

Bachelor Thesis

One-sided, budget-balanced reserve-price mechanisms for ridesharing

Philipp Klein

Date of Submission: 27. Juli 2026
Advisor: Prof. Dr. Marie Schmidt



Julius-Maximilians-Universität Würzburg
Lehrstuhl für Informatik I
Algorithmen und Komplexität

Abstract

This thesis proposes a budget-balanced mechanism for ridesharing. Travelers report their trip data, their car model and open seats, and a willingness-to-pay (WTP) for being taken along as a rider, while their roles as drivers or riders are assigned endogenously by the mechanism. Since driver cost are computed from the publicly verifiable car model, only the riders' private WTP is strategically relevant. This makes the problem one-sided and therefore escapes the assumptions of the Myerson-Satterthwaite impossibility theorem. Every feasible driver-rider pair is assigned a one-sided reserve price, which covers the driver's detour cost plus half the distance cost of the otherwise shared part of the trip. The allocation is computed in polynomial time as a maximum-weight matching on a graph over the travelers. A matched rider pays the larger of the reserve price and the Clarke-Pivot-Payment directly to their driver. The mechanism is strongly budget balanced and individually rational by construction and efficient under the reserve price constraint. It is not incentive compatible, as a rider can profit from underreporting when drivers differ in their cost per kilometer. Assuming a uniform cost per kilometer restores incentive compatibility, but puts efficiency and individual rationality into question instead.

Zusammenfassung

Diese Bachelorarbeit stellt einen budget-balanced Mechanismus für Ridesharing vor. Teilnehmer geben ihre Fahrtdaten, ihr Automodell und angebotene Anzahl an Sitzplätzen, sowie ihre Zahlungsbereitschaft als Mitfahrer an. Der Mechanismus bestimmt endogen wer Fahrer und Mitfahrer ist. Da die Kosten eines Fahrers aus dem überprüfbaren Automodell berechnet werden, ist nur die private Zahlungsbereitschaft der Mitfahrer strategisch relevant. Somit wird das Problem einseitig und die Annahmen des Unmöglichkeitstheorem von Myerson und Satterthwaite gelten hier nicht mehr. Jedem zulässigen Fahrer-Mitfahrer-Paar wird ein einseitiger Reserve Price zugewiesen, welcher die Detourcost des Fahrers sowie die halben Distanzkosten des ansonsten gemeinsam gefahrenen Teils der Strecke abdeckt. Die Zuordnung wird in Polynomialzeit als Maximum-Weight-Matching auf einem Graphen über den Teilnehmern berechnet. Ein vermittelter Mitfahrer zahlt das Maximum aus Reserve Price und Clarke-Pivot-Payment direkt an seinen Fahrer. Der Mechanismus ist per Konstruktion strongly budget-balance und individuell rational sowie effizient gegeben die Einschränkungen aus dem Reserve Price. Er ist nicht incentive

compatible, da ein Mitfahrer von einer manipulierten niedrigeren Zahlungsbereitschaft profitieren kann, wenn sich die Fahrer in ihren Kosten pro Kilometer unterscheiden. Die Annahme einheitlicher Kosten pro Kilometer stellt die Incentive Compatability wieder her, stellt dafür aber Effizienz und individuelle Rationalität in Frage.

Contents

1	Introduction	6
2	Background	8
2.1	Taxonomy of the Ridesharing Problem	8
2.2	Mechanism Properties	10
2.3	Myerson-Satterthwaite Theorem	11
3	Related literature	12
3.1	Matching and Optimization	12
3.2	Auction and VCG Based Mechanisms	12
3.3	Deficits, Reserve Prices, and Budget Balance	13
3.4	Cost Sharing and Stability	14
4	VCG mechanism solution	15
4.1	The Model and Mechanism	15
4.2	Payments and resulting Deficits	16
5	The Model	18
5.1	Commuters and Types	18
5.2	Travel Distances, Times, and Cost	18
5.3	Valuations and Utilities	19
5.4	Model Discussion	19
5.4.1	One-Sidedness and Public Driver Cost	20
5.4.2	The Linear Time Cost	20
5.4.3	Truthful Reporting of the Remaining Type Components	20
5.4.4	Eliciting a Willingness-to-Pay Instead of Cost	21
5.4.5	Time Windows and Long Detours	21
5.5	Reserve Price Calculation	22
5.5.1	Budget Balance and Cost Sharing	23
5.5.2	The Reserve Price Case by Case	23
5.5.3	Reserve Prices and Efficiency	24
6	The Mechanism	25
6.1	Feasible Pairs	25
6.2	The Matching Graph and the Efficient Allocation	26
6.3	Payments	27
6.4	Properties	27

7 Conclusion	31
Bibliography	34

1 Introduction

Ridesharing means that two or more travelers with similar routes make a trip together in one car instead of getting there separately. The car and the cost are shared, potentially one vehicle is taken off the road and therefore emissions and congestion get reduced. A traveler who would otherwise use public transportation may get to their destination faster and more comfortably and could potentially use the travel time more productively. A traveler who'd use their own car otherwise may enjoy the same benefits, besides getting there faster. The potential for more ridesharing is considerable, since most car trips are made with empty seats. In the United States roughly three quarters of all commute trips are single occupant vehicle trips, and the numbers for Europe are similar [SASG18]. Commuting is the setting in which this potential is easiest to realize. Commuters travel to and from work at roughly the same time on every working day, so their trips are known well in advance. Participation nevertheless stays low.

Ridesharing can be viewed as a mechanism design problem. Participants hold private information about where they want to go, when they need to be there and what a ride is worth to them, and they will report whatever serves them best. A mechanism collects these reports, decides who shares a ride with whom and determines the payments. Four properties are commonly asked of such a mechanism, namely efficiency, incentive compatibility, individual rationality and budget balance. They are defined in Section 2.2.

The obvious candidate is the VCG mechanism with Clarke-Pivot-Payments, and it is the starting point of this thesis. It maximizes social welfare, leaves no participant with a reason to misreport their type and never makes participation worse than staying away. However, it will usually result in a large budget deficit. Drivers are paid the full value they create for their riders, while riders only pay for the externality they impose on the other participants. A rider who causes no detour and does not displace another rider pays nothing at all. So the mechanism needs to be subsidized to pay drivers in full. This is discussed in detail in Chapter 4.

The Myerson-Satterthwaite theorem shows that this is not a flaw of the VCG mechanism [MS83]. In a two-sided market with private information on both sides, no mechanism satisfies all four properties at once. One of them has to be given up, and VCG gives up budget balance. For a platform which is meant to actually be run in the real world, that is the wrong choice. The price a rider pays should be equal to the payment the driver receives, so that riders can just pay drivers directly and everything is taken care of. The question of this thesis is how close one can stay to the VCG solution once budget balance is treated as a hard constraint.

The approach taken here is to model the problem differently and floor VCG payments with a one-sided reserve price. The impossibility result needs private information on both sides of the market, so the model of this thesis removes it from one side. Drivers report their car model, from which their cost per kilometer can be computed, which makes the driver side public. Riders report a willingness to pay for being taken along, which is the only strategically relevant piece of private information left.

Every feasible pair of a driver and a rider is then assigned a reserve price. It covers the detour cost the driver incurs and adds a share of the otherwise shared route between driver and rider. A rider is only eligible for a match if their willingness to pay reaches that reserve price. The allocation is computed as a maximum weight matching on a graph whose nodes are the commuters. The the roles of driver and rider are not fixed in advance but determined by the mechanism. A matched rider pays the larger of the reserve price and the Clarke pivot payment directly to their driver.

The resulting mechanism is strongly budget balanced and individually rational by construction, and it maximizes social welfare over those matchings which the reserve prices admit. Those properties and to what extent incentive compatibility can be achieved are discussed in Section 6.4.

The remainder of this thesis is structured as follows. Chapter 2 classifies the ridesharing problem, defines the four mechanism properties and states the impossibility result which motivates the model. Chapter 3 discusses the related literature. Chapter 4 shows how the VCG mechanism with Clarke-Pivot-Payments solves the two-sided ridesharing problem and where its deficit comes from. Chapter 5 introduces the one-sided model together with the reserve price and discusses the modeling choices behind both. Chapter 6 defines the mechanism itself and evaluates it against the four properties. Chapter 7 concludes.

2 Background

This chapter will give a short introduction to some relevant background knowledge.

