LP problem to be solved may not have a unique solution. To check whether the allocation is the Nucleolus of the game or not, the dual variables of the binding constraints are used. These can easily be retrieved from AIMMS as well. As long as the solution of the LP problem is not unique, the binding and active constraints are fixed and AIMMS is called again to solve another LP problem.

6.3 Penalty Methods

When some (or all) of the customers in a route want to be served within specific time windows, they can be penalized for this undesired behaviour. This can be done by using one of the approaches described in Section 5.2Penalizing Time Windowssection.5.2. All of these methods require that an allocation of emissions for the route without considering time windows is available. Therefore, one should first determine a solution by using one of the allocation methods described above. Because the penalties do not depend on this allocation but only require that there is one to add the penalty to, there will be four different allocations of the additional emissions generated by applying the penalty methods to a route. Only in case a penalty is of such size that the total emission allocated to a customer becomes zero, the sizes of the penalties may differ among the allocations determined for the route without time windows. When such an allocation is known, the optimal route taking into account the time windows should be determined by the VRPTW program in RESPONSE. If there is no difference in emissions between this route and the one without considering any time windows, there is no need for time window penalties. When the time windows are such that it is no longer possible to take the trip desired by the LSP, the penalties as defined in Section 5.2Penalizing Time Windowssection.5.2 should be determined. This is done in the following way.

The first approach only requires a check whether a customer has at least one time window. The implementation of the second penalty method consists of a simple calculation to determine the total time window length per customer. As they require the individual impact of a customer on the trip duration, implementing the last two approaches is more challenging. In Chapter 5Methodologychapter.5 it has been explained that the impact of a customer is determined by shifting it to other positions within the route while ignoring its time window(s). By departing as early as possible, it can be checked whether such a sequence is feasible w.r.t. the time windows of the other customers. If it is, the minimum duration of the route that appears should be calculated to see whether this sequence would have been a better option for the LSP. If it is, then the customer is said to have impact. Because the program was not yet able to determine the minimum duration of a specific sequence of customers, an algorithm to find the departure time that minimizes the duration was suggested in the previous chapter. Due to the loops inside the algorithm, its implementation is presented as a flow chart in Figure 6.2Flow chart describing the implementation of the algorithm used to determine the optimal departure timefigure.6.2.

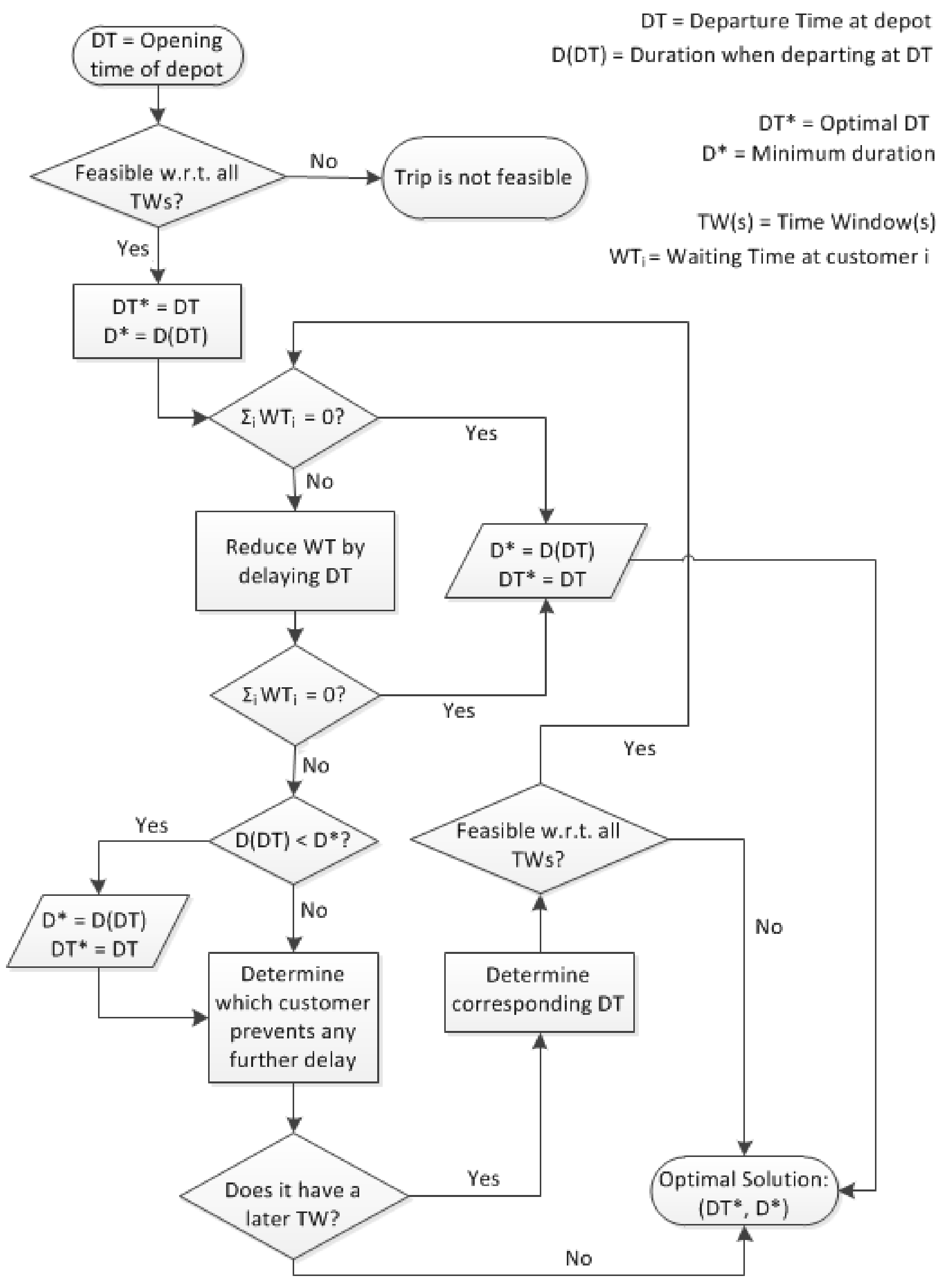

Figure 6.2: Flow chart describing the implementation of the algorithm used to determine the optimal departure time

Chapter 7

Test Cases

The allocation methods are applied to a business case of a Dutch retailer which is available at TNO and to some hypothetical cases. The inclusion of a business case will show that the research is not only scientifically relevant, but that results can also be used in practice. The hypothetical cases will be used to investigate what kind of results appear in 'extreme' scenarios in terms of locations and order sizes for different numbers of customers. For both types of cases, the emission of a route is calculated by applying a truck emission function used by TNO. As the results are dependent on this function, its behaviour is discussed in the first section of this chapter.

If a method works well under the 'normal' conditions of the business case, it can be considered as satisfactory because of its practical use. However, if a method has acceptable results in all test cases it is even better. Such a method is more robust and thus generally applicable. Its use can be extended to any other LSP. Also, the method does not have to be revised when properties such as the location and order size of customers change. It is possible that in some situations one method is best, while in others another one is preferred. In this case it is interesting to know which factors are important when selecting an allocation method. For example, it could be that a method is only preferred when the number of customers in a route is relatively low. A reason for this can be that the computational effort may be too large for routes with more customers.

7.1 Truck Emission Function

The function that will be used to calculate truck emissions is retrieved from Ligterink et al. [19]. According to their model, the emissions are dependent on the velocity, total vehicle mass and specific power of the truck. The specific power is the rated engine power in kilowatt divided by the total vehicle weight. They base their model on a sample of heavy-duty trucks that more or less represents the Dutch fleet composition. Because truck characteristics differ among countries and the business case that will be studied in this thesis includes trips performed by heavy-duty trucks within the Netherlands, this is an important advantage of this study in comparison with others. The CO_2-function is a fleet-average estimate and as every truck has its own characteristics; deviations of 20% may exist. However, they conclude that their method provides better predictions and more accurate emissions than previously developed models. This is due to the fact that they monitored real-world emissions in common traffic situations by using accurate on-road emission measurements. Before, emissions were derived using engine tests and chassis dynamometer tests, in combination with extensive modeling. Apart from this technical advantage, the function is capable of modelling emissions with varying payload. In a distribution (or collection) route, the vehicle weight is likely to decrease (respectively increase) along the way. According to Ligterink et al., their model is unique in the sense that it describes the emission of heavy-duty trucks representative for the European situation while it simultaneously allows for a continuous variation in vehicle weight.

To illustrate the effect of varying payload on the amount of CO_2 per kilometer, the function is plotted in Figure 7.1Truck emission functionfigure.7.1. The weight of the vehicle itself is excluded from the payload which is displayed on the horizontal axis. Because the influence of the payload on the amount of emissions is dependent on the velocity of the truck, the function is plotted for multiple speeds. Assuming that truck drivers respect the maximum speed of 80 km/h on Dutch highways, it is not necessary to investigate higher speeds.

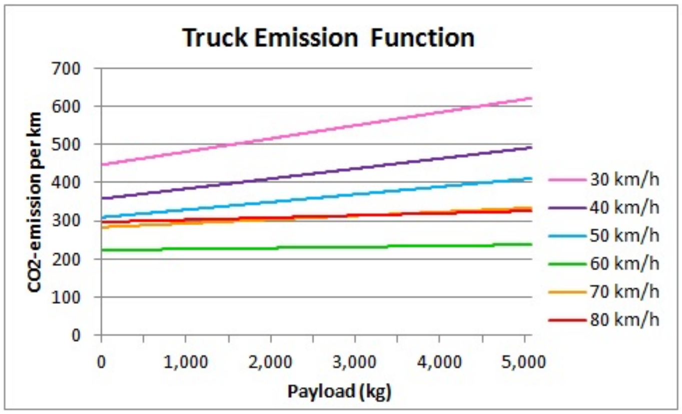

Figure 7.1: Truck emission function

The figure shows that a truck emits a lot of CO_2 when its speed is low. The most efficient driving speed in terms of CO_2 emissions is about 60 km/h. Moreover, the graph has a very low gradient at this velocity. This implies that the influence of the truck load on the amount of CO_2 emitted is almost negligible. In other words, a full truck emits only slightly more CO_2 than an empty truck. When a truck drives at 30 km/h instead, the vehicle weight does play a significant role.

When the distances between consecutive stops are small, a relatively large part of the trip takes place within the urban driving cycle. When the customers are not situated close to each other, the highway will be used a lot instead. Because the influence of the payload is larger at lower speeds, the vehicle weight is expected to have a larger impact in trips with small driving times.

7.2 Business Case

TNO has provided a data set of a Dutch retailer consisting of thirteen routes that the company used to supply its shops on a weekly basis. Moreover, the corresponding weekly demands of every shop are known for a period of 52 weeks. As all customers and depots have one or more time window, the dataset can be used to test both the allocation methods as well as the penalty methods.

In order to test the game theoretic allocation methods on the business data, all time windows are ignored. As the routes in the data file all took into account these time windows, they should be re-optimized to find the fastest route without considering time windows. For each of these routes, the allocation will be computed for every week using the five methods described in Chapter 5Methodologychapter.5. This will result in a general overview. To investigate the behaviour of the methods on a smaller level, two of the routes are explored in more detail. The difference in emissions between the fastest route without considering time windows and the one that is actually driven will be used to penalize the customers for the inconvenience caused by their time window(s).

7.2.1 Data Description

The data set of the retailer consists of thirteen routes that were used to visit its shops located in the Netherlands on a regular basis (i.e. once or twice a week). Every route includes 9 to 11 shops, all of

Table 7.1: Length and number of customers per route

Route nr.	1	**2**	3	4	5	**6**	7	8	9	10	11	12	13
# Customers	11	**9**	11	10	11	**11**	9	10	10	11	10	10	10
Length (km)	285	**205**	252	230	520	**620**	212	351	468	247	459	333	474

them having multiple time windows per week. While some of these time windows are only needed to avoid that service is scheduled when there is no personnel present at the shops (e.g. during the night), others are more specific. A few customers have two time windows per day, but the majority of them (about 94%) has only one. The average length of a customer's time window(s) is 9 hours and 43 minutes per day, the minimum duration is 2 hours and the corresponding maximum is 13 hours. Every route starts and ends at one of two distribution centers, which also dictate a time interval within which the routes have to be driven. However, these are very wide and are therefore unlikely to have any effect on the duration of routes.

Besides the routes, the data also include the weekly demand of every shop in 2009. This weekly demand is defined as an amount of loading units (LUs), which in this case are standard carton boxes as used by this specific retailer. In general, this demand was spread out equally over two visits per week. In case the demand was very low in all routes, it was decided to serve every shop only once a week. As this decision was made on a global level instead of for every shop or route individually, some of the routes were performed twice a week even though their total weekly demand would have filled even less than half a truck load. The extent to which the truck was filled in these trips is derived from dividing the total demand in units by the maximum volume of the truck, which is 507 units. Although the maximum allowable weight of the truck could be restrictive too, it is not very likely in the retail industry. The exact weight of the loading units is unknown and dependent on the products in the boxes. Because this information is needed to calculate the emissions properly, an estimation of the average weight of a box has been made. Taking into account the size of the boxes, it seems reasonable to assume an average weight of 10 kilograms.

