Brown University

Treating Agents as Workers

Author:
Roman Blum

Advisor:
Professor Amy Greenwald

February 8, 2019

Chapter 1

Introduction

We study centralized marketplaces (e.g., matching markets) where the market sells goods and agents buy them. The twist we introduce in this work is that when agents end up with goods they don't like, the market pays them to compensate for their inconvenience. One motivating example is bike-sharing systems, in which users pay to ride bikes from one place to another. One problem with these systems is that there are not always bikes readily available when and where users want them. Instead of using an external mechanism, like trucks, to redistribute bikes in anticipation of user needs, it should be possible to pay some users, who are internal to the system, to take care of the load balancing. The goal of this work is to propose a framework for calculating appropriate payments when agents are not strictly consumers, but might also be workers.

The mathematical model of marketplaces that we adopt is that of auctions. Auctions are processes where goods are sold at endogenously determined prices. Auctions are fully defined by how bids are reported, how winners are determined, and how payments are computed. In this work, we consider **sealed-bid** auctions, in which all **agents** simultaneously report their bids to the auctioneer, so that no agent's bid can depend on another's. Typically, an agent's bid depends on its true valuation: i.e., their value for the goods up for auction. Depending on these bids, an auctioneer then assigns goods to agents via an **allocation rule**. We consider allocation rules that maximize **social welfare**, the sum of all agents' values for goods received.

In a **direct** mechanism,[1] the auctioneer asks the agents for bids in the

[1]Auctions are a special case of mechanisms, which collect reports from agents, and

form of values. We say that a direct auction is **dominant strategy incentive compatible** (DSIC) when every agent can maximize their utility by bidding truthfully (i.e., reporting their true valuation). DSIC auctions can be useful in practice as they drastically simplify the reasoning process for participating agents. We say an auction satisfies **individual rationality** (IR) if all agents are guaranteed non-negative utility.

In a seminal paper, Myerson [1981] derived a payment formula that ensures the DSIC and IR properties in single-parameter auctions, where an agent's valuation is a single number, canonical examples being the single good auction [Vickrey, 1961] and sponsored search auctions [Edelman et al., 2007]. However, there are important auction environments where agents' valuations cannot be described by only a single parameter. Multiple parameters are necessary, for example, in auctions for multiple goods. Vickrey [1961], Clarke [1971], Groves [1973] designed an auction—the VCG mechanism—which is DSIC and IR in the multi-parameter case. The most popular VCG payment rule is based on the **Clarke pivot rule**. This mechanism charges agents their **externality**: i.e., the amount by which their participation in the auction impacts the other agents. This mechanism can be viewed as a generalization of Myerson's auction to the multi-parameter setting.

Recall that the auctioneer's objective is welfare maximization. Assuming only the usual allocation constraints (i.e., the auctioneer cannot allocate more goods than the number that are up for auction), welfare cannot be maximized by allocating goods to agents with negative valuations. But the problems of interest here have additional allocation constraints. One example is: *all* goods must be allocated (e.g., all employees must be allocated a carpool to work). Another example is: movable goods, like bikes and rental cars, can only be allocated if there is enough space at their destination for them to be returned (packing constraints). In these sorts of problems, allocating agents goods they really don't want—i.e., for which they have *negative* value—is sometimes necessary for the greater good.

Negative valuations are a defining feature of this work. We allow for their possibility because we aim to model auction settings in which some agents emerge as workers, who don't pay for their allocation, but are instead paid by the auctioneer. When we allow for negative valuations VCG with Clarke's pivot rule and Myerson's auction maintain their DSIC property, but

make decisions based on those reports. In mechanisms, like in auctions, the agents' reports reflect their private information about the setting.

they lose their IR guarantees: agents are no longer guaranteed non-negative utilities. This observation leads to the central technical problem studied in this work: *Can we design an auction that is DSIC and IR, even when valuations can be negative?* Intuitively, this goal seems achievable, simply by paying agents when they are allocated goods they don't want. In other words, whenever agents receive a negative allocation, they should also receive a negative payment. But how negative should that payment be?