2.1 Taxonomy of the Ridesharing Problem

The ridesharing problem has many different facets and can be modeled in many different ways. In the following, the different variants of the problem will be presented in detail. The classification follows the survey of Agatz et al. [AESW12]. The last paragraph then places this thesis within the taxonomy.

Static and dynamic settings. In a static setting all participants submit their data before a fixed cut-off point. The mechanism computes all matchings and payments at once and puts them into effect afterwards. In a dynamic setting participants enter their data over time and either receive offers within a short time or quickly become visible to others who are looking for a ride. Dynamic systems are usually implemented with a rolling horizon approach, where the matching problem is solved repeatedly for the requests which arrived within a short time frame [AESW12]. In particular, for commuter ridesharing, trips are known well in advance, which makes the static setting a natural choice for a commuter platform.

Fixed roles and role flexibility. Many models take the set of drivers and the set of riders as given and mutually exclusive. The allocation problem is then bipartite and comparatively easy to solve. Under role flexibility a participant who owns a car can also be assigned the role of a rider, so the role is endogenously decided by the mechanism instead of being fixed beforehand. This generally requires more computational work. For commuting it is a very relevant scenario. When two commuters with similar routes both own a car, a match between them means that only one of the two cars is used.

Number of participants per rideshare. A further dimension is how many drivers and riders a single match may contain.

- **One-to-one.** Every driver is matched with at most one rider and every rider with at most one driver.
- **One-to-many.** A driver can take along several riders, limited by the number of free seats in their car, while a rider can only be matched to a single driver.
- **Many-to-one.** A driver serves at most one rider, but a rider can be matched to several drivers in sequence to reach their goal.

- **Many-to-many.** A combination of the two previous cases.
- **Multimodal.** The mechanism additionally takes other publicly available travel options for a route and their cost into account, such as public transportation or taxis.

Ride-back. In particular in commuter ridesharing, the participating riders usually need a ride to work in the morning and then also a ride back home in the evening. The two directions can be treated as separate problems. This makes it pretty infeasible to have less cars on the roads and save emissions by bringing together commuters who'd both take a car on their own otherwise. But a commuter who'd otherwise get to work by public transportation can of course still make one of their rides that way. The more appropriate mechanistic solution in this scenario is to guarantee a ride-back and match a rider in the morning only if a suitable return ride exists as well. Guaranteeing the return trip makes participation considerably more attractive, but it couples the two matching problems. Hence, it rules out matches which would otherwise be feasible for one-direction trips.

Types and their elicitation. From the mechanism design perspective one central question is which information the participants hold privately and how it is elicited. Typical components of a participant's type are their origin, their destination, a time window, the number of free seats and a monetary valuation of the ride. But it can go further and include specific valuations such as cost per km driven, cost per minute spent driving/riding or personal caps on detour time and distance. Most models in the literature are two-sided, meaning that drivers and riders both hold a private valuation which has to be elicited. Usually the participants just state their type to the mechanism. Therefore it is crucial for a mechanism to incentivize truthful reporting or at least limit the possible gains by a manipulated type report.

Objective, computation and payments. The remaining dimensions concern what the matching optimizes for, how it is computed and how the mechanism-determined payments are made. Common objectives are minimizing the total travel distance or the total cost, maximizing the number of matched participants, or maximizing social welfare. The computational approaches range from polynomial time matching algorithms on a suitably constructed graph to integer programs and heuristics. Payments can be made directly from rider to driver, but this is only sufficient if the rider's price equals the driver's payment. Otherwise the mechanism operator must collect the money from the riders and distribute it among the drivers and potentially subsidize the mechanism.

Classification of this thesis. The mechanism developed in this thesis is static and allows for role flexibility. All matchings are one-to-one and there is no ride-back. The problem is modeled as one-sided, since the driver cost are assumed to be publicly known and only the riders' willingness to pay has to be elicited. The allocation maximizes social welfare subject to a reserve price constraint and is computed as a maximum weight matching in

polynomial time. Payments are only made directly from riders to drivers, so no operator is involved and no deficit can arise. The properties which these design choices are meant to achieve are defined in the following section.

2.2 Mechanism Properties

A mechanism collects a report $\hat{\theta}_i$ from every participant i , chooses an allocation π and determines a payment p_i for each participant. Four properties are commonly used to judge such a mechanism, namely budget balance, individual rationality, incentive compatibility and efficiency. This thesis' mechanism will also be evaluated with regard to these properties in Section 6.4. Throughout, utilities are assumed to be quasi-linear, so a participant's utility is their valuation of the allocation minus their payment,

$$u_i = v_i(\pi) - p_i.$$

A payment p_i is positive if i pays and negative if i receives money.

Budget balance. Budget balance concerns the money which flows into and out of the mechanism. A mechanism is strongly budget balanced if the payments of all participants sum to zero, that is $\sum_i p_i = 0$. Everything participants receive is paid for by other participants. The mechanism operator neither keeps nor adds money.

Everything paid by some participants is received by others and the mechanism operator neither keeps nor adds money. Weak budget balance only requires $\sum_i p_i \geq 0$, so the operator is allowed to retain a surplus but must not subsidize the participants. If $\sum_i p_i < 0$, the mechanism runs a deficit and has to be financed from the outside.

Individual rationality. A mechanism is individually rational if no participant is made worse off by taking part in it. The utility of not participating is assumed to be zero, so individual rationality means $u_i \geq 0$ for every participant i . If a participant can have negative utility as a result of the mechanism, they are of course less likely to take part, everything else being equal.

Incentive compatibility. A mechanism is called (dominant-strategy) incentive compatible if all participants are best off by truthfully reporting their type, regardless of what others report. Formally, for every participant i with true type θ_i , every possible misreport θ'_i and all possible reports by the other participants without i , θ'_{-i} ,

$$u_i(\theta_i, \theta'_{-i}) \geq u_i(\theta'_i, \theta'_{-i}).$$

A mechanism's allocation of drivers to riders is always made to fulfill an optimization goal. So for a mechanism to actually work and truly maximize this goal, the elicited information it works with must be true. Therefore a mechanism which invites misreporting cannot be relied upon to produce the allocation it was designed for.

Efficiency. A mechanism is efficient if it always finds the optimal allocation to its optimization goal, given any feasible input.

So if the goal is social welfare maximization, an efficient mechanism always finds the allocation π^* which maximizes the sum over all participants' valuations. Formally that is

$$\pi^* \in \arg \max_{\pi} \sum_i v_i(\pi).$$

In this scenario all matchings of an efficient allocation have non-negative social welfare, otherwise they would contribute negatively to the above sum.

Clearly, all of these properties are desirable. The difficulty of mechanism design lies in the fact that they cannot be obtained at the same time if both riders and drivers hold private information, which is the subject of the following section.

2.3 Myerson-Satterthwaite Theorem

The four properties of the previous section are of course individually implementable, but sadly they cannot generally be obtained together. The classic impossibility result is due to Myerson and Satterthwaite [MS83]. They study bilateral trade, where a single seller owns one indivisible good and a single buyer is interested in it, and where the seller's cost and the buyer's valuation are private information. For such a market they show that no mechanism is efficient, incentive compatible, individually rational and budget balanced at the same time.

Bilateral trade is a special case of the ridesharing problem, since a single driver matched to a single rider is exactly a seller and a buyer negotiating over one seat. The impossibility therefore carries over.

Theorem 2.1 ([ZZG⁺14, MS83]). *It is impossible to have a truthful, efficient and individually rational ridesharing mechanism without outside subsidies.*

A mechanism designer has to give up one of the four properties. The VCG mechanism with Clarke-Pivot-Payments (see Chapter 4) is efficient, truthful and individually rational but gives up budget balance.

The theorem is based upon on private information being held on both sides of the market. This is the assumption which the model of this thesis effectively drops. If the driver cost are publicly known and only the riders hold private valuations, then the problem is one-sided and the impossibility no longer applies directly. Whether all four properties can then actually be achieved is a separate question, which Chapter 6 takes up.

3 Related literature

Research on ridesharing has developed along two fairly separate lines. The first line treats ridesharing as an optimization problem and asks how a platform can compute good matches under spatial and temporal constraints. The second line treats it as a market and asks how participants can be brought to truthfully state their type so that computations can be made to actually fulfill their optimization goal. Our mechanism sits in the second line but of course the allocation step is still an optimization problem. We therefore review both.