7.2.2 Customer Level

As all routes are quite comparable in terms of number of customers and their order sizes, it does not add a lot of value to investigate all 13 routes extensively on customer level. This would also be too time consuming. Table 7.1Length and number of customers per routetable.7.1 shows that the routes do differ quite a lot in length; the longest trip is more than three times as long as the shortest one. To investigate the effect of the route length, these routes (numbers 2 and 6) are examined in more detail.

To have a better understanding of these two routes, they are depicted in Figure 7.2Longest and shortest route in the business casefigure.7.2. It can be observed that the 9 customers served in the shortest route are located relatively close to each other and to the depot. In the longest route however, the distance between the shops and especially the distance from the depot to the first and last customer visited is particularly large. From the information of the retailer it is known that the customers in the shortest route are situated in a huge urban area, whereas the shops in the longest route are located in a remote rural area. As explained in Section 7.1Truck Emission Functionsection.7.1, the payload has a smaller effect on the amount of CO_2-emissions for higher speeds. Because route 6 includes a larger percentage of kilometers to be driven on the highway, the impact of the payload is expected to be lower than it is for route 2.

7.3 Hypothetical Cases

The hypothetical cases are included to investigate to what extent the allocation methods give desirable results in 'extreme' situations. Imagine for example that 9 out of 10 customers are located very close to the depot (e.g. within a region of 10 kilometers) and that the 10^{th} customer is located

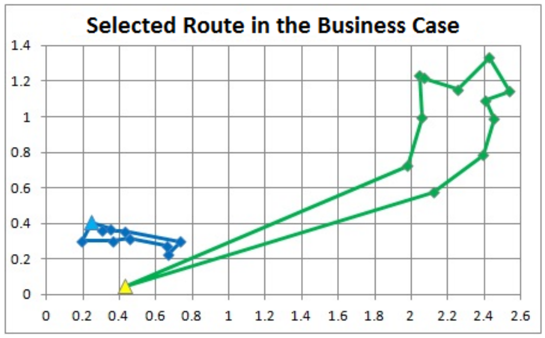

Figure 7.2: Longest and shortest route in the business case

much further way (e.g. more than 100 kilometers away). How large should the emission allocated to this customer be? Comparable scenarios can be made in terms of order sizes and numbers of customers. In the hypothetical cases these properties have been varied for all customers except for the selected customer, whose location and demand is kept fixed. In this way, the influence of the other customers on its allocated emission can be measured. A more detailed description of the 90 test cases that were constructed in this way is given below.

7.3.1 Number of Customers

In order to observe the effect of the number of customers on the amount of emissions allocated to a specific customer, three different numbers of other customers are considered. A small route includes 4 such other customers, an average sized route has 9 others and a large one includes 14 of them. This brings the total number of customers to 5, 10 and 15 respectively.

7.3.2 Locations

To vary the distances between the depot (D), the selected customer (S) and the other customers (O), the grid shown in Figure 7.3Grid with blocks used in hypothetical casesfigure.7.3 is used. In every situation, the depot and the customers are located in block A, B and/or C. Which of them is situated where differs per situation; an overview can be found in Table 7.2Scenarios in the hypothetical casestable.7.2. The scenarios are divided into three types based on the distance between the depot and the selected customer. For situations of the same type, the stand-alone emission of the selected customer is identical. The depot is always in the middle of a block, so it is at (5,5) in all situations except for the 7^{th} one, where it is at (50,50). The locations of the customers are randomly generated on a grid of 10 by 10 and are only determined once. In order to have the same situation in every block, the x-values and y-values are increased with 45 or 90 when customers are in block B or C respectively. If the customers would have a different location within such a block in every test case, the ceteris paribus condition would not hold. In case they are all situated in the same block (as is the case in scenario 1), their locations are as shown in Figure 7.4Locations of the customers and depot within blocks A, B and/or Cfigure.7.4.

The network used to connect the locations includes a direct link between all customers and between the depot and every customer. To simplify matters, the truck is assumed to drive at an average speed on each of these links. In reality, the driving speed will depend on the road type. Within a city center, a truck will typically drive slower than on the highway. To incorporate these differences, the truck is assumed to have an average speed of 35 kilometers per hour over links connecting two points within the same region and an average of 60 kilometers per hour over links connecting two points in different regions.

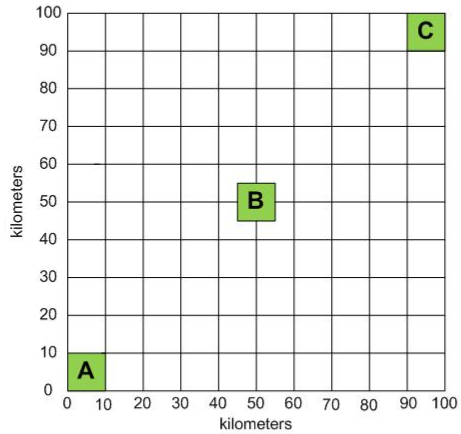

Figure 7.3: Grid with blocks used in hypothetical cases

Table 7.2: Scenarios in the hypothetical cases

		Block		
	Scenario	A	B	C
Type I	1	D,S,O	-	-
	2	D,S	O	-
	3	D,S	-	O
Type II	4	D,O	S	-
	5	D	S,O	-
	6	D	S	O
	7	O	D	S
Type III	8	D,O	-	S
	9	D	O	S
	10	D	-	S,O

7.3.3 Order Size

To observe the effect of the order size of other customers on the allocated emission of a single customer, only the order size of the others changes over the different test cases. The selected customer has a constant demand. Its order size has been chosen such that it is possible for all the other customers in the route to demand 5 times as much. Based on the business data, one may conclude that the demand of the selected customer is extremely low compared to the others in that case. As the maximum load of the truck is assumed to be 507 units and the maximum number of other customers is 14, the order size of the selected customer is equal to $507/(14 \text{ x } 5+1)$, which is about 7.14 units. A disadvantage of this approach is that when the selected customer is supposed to demand a relatively large amount, the truck will be almost empty. For the selected customer to demand a relatively large amount, the others have an order size which is 5 times as small, which is 1.43 units. When there are only 4 other customers in the route, the truck load is only 2.5% of its maximum value in this case. In the business data, the minimum truck load at the beginning of a route was observed to be about 16%.

7.4 Comparing Allocations

The test cases described above will be used to test the allocation and penalty methods on the most desired properties: fairness, robustness and if time windows are taken into account also the stimulance to broaden them. The computation times of the methods will also be discussed in

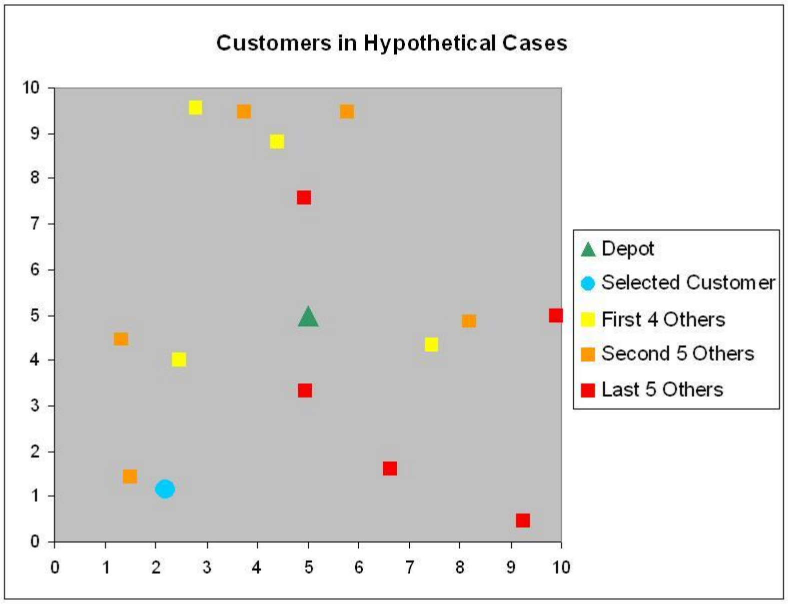

Figure 7.4: Locations of the customers and depot within blocks A, B and/or C

Chapter 8Resultschapter.8.

7.4.1 Fairness

In Chapter 5Methodologychapter.5, the fairness of the allocation methods was already discussed by using the properties defined in the Literature Review. As the Star-method and Shapley value do not guarantee a core solution, it is interesting to check how often and in what cases their allocation belongs to the core. The other three methods are stable by definition. For the penalty methods, one may assume that it is fair to allocate the additional emissions to the time windows that are responsible for the increase. Because the impact methods try to penalize only those time windows which are responsible for the increase in emissions, they are based on the 'polluter-pays' principle and can therefore be regarded as fair.

7.4.2 Robustness

For the business case, the variation of the allocated emissions over the 52 weeks is an indication of the actual variation in the allocated values. Because the demand of every customer changes over the weeks, the results do not satisfy the ceteris paribus condition. Therefore, the robustness of the methods will also be measured in two other ways. Both the effect of modifying a customer's own demand while the order sizes of other customers in the route remain constant and vice versa will be investigated for the shortest and longest route. In these routes, the 'average' customer and the 'average' week in terms of order size over the 52 weeks are taken as a reference point. When the demand of a customer is held constant, it is equal to its order size that was measured in this 'average' week. For the cases in which the demand of the single customer is constant, the demand of the other customers ranges from one to the maximum value for which the orders still fit in the truck. In the other cases, the demand of the selected customer ranges from one to the value for which the truck load is completely filled with orders.

The robustness of the penalty methods will be measured by fixing one of the customers' time windows and varying the time windows of the others by increasing and decreasing their length around their mid-point. As the effect from this may be dependent on the selected customer, it will be investigated for three of them in both the shortest and the longest route in the business case.

7.4.3 Stimulance to broaden Time Windows

Because the second penalty method assigns higher penalties to smaller time windows, it is clear that this method stimulates the customers to broaden their time windows. Moreover, as the first method penalizes every customer with a time window of any length in the same way, one does not have to test whether this method satisfies the desired property. The extent to which the third and fourth method stimulate customers to broaden their time windows is less easily determined beforehand. Therefore, it is suggested to vary the length of a single customers' time window while fixing the time windows of the other customers. The penalty methods are said to stimulate the customers to widen their time windows if the penalty of this single customer decreases as its time window is widened.

Chapter 8

Results

This chapter discusses the results of the application of the allocation and penalty methods to the test cases described in Chapter 7Test Caseschapter.7. Before describing the actual results, the way of reporting them is discussed. In Section 2 the results of applying the allocation methods to both the business case and the hypothetical cases are discussed. The results that were obtained from the penalty methods are described in the final section.

8.1 Reported Results

The amount of emissions allocated to a customer can be expressed per order, but also per loading unit. Which of these two is most appropriate depends on the relationship between the payload and the amount of emissions. If the emissions largely depend on the vehicle weight, emissions per loading unit may be more stable than emissions per order. Then, it is legitimate to display allocations using emissions per load unit. However, the emission function used in this research is relatively insensitive to changes in the payload (see Section 7.1Truck Emission Functionsection.7.1). This implies that comparing the allocation methods using emissons per order seems to be more appropriate.

To illustrate this difference between emissions per order and per loading unit, the values are displayed for one of the customers in the business case in Figure 8.1Emission per order versus emission per loading unit (LU)figure.8.1. The first graph indicates that the emission per order is relatively constant; the emission for a demand of 10 units is almost equal to the emission allocated to a demand of 50 units. One should notice that this graph does not satisfy the ceteris paribus condition; the demand of the other customers in the route changes as well (mostly in rather the same way). The graph does imply that for larger orders, the emission per loading unit is lower. This decrease is exactly what the second graph illustrates. For other customers, the same shapes of these graphs were observed. The only thing that differs among customers in the same route is the order in which the values of the different methods are plotted. This is due to the fact that all emissions allocated should sum up to the total emission of the route (this is the efficiency property). In this example, the Nucleolus allocates the largest and the Star-method the smallest amount of emissions to the customer considered. In order to allocate the same amount of emissions in total, the Star-method should allocate more emissions than the Nucleolus to at least one of the other customers.

From these two graphs one may already suspect the location of a customer to be much more important for the amount of emissions than the size of its demand. Because the emissions per unit decrease in the order size, the emissions should stimulate customers to lower their order frequency by increasing their order size. To be able to judge solution methods on their fairness and robustness it seems most appropriate to compare (the statistics based on) the allocated emissions per order.

To compare the robustness of methods, the variation in the allocated emissions should be reported. The use of standard deviations is not recommended because they are scale dependent.