To guarantee IR, Myerson's payment formula includes an offset term that ensures that the payment of the lowest value does not exceed the value of the corresponding allocation. In our setting, this offset term gains new relevance. We thus adjust Myerson payments via offsets to ensure IR. In doing so, there is some flexibility, so at the same time, we also seek to maximize the auctioneer's revenue: i.e., to minimize payments. We find these offsets using a straightforward binary search.

We also revisit the design of a VCG pivot rule. The Clarke pivot rule is appropriate when agents can impose only *negative* externalities on other agents. But in our setting, agents who work for the greater good impose *positive* externalities. Hence, a slight tweak to the Clarke pivot rule is necessary [Nisan, 2007a]. While this tweak ensures the mechanism is DSIC, it does not ensure IR. As in the single-parameter setting, offsets are necessary to achieve IR. We leave for future research the computational problem of efficiently finding revenue-maximizing offsets in the multi-parameter case.

The answer to the aforementioned question is, thus: *we provide generalizations of both auction mechanisms—Myerson, for the single-parameter setting, and VCG, for bounded valuations in the multi-parameter setting—which are both DSIC and IR, even when valuations can be negative.*

We motivate our problem setup using an auction for comparable goods in which agents have unit demands. Agents are said to exhibit **unit-demand valuations** over k distinct goods when each agent is characterized by k values, one per good, and their value for any set of goods is their maximum value over all goods in that set. Prior to the auction, the auctioneer and the agents agree to a public penalty function over the space of goods. Then agents publicly announce a preferred good, and privately report their bid for that good. The auctioneer then allocates goods, with some agents receiving goods far from their preferred good. Agent's valuations for these non-preferred goods are determined through the penalty function, and can be negative.

We study two problems that fit into this framework: carpool assignment [Hartman et al., 2014] and bike sharing. In both, agents have unit-demands defined over a physical space, so the goods are comparable, making it is easy to construct a public penalty function based on distance. As always, the goal is to maximize welfare while satisfying allocation constraints. We find that in both these problems, it is sometimes the case that allocating agents goods they don't want is necessary to construct the welfare-maximizing allocation (i.e., to serve the greater good). We present concrete examples of this phenomenon, and compute revenue-maximizing DSIC and IR payments.

Chapter 2

Background Information

2.1 Myerson's Lemma

We start in the single-parameter setting. For every agent $i \in N$, let v_i denote their valuation. Similarly, let $\mathbf{v}_{-i}$ be the valuations of every bidder besides i. Then $x_i(v_i, \mathbf{v}_{-i})$ is bidder i's allocation for bidding v_i when other agents bid $\mathbf{v}_{-i}$. Similarly, $p_i(v_i, \mathbf{v}_{-i})$ is bidder i's payment. Let $\underline{v}_i$ and $\overline{v}_i$ represent the minimal and maximal possible valuations for agent i. We assume quasilinear utility where, $u_i(v_i, \mathbf{v}_{-i}) = v_i x_i(v_i, \mathbf{v}_{-i}) - p_i(v_i, \mathbf{v}_{-i})$.

The two critical properties we expect of payment rules are individual rationality and incentive compatibility. Individual rationality guarantees non-negative utility, whereas incentive compatibility states agents maximize utility by bidding truthfully. Formally,

Definition 2.1.1. A payment rule is **individually rational** if and only if for all i, v_i, and v_{-i}: $u_i(v_i, \mathbf{v}_{-i}) \geq 0$

Definition 2.1.2. Incentive Compatibility A payment rule is **incentive compatible** if and only if for all i, v_i, and v_{-i}:

$$u(v_i, \mathbf{v}_{-i}) \geq v_i x_i(t_i, \mathbf{v}_{-i}) - p_i(t_i, \mathbf{v}_{-i}), \quad \forall t_i \in T_i$$