3.1 Matching and Optimization

Agatz et al. [AESW12] give the standard taxonomy, which differentiates settings by whether a driver serves one or several riders and by whether a rider may transfer between drivers. The one-to-one scenario reduces to a maximum weight matching problem and is solvable in polynomial time, while the remaining cases become variants of the pickup and delivery problem and are NP-hard. In an earlier simulation study for Metro Atlanta, Agatz et al. [AESW12] show that a rolling horizon strategy, which collects requests over a short window before reoptimizing, outperforms immediate matching. Stiglic et al. [SASG18] study which kind of participant flexibility actually produces matches and find that the detour a driver is willing to accept matters considerably more than the width of a time window. Armant and Brown [AB14] and Chen et al. [CMSQ16] extend the formulation with role switching and with meeting points, and Sun et al. [SZQ25] show how sensitive such models are to the assumed value of time. Furuhata et al. [FDO⁺13] survey operating ridesharing agencies and classify their pricing practice into catalog prices, rule based prices, and prices negotiated between the participants.

What these models have in common is that they take reported trip information at face value. Participants state their type and the platform assumes that these announcements are truthful. Tafreshian et al. [TMY20] state this explicitly in their review and note that strategic behavior is largely absent from the ride matching literature.

3.2 Auction and VCG Based Mechanisms

A smaller body of work applies mechanism design directly to ridesharing. Kleiner et al. [KNZ11] propose a system of parallel auctions in which a rider states a price and a driver decides whether to accept the resulting match. Kamar and Horvitz [KH09] evaluate several VCG based payment schemes on real commuting logs and show the tension between truthfulness and the payments the platform has to make. Shi et al.

[SHC22] design a second-price mechanism in which drivers report a private cost per unit distance and riders are charged a fare that the platform fixes in advance. Their mechanism is one-sided in the sense that only one side of the market holds private information. It is worth noting that their asymmetry runs opposite to ours. Shi et al. keep the driver cost private and treat riders as price takers, whereas we assume the driver cost to be publicly known and let the riders report a willingness-to-pay.

3.3 Deficits, Reserve Prices, and Budget Balance

The obstacle that all of this work runs into is the impossibility result of Myerson and Satterthwaite [MS83], as discussed in Section 2.3. Parkes et al. [PKE01] respond to this in the setting of combinatorial exchanges by imposing budget balance and individual rationality and then accepting a controlled loss of truthfulness. The ridesharing literature has mostly made the opposite choice and kept truthfulness while accepting a deficit.

The starting point for this thesis is the work of Zhao et al. [ZZG⁺14]. They model commuters who report a solo travel cost, a time window, and a seat count, and whose roles as driver, rider, or unmatched are decided by the mechanism rather than declared in advance. Plain VCG with Clarke pivot payments is truthful and individually rational in their model but can run an arbitrarily large deficit, since a driver is paid the full welfare he creates while a rider pays only the detour cost he imposes. They then propose two alternatives. A serial dictatorship with fixed prices controls the deficit completely but assigns roles without using cost information and is correspondingly inefficient. Their VCG mechanism with two-sided reserve prices sets a floor on what a rider pays and a ceiling on what a driver receives. It is truthful and individually rational whenever the riders' reserve price is no smaller than the drivers' reserve price. The deficit is then bounded rather than eliminated.

The same idea appears elsewhere with different parameters. Lloret-Batlle et al. [LBMN17] let commuters report a value of time and a value of distance and use reserve levels on both to classify participants into roles endogenously, which squeezes the VCG deficit from both sides. Li et al. [LCLZ25] set reserve prices per unit distance on the willingness to pay of riders and the willingness to accept of drivers in a ride sourcing market. Masoud and Tafreshian [Mas21] and [TM22] design truthful subsidy schemes, which keep efficiency and truthfulness by letting an outside party pay the difference.

Our mechanism uses the reserve price idea but implements a one-sided reserve price. Then the payments in our model flow directly from rider to matched driver and budget balance is a structural property. The reserve price serves as a floor on the rider payment, and it contains a share of the base cost of the driver so that a rider who causes no detour still contributes to the trip. This is the case in which Clarke pivot payments collapse to zero and in which the two-sided constructions above still rely on a subsidy or on a bounded deficit.

3.4 Cost Sharing and Stability

A further group of papers shares our concern that the money should come from the participants themselves. Hu et al. [HUV21] develop a cost sharing framework for a commuter driver who wants to recover cost rather than make a profit, and they show that a reduced burden for the first passenger and a fair split of the direct cost of the driver cannot both be had. Their mechanisms are budget balanced by construction, but the matching is taken as given and lies outside the model. Peng et al. [PSJ⁺20] compute a stable matching for commuters and let riders offer compensation in order to make themselves more attractive. This makes the riders the strategic side and the drivers reactive. Tafreshian and Masoud [TM20] reformulate role assignment and matching as a roommate problem with transferable utility and show that a stable outcome need not exist initially, because a small external subsidy can restore one.

4 VCG mechanism solution

This chapter shows how the ridesharing problem is solved by the Vickrey-Clarke-Groves (VCG) mechanism with Clarke-Pivot-Payments (CPP) when the problem is modeled two-sided, as done by Zhao et al. [ZZG⁺14]. The mechanism is efficient, individually rational and incentive compatible but will usually result in a large budget deficit.

For consistency with the rest of this thesis, everything is presented for the one-to-one case, while Zhao et al. allow several riders per driver.

4.1 The Model and Mechanism

Every traveler i holds a private type θ_i which contains their origin, their destination, a time window, the number of free seats and the cost c_i of making the trip on their own. The crucial difference to the model of this thesis is that c_i is private on both sides of the market. A driver knows what their drive costs them and a rider knows what their solo trip (without a car) costs them, and neither is observable by the mechanism. Based on the reported types the mechanism assigns every traveler a role, matches drivers with riders and determines the payments.

An allocation π is valued by a rider r at the cost they save by being taken along and by a driver d at the negative of the detour they take on,

$$v_r(\pi) = c_r, \quad v_d(\pi) = -\text{detourcost}_{d,r},$$

and zero for unmatched travelers. The mechanism selects the allocation which maximizes social welfare, the sum over all participants' valuations. Formally that is,

$$\pi^*(\theta) \in \arg \max_{\pi} V(\theta, \pi), \quad V(\theta, \pi) = \sum_i v_i(\pi).$$

The payment of traveler i is the Clarke-Pivot-payment

$$x_i^{\text{Clarke}}(\theta) = V(\theta_{-i}, \pi_{-i}^*) - (V(\theta, \pi^*) - v_i(\pi^*)), \quad (4.1)$$

where π_{-i}^* is the efficient allocation if i were not to participate in the mechanism. So the first term is the social welfare if i were not to participate. The second term is the welfare of the efficient allocation with i minus i 's own valuation of that allocation. The difference between these two terms is then the externality i imposes on everybody else, and that is what i pays. A positive x_i means i pays and a negative x_i means i receives money. Riders pay for the cost they cause and drivers are paid for the cost they save their passengers. With this payment rule the mechanism is efficient, individually rational and incentive compatible [ZZG⁺14].

4.2 Payments and resulting Deficits

In the following three crucial cases are presented to see how the payments work and understand why a large budget deficit is effectively inevitable in the real world.

A single feasible pair without a detour. Suppose the only mechanism participants are a driver d and a rider r who have the same origin and destination, so $\text{detourcost}_{d,r} = 0$. The welfare of the match is c_r . If r were not to participate, no matches could be made and $V(\theta_{-r}, \pi_{-r}^*) = 0$. The second term of the CPP ($V(\theta, \pi^*) - v_i(\pi^*)$) equals zero as well, since the driver does not incur a detour. So the rider's payment $x_r^{\text{Clarke}}(\theta) = 0 - 0 = 0$. The driver, on the other hand, receives $x_d^{\text{Clarke}}(\theta) = -c_r$, the full value the ride has for the rider. Again, the CPP's first term equals zero, since no matches can be made when d does not participate. The second term equates to c_r . Therefore the mechanism has to pay out c_r , but does not receive a payment.

A single feasible pair with a detour. Now let $\text{detourcost}_{d,r} > 0$, so the welfare of the match is $c_r - \text{detourcost}_{d,r}$. The same calculation gives $x_r = \text{detourcost}_{d,r}$ and again $x_d = -c_r$. The rider pays the detour cost they cause and the driver still receives the rider's full valuation. The mechanism again pays the driver c_r , this time against an income of $\text{detourcost}_{d,r}$. The detour cost lie within $]0; c_r]$, since a detour by definition costs something and the match will only be made if its social welfare is non-negative. So the mechanism will in most cases again incur a deficit, which will be large for relatively long trips and short detours.