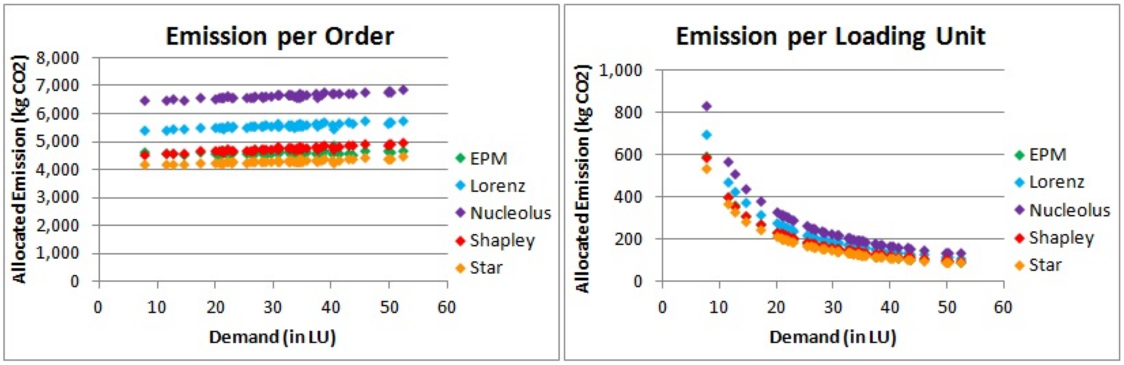

Figure 8.1: Emission per order versus emission per loading unit (LU)

This implies that a specific standard deviation can be considered as small and large at the same time. A value of 500 kg CO_2 will be considered as small for customers with an average allocated emission of 50,000 kg but is relatively large for customers that only receive 1,000 kg of CO_2 on average. Therefore the Coefficient of Variation (CoV), which is equal to the standard deviation divided by the corresponding average value, is suggested to normalize the standard deviations. In the previous example, the customer with an average allocated emission of 50,000 kg will have a CoV of $500/50,000 = 0.01$ or 1% and the others receiving only 1,000 kg of CO_2 on average will have a CoV of $500/1,000 = 0.5$ or 50%.

8.2 Allocation methods

The allocation methods have been tested on the business case (without considering its time windows) and on the hypothetical cases. For a description of these cases the reader is referred to Chapter 7Test Caseschapter.7. First, the results of the business case are described on route and on customer level and afterwards the results of the hypothetical cases are discussed. To summarize these results, the final section presents a comparison of the allocation methods using the desired properties which were defined in the Problem Definition.

8.2.1 Business Case

The results of the business case are discussed in the following way. First, a general overview of the results of all thirteen routes is given. Afterwards, the shortest and longest route are explored in more detail, especially to investigate the fairness and robustness at customer level.

Overall Results

For all thirteen routes, the five allocation methods described in Chapter 5Methodologychapter.5 have been applied to 52 weeks of varying demands. Because the total amount of emissions to be shared in a route is the same for all methods, the average allocated emissions are also the same. It is more interesting to view the overall results of the fluctuation in the allocated emissions to customers. For every allocation method, the CoVs of the allocated emissions to all customers are summarized in a boxplot in Figure 8.2Boxplot of the CoVs of the allocated emissions to customers in all routesfigure.8.2.

In this figure, the minimum and maximum CoV are indicated by the vertical lines. The boxes in between those lines indicate the middle 50% of the observations and the boundary between these boxes is the median CoV value. The average CoV values are indicated by the black diamonds. Whereas these average values are relatively close to each other, the maximum values are not. Due to their high value compared to the others, the maximum CoV values of the EPM and Lorenz

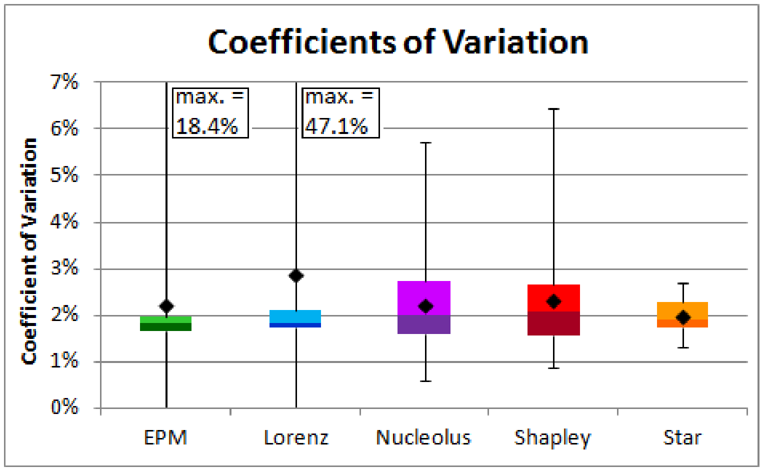

Figure 8.2: Boxplot of the CoVs of the allocated emissions to customers in all routes

allocation are displayed next to their boxplots. Especially for the Lorenz allocation, the maximum value is very high. This value was observed for a customer that did not receive the lowest or highest emission allocated in the route in any of the weeks. The large fluctuation measured strengthens the remark made in Chapter 5Methodologychapter.5 that the emissions allocated to customers that do not receive either of these two values are chosen rather arbitrary. Even though the average CoV value of the Lorenz allocation does indicate that the large fluctuation is an exception, the maximum CoV value will be a large disadvantage for the method in terms of robustness. Although its maximum was not as high as for the Lorenz, the EPM suffers from the same problem. The Shapley value and Nucleolus show acceptable CoV values and the Star-method generates solutions with an even smaller CoV value on average.

The EPM, Nucleolus and Lorenz allocation always generate a stable solution. For the allocations of the Shapley value and Star-method it has been checked whether they satisfy all the pseudo-core restrictions. For every route, the percentage of the 52 allocations which is in the pseudo-core is shown in the second and third column of Table 8.1Overall results of the business case without time windowstable.8.1.

Table 8.1: Overall results of the business case without time windows

Route nr.	% of **Star** allocations in pseudo-core	% of **Shapley** allocations in pseudo-core	Average # LP problems solved by **Nucleolus**	Average profit % **EPM**	% of 'equal split allocations' **Lorenz**
1	0%	98.1%	6.5	62.1%	0%
2	0%	86.5%	6.4	63.8%	0%
3	0%	100%	1	77.2%	15.4%
4	53.8%	100%	2.2	71.0%	0%
5	100%	100%	1	78.5%	0%
6	100%	100%	1	85.0%	100%
7	0%	100%	1	75.9%	0%
8	0%	100%	3	77.7%	0%
9	100%	100%	1	82.3%	100%
10	0%	98.1%	3.7	70.0%	0%
11	0%	100%	2.7	77.8%	0%
12	0%	100%	1	81.0%	0%
13	100%	100%	1	82.3%	100%
Average	34.9%	98.7%	2.4	75.7%	24.3%

Whereas the Shapley value almost always results in a stable solution, the allocation of the Star-method belongs to the pseudo-core in only 34.9% of the cases. For none of the routes, the percentage of Shapley allocations in the pseudo-core is lower than the percentage of Star-allocations. This implies that in terms of stability, the Shapley value outperforms the Star-method.

Besides the percentages of allocations that are in the pseudo-core, the average number of LP problems that have been solved to find the Nucleolus is also displayed in Table 8.1Overall results of the business case without time windowstable.8.1. For approximately half of the routes, solving a single problem turned out to be sufficient in all of the weeks. The maximum number of problems required to find a unique solution was 9 and this value was only found in the first two routes.

The fifth column shows the average percentual profit taken over all customers when the EPM is used to allocate emissions. The percentual profit is computed by comparing the emission allocated to a customer with its stand-alone emission. High percentages indicate that on average customers benefit a lot from the cooperation. Generally, the gains from cooperation will be larger when customers are situated relatively close to each other. This is because the amount of emissions is largely dependent on the amount of driven kilometers. When compared to the location of the depot, the customers are located relatively close to each other, they can share the emissions of driving to and from their area. In that case, the customers are said to have a lot of synergy.

For completeness, the percentage of equal split allocations found by the Lorenz allocation is shown in the last column of Table 8.1Overall results of the business case without time windowstable.8.1. In nearly a quarter of the routes, allocating an equal amount of emissions to every customer results in a core-solution.

The above results have been combined in one table to be able to investigate the (dis)similarities between the allocation methods. First of all, there are only three routes for which the percentage of the Shapley allocations that lies in the pseudo-core is below 100%. For these three routes, none of the Star and equal split allocations are in the pseudo-core. Moreover, it are exactly these three routes for which the largest average amount of LP problems were required to solve the Nucleolus. Furthermore, these routes also correspond to the three lowest average profit percentages generated by the EPM. As said, a low average of percentual profits indicates a low level of synergy among the players. Because the probability that an allocation belongs to the pseudo-core will generally be lower for smaller pseudo-cores and the lowest percentage of core allocations is found for these three routes, one may conclude that the pseudo-core is smaller than for the other routes.

Customer Level

The overall results presented in the previous section cannot be used to compare the allocation methods with respect to their fairness and robustness on customer level. Therefore, consider the longest and shortest route of the business case in Figure 8.3Longest and shortest route in the business casefigure.8.3. Notice that the customers have been numbered in order to be able to distuingish them from each other when discussing their results.

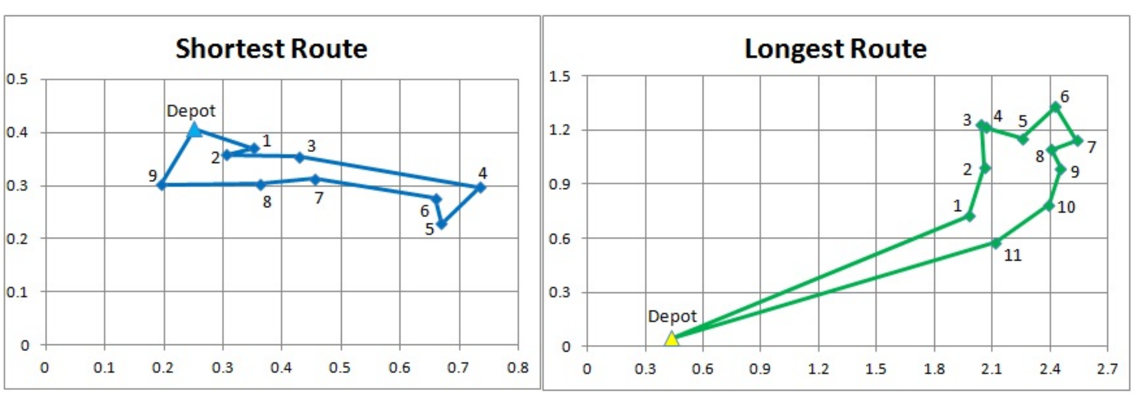

Figure 8.3: Longest and shortest route in the business case

In Chapter 7Test Caseschapter.7 it has already been mentioned that in the longest route the distance to be driven from and to the depot is relatively large. Compared to this distance, the customers are located relatively close to each other. Due to this high level of synergy, the average percentual profit over all the customers is very large (see Table 8.1Overall results of the business case without time windowstable.8.1). Moreover, the Star-allocation, Shapley value and equal split allocation all belong to the core. For the shortest route, the average percentual profit is only 63.8%. From Figure 8.3Longest and shortest route in the business casefigure.8.3 one can also observe that the players have less synergy with each other. The Star and equal split allocation are never in the core and the Shapley value is not always stable.

In Figures 8.4Average allocated (left axis) and stand-alone emissions (right axis) of customers in the longest routefigure.8.4 and 8.5Average allocated (left axis) and stand-alone emissions (right axis) of customers in the shortest routefigure.8.5 the average stand-alone emission of every customer is plotted on the right axis. The average emissions allocated to them can be read from the left axis. Because the Star-method allocates an equal percentual profit of the stand-alone emission to customers, its graph has the same shape as the stand-alone graph in both figures.