We know present Myerson's Lemma:

Lemma 2.1.3 (Myerson [1981]). *An auction satisfies incentive compatibility and individual rationality if and only if the following two conditions hold* $\forall i \in N, \forall v_i \geq t_i \in T_i, \forall \mathbf{v}_{-i} \in T_{-i}$*:*

1. *The allocation rule is monotone:*

$$x_i(v_i, \mathbf{v}_{-i}) \geq x_i(t_i, \mathbf{v}_{-i})$$

2. *Payments are computed as follows:*

$$p_i(v_i, \mathbf{v}_{-i}) = p_i(\underline{v}_i, \mathbf{v}_{-i}) + v_i x_i(v_i, \mathbf{v}_{-i}) - \int_0^{v_i} x_i(z, \mathbf{v}_{-i}) \mathrm{d}z$$

Therefore if we have a monotonic allocation rule, we can use Myerson's payment rule to find individually rational and incentive compatible payment function.

2.2 Vickrey-Clarke-Groves Mechanisms

As we enter the multi-parameter setting, we see allocations as "bundles" of goods. We define A to be the set of all possible allocations. For agent i and outcome $a \in A$, let a_i be agent i's allocation and $v_i(a)$ be agent i's valuation of the goods in a_i. Instead of an allocation function x, let $f : V_1 \times \ldots \times V_n \rightarrow A$ be the social choice function. We say the social welfare of an outcome $a \in A$, is $\sum_i v_i(a)$.

A Vickrey-Clarke-Groves (VCG) mechanism is defined as one where:

- $f(\mathbf{v}) = \arg\max_{a \in A} \sum\limits_j v_j(a)$, i.e. social welfare is maximized. And
- $p_i = h_i(\mathbf{v}_{-i}) - \sum\limits_{j \neq i} v_j(f(v_i, \mathbf{v}_{-i}))$, where $h_i(\mathbf{v}_{-i})$ is a function which does not depend on v_i. We call h_i a **pivot rule**.

Payments are *independent* from the agent's valuation. Therefore they have no motivation to lie and the general VCG mechanism is incentive compatible. The most common choice of h_i is the Clarke pivot rule defined as follows:

$$h_i(\mathbf{v}_{-i}) = \max_{a \in A} \sum_{j \neq i} v_j(a) \tag{2.1}$$

The Clarke rule has the following two properties; payments are non-negative, and if valuations are non-negative the mechanism guarantees individual rationality.

Chapter 3

Negative Valuations

Two assumptions that are typical of valuations in DSIC auctions are: 1. **monotonicity**, so that increasing a bid never decreases an allocation, and 2. **normalization**, meaning the value of the empty allocation is fixed at zero. Normalization implies that that **surplus**—the product of an agent's value and their allocation—is never negative. we retain the monotonicity assumption, but dispense with normalization. We can ask, what happens to our established mechanisms when we remove the non-negativity assumption on valuations? We now address the theoretical issues that arise when of allocating goods to agents with negative valuations. First, we reproduce the proof of Myerson's Lemma to confirm the DSIC property still holds. Then, working through a small example, we show how to find an offset term that guarantees IR, while at the same time maximizing revenue.

Turning our attention to the multi-parameter setting, we show that Clarke's pivot rule is not individually rational under negative valuations. We instead introduce a slightly modified pivot rule, which ensures the DSIC property in the presence of both negative and positive externalities. To ensure IR as well as DSIC, we propose a similar trick as in the single-parameter setting, namely computing offset terms. While theoretically sufficient, this solution is not yet practical for general, multi-parameter auction settings.

3.1 Negative Allocations and Valuations Under Myerson

We now present a proof of Myerson's lemma to confirm it has no dependence on v_i or x_i being non-negative. Restating the lemma:