Competition between two riders. Let one driver d be feasible for two riders r_1 and r_2 , and denote the welfare of the respective pair by $w_i = c_i - \text{detourcost}_{d,r_i}$. Assume $w_1 > w_2$, so the mechanism matches r_1 . Without r_1 the driver would take r_2 along and the social welfare is w_2 . On the other hand, the social welfare of the matching with r_1 but without r_1 's valuation is the negative of the caused detour cost $-\text{detourcost}_{d,r_1}$. Since the second term is subtracted from the first, r_1 's payment is

$$x_{r_1} = w_2 + \text{detourcost}_{d,r_1}.$$

The driver still just receives the social welfare created by him $x_d = -c_1$. The budget deficit is then exactly $w_1 - w_2$, the difference in social welfare between the pairings. This is the case since the mechanism pays out c_1 and gets x_{r_1} , so the deficit is:

$$\begin{aligned} c_1 - x_{r_1} &= c_1 - (w_2 + \text{detourcost}_{d,r_1}) \\ &= c_1 - \text{detourcost}_{d,r_1} - (c_2 - \text{detourcost}_{d,r_2}) \\ &= w_1 - w_2 \end{aligned}$$

Therefore the mechanism's gap between income and payout is much smaller under competition, as x_{r_1} now includes the entire social welfare of the other pairing in addition to the caused detour by r_1 .

Note that under competition the payment determination is effectively a sealed-bid second-price auction. The riders 'bid' w_i , the welfare their pairing would generate, and the winner pays the second highest bid plus their own detour cost.

Summation. When there is no competition the rider pays nothing or the detour cost, while the driver receives c_r . Consequently the mechanism has a deficit of $c_r - \text{detourcost}_{d,r}$. That is the entire social welfare the match generates! Under competition the deficit is limited to $w_1 - w_2$, the difference in social welfare between the pairings.

5 The Model

This chapter first defines this thesis' base model. The modeling choices and their implications are then discussed in Section 5.4.

5.1 Commuters and Types

Let C be the set of all commuters and Θ the set of all types. All commuters $i \in C$ have a private type $\theta_i \in \Theta$, with

$$\theta_i = (o_i, d_i, [est_i, lat_i], WTP_i, car_i, seats_i).$$

Here o_i and d_i are the origin and destination locations of i . The interval $[est_i, lat_i]$ is the commuter's time window. It indicates when i could leave o_i at the earliest, the earliest start time est_i , and when they have to arrive at d_i at the latest, the latest arrival time lat_i . What a commuter is willing to pay to be taken along for the trip from o_i to d_i within $[est_i, lat_i]$ is denoted by WTP_i . If a commuter only wants to offer their services as a driver and is not interested in potentially becoming a rider, then $WTP_i = 0$. A commuter who is willing to be a driver must declare their exact car model car_i and the number of passengers they are willing to take along $seats_i$. If a commuter does not have a car or is not interested in driving themselves, car_i is set to *none* and $seats_i = 0$.

5.2 Travel Distances, Times, and Cost

Given car_i , the assumption is made that a driver d 's cost per km c_d^{km} is publicly known. Furthermore, it is assumed that any additional time a commuter spends on a ride beyond the duration of their solo trip from o_i to d_i is accurately priced per minute at the minimum wage. This is captured by the constant $timecost = 0.25/\text{min}$, for a minimum wage of 15/h.

From a commuter i 's origin o_i , destination d_i , and time window $[est_i, lat_i]$, the mechanism can compute the best route from o_i to d_i . The route's length in km is given as the travel distance td_i and its duration in minutes as the travel time tt_i . For a rideshare, the entire trip's distance and time are denoted by $td_{d,r_1,\dots,r_n}$ and $tt_{d,r_1,\dots,r_n}$, where d is the rideshare's driver and r_1 through r_n are all the riders participating in the rideshare. The corresponding route starts and ends at the driver's origin o_d and destination d_d .

These definitions allow multiple riders to be picked up by a single driver. The remainder of the model is only defined for the one-to-one scenario, where all matchings are between exactly one driver and one rider.

The detour cost a driver d incurs when matched with rider r is the sum of the additional distance and time cost incurred by d as a result of the rideshare, compared to d 's own solo route:

$$\begin{aligned} \text{detourcost}_{d,r} &= \text{detour distance cost}_{d,r} + \text{detour time cost}_{d,r} \\ &= (td_{d,r} - td_d) \cdot c_d^{km} + (tt_{d,r} - tt_d) \cdot \text{timecost} \end{aligned}$$

The additional kilometers driven by d when taking r along, instead of driving on their own, are given by subtracting d 's solo travel distance td_d from the length of the ridesharing tour $td_{d,r}$. They are priced with the publicly known c_d^{km} , the cost d incurs per km with their specific car. The extra time spent traveling is calculated analogously and priced with timecost , the minimum wage per minute.

5.3 Valuations and Utilities

For a matching π between riders and drivers given by the mechanism, the valuations and utilities are as follows. All utility functions are assumed to be quasi-linear.

A rider r values a matching π according to the value function $v_r(\pi)$:

$$v_r(\pi) = \begin{cases} WTP_r, & \text{if } r \text{ is matched,} \\ 0, & \text{if } r \text{ is unmatched.} \end{cases}$$

A matched rider r 's utility then is

$$u_r = WTP_r - p_r,$$

where p_r is the price rider r has to pay to their driver, which is determined by the mechanism. An unmatched rider has a utility of zero.

On the other hand, for a driver d , matched with a rider r in π , their valuation of this matching is:

$$v_d(\pi) = \begin{cases} -\text{detourcost}_{d,r}, & \text{if driver } d \text{ is matched with rider } r, \\ 0, & \text{if driver } d \text{ is unmatched.} \end{cases}$$

As defined at the end of the last section, $\text{detourcost}_{d,r}$ is the specific detour distance and time cost d incurs for taking r along. Therefore, a driver d 's utility when matched to rider r is

$$u_d = p_r - \text{detourcost}_{d,r},$$

where again p_r is the mechanism-determined price r has to pay d . An unmatched driver derives a utility of zero.

5.4 Model Discussion

The model just defined makes a number of assumptions and modeling choices which deserve a closer look. This section discusses them one by one.

5.4.1 One-Sidedness and Public Driver Cost

In the literature, the ridesharing problem is usually modeled as a two-sided problem, where both drivers and riders hold private information which needs elicitation. We call this model one-sided, because drivers are not directly asked for their cost. Instead, they report their specific car model car_i as part of their type. A matched rider can easily verify this. Therefore it is assumed that car_i is always stated truthfully. However, it should be noted that in practice it actually is not that easy for a rider to verify the exact car model. While the brand and the general model are easy to identify, the exact motorization is not necessarily knowable just from seeing the car.

Given the truthfully stated car model car_i , the assumption is made that the actual cost per km of driver d , denoted by c_d^{km} , can be accurately calculated. To do this, the current fuel cost for the specific car needs to be looked up. But even more can be incorporated into c_d^{km} . There are insurance and maintenance cost, taxes, and even the value depreciation could be part of it. For example, the ADAC provides extensive lists of all of these cost factors for most car models ². The mechanism operator can then decide which of these factors should actually be a part of c_d^{km} . In practice, the fuel cost per km of course also differs depending on the speed driven. For the purpose of this mechanism, it is assumed that c_d^{km} accurately reflects the cost per km of driver d .

5.4.2 The Linear Time Cost

Time cost is not always modeled in the literature. By using a constant time cost in this model, we essentially have a linear time cost function. Therefore it is assumed that everybody has the same time cost and that every additional minute is worth the same. A non-linear function, where the cost for an additional single minute increases over time at some point, is probably better suited to accurately model the time cost. Longer detours would then become more expensive. More about the time cost function is said in Section 5.4.5 at the end of this discussion.

5.4.3 Truthful Reporting of the Remaining Type Components

What about the other type features, besides WTP_i ? We now show how participants cannot gain anything by manipulating those.

What could a rider gain from manipulating o_i and/or d_i ? In general, this means they will not be picked up or dropped off where they actually need it. So at first glance, it does not make sense to falsify this information. However, one could imagine that a rider states an origin and/or destination which is closer to important points in the traffic network, like a central station, instead of their actual locations. This would be done in the hope of decreasing the detour cost and increasing the likelihood of being matched. The rider then incurs the cost for getting to such a point, or for getting to their destination from one, and their WTP would also change accordingly. So this is not so much a manipulation where a rider gains something illicitly, but rather just a different

²<https://www.adac.de/rund-ums-fahrzeug/auto-kaufen-verkaufen/autokosten/uebersicht/>

demand towards the mechanism. A driver could also declare only a part of their actual tour as their type. Again, we do not view this as a manipulation, but just as a different preference by the driver. If they only want to offer a rideshare for a certain part of their trip, that is fine.