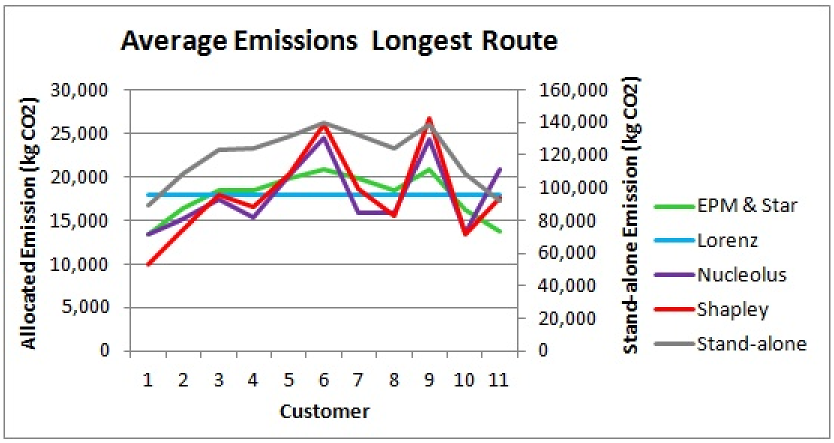

Figure 8.4: Average allocated (left axis) and stand-alone emissions (right axis) of customers in the longest route

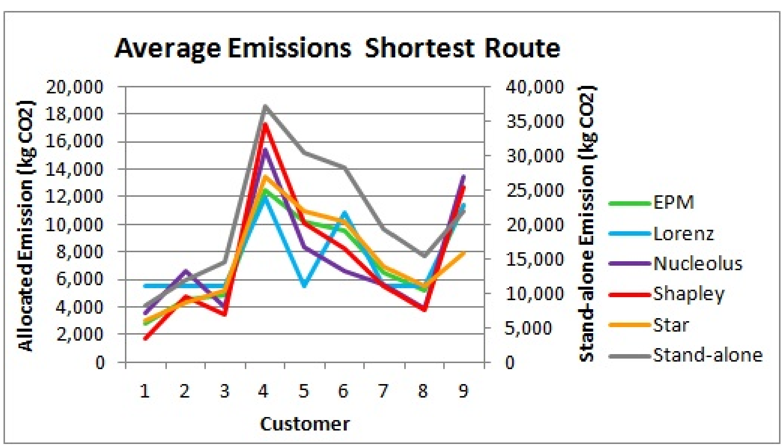

Figure 8.5: Average allocated (left axis) and stand-alone emissions (right axis) of customers in the shortest route

Let us first consider the longest route. The average emissions allocated by the Lorenz allocation form a horizontal line because every customer gets an equal amount of emissions in every week. As already expected, the results of the EPM and Star-Method are identical. The Shapley value and Nucleolus allocate the emissions in a similar way, but fluctuate a lot more than the others. This indicates that these methods consider a more uneven division of emissions as being fair. Except for the Lorenz allocation, all methods allocate the largest amount of CO_2 to the two customers with the largest stand-alone emissions (nr. 6 and 9). Only the eleventh customer seems to be treated unjustly by the Nucleolus. Its stand-alone emission is far below average and the amount of CO_2 it receives from this method is above average.

Because the customers are situated much closer to the depot, the average emissions allocated to customers in the shortest route are a lot smaller. All of the methods show a peak at customer 4, which has the largest stand-alone emission on average. In Figure 8.3Longest and shortest route in the business casefigure.8.3, one can observe that customers 4, 5 and 6 are located far away from the depot. However, as they are situated relatively close to each other, they can share the emissions of driving to and from their area. Customer 9 is located less far away from the depot, but it merely has any synergy with the other customers. The Star-method is the only method that does not take into account the benefits from neighbouring customers. Because it does not consider the stand-alone emissions of subsets and only looks at the individual properties of customers, it allocates too much CO_2 to customers 4, 5 and 6 and too less CO_2 to customer 9. This can be observed from Figure 8.5Average allocated (left axis) and stand-alone emissions (right axis) of customers in the shortest routefigure.8.5 by noticing that for these points, the gap between the graph of the EPM and the graph of the Star-method is the largest. The fact that the Star allocation does not belong to the core in any of the weeks implies that some of the customers would actually deserve a larger percentual profit than others. Generally, this implies that the amount of synergy is unevenly distributed among the customers. When all of them would be equally far away from (or close to) each other, they would all deserve a low (or high) benefit. However, when some are clustered together while others are not, then the ones within the cluster deserve a larger profit than the others. In that case, the Star-method cannot be regarded as a fair approach. Just like for the longest route, the Nucleolus and especially the Shapley value fluctuate heavily and allocate a lot of CO_2 to customers 4 and 9.

8.2.2 Hypothetical Cases

The hypothetical cases have been included to ivestigate the effect of extreme situations in terms of demand size, location and number of customers on the allocated emission of a single selected customer. The demand of this customer is the same over all test instances and in every of the three types of scenarios it is situated equally far away from the depot. The overall results of the hypothetical cases can be found in Appendix AResults Hypothetical Casesappendix.A. This section gives a description of the most important results from these hypothetical test cases. First, an overview of the results is presented and then the influence of the three variables that have been varied is discussed in different subsections.

Overall Results

In all three tables in Appendix AResults Hypothetical Casesappendix.A, the third column shows whether the problem has a non-empty pseudo-core. Although it was expected and observed in the business case that problems without time windows have a non-empty core, this is apparently not true in 'extreme' situations. For some scenarios, the core is empty when the selected customer has a relatively low demand. In the 7^{th} scenario, the core is even empty for every number of customers and for every order size of the other customers. However, when one takes a closer look at this scenario this can be explained easily. The depot is situated in region B, the selected customer is in C and all the others are in A. As for each number of customers the selected customer is visited first, the demand of all the others is transported all the way to C and then to A. Therefore, a lot of CO_2 is emitted. The fact that these routes are constructed by the VRPTW program is because it only minimizes the duration of the routes. If the amount of emissions would have been minimized, the routes in this scenario would probably be such that the selected customer is visited last. Note that in that case, it is not guaranteed that the pseudo-core is non-empty. Because the EPM and Lorenz allocation are not able to generate a solution when the pseudo-core is empty, the results of these problems will not be analyzed in further detail.

Number of Customers

For the overall results of the test cases with 5, 10 and 15 customers, the reader is referred to the tables depicted in Appendix AResults Hypothetical Casesappendix.A. The differences between these tables can most easily be read from the summaries in the bottom lines of the tables. It turns out that for routes including more customers, the likelihood of an empty core is larger. However, this increase is only observed for test cases in which the selected customer has a relatively low demand compared to the other customers. As the average number of LP problems solved to find the Nucleolus is only 3.5 for routes with 5 customers and 18.5 for those including 15 customers, the computational effort of this method is clearly affected by the number of customers.

The effect of the number of customers on the amount of CO_2 allocated to the selected customer is displayed in Figure 8.6Average emission allocated to the selected customer for 5, 10 and 15 customersfigure.8.6. The graph shows the average emission allocated to the customer over all scenarios for which the core is non-empty and in which every customer demands an equal amount of loading units. It turns out that every allocation method allocates a smaller amount of CO_2 to the selected customer when more customers are added to the route. However, the impact is a lot stronger for the Star-method than for the other methods. As mentioned earlier, the Star-method is the only method that does not account for the degree of synergy of customers. Because the other customers are always situated in the same block, they generally have much more synergy than the selected customer has with the other customers.

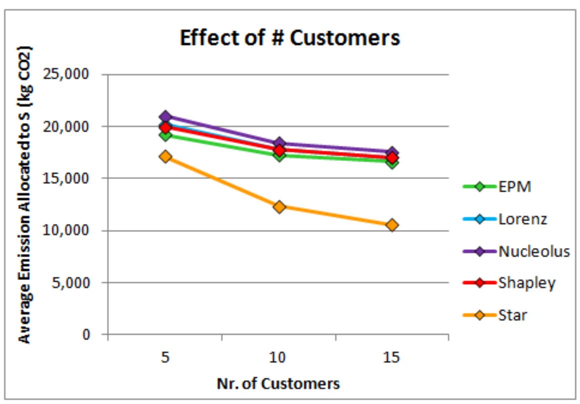

Figure 8.6: Average emission allocated to the selected customer for 5, 10 and 15 customers

Locations

The hypothetical cases have been classified in 3 categories based on the distance between the selected customer and the depot. In Chapter 7Test Caseschapter.7 it has been explained that within each of these types of scenarios, the location of the other customers is varied among the blocks. For all cases in which the distance between the selected customer and the depot is equal, the stand-alone emission of the selected customer is the same. In Figure 8.7Average emission allocated to the selected customer for different locations of other customersfigure.8.7, the effect of the locations of the other customers on the amount of emissions allocated to the selected customer is displayed for the three categories.

First of all, observe from the scales used on the y-axes that the emissions allocated to the customer are much higher when it is located further away from the depot. This is easily explained by the fact that the distances between the blocks in which the customers and depot are located

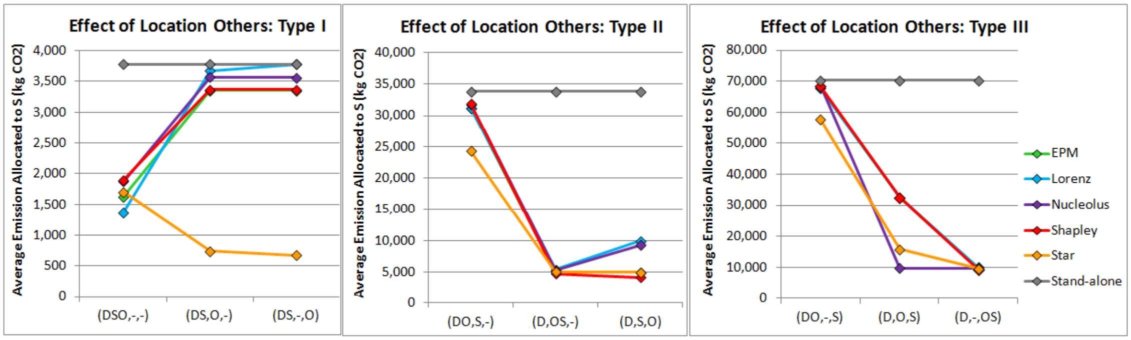

Figure 8.7: Average emission allocated to the selected customer for different locations of other customers

are quite large. The stand-alone emission of the third type of scenarios is also much higher than the stand-alone emission of the second type, which is again much higher than the stand-alone emission of the first type.

For this first type of scenarios, the graph displays the same pattern for all the allocation methods except for the Star-method. When the depot and all customers are located in the same block, the Star-method allocates approximately the same amount of CO_2 to the customer as the other methods. However, as the other customers are moved to block B and C, the emission allocated by the Star-method decreases. Again, this is due to the fact that the Star-method is the only method that does not take into account the stand-alone emissions of subsets. In this case, the difference with the other allocation methods is particularly large because, as can be observed from Figure 7.4Locations of the customers and depot within blocks A, B and/or Cfigure.7.4, the selected customer is located to the south-west of the depot whereas the others are situated to the north-east of the depot. In this case, the other customers have much more synergy and should therefore receive a larger percentual emission reduction than the selected customer. In the other scenarios for which the selected and other customers are situated in a different block, the Star-methodology allocates a relatively low amount of emissions to the selected customer as well.

In the second and third graph, the differences between the allocation mehods are smaller. For these two types of scenarios, the emission allocated to the selected customer is especially large when the other customers are located closer to the depot. This seems to be reasonable because in that case the truck needs to travel a lot of kilometers for the selected customer only.

Order Size

In the tables presented in Appendix AResults Hypothetical Casesappendix.A, the second column shows whether the demand of the selected customer is relatively high, average or relatively low. Because the selected customer has a constant demand while the demand of the other customers fluctuates, one may also say that the demand of the other customers is low, average or high respectively. Besides the fact that for cases in which the demand of the selected customer is relatively low the pseudo-core turns out to be empty more often, only minor differences can be observed. This is illustrated by Figure 8.8Average emission allocated to the selected customer for different demand sizes of other customersfigure.8.8. It can be concluded that the order size does not have a lot of influence on the allocations.

8.2.3 Properties of the Allocation Methods

Below, the extent to which the allocation methods satisfy the requirements of a good solution approach is discussed.

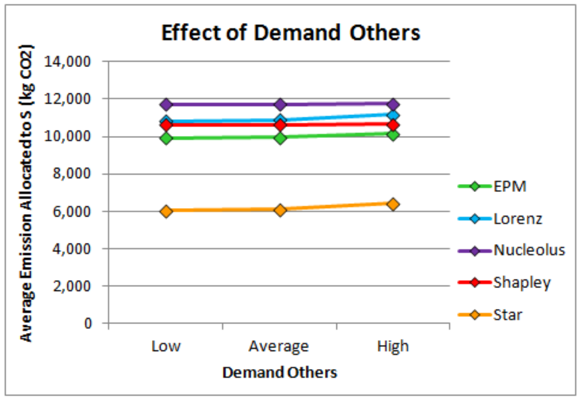

Figure 8.8: Average emission allocated to the selected customer for different demand sizes of other customers

Fairness

In Chapter 5Methodologychapter.5, the allocation methods were already compared with each other in terms of the desired fairness properties described in Chapter 3Literature Reviewchapter.3. Even though the Star-method and Shapley value satisfy a lot of these properties, they cannot guarantee a core-solution. Because the Shapley value was able to generate a solution in the core for nearly all the instances of the business case, this does not seem to be a large disadvantage for this method. Due to the fact that the Star-method does not take into account the level of synergy customers have, its allocation was more often outside than inside the core. Therefore, this method will not be preferred in terms of fairness.

Robustness

As has been discussed in Chapter 7Test Caseschapter.7, the robustness of the allocation methods is analyzed in three different ways in the business case. First of all, the actual variation in the allocated emissions is discussed. Then, the ‘average’ customer’s demand is fixed at its average value while the demand of the other customers is increased from one loading unit to a full truck load. Afterwards, the demand of the others is fixed while the demand of the selected customer is varied.