Manipulating the time window also does not make sense for any participant. Extending the window beyond one's actually available time in the hope of increasing the likelihood of a match leads to unwanted matches. Any match which is only possible because of the extended window means the participant will have to leave earlier than they can or arrive later than they are allowed to. Shortening the time window as a rider just leads to forgone matches. For a driver, a shortened time window also leads to forgone matches and it limits the possible detour time. A driver who wants to limit their detour time in this way, again, is not really manipulating, but rather truthfully stating how much spare time they offer for detours.

Finally, the number of seats in a driver's car $seats_i$ might be understated by a driver. One more time, we view this as a truthful statement with regard to the actual preferences of the driver. If a driver does not want to offer all of their actually available seats, that is fine. Overreporting seats, of course, cannot benefit a driver.

So we assume that all type information besides WTP_i is always truthfully stated. The ridesharing problem then is a one-sided problem in this model, where only the riders' WTP needs to be truthfully elicited.

5.4.4 Eliciting a Willingness-to-Pay Instead of Cost

Most papers in the literature ask riders for their cost to get from origin to destination if they are not participating in the mechanism. Here, a different approach is taken by asking riders for their willingness-to-pay (WTP) to get from origin to destination. This is done to take into account how a rider's actual valuation is probably not the same between taking public transportation and a rideshare. A rideshare might be quicker, is probably more comfortable, and could allow the time spent traveling to be used more productively. It also allows a potential driver who is open to being a rider to state their WTP for being driven. Again, this is probably different from the cost they would incur by driving themselves, as the time spent traveling can now be used differently.

5.4.5 Time Windows and Long Detours

Using a time interval like $[est_i, lat_i]$ to determine time feasibility is the standard in the ridesharing literature. However, it should be noted that this is somewhat problematic. Here is an example to show why. Assume a driver's ride, for example to get to work, takes 30 minutes. They have a flexible work schedule so that they can arrive at work between 8 and 9 am. Therefore, their time interval is [7:30 am, 9:00 am]. But this now leaves an entire hour of potential detour time the driver could be burdened with. This is twice the duration of their own trip and could be even larger for a more flexible time window! It is fair to assume that most drivers with such a truthfully stated time interval do not actually want to incur detour time which exceeds the duration of their own trip,

even if they are being paid for it.

To avoid such long detour times, there are a couple of options, but all of them have downsides, too. The mechanism could implement a global cap on detour time, so that any matching exceeding the cap is not feasible. Of course, this will forgo matches where a driver actually would be willing to take a longer detour, so efficiency is lost. A different option is for participants to state their individual maximum detour time as part of their type. However, in reality such a rigid cap on the detour time is probably not accurate. A driver might really be fine with a detour which takes a minute longer than their cap, if it pays enough. So, again, efficiency might be lost.

Finally, the time cost function might be best suited to alleviate the time window problem. As stated in Section 5.4.2, a non-linear time cost function, where the cost per minute rises with every additional minute, would probably be a more accurate reflection of the detour time cost a driver incurs. This would make longer detours more expensive for riders and thereby compensate drivers appropriately for long detours. The time cost function could also depend on a driver's stated time window, so that there is an individual time cost function for every driver. Extra detour time could then, for example, only become more expensive once a certain threshold is crossed, like half the duration of the driver's own route. Using a different, more expensive time cost function could again forgo some matches which the participants would be fine with, so efficiency could be lost.

5.5 Reserve Price Calculation

All proposed mechanisms in this thesis use a one-sided reserve price. For every combination of a driver and one or more riders, whose location and time constraints make a matching feasible, the mechanisms calculate a unique reserve price. For the combination to then be eligible for matching, every participating rider's WTP_{r_i} must be at least as high as their reserve price. The reserve price is the minimum price a rider will have to pay and at the same time the minimum payment a driver will receive. In the one-to-one scenario, rider r can therefore only be matched with driver d if $WTP_r \geq res_{d,r}$. The reserve price $res_{d,r}$ is calculated as follows:

$$res_{d,r} = detourcost_{d,r} + 0.5 \cdot c_d^{km} \cdot \min\{td_r, td_d\}$$

The reserve price has two components. The first is the detour cost $detourcost_{d,r}$, which d incurs as a result of being matched with r . The second makes the rider pay for half of the driver-dependent distance cost of the distance they share with the driver beyond the detour. The rider pays for half the distance cost of their own route td_r , unless this route is longer than the driver's base route td_d , in which case half the distance cost of td_d has to be paid. The following discussion motivates the reserve price and shows case by case how well $\min\{td_r, td_d\}$ reflects the actually shared distance.

5.5.1 Budget Balance and Cost Sharing

To have a (strongly) budget-balanced mechanism, all payments go directly from riders to drivers. What a rider pays is exactly what the matched driver gets. As shown in Section 4.2, the VCG mechanism with Clarke-Pivot-Payments results in a rider payment of zero if the rider does not cause a detour and does not displace another potential rider, while the driver gets paid the rider's valuation. With only the direct payments from riders to drivers, a payment of zero clearly would not be fair. Of course, the driver would drive anyway, but their cost should be split with the rider. This is where the second component of the reserve price comes in. The first component makes the rider pay for the detour cost they cause, just like a Clarke-Pivot-Payment would. The second component makes the rider additionally share some of the cost of the ride the driver would take anyway.

5.5.2 The Reserve Price Case by Case

To illustrate how the two components work together, we now go through all relevant cases of how the driver's and the rider's routes can relate to each other.

First, consider the two cases without a detour, so $\text{detourcost}_{d,r} = 0$.

If both driver and rider have the exact same origin and destination, they also have the same travel distance $td_r = td_d = td_{d,r}$ and the reserve price becomes $res_{d,r} = 0.5 \cdot c_d^{km} \cdot td_{d,r}$. Hence, r must at least be willing to pay for half of the entire ride's distance cost to be eligible for matching. Conversely, the driver recovers half of the distance cost which they would also have when not participating in the mechanism.

The second case is a rider whose origin and destination are on the driver's route, but $td_r < td_d$ and $\min\{td_r, td_d\} = td_r$. Then we have $res_{d,r} = 0.5 \cdot c_d^{km} \cdot td_r$. Now the rider only pays for half the distance cost of their own, shorter route and not of the driver's longer route. In this way, it is avoided that a rider has to pay for distance they are not a part of. Note that without a detour, the rider's route can never be longer than the driver's, so these two cases are exhaustive.

Now consider the cases with $\text{detourcost}_{d,r} > 0$. If there is a detour, the entire rideshare's distance is longer than the distance the driver would drive on their own, $td_{d,r} > td_d$. Recall that the detour distance cost is $(td_{d,r} - td_d) \cdot c_d^{km}$. So any additional travel distance the driver incurs is already accounted and paid for by the first component of the reserve price. What remains is the distance cost of the driver's base route td_d and the question to what extent the rider should share it.

If $td_r \geq td_d$, then $\min\{td_r, td_d\} = td_d$. In addition to the detour cost, the rider then pays for half of the driver's entire base route. The additional kilometers by which td_r might exceed td_d are already covered by the detour cost.

If $td_r < td_d$, then $\min\{td_r, td_d\} = td_r$. The rider then pays for half the distance cost of their own route, since besides the detour distance, the car is shared for at most td_r kilometers of the driver's base route. In most cases, $\min\{td_r, td_d\}$ accurately reflects the shared distance which is not already part of the detour. This is however not the case if the rider's origin or destination are further away from the driver's destination than the

driver's origin. To show why this is the case, we go through the most extreme scenario in the following. Consider a rider whose trip runs contrary to the driver's direction of travel. Suppose the driver, after dropping the rider off at d_r , is still further away from their destination than they were at their own origin. Then the rider's route has brought the driver no closer to d_d and can be regarded as being entirely a part of the detour. The rider then effectively pays twice for their own route. They pay its full distance cost once as part of the detour cost and then again half of its distance cost through the second component of the reserve price. If for example only the rider's origin is further away from d_d , then the rider will only pay twice for a fraction of his own route. This can still be considered fair in the following sense. The reserve price in this scenario depends only on the caused detour cost and the length of the rider's own route, not on the direction of that route. A different rider with the same detour cost and the same route length td_r , whose trip happens to bring the driver closer to their destination, pays exactly the same reserve price. Also, from the driver's perspective, both of these riders cause equal cost and therefore they should have the same reserve price. So from the perspective of the individual rider, the double payment might seem unfair. Seen as a pricing rule across all riders, however, it treats equal riders equally and charges more where more inconvenience is caused. In this sense, the double payment is less a flaw of the reserve price and more a consequence of a uniform pricing rule. It could even be considered a feature in the sense that riders have the same reserve price for the same detour cost and td_r .