Actual Variation due to Changes in Demand of all Customers
For the customers in the longest and shortest route, the COV values are plotted in Figure 8.9CoV of the emissions allocated to customers in the shortest and longest routefigure.8.9.

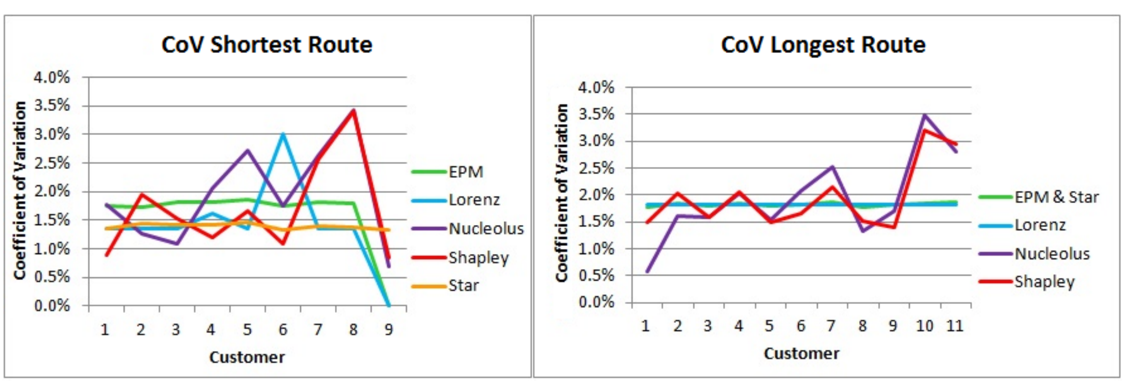

Figure 8.9: CoV of the emissions allocated to customers in the shortest and longest route

In general, the CoV of the allocated emission varies most for the Shapley value and Nucleolus. Because these methods have a larger maximum value than the others, they would not be preferred based on these figures. However, as the maximum standard deviation is only between 3 and 3.5% of the average allocated emission, the variation is still acceptable. For both routes, the variation of the emissions allocated by the Star-method is roughly the same for every customer. Because its values are also relatively low, the Star-method would be preferred based on these graphs. The CoV values of the EPM and Lorenz are also relatively stable, but are very small for the last customer in the shortest route. Actually, the CoV is approximately zero because these methods allocate almost the same amount of emissions to this customer in every week. Therefore, its standard deviation and thus also its CoV value is negligible for these methods. The reason why the Lorenz allocation allocates a fixed amount of CO_2 to this customer is that the customer never deserves the benefit of the equal split allocation due to its low amount of synergy with the others. Therefore, the emission it receives is the smallest amount that still belongs to the pseudo-core. The same reasoning applies to the EPM. In Figure 8.9CoV of the emissions allocated to customers in the shortest and longest routefigure.8.9, a peak in the graph of the Lorenz allocation can be observed for customer 6 in the shortest route. As this customer also receives a high average emission from the Lorenz allocation (see Figure 8.5Average allocated (left axis) and stand-alone emissions (right axis) of customers in the shortest routefigure.8.5), the standard deviation is very high. This large variation is probably because the Lorenz allocation only minimizes the difference between the smallest and largest amount of emissions allocated to customers. For customers that receive neither of these amounts of emissions, the amount is picked somewhat arbitrarily on the interval from the minimum amount it should get to the largest amount that is allocated within the route. Therefore, the results of this customer show the disability of the Lorenz allocation to guarantee a unique solution.

Varying the Demand of Other Customers
The influence of the demand of other customers on a customer's emission can be measured by fixing the demand of that customer and letting the demand size of the others vary. To this end, the 'average' customer in terms of demand size has been selected in both the shortest and longest route. The demand of this customer is its average demand and the others have a demand from zero to a full truck load. The emissions allocated to the selected customers are displayed in Figure 8.10Varying the demand of other customers in the shortest and longest routefigure.8.10.

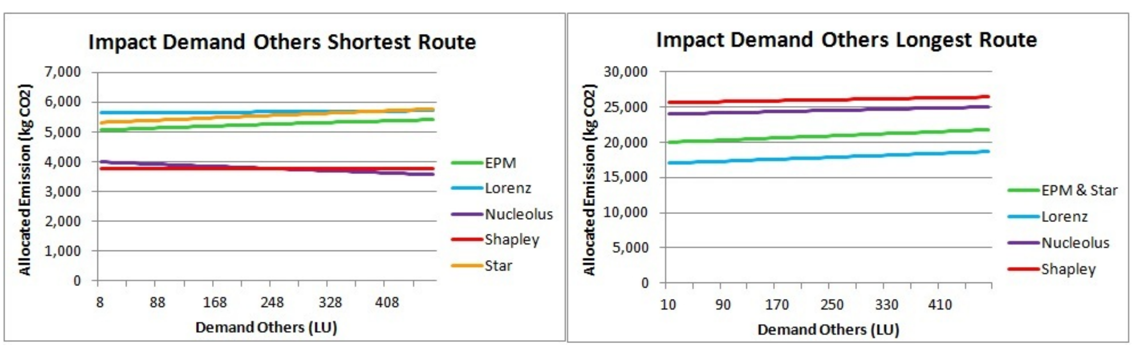

Figure 8.10: Varying the demand of other customers in the shortest and longest route

Because these graphs are all relatively horizontal, the total percentual changes in the allocated emission due to the increase of the demand may be more useful. Therefore, these are depicted in Table 8.2Total increase of the allocated emission due to the increase of the demand of otherstable.8.2.

Because the selected customer does not change its own behaviour, the variation in the emission it receives is preferably small. Compared to the other methods, the Shapley value is relatively insensitive to changes in the demand size of others. When considering this type of robustness only, the Shapley value would therefore be preferred.

Table 8.2: Total increase of the allocated emission due to the increase of the demand of others

	EPM	Lorenz	Nucleolus	Shapley	Star
Shortest Route	6.8%	1.4%	-10.7%	0.3%	8.6%
Longest Route	8.7%	9.5%	4.2%	3.1%	8.7%

Varying the Demand of the Selected Customer
It is also interesting to investigate the effect of a customer's own order size on its allocated emission. The emissions allocated to the selected customer when varying its own demand and keeping the demands of the others fixed are displayed in Figure 8.11Varying the demand of the selected customer in the shortest and longest routefigure.8.11.

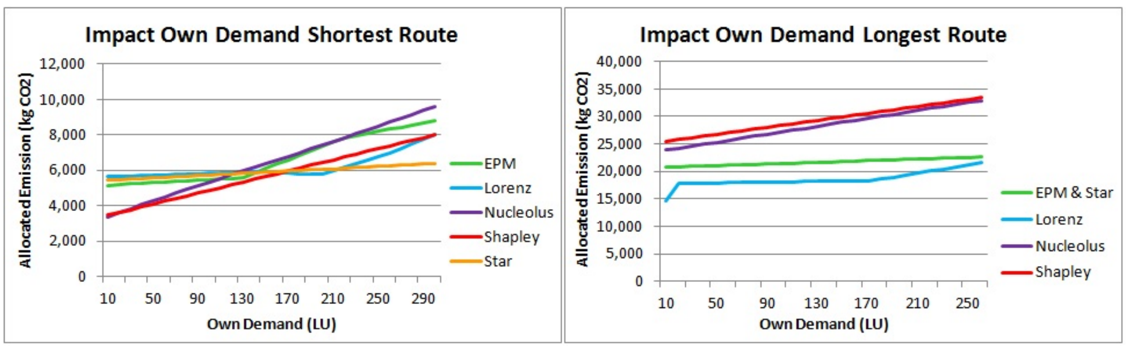

Figure 8.11: Varying the demand of the selected customer in the shortest and longest route

Because the total amount of CO_2 increases as the customer increases its demand, an increase in the emissions allocated is justified. However, a decrease is not regarded as fair as other customers are harmed by the increase while the customer responsible for the increase in the total emissions benefits from it. In Figure 8.11Varying the demand of the selected customer in the shortest and longest routefigure.8.11, this is only observed once. The emissions allocated by the Lorenz allocation decrease slightly for demand sizes between 150 and 190 units in the shortest route. Even though the total decrease is only 2.7%, it is undesired and probably due to the fact that the Lorenz allocation does not guarantee a unique solution. The Shapley value and especially the Nucleolus allocate the extra emissions to a large extent to the customer that is responsible for the increase.

Computation Times

The practical use of a method is strongly related to its computation time. As time is money, LSPs are not likely to be willing to wait more than e.g. an hour for a method to come up with a solution. One should keep in mind that informing customers about the CO_2 emissions of their cargo is (at least at the moment) only an extra service provided by these companies. If it takes too much effort to calculate the emissions, they will generally lose interest in reporting them.

Generally, complex methods require a larger computational effort than simple ones. Moreover, this difference in computation time often increases as problems are extended. For the allocation of emissions, the number of customers in a route turns out to be an important factor in this respect. As can be observed from Figure 8.12Computation times of all allocation methods except for the Nucleolus for different numbers of customersfigure.8.12, the computation times of the relatively simple Star-method are the only values that do *not* increase in the number of customers. Especially for larger problem instances, the methods based on game theoretical concepts have a relatively large computation time.

The computation time of the Nucleolus is not displayed in this graph, because it is not only

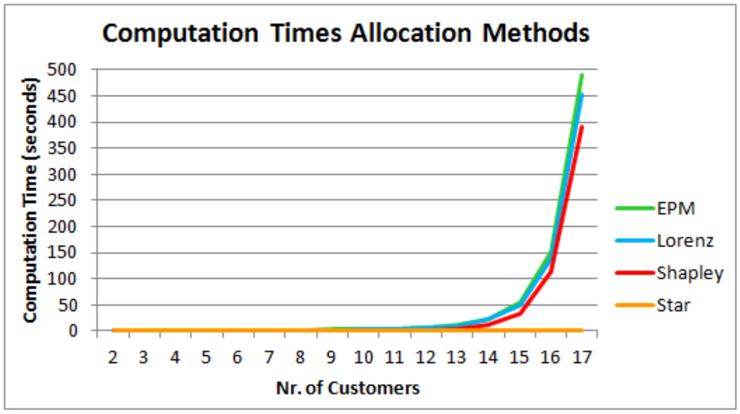

Figure 8.12: Computation times of all allocation methods except for the Nucleolus for different numbers of customers

dependent on the number of customers on the route but also on the number of LP problems to be solved. As the Lorenz allocation and EPM are determined by solving a single problem and the Nucleolus is found by solving at least one problem, a lower bound on its computation time is approximately equal to the time required to apply the EPM or the Lorenz allocation. For the three different numbers of customers that were investigated in the hypothetical cases, the relationship between the number of required LP problems and the computational effort is depicted in Figure 8.13Computation times of the Nucleolus for 5, 10 and 15 customers and different numbers of LP problemsfigure.8.13.

For each of the three figures, a regression line is added to illustrate the strong relationship between the number of LP problems and the computation time. As solving an LP problem requires computational effort, this correlation was already expected. However, the regression line fits the data very good in the second and third graph. This is also indicated by the high values of R^2. The fact that the value of R^2 is lower for problems with only five customers is probably due to the fact that the computation times are very small for these problems. On such a small scale, differences in computation time may arise due to e.g. other processes performed by the computer.

It is clear that all of the methods can generate a solution with 10 customers within a short time period. For the routes in the business case, the largest average computation time was 4.85 seconds per trip. Of course, the user can decide for himself which computation time is still acceptable and which is not. As applying one of the advanced methods to routes including 18 customers or more takes more than an hour, the Star-method is likely to be preferred in such cases. Obviously, these times will differ among different (types of) computers. In this research, the computations were performed on an Intel(R) Core(TM) i5 CPU 650 @ 3.20 GHz with 3.42 GB of RAM.

8.3 Penalty Methods

The penalty methods have been applied to all the routes of the business case. First, an overview of the results is given and then the shortest and longest route are investigated in more detail again. The extent to which the penalty methods have the desired properties on customer level is explored in the final section.

8.3.1 Overall Results

The overall results of the thirteen routes while taking into consideration the time windows of customers are summarized in Table 8.3Overall results of the test cases with time windowstable.8.3.

In the second column, the percentage of the 52 weeks for which the emissions of the route with time windows are larger than for the route without considering time windows is displayed.

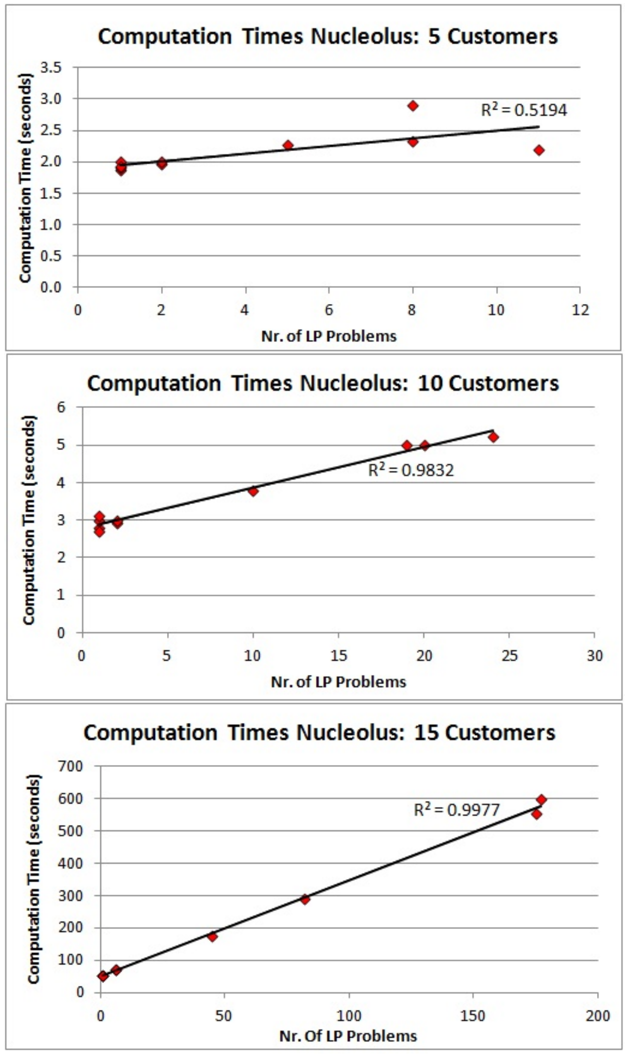

Figure 8.13: Computation times of the Nucleolus for 5, 10 and 15 customers and different numbers of LP problems

It turns out that the time windows of the customers in route 12 do not cause any inconvenience for the LSP. In all of the 52 weeks, the LSP can drive the same route as it would have taken when time windows would not be considered.

The third column shows the average number of customers that have impact according to the third and fourth penalty method. Only in the first and last route, this is not a constant value for different values of demand. For these two routes, the average penalty to be shared (last column) is also much higher than for the other routes. In routes 7 and 11, none of the individual customers can be held responsible for the extra emissions caused by their time windows. Apparently, the combination of time windows of a group of customers is such that the fastest way to serve them is not possible. In that case, the impact methods divide the extra emissions equally among the customers.

8.3.2 Longest and Shortest Route

For simplicity, statistics on the shortest and longest route are again studied in more detail. The routes that are taken when time windows are taken into account are displayed in Figure

Table 8.3: Overall results of the test cases with time windows

Route nr.	CO_2 with TWs > CO_2 without TWs	Average # Customers with Impact	Average Penalty to be shared (kg CO_2)
1	100%	2.9	10,019
2	100%	2	3,833
3	100%	2	2,932
4	100%	2	2,863
5	100%	1	5,065
6	100%	5	8,394
7	98.1%	0	595
8	100%	3	707
9	100%	1	4,872
10	100%	2	6,615
11	98.1%	0	2,632
12	0%	-	-
13	100%	3.9	14,641
Average	92.0%	2.1	5,264

8.14Longest and shortest route with time windowsfigure.8.14. The customers marked in black are the ones that are considered to have impact. In the shortest route, customer 2 (customer 9 in Figure 8.3Longest and shortest route in the business casefigure.8.3) could better be served at the end of the day but it only wants to be served during the morning of the day. For the other customer marked in black (nr. 6), the duration of the trip would have been smaller if it would be able to visit it between customers 3 and 4. In the longest route, customers 5 to 9 all have impact because the sequence in which they are served differs from the sequence in which they were visited without considering time windows (also displayed in Figure 8.3Longest and shortest route in the business casefigure.8.3).

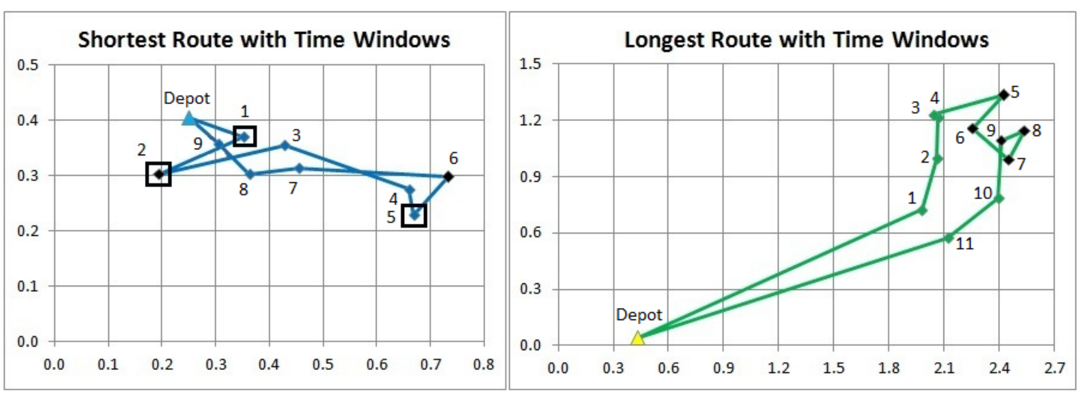

Figure 8.14: Longest and shortest route with time windows

8.3.3 Properties of the Penalty Methods

The extent to which the penalty methods possess the desired properties described in Chapter 2Problem Definitionchapter.2 is discussed in the following paragraphs.

Fairness

Unlike the game theoretic allocation methods, the methods to penalize time windows were not already defined in the literature. There are also no fairness criteria or whatsoever available to characterize the penalty methods in terms of their fairness. Because the third and fourth method only penalize these customers that have impact, they are based on the 'polluter-pays principle' and can therefore be regarded as fair.

Robustness

The robustness of the penalty methods have been measured by keeping one of the customers' time window(s) fixed while varying the size of the time windows of the others. In the first paragraph, varying the length of the time windows of all the customers in a route is shortly discussed. Whenever an increase in the length of time windows leads to a larger duration of the trip, the value of the objective of the VRPTW is worsened while the number of possibilities to visit the customers has increased (or at least not decreased). In that case, the route is not useful for testing the robustness of the penalty methods.

Varying the Time Window Length of all Customers
The time windows of all customers are varied to see whether the VRPTW program is capable of selecting faster routes in case they are broadened. It turns out that for the shortest route, the program is indeed able to find better routes when time windows are widened. It is disappointing that this is not true for the longest route. Taking the original time windows resulted in a better solution than taking 110% of their width. The routes constructed using 90 and 120% of the time windows lengths resulted in an identical solution, which was worse than the ones using 100 en 110%. Apparently, the VRPTW program is unable to select the fastest possibility to visit the customers in the longest route. Therefore, it is not legitimate to investigate the robustness of the penalties by varying the time windows of the selected or the other customers.

Varying the Time Window Length of other Customers
The influence of the time windows of other customers can be measured by varying them while fixing the time window of a selected customer. Because the influence may depend on the selected customer, it will be investigated for three of them in the shortest route. This is done by increasing and decreasing their time window width around their mid-point. Note that the three selected customers are chosen based on their position within the original route without time frames (first, middle, last). Because the time window and demand of the selected customer does not change, the influence on its penalty is preferably low. The time windows of the three customers framed in Figure 8.14Longest and shortest route with time windowsfigure.8.14 are fixed at their original value while the time windows of the other customers in the route are varied from 40 to 160% of their original size. Their mid-point remains the same over all instances. The total penalty to be shared in the routes that were constructed using these time windows and the penalties allocated to these three customers are displayed on the right and left axis respectively in Figure 8.15Varying the length of the time windows of other customersfigure.8.15.

From the figures it can be observed that methods 1 and 2 show similar behaviour for all three customers. That is, method 1 has the same shape as the graph of the total penalty to be shared as it is equal to the total penalty divided by the number of customers. Therefore there is a direct relation between the total penalty to be shared and the penalty allocated by method 1. For method 2 it can be observed that for the smallest time windows of other customers, the penalty to the selected customer increases for a non-increasing total penalty to be shared.

For example, consider customer 1. While the time windows of the other customers increase from 40% to 80%, the penalty allocated to the selected customer increases. This can be explained by the fact that the size of the selected customers' time window relative to the others decreases. For a further extension of the time windows the total penalty to be shared decreases that much that the penalty assigned to the selected customer also decreases.

For method 3 and 4 significant differences can be observed between customer 2 and customers 1 and 5. As described in Chapter 5Methodologychapter.5 these methods allocate emissions to the customers of which the time windows are considered to have a negative impact on the total duration of the route. Customers 1 and 5 have no impact for this route and therefore both methods

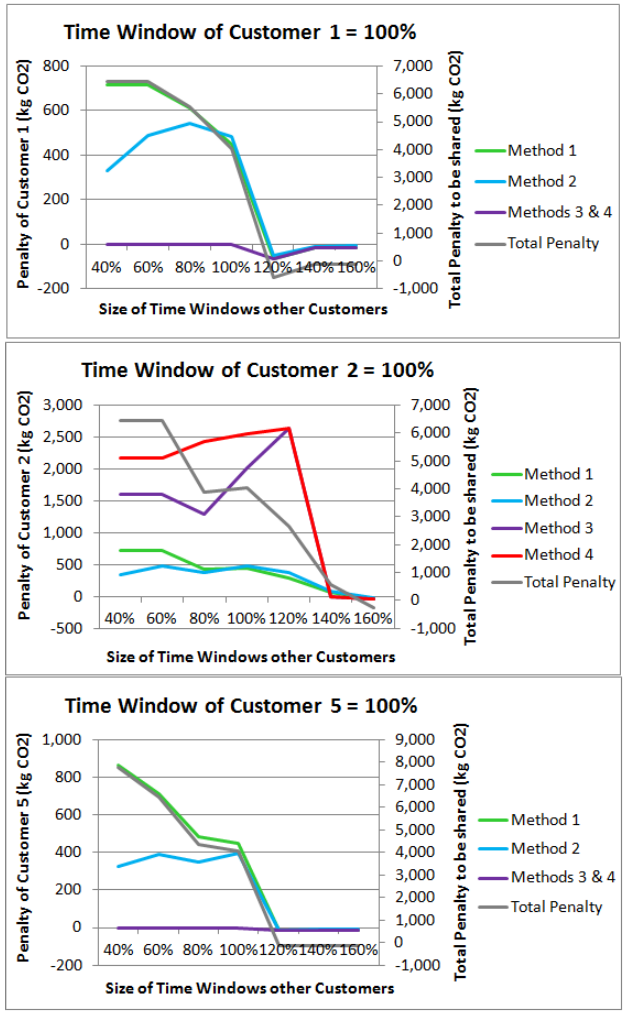

Figure 8.15: Varying the length of the time windows of other customers

allocate no emissions to this customer. Note that for a negative total penalty, the impact is not calculated and the negative penalty is shared equally among all customers. In that case methods 3 and 4 give the same results as method 1. Customer 2 does have a negative impact on the route duration which can clearly be seen from the allocated penalties calculated by method 3 and 4. For increasing time windows of the other customers, the penalty for both method 3 and 4 increases towards a maximum (for 120% of the time windows of other customers). At that maximum, customer 2 is allocated the total penalty as it is the only customer in the route having a negative impact. A further increase in the time windows of the other customers enables the VRPTW program to find a route in which the time window of customer 2 has no negative impact anymore.

For methods 1 and 2, a change in the time windows of the other customers only leads to small changes in the penalty assigned to the selected customers. This is because these methods follow the trend of the total penalty to be shared. An increase in the amount of emissions will lead to an increase of the penalty for every customer. This implies that penalty methods 1 and 2 can be considered as robust. The opposite is true for methods 3 and 4. The penalties generated by these methods are highly dependent on the specific route that is constructed and they do not follow the trend of the total penalty to be shared. Due to small changes in the time windows of other customers, the penalty of the selected customer can fluctuate heavily.

Stimulance to broaden Time Windows

As discussed in Chapter 7Test Caseschapter.7, the extent to which the third and fourth penalty method stimulate the customers to broaden their time windows is measured by varying the length of the time window of a single customer while fixing the size of all the others. This is exactly the opposite of what has been done in the previous section. Preferably, a broader (smaller) time window generates a lower (higher) penalty for the selected customer. Unfortunately, the routes that were constructed by the VRPTW program were not faster for broader time windows of the selected customer. The time required to perform the route sometimes increased and sometimes decreased as the time window of the customer was widened. Clearly, the program used to construct the routes should first be improved to be able to observe the influence of changes in time windows.

Computation Times

Just like for an allocation method, it is important that the computational effort of a penalty method is low. Because the computation times for the first and second and for the third and fourth were of similar size, they are combined in the two lines depicted in Figure 8.16Computation times of the penalty methods for different numbers of customersfigure.8.16.