5.5.3 Reserve Prices and Efficiency

Introducing a reserve price in a ridesharing mechanism always means some feasible matches might be forgone. There could be a driver who is willing to accept less than the reserve price, while the only feasible rider for them has a true WTP below the reserve price. The feasibility condition $WTP_r \geq res_{d,r}$ does not allow such a match to be made. So a reserve price inevitably means that the efficiency of the mechanism cannot be guaranteed.

6 The Mechanism

The mechanism combines an efficient, reserve-price-constrained allocation with Clarke pivot payments, analogous to the VCG mechanism with reserve prices of Zhao et al. [ZZG⁺14]. In contrast to their setting, the problem here is one-sided. The only privately held and strategically relevant type information is a traveler's willingness to pay WTP_i and the reserve price is one-sided as well. All remaining components of θ_i , namely o_i , d_i , $[est_i, lat_i]$, car_i and $seats_i$, are assumed to be reported truthfully (see Section 5.4.3). As in the model above, the mechanism is defined for the case of a single rider per driver.

Each traveler $i \in T$ submits a report θ_i of their type to the mechanism. The mechanism then determines a matching and the payments in the following steps.

6.1 Feasible Pairs

The mechanism first determines which travelers could share a ride in which role. An *ordered* pair (d, r) of commuters $d, r \in T$, $d \neq r$, where d acts as the driver and r as the rider, is called *feasible* if and only if all of the following conditions hold:

1. d offers a car: $car_d \neq \text{none}$ and $seats_d \geq 1$,
2. r is interested in riding along: $WTP_r > 0$,
3. the pair is time- and location-feasible: the rideshare route from o_d through o_r and d_r to d_d admits a departure time such that d departs no earlier than est_d , r is picked up no earlier than est_r , and r and d arrive at d_r and d_d no later than lat_r and lat_d respectively,
4. the rider's willingness-to-pay covers the reserve price: $WTP_r \geq res_{d,r}$.

The set of all feasible ordered pairs is denoted by $F \subseteq T \times T$. Note that conditions 1 and 2 do not exclude one another. A participant i with $car_i \neq \text{none}$, $seats_i \geq 1$ and $WTP_i > 0$ may appear in feasible pairs both as a driver and as a rider. Which role such a traveler ultimately takes is not fixed in advance but determined endogenously by the mechanism.

The social welfare generated by realizing a feasible pair $(d, r) \in F$ is the sum of the two participants' valuations,

$$v_{d,r} = WTP_r - \text{detourcost}_{d,r}. \quad (6.1)$$

Notice that, by condition 4 and the definition of the reserve price, the following holds for every feasible pair (d, r) :

$$v_{d,r} = WTP_r - \text{detourcost}_{d,r} \geq \text{res}_{d,r} - \text{detourcost}_{d,r} = 0.5 \cdot c_d^{km} \cdot \min\{td_r, td_d\} > 0.$$

So every feasible pair generates positive social welfare, since c_d^{km} and all travel distances are positive numbers.

6.2 The Matching Graph and the Efficient Allocation

The allocation problem is represented as a matching problem on an undirected, weighted graph $G = (T, E)$, with one node for each participant. An edge $\{i, j\} \in E$ exists if and only if at least one of the two orientations of the pair is feasible, i.e. $(i, j) \in F$ or $(j, i) \in F$. Since it is an undirected graph, note the use of set notation for the edges, so that the edge $\{i, j\}$ is the same edge as $\{j, i\}$. The weight of an edge is the social welfare of its best feasible orientation:

$$w_{\{i,j\}} = w_{\{j,i\}} = \max\{v_{d,r} \mid (d, r) \in F \text{ with } \{d, r\} = \{i, j\}\}. \quad (6.2)$$

For every edge, the mechanism stores its social welfare maximizing orientation, which fixes the roles of i and j in case the edge is selected. If both orientations are feasible and attain the same welfare, one of them is chosen randomly.

Now a maximum-weight matching on G can be calculated to find the social welfare maximizing allocation of riders to drivers. In a graph a matching $M \subseteq E$ contains at most one edge incident to each node. So here every participant is assigned at most one other participant and the stored orientation of an edge assigns the roles of rider and driver for the pair.

Given any reported type profile $\theta = (\theta_1, \dots, \theta_{|C|})$, the mechanism computes the *efficient allocation* $\pi^*(\theta)$ as a maximum-weight matching on G :

$$\pi^*(\theta) \in \arg \max_{M \subseteq E} \sum_{e \in M} w_e. \quad (6.3)$$

So the efficient allocation is a matching which maximizes social welfare. Note the use of the operator $\in$. This is necessary, since there can be more than one allocation which achieves the maximum social welfare. Ties between matchings of equal total weight are again resolved randomly. A maximum-weight matching on an undirected graph can be computed in polynomial time. Edmonds' blossom algorithm [Edm65] solves this in $O(EV^2)$ and an implementation running in $O(V^3)$ is due to Gabow [Gab74].

The total social welfare of a matching M is denoted by

$$V(\theta, M) = \sum_{e \in M} w_e,$$

so that $\pi^*(\theta)$ maximizes $V(\theta, \cdot)$ over all feasible matchings M in G .

6.3 Payments

Payments flow exclusively and directly from matched riders to their respective drivers. Unmatched commuters pay and receive nothing. A rider r 's payment when matched with driver d is then the larger of the pairing-specific reserve price and the CPP:

$$p_r = \max\{res_{d,r}, x_r^{Clarke}(\hat{\theta})\}. \quad (6.4)$$

The reserve price has been defined in Section 5.5 and the CPP in Section 4.1. As the reserve price incorporates both the detour cost and a cost-share of the distance traveled together besides the detour, it will always be larger than the CPP, if there is no competition. In this way the zero- or detour-only-payments of the CPP are avoided. If there is competition for a seat, either value could be larger, depending upon detour costs and WTPs. If the CPP is larger than the reserve price, we are again in the aforementioned second-price auction. The riders' $WTP_i - res_{d,i}$ are effectively the bids and the winner pays the social welfare generated by the other pairing plus their own detour cost.

6.4 Properties

The previously presented mechanism properties (see Section 2.2) are now being evaluated for the proposed mechanism in this section.

Budget balance. Every payment p_r is transferred directly from a matched rider to their driver. Hence, the mechanism never runs a deficit or a surplus. It is strongly budget balanced by construction.

Individual rationality. For a rider r matched with a driver d , $res_{d,r} \leq WTP_r$, by the feasibility condition 4 from Section 6.1. Furthermore $x_r^{Clarke} \leq WTP_r$, because otherwise the matching would result in negative social welfare and not be made in the first place. Therefore r 's payment $p_r = \max\{res_{d,r}, x_r^{Clarke}(\hat{\theta})\} \leq WTP_r$. Consequently r always derives non-negative utility from a matching, $u_r = WTP_r - p_r \geq 0$.

On the other hand, a driver d matched with a rider r has a utility $u_d = p_r - detourcost_{d,r}$. This is also always non-negative, because $p_r \geq res_{d,r} \geq detourcost_{d,r}$. Of course this only the case given the assumption that the detour cost are actually correctly modeled. Unmatched travelers obtain a utility of zero.

Constrained efficiency. Due to the reserve price, the mechanism will forgo social welfare increasing matchings between a rider r with $detourcost_{d,r} \leq WTP_r < res_{d,r}$ and a driver d who would accept this offer and is otherwise unmatched. This is inevitable when a reserve price is implemented. However, given any type profile and the resulting possible pairings between riders and drivers, which do fulfill the feasibility constraints, the mechanism will always correctly compute the (or one of the) social welfare maximizing

allocation(s). As shown in Section 6.2, the allocation problem can be modeled as max-weight matching and algorithms solving this problem exist and can even be computed in polynomial time.

Incentive compatibility. The mechanism is not incentive compatible. A rider can gain utility by underreporting their willingness-to-pay under certain conditions. We show this with a minimal example, which is depicted in Figures 6.1 and 6.2. The figures also illustrate how the allocation problem is solved as a matching problem on the graph defined in Section 6.2.

There are two drivers d_1, d_2 and two riders r_1, r_2 . All pairs are time- and location-feasible and do not require a detour, so all detour cost are zero. When there are no detours, the reserve price is rider-independent and only covers half of a driver's distance cost. Assume the reserve prices are $res_{d_1, r} = 2$ and $res_{d_2, r} = 1$ for both riders r . This difference in reserve price is then due to differing c_d^{km} . So here d_1 has a car with higher fuel cost, insurance cost, etc. The riders' true valuations are $WTP_{r_1} = 2$ and $WTP_{r_2} = 1.5$.




Fig. 6.1: The matching graph for truthful reports. The nodes are the participants. Edges exist between two nodes if the pairing is feasible. Edge weights are a pairing's generated social welfare. The green edges indicate the social welfare maximizing allocation.

First, consider the case where both riders report truthfully, shown in Figure 6.1. Since all detour cost are zero, the weight of an edge, the social welfare generated by its pairing, is simply the reported WTP of its rider. Rider r_2 is not feasible for driver d_1 , because $WTP_{r_2} = 1.5 < 2 = res_{d_1, r}$. This leaves the edges $\{r_1, d_1\}$ and $\{r_1, d_2\}$, both with weight 2, and the edge $\{r_2, d_2\}$ with weight 1.5. The maximum-weight matching pairs r_1 with d_1 and r_2 with d_2 , for a total social welfare of 3.5. Now consider r_1 's payment. If r_1 were not to participate, the mechanism would still match r_2 with d_2 and the social welfare would be 1.5. In the chosen allocation, the social welfare of everybody but r_1 is also 1.5. So r_1 causes no detour and displaces nobody, and their Clarke-Pivot-Payment is $x_{r_1}^{Clarke} = 1.5 - 1.5 = 0$. This is exactly the situation the reserve price was introduced

for and r_1 pays

$$p_{r_1} = \max\{res_{d_1, r_1}, x_{r_1}^{Clarke}\} = \max\{2, 0\} = 2.$$

Their utility is $u_{r_1} = 2 - 2 = 0$.




Fig. 6.2: The matching graph with the manipulated report $WTP_{r_1}' = 1.6$. The edge between r_1 and d_1 vanishes. There is a new social welfare maximizing allocation.

Now let r_1 instead report $WTP_{r_1}' = 1.6$, shown in Figure 6.2. The pair (d_1, r_1) is no longer feasible, because $1.6 < 2 = res_{d_1, r_1}$, so the edge $\{r_1, d_1\}$ disappears from the graph. Both riders now compete for d_2 . The maximum-weight matching pairs r_1 with d_2 , since $1.6 > 1.5$, and leaves r_2 unmatched. Without r_1 , the mechanism would again match r_2 with d_2 for a social welfare of 1.5. In the actually chosen allocation the social welfare without i is zero. The Clarke-Pivot-Payment is therefore $x_{r_1}^{Clarke} = 1.5 - 0 = 1.5$. This time it exceeds the reserve price and r_1 pays

$$p_{r_1} = \max\{res_{d_2, r_1}, x_{r_1}^{Clarke}\} = \max\{1, 1.5\} = 1.5.$$

Their true utility is $u_{r_1} = WTP_{r_1} - p_{r_1} = 2 - 1.5 = 0.5$. So r_1 gains 0.5 utility by underreporting and incentive compatibility is broken.

The example also shows where the manipulation comes from. The reserve price a rider has to pay depends on c_d^{km} , the driver's cost per km. By underreporting, r_1 makes the match with the expensive reserve price infeasible and forces a matching with the cheaper driver d_2 . There the payment is determined by the competition with r_2 , in the second-price auction manner described previously. This price is lower than the reserve price of the forgone match with d_1 .

Under the VCG mechanism with Clarke-Pivot-Payments alone, such a manipulation can never be profitable, since the mechanism is incentive compatible. More specifically, in the example the riders would not have pay anything, since no detours are caused and they are not in competition. The reserve prices breaks the incentive compatibility here, because the minimum payment now varies with the selected pairing and a rider's report

influences which pairing is selected. It should be noted that successfully executing such a manipulation requires prior knowledge of the other participants' types and reserve prices. So while truthful reporting is no longer a dominant strategy, a rider cannot systematically profit from misreporting without this information. How large this possible manipulation gain is can be seen in the calculation of r_1 's utility in the manipulation scenario. It is simply the difference between the two riders' truthful valuations. Notice how it is inconsequential which manipulated *WTP* r_1 states, as long as it is smaller than the more expensive reserve price and larger than the other rider's *WTP*.

Incentive compatibility with constant c^{km} . As written in the last paragraph, incentive compatibility breaks down because of differing driver distance cost c_d^{km} . Hence, if we make a further simplifying assumption and assume every driver has the same cost per km, a constant c^{km} , the above incentive compatibility-breaking example vanishes. And incentive compatibility is achieved, since CPPs are incentive compatible and if a rider otherwise has to pay the reserve price a manipulated report can't lead to any gain. A platform operator could for example take the average or median across all cars for c^{km} . Of course this assumption is not realistic. Actually, drivers who do have a car with an unusually high cost per km could now theoretically end up with negative utility. Imagine a scenario where the driver has a rather short ride themselves but gets matched to make a very long detour with a rider, who did not have any competition for the seat. The reserve price will now underprice the detour, while the second component only contains half the cost of the driver's short route, which is again undervalued. But this scenario is quite unlikely. Furthermore some efficiency would in reality be lost, since drivers of cars with actually cheaper cost per km than c^{km} , can now go unmatched when they would have been matched with c_d^{km} .

Now we can answer the question left open in Section 2.3. Making the problem one-sided escapes the assumptions of the Myerson-Satterthwaite impossibility theorem, but it does not automatically restore all four properties. The proposed mechanism with driver-dependent c_d^{km} is (strongly) budget balanced, individual rational and efficient, given the reserve price constraint. But sadly incentive compatibility breaks down in the explained scenario. By further (unrealistically) simplifying the model with constant, driver-independent c^{km} , incentive compatibility is achieved. And given the new assumption the other properties remain intact.

7 Conclusion

This thesis asked how close a ridesharing mechanism can stay to the VCG solution once budget balance is treated as a hard constraint. The Myerson-Satterthwaite theorem rules out having efficiency, incentive compatibility, individual rationality, and budget balance together in a two-sided market. The answer given here was to change the modeling and install a one-sided reserve price to floor the CPP payment rule. The driver cost are computed from the publicly verifiable car model, which makes the driver side of the market public. The riders' willingness-to-pay remains as the only strategically relevant private information. On top of this one-sided model, every feasible pair of a driver and a rider was assigned a reserve price. It covers the driver's detour cost and adds half of the distance cost of the shared part of the trip besides the detour. Hence, a rider who causes no detour and does not displace another rider has to pay something. A matched rider pays the larger of the reserve price and the Clarke-Pivot-Payment directly to their driver.

The resulting mechanism is strongly budget balanced and individually rational by construction. It maximizes social welfare over the matchings which the reserve prices admit and the allocation is computable in polynomial time as a maximum-weight matching. Incentive compatibility, however, does not fully survive. A rider can profitably underreport their willingness-to-pay in order to force the allocation away from a match with an expensive reserve price and towards a cheaper driver. The manipulation rests entirely on drivers differing in their cost per km. Assuming a constant, driver-independent cost per km achieves incentive compatibility and keeps the other properties intact. However this a very strong assumption and in reality a driver with an expensive car may then no longer be fully compensated for a detour. A driver with a cheap car might now go unmatched. Therefore efficiency and individual rationality come into question instead. So even in the one-sided model, the four properties are here only achieved by a very simplifying assumption.

Several directions for future work follow from this. The most natural extension is the one-to-many scenario, where a driver takes several riders along. The model of Chapter 5 already defines travel distances and times for tours with multiple riders, but the reserve price needs adjustment. In the one-to-one case, the rider pays half of the shared distance cost precisely because the car is shared between exactly two people. With several riders, every kilometer of the tour can be shared by a different subset of the passengers, so the shared distance and the fair cost share differ per rider and depend on the pick-up and drop-off order and location. The detour cost can also no longer be attributed to a single rider, since the detour caused by one rider depends on which other riders are part of the tour. On top of that, the allocation problem stops being a matching problem

and becomes NP-hard, as classified in Section 2.1. A second extension is the commuter scenario with ride-backs.

Finally, two more directions seem promising. The linear time cost could be replaced by the non-linear, driver-individual time cost functions discussed in Section 5.4.5, which would price long detours more accurately and alleviate the time window problem. And a simulation on realistic travel data could quantify what the reserve price actually costs. How many welfare-increasing matches does it actually forgo? How does this compare to the deficit a subsidized VCG mechanism incurs on the same data?

AI Disclaimer

Anthropic's Fable 5 and Claude 4.8 were used to enhance the visual structure of the text, create the graphics and enhance the language in spots for better readability. All scientific contributions, ideas and content are the author's own unless otherwise quoted.

Bibliography

- [AB14] Vincent Armant and Kenneth N. Brown: Minimizing the driving distance in ride sharing systems. In *Proceedings of the 26th IEEE International Conference on Tools with Artificial Intelligence (ICTAI)*, pages 568–575, 2014, 10.1109/ICTAI.2014.91.
- [AESW12] Niels Agatz, Alan Erera, Martin Savelsbergh, and Xing Wang: Optimization for dynamic ride-sharing: A review. *European Journal of Operational Research*, 223(2):295–303, 2012, 10.1016/j.ejor.2012.05.028.
- [CMSQ16] Wenyi Chen, Martijn Mes, Marco Schutten, and Job Quint: A ride-sharing problem with meeting points and return restrictions. Beta working paper 516, University of Twente, 2016.
- [Edm65] Jack Edmonds: Maximum matching and a polyhedron with 0,1-vertices. *Journal of Research of the National Bureau of Standards B*, 69B(1–2):125–130, 1965.
- [FDO⁺13] Masabumi Furuhata, Maged Dessouky, Fernando Ordóñez, Marc Etienne Brunet, Xiaoqing Wang, and Sven Koenig: Ridesharing: The state-of-the-art and future directions. *Transportation Research Part B: Methodological*, 57:28–46, 2013, 10.1016/j.trb.2013.08.012.
- [Gab74] Harold N. Gabow: *Implementation of algorithms for maximum matching on nonbipartite graphs*. PhD thesis, Stanford University, Stanford, CA, 1974.
- [HUV21] Shichun Hu, Maged M. Dessouky, Nelson A. Uhan, and Phebe Vayanos: Cost-sharing mechanism design for ride-sharing. *Transportation Research Part B: Methodological*, 150:410–434, 2021.
- [KH09] Ece Kamar and Eric Horvitz: Collaboration and shared plans in the open world: Studies of ridesharing. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI)*, pages 187–194, 2009.
- [KNZ11] Alexander Kleiner, Bernhard Nebel, and Vittorio Amos Ziparo: A mechanism for dynamic ride sharing based on parallel auctions. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 266–272, 2011.

- [LBMN17] Roger Lloret-Batlle, Neda Masoud, and Daisik Nam: Peer-to-peer ridesharing with ride-back on high-occupancy-vehicle lanes: Toward a practical alternative mode for daily commuting. *Transportation Research Record*, 2668(1):21–28, 2017.
- [LCLZ25] Ruijie Li, Haiyuan Chen, Xiaobo Liu, and Kenan Zhang: The VCG pricing policy with unit reserve prices for ride-sourcing is 3/4-individual compatibility. *Transportation Research Part C: Emerging Technologies*, 171:104991, 2025, 10.1016/j.trc.2024.104991.
- [Mas21] Neda Masoud: Design and operation of efficient and budget-balanced shared-use mobility systems. Report 12, Center for Connected and Automated Transportation, University of Michigan, 2021.
- [MS83] Roger B. Myerson and Mark A. Satterthwaite: Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29(2):265–281, 1983, 10.1016/0022-0531(83)90048-0.
- [PKE01] David C. Parkes, Jayant R. Kalagnanam, and Marta Eso: Achieving budget-balance with Vickrey-based payment schemes in exchanges. In *Proceedings of the 17th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1161–1168, 2001.
- [PSJ⁺20] Zixuan Peng, Wenxuan Shan, Peng Jia, Bin Yu, Yonglei Jiang, and Baozhen Yao: Stable ride-sharing matching for the commuters with payment design. *Transportation*, 47(1):1–21, 2020, 10.1007/s11116-018-9960-x.
- [SASG18] Mitja Stiglic, Niels Agatz, Martin Savelsbergh, and Mirko Gradisar: Enhancing urban mobility: Integrating ride-sharing and public transit. *Computers & Operations Research*, 90:12–21, 2018.
- [SHC22] Bing Shi, Xizi Huang, and Zhi Cao: An incentive-compatible and efficient mechanism for matching and pricing in ride-sharing. In *Knowledge Science, Engineering and Management (KSEM 2022)*, volume 13370 of *Lecture Notes in Artificial Intelligence*, pages 191–204. Springer, 2022, 10.1007/978-3-031-10989-8_16.
- [SZQ25] Min Ci Sun, Cheng Zhang, and Luca Quadrifoglio: The critical role of the value of time in peer-to-peer ridesharing models. *Transportation*, 2025, 10.1007/s11116-025-10677-8.
- [TM20] Amirmahdi Tafreshian and Neda Masoud: Using subsidies to stabilize peer-to-peer ridesharing markets with role assignment. *Transportation Research Part C: Emerging Technologies*, 120:102770, 2020, 10.1016/j.trc.2020.102770.
- [TM22] Amirmahdi Tafreshian and Neda Masoud: A truthful subsidy scheme for a peer-to-peer ridesharing market with incomplete information.

Transportation Research Part B: Methodological, 162:130–161, 2022, 10.1016/j.trb.2022.05.010.

- [TMY20] Amirmahdi Tafreshian, Neda Masoud, and Yafeng Yin: Frontiers in service science: Ride matching for peer-to-peer ride sharing: A review and future directions. *Service Science*, 12(2-3):44–60, 2020, 10.1287/serv.2020.0258.
- [ZZG⁺14] Dengji Zhao, Dongmo Zhang, Enrico H. Gerding, Yuko Sakurai, and Makoto Yokoo: Incentives in ridesharing with deficit control. pages 1021–1028, 2014.

Titel der Bachelorarbeit:

One-sided, budget-balanced reserve-price mechanisms for ridesharing

Thema bereitgestellt von (Titel, Vorname, Nachname, Lehrstuhl):

Prof., Marie, Schmidt, Lehrstuhl für Informatik I - Algorithmen und Komplexität

Eingereicht durch (Vorname, Nachname, Matrikel):

Philipp, Klein, 2695600

Ich versichere, dass ich die vorstehende schriftliche Arbeit selbständig verfasst und keine anderen als die angegebenen Quellen und Hilfsmittel benutzt habe. Die benutzte Literatur sowie sonstige Hilfsquellen sind vollständig angegeben. Wörtlich oder dem Sinne nach dem Schrifttum oder dem Internet entnommene Stellen sind unter Angabe der Quelle kenntlich gemacht.

Weitere Personen waren an der geistigen Leistung der vorliegenden Arbeit nicht beteiligt. Insbesondere habe ich nicht die Hilfe eines Ghostwriters oder einer Ghostwriting-Agentur in Anspruch genommen. Dritte haben von mir weder unmittelbar noch mittelbar Geld oder geldwerte Leistungen für Arbeiten erhalten, die im Zusammenhang mit dem Inhalt der vorgelegten Arbeit stehen.

- ☒ Mit dem Prüfungsleiter bzw. der Prüfungsleiterin wurde abgestimmt, dass für die Erstellung der vorgelegten schriftlichen Arbeit Chatbots (insbesondere ChatGPT) bzw. allgemein solche Programme, die anstelle meiner Person die Aufgabenstellung der Prüfung bzw. Teile derselben bearbeiten könnten, entsprechend den Vorgaben der Prüfungsleiterin bzw. des Prüfungsleiters eingesetzt wurden. Die mittels Chatbots erstellten Passagen sind als solche gekennzeichnet.

Der Durchführung einer elektronischen Plagiatsprüfung stimme ich hiermit zu. Die eingereichte elektronische Fassung der Arbeit ist vollständig. Mir ist bewusst, dass nachträgliche Ergänzungen ausgeschlossen sind.

Die Arbeit wurde bisher keiner anderen Prüfungsbehörde vorgelegt und auch nicht veröffentlicht. Ich bin mir bewusst, dass eine unwahre Erklärung zur Versicherung der selbständigen Leistungserbringung rechtliche Folgen haben kann.

Würzburg, 27.07.26,



Ort, Datum, Unterschrift