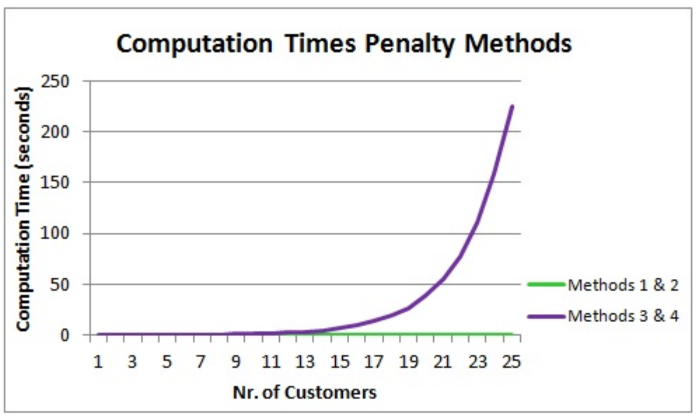

Figure 8.16: Computation times of the penalty methods for different numbers of customers

Due to the simplicity of penalty methods 1 and 2, their computation times remain approximately zero for all numbers of customers. The computation times of the penalty methods based on the impact of the time windows are dependent on the number of customers. This has to do with the calculation of the impact of a customer. It is calculated by shifting the customer to all other positions within the route. For each of these positions, it is checked whether this sequence is possible and if so, whether it takes less time than the original trip. Because the impact needs to be calculated for every customer and the possible shifts increase in the number of customers, it is quite logical that penalty methods 3 and 4 require more time for larger numbers of customers. However, for the routes in the business case the results were obtained quickly using these methods. Even for routes including 21 customers the computation time is less than a minute and for routes in which 25 customers are served the computation time is less than 4 minutes.

One should notice that the shifting of customers is a simplification of re-optimizing the VRPTW. Partly, this has been done to save computation time. When the VRPTW would have been re-optimized for every customer, penalty methods 3 and 4 would require a larger computational effort. It depends on the speed of the VRPTW program whether the penalty methods based on the impact would still be useful for practical problems.

Chapter 9

Conclusions

In the previous chapter, a large amount of results have been discussed. This chapter will provide an overview of the most important observations. By using these results and the information presented in the earlier chapters, the second and third research question will be answered.

First of all, it was concluded that in a distribution route the amount of CO_2 emissions is merely influenced by the weight of the truck and its load. The specific route that is driven is much more important. Because the influence of the payload is small, the total amount of emissions will be lower if customers would decrease their order frequency by increasing the size of their orders. If the LSP can stimulate customers to increase their order size, less CO_2 will be emitted.

The practical use of the allocation methods was investigated by applying them to a business case. For ten out of thirteen routes in the business case, the allocation generated by the Shapley value was in the pseudo-core for all of the 52 instances of demand sizes. For the other routes, only a few exceptions were observed. This implies that even though the Shapley value does not guarantee a stable solution, it does generate a solution in the pseudo-core in almost any 'normal' situation. As applying the Star-methodology resulted in a solution in the pseudo-core in only 34.9% of the cases, this cannot be concluded for this allocation method. The other three allocation methods are stable by definition.

Because the average profit percentage of the EPM taken over all routes was rather high (75.7%), it was concluded that the customers in the routes generally have a lot of synergy. The difference in the amount of synergy between the routes was illustrated by the comparison of the shortest and longest route in the business case. Whereas in the shortest route, the average percentual profit was only 63.8%, it was as high as 85% for the longest route. Because the distance between the depot and the group of customers served was particularly large for the longest route, the customers could share the emissions of driving to the rural area in which they were situated. In the smallest route, not every customer was situated at the same side of the depot and therefore not all of them could benefit from the locations of the others. In this route, the amount of synergy with other customers differed quite a lot among the customers. Because the Star-method does not consider the stand-alone emissions of subsets, it allocates an amount of emissions that is too low to customers that have a relatively low amount of synergy with other customers. This is the reason that its solution does not belong to the pseudo-core in the smallest route but does in the longest route.

From the hypothetical cases, it was observed that in 'extreme' situations the core is not always non-empty. Cases for which the pseudo-core turned out to be empty were left out from the analysis. From the cases that were analyzed it was concluded that the allocated amount of emissions decreases if the number of customers in the route increases. The demand size of the other customers does not have a lot of effect on the amount of emissions allocated to the selected customer. Instead, the locations of the other customers largely influenced the amount of CO_2 that was allocated to the selected customer for all of the allocation methods. The main difference between the approaches observed in the hypothetical cases was that the Star-method allocates a relatively low amount of emissions to the selected customer because it does not consider the larger amount of synergy the other customers have.

The extent to which the allocation methods satisfy the desired properties defined in Chapter 2Problem Definitionchapter.2 determines which method should be preferred. To summarize the results and findings from earlier chapters in terms of these properties, the allocation methods are presented as a ranking for each of these properties in Table 9.1Ranking of the allocation methodstable.9.1. For every property, a value of 1 indicates that the method should be preferred and a value of 5 implies that it should not.

Table 9.1: Ranking of the allocation methods

	EPM	**Lorenz**	**Nucleolus**	**Shapley**	**Star**
Fairness	3	3	2	1	5
Robustness	4	5	3	2	1
Understandability	3	3	5	2	1
Computational effort	3	3	5	2	1

In Chapter 5Methodologychapter.5, the *fairness* of the allocation methods was discussed by using the fairness criteria defined in Chapter 3Literature Reviewchapter.3. From this discussion, it was concluded that due to their inability to guarantee a unique solution the Lorenz allocation and EPM satisfy a relatively small amount of these fairness properties. Although the Star-method does satisfy a lot of them, the fact that its allocation is often outside the core is of such importance that the method would only be preferred for cases in which the level of synergy is relatively equal for all sets of customers. Because in that case its allocation will be equal to the allocation of the EPM, in terms of fairness one would never be worse off by taking the EPM instead of the Star-method. As choosing between the EPM and Lorenz allocation is a matter of taste, they receive the same number in the ranking. Because the Shapley value satisfies almost all the fairness properties and its allocation is very often in the core, it is preferred in terms of fairness. The Nucleolus is a good alternative, but does not satisfy the monontonicity and additivity properties.

The *robustness* of the allocation methods was measured in three ways, namely by observing the fluctuation in allocated emissions due to the actual variation in order sizes, by fixing the selected customers' demand and by fixing the demand of the other customers. The observed CoVs of the allocated emissions showed that in some routes the Lorenz allocation and EPM suffer from the fact that they are unable to generate a unique solution. Even though for most customers they provide an allocated emission which is relatively stable over the different instances of demand sizes, there are a few customers for which the amount fluctuates quite a lot. Because the maximum CoV value was especially high for the Lorenz allocation, it is considered to be less robust than the EPM. The lowest CoV values were observed for the Star-method. When the demand of a selected customer is fixed while the demand of the others change, the effect on its allocated emission is preferably small. In this respect, the Shapley value would have been preferred. When the demand of the selected customer is increased, it is important that the emission allocated to this customer does not decrease. The fact that this was observed for the Lorenz allocation is probably also because the method does not guarantee to find a unique solution.

For completeness, the *understandability* property is also added to the ranking of the allocation methods in Table 9.1Ranking of the allocation methodstable.9.1. In Chapter 5Methodologychapter.5, the motivation for this ranking has already been discussed.

Finally, a ranking of the methods in terms of their *computational effort* is presented. For the business case, the computation time of all the allocation methods was at most a few seconds per route and therefore acceptable in all cases. However, when the number of customers in a route increases the computation times will rise sharply for all allocation methods except for the Star-method. For the Shapley value, EPM and Lorenz allocation, a route with 16 customers will require only 5 minutes of computation time but a route including 18 customers will require more than an hour. Because the Nucleolus is found by solving at least one LP problem, its computational effort may even be larger. Which amount of computation time is regarded as acceptable will depend on the application.

Even though the Shapley value shows the lowest value on average in Table 9.1Ranking of the allocation methodstable.9.1, it may not be preferred in any case. The choice of an allocation method should always depend on the application. For example, if the situation is such that customers do not agree with a solution outside the core, then a stable method would be more appropriate. In case the number of customers on a route and/or the number of routes to be evaluated are very large, the Star-method may be preferred due to its low computational effort.

The application of the penalty methods to the business case showed that the methods based on the impact that customers' time windows have are highly dependent on the specific route constructed by the VRPTW program. Because the penalties of the relatively simple methods are closely related to the total penalty to be shared, they are much more robust than the other two. The extent to which the impact methods stimulate customers to widen their time windows could not be measured because the routes constructed by the VRPTW did not improve as the time window of a single customer was widened. It was already explained that because method 2 penalizes customers based on the length of their time windows, this method stimulates customers to broaden them. As method 1 assigns an equal penalty to every customer with at least one time window, this method clearly does not. The computation times of the penalty methods were higher for methods 3 and 4 but were of such level that they should not be decisive when determining which of the approaches should be preferred. Because it seems to be fair to assign penalties to the customers who are responsible for the extra amount of CO_2, but the high amount of fluctuation that was observed for the impact methods is not desired, it is suggested to combine method 2 with method 3 or 4. In this case, customers are penalized based on the 'polluter-pays' principle, they are stimulated to increase the length of their time windows and their penalties are relatively robust.

Chapter 10

Further Research

In this final chapter of the thesis, several recommendations for further research will be made. Some of these suggestions were already mentioned in earlier chapters, but will be explained in more detail below.

First of all, due to the large computational effort of some operations, a few simplifications were needed in this research. One of these simplifications is the use of estimations of the stand-alone emissions of subsets. This resulted in the use of the so-called 'pseudo-core' instead of the actual core. Even though the differences are assumed to be small, it would be interesting to test the methods using the actual values of the stand-alone emissions.

Another suggestion for further research is to modify the mathematical definition of the EPM and Lorenz allocation such that they will always find a unique solution. This will probably reduce the high CoV values of the allocated amount of emissions for these methods and may therefore imply a large improvement. A unique solution may be achieved by extending the methods with a similar technique that is used to find a unique Nucleolus. For example, to find a unique Lorenz allocation, one may first minimize the absolute difference in allocated emissions over all pairs of customers. This is the way the Lorenz allocation is obtained currently. However, multiple allocations may have the same minimum difference. Therefore, a step to minimize the difference between the smallest amount of emission and the second largest emission may be added. As long as the solution is not unique, another LP problem can be solved to minimize the differences between customers. Note that this is only one of the several possibilities to find a unique Lorenz allocation. Because the largest difference between the customers has already been minimized in the first step, the fairness of distributional equality has already been achieved as much as possible and the additional steps are primarily needed to guarantee uniqueness of the solution. The same reasoning can be applied to the achievement of a unique EPM allocation.

For the penalty methods, it would be interesting to investigate to what extent the results of the 'actual' impact of a customer would differ from the results presented in Chapter 8Resultschapter.8. In this research, two simplifications have been made in order to keep the computation time of the impact methods at a reasonable level. First of all, only the impact of the time window(s) of a single customer has been investigated. In some cases, the time windows of a subset of two or more customers could be responsible for the fact that the LSP cannot take the fastest route. In two of the routes in the business case, no single customer was found to have impact. However, the time windows of at least one subset of customers including two or more customers was such that the amount of emissions in the route was higher than without taking into account any of the time windows. Especially for such types of routes, it would have been better to consider the impact of subsets as well. The other simplification made in determining the impact of a customer's time window was made because the repetitive optimization of routes takes a lot of computation time. Therefore, it was suggested to only shift the customer whose time window is ignored to other positions within the route and then check whether this option would have resulted in a faster trip. In some cases, modifying the positions of two or more customers might have resulted in a faster trip.

In the business case used to test both the allocation and the penalty methods, the same routes were used to visit the same set of customers in every week. It would also be interesting to test the methods on a data set in which the same customers are served in different routes. Then, the effect of the route on the amount of emissions can be investigated. From the hypothetical cases one could observe that the number and especially the locations of the other customers are important for the amount of emissions allocated to a customer. However, these cases were rather ‘extreme’ in the sense that the selected customer was situated very far away from the others in some of the scenarios. It would be interesting to see the extent to which an allocated amount of emissions to a customer would be influenced by the actual variation in routes that are chosen by an LSP.

Because the customers of the business case only varied their demand among the different weeks and did not change their time windows, the influence of their actual variation could not be investigated. The lengths of the original time windows of the customers were modified, but their positions on the day remained the same over the instances. Investigating the effect of both varying the size and shifting the position of the customers’ time windows is an interesting topic. Due to time limitations this investigation has not been performed in this thesis.

Bibliography

[1] J. Albrecht. The diffusion of cleaner vehicles in co_2 emission trading designs. *Transportation Research Part D: Transport and Environment*, 5(5):385401, September 2000.

[2] J. Arin, J. Kuipers, and D. Vermeulen. Geometry and computation of the lorenz set. *International Journal of Game Theory*, 6(2):223–238, 2004.

[3] O.N. Bondareva. Some applications of linear programming methods to the theory of cooperative games. *Problemy Kibernetiky*, 10:119–139, 1963.

[4] M. Borderé. Een empirische vergelijking van 'gain sharing' mechanismes in de logistiek. Master's thesis, Universiteit Gent, 2009.

[5] B. Dutta and D. Ray. A concept of egalitarianism under participation constraints. *Econometrica*, 57:615–635, 1989.

[6] S. Engevall, M. Göthe-Lundgren, and P. Värbrand. The travelling salesman game: An application of cost allocation in a gas and oil company. *Annals of Operations Research*, 82:453–471, 1998.

[7] S. Engevall, M. Göthe-Lundgren, and P. Värbrand. The heterogeneous vehicle-routing game. *Transportation Science*, 38(1):71–85, February 2004.

[8] J. Fuglestvedt et al. Climate forcing from the transport sectors. *PNAS*, 105(2):454–458, January 2008.

[9] S. Dasgupta et al. The impact of sea level rise on developing countries: A comparative analysis. Technical report, World Bank, February 2007.

[10] S. Solomon et al. Irreversible climate change due to carbon dioxide emissions. *PNAS*, 106(6):1704–1709, February 2009.

[11] M. Frisk, M. Göthe-Lundgren, K. Jörnsten, and M. Rönnqvist. Cost allocation in collaborative forest transportation. Technical Report 15, NHH Dept. of Finance & Management Science, November 2006.

[12] D. Gately. Sharing the gains from regional cooperation: A game theoretic application to planning investment in electric power. *International Economic Review*, 15(1):195–208, 1974.

[13] D. Gillies. *Contributions to the Theory of Games IV*, chapter Solutions to general non-zero-sum games, pages 47–85. Princeton University Press, 1959.

[14] J.H. Grotte. Computation of and observation on the nucleolus, and the central games. Master's thesis, Cornell University, September 1970.

[15] T. Hokari. Population monotonic solutions on convex games. *International Journal of Game Theory*, 29:327–338, 2000.

[16] J. Homberger and H. Gehring. Two evolutionary metaheuristics for the vehicle routing problem with time windows. *Information Systems and Operational Research*, pages 297–318, 1999.

[17] J.L. Hougaard. *An Introduction to Allocation Rules*, chapter Cost Allocation as Cooperative Games, pages 61–96. Springer, first edition, 2009.

[18] R. Jain, D.M. Chiu, and W. Hawe. A quantitative measure of fairness and discrimination for resource allocation in shared systems. Technical report, Digital Equipment Corporation, September 1984.

[19] N. Ligterink, L. Tavasszy, and R. de Lange. A speed and payload dependent emission model for heavy-duty road freight transportation. TNO and Delft University of Technology, 2010.

[20] S.C. Littlechild and K.G. Vaidya. The propensity to disrupt and the disruption nucleolus of a characteristic function game. *International Journal of Game Theory*, 5:151–161, 1976.

[21] M. Maschler and B. Peleg. A characterization, existence proof and dimension bounds for the kernel of a game. *Pacific Journal Mathematics*, 18:289–328, 1966.

[22] M. Maschler, B. Peleg, and L.S. Shapley. Geometric properties of the kernel, nucleolus, and related solution concepts. *Mathematics of Operations Research*, 4(4):303–338, 1979.

[23] N. Megiddo. On the nonmonotonicity of the bargaining set, the kernel, and the nucleolus of a game. *SIAM Journal on Applied Mathematics*, 27:355–358, 1974.

[24] J. Von Neumann and O. Morgenstern. *Theory of games and economic behavior*. Princeton University Press, second edition, 1947.

[25] D. Schmeidler. The nucleolus of a characteristic function game. *SIAM Journal of Applied Mathematics*, 17(6):1163–1170, November 1969.

[26] L. S. Shapley. A value for n-person games. *Annals of Mathematics Studies*, 28:307–317, 1953.

[27] L. S. Shapley. On balanced sets and cores. *Naval Research Logistics*, 14:453–460, 1967.

[28] L. S. Shapley. Cores of convex games. *International Journal of Game Theory*, 1:11–26, 1971.

[29] M. Shubik. Incentives, decentralized control, the assignment of joint costs, and internal pricing. *Management Science*, 8(3):325–343, 1962.

[30] A. Sichwardt. Co_2 allocation in road transportation for alwex transport ab. Master's thesis, Linnaeus University, 2011.

[31] Y. Sprumont. Population monotonic allocation schemes for cooperative games with transferable utility. *Games and Economic Behavior*, 2(4):378394, December 1990.

[32] J. L. Sullivan, R. E. Baker, B. A. Boyer, R. H. Hammerle, T. E. Kenney, L. Muniz, and T. J. Wallington. Co2 emission benefit of diesel (versus gasoline) powered vehicles. *Environmental Science and Technology*, 38(12):32173223, May 2004.

[33] S.H. Tijs. *Game Theory and Mathematical Economics*, chapter Bounds for the core and the tau-value, pages 123–132. North-Holland, 1981.

[34] S.H. Tijs and T.S.H. Driessen. Game theory and cost allocation problems. *Management Science*, 32(8):1015–1028, August 1986.

[35] F. Tricoire, M. Romauch, K. Doerner, and R. Hartl. Heuristics for the multi-period orienteering problem with multiple time windows. *Computers and Operations Research*, 37:351–367, 2010.

[36] H.P. Young. Monotonic solutions of cooperative games. *International Journal of Game Theory*, 14(2):65–72, 1985.

Appendix A

Results Hypothetical Cases

A.1 Cases with 5 Customers

Table A.1: Results of the hypothetical test cases including 5 customers

Scenario	Demand Selected Customer	Non-empty pseudo-core?	Star allocation in pseudo-core?	Shapley allocation in pseudo-core?	# LP problems solved by Nucleolus	Proft % Sel. Cust. in EPM	Lorenz = 'equal split allocation'?
1	Rel. High	Yes	No	Yes	11	42.5%	Yes
(DSO,-,-)	Average	Yes	No	Yes	11	42.1%	Yes
	Rel. Low	Yes	No	Yes	10	40.2%	Yes
2	Rel. High	Yes	No	Yes	2	13.8%	No
(DS,O,-)	Average	Yes	No	Yes	2	12.5%	No
	Rel. Low	Yes	No	No	2	6.1%	No
3	Rel. High	Yes	No	Yes	2	13.9%	No
(DS,-,O)	Average	Yes	No	Yes	2	12.6%	No
	Rel. Low	Yes	No	No	2	6.1%	No
4	Rel. High	Yes	No	No	8	6.3%	No
(DO,S,-)	Average	Yes	No	No	8	6.3%	No
	Rel. Low	No	-	-	-	-	-
5	Rel. High	Yes	Yes	Yes	1	77.0%	Yes
(D,SO,-)	Average	Yes	Yes	Yes	1	76.9%	Yes
	Rel. Low	Yes	Yes	Yes	1	76.3%	Yes
6	Rel. High	Yes	Yes	Yes	1	76.3%	Yes
(D,S,O)	Average	Yes	Yes	Yes	1	76.1%	Yes
	Rel. Low	Yes	Yes	Yes	1	75.6%	Yes
7	Rel. High	No	-	-	-	-	-
(O,D,S)	Average	No	-	-	-	-	-
	Rel. Low	No	-	-	-	-	-
8	Rel. High	Yes	No	No	8	3.0%	No
(DO,-,S)	Average	Yes	No	No	8	3.0%	No
	Rel. Low	No	-	-	-	-	-
9	Rel. High	Yes	No	Yes	1	53.1%	No
(D,O,S)	Average	Yes	No	Yes	1	53.1%	No
	Rel. Low	Yes	No	Yes	1	53.1%	No
10	Rel. High	Yes	Yes	Yes	1	78.5%	Yes
(D,-,SO)	Average	Yes	Yes	Yes	1	76.1%	Yes
	Rel. Low	Yes	Yes	Yes	1	77.9%	Yes
Overall % or Avg.		83.3%	36%	76%	3.5	42.4%	48%

A.2 Cases with 10 Customers

Table A.2: Results of the hypothetical test cases including 10 customers

Scenario	Demand Selected Customer	Non-empty pseudo-core?	Star allocation in pseudo-core?	Shapley allocation in pseudo-core?	# LP Problems solved by Nucleolus	Profit % Sel. Cust. in EPM	Lorenz = 'equal split allocation'?
1	Rel. High	Yes	No	Yes	18	66.9%	No
(DSO,-,-)	Average	Yes	No	Yes	19	66.4%	No
	Rel. Low	Yes	No	Yes	35	64.1%	No
2	Rel. High	Yes	No	Yes	2	13.5%	No
(DS,O,-)	Average	Yes	No	No	2	10.6%	No
	Rel. Low	No	-	-	-	-	-
3	Rel. High	Yes	No	Yes	2	13.5%	No
(DS,-,O)	Average	Yes	No	No	2	10.6%	No
	Rel. Low	No	-	-	-	-	-
4	Rel. High	Yes	No	Yes	18	8.3%	No
(DO,S,-)	Average	Yes	No	Yes	20	8.3%	No
	Rel. Low	No	-	-	-	-	-
5	Rel. High	Yes	Yes	Yes	1	88.1%	Yes
(D,SO,-)	Average	Yes	Yes	Yes	1	87.9%	Yes
	Rel. Low	Yes	Yes	Yes	1	87.2%	Yes
6	Rel. High	Yes	Yes	Yes	1	88.5%	Yes
(D,S,O)	Average	Yes	Yes	Yes	1	88.4%	Yes
	Rel. Low	No	Yes	Yes	1	87.7%	Yes
7	Rel. High	No	-	-	-	-	-
(O,D,S)	Average	No	-	-	-	-	-
	Rel. Low	No	-	-	-	-	-
8	Rel. High	Yes	No	Yes	17	4.0%	No
(DO,-,S)	Average	Yes	No	Yes	17	4.0%	No
	Rel. Low	No	-	-	-	-	-
9	Rel. High	Yes	No	Yes	1	53.5%	No
(D,O,S)	Average	Yes	No	Yes	1	53.5%	No
	Rel. Low	Yes	No	Yes	1	53.5%	No
10	Rel. High	Yes	Yes	Yes	1	89.0%	Yes
(D,-,SO)	Average	Yes	Yes	Yes	1	88.9%	Yes
	Rel. Low	Yes	Yes	Yes	1	88.2%	Yes
Overall % or Avg.		76.7%	39.1%	91.3%	7.1	53.2%	39.1%

A.3 Cases with 15 Customers

Table A.3: Results of the hypothetical test cases including 15 customers

Scenario	Demand Selected Customer	Non-empty pseudo-core?	Star allocation in pseudo-core?	Shapley allocation in pseudo-core?	# LP problems solved by Nucleolus	Profit % Sel. Cust. of EPM	Lorenz = 'equal split allocation'?
1	Rel. High	Yes	No	No	39	70.3%	No
(DSO,-,-)	Average	Yes	No	No	48	69.8%	No
	Rel. Low	No	-	-	-	-	-
2	Rel. High	Yes	No	Yes	2	13.1%	Yes
(DS,O,-)	Average	Yes	No	Yes	2	8.6%	No
	Rel. Low	No	-	-	-	-	-
3	Rel. High	Yes	No	Yes	2	13.1%	No
(DS,-,O)	Average	Yes	No	No	2	8.6%	No
	Rel. Low	No	-	-	-	-	-
4	Rel. High	Yes	No	No	65	8.2%	No
(DO,S,-)	Average	Yes	No	No	85	8.2%	No
	Rel. Low	No	-	-	-	-	-
5	Rel. High	Yes	Yes	Yes	1	91.3%	Yes
(D,SO,-)	Average	Yes	Yes	Yes	1	91.1%	Yes
	Rel. Low	Yes	Yes	Yes	1	90.2%	Yes
6	Rel. High	Yes	Yes	Yes	1	92.1%	Yes
(D,S,O)	Average	Yes	Yes	Yes	1	91.9%	Yes
	Rel. Low	Yes	Yes	Yes	1	91.1%	Yes
7	Rel. High	No	-	-	-	-	-
(O,D,S)	Average	No	-	-	-	-	-
	Rel. Low	No	-	-	-	-	-
8	Rel. High	Yes	No	Yes	65	4.0%	No
(DO,-,S)	Average	Yes	No	No	85	4.0%	No
	Rel. Low	No	-	-	-	-	-
9	Rel. High	Yes	No	Yes	1	54.8%	No
(D,O,S)	Average	Yes	No	Yes	1	54.8%	No
	Rel. Low	Yes	No	Yes	1	54.8%	No
10	Rel. High	Yes	Yes	Yes	1	92.3%	Yes
(D,-,SO)	Average	Yes	Yes	Yes	1	92.1%	Yes
	Rel. Low	Yes	Yes	Yes	1	91.4%	Yes
Overall % or Avg.		73.3%	40.9%	68.2%	18.5	54.4%	45.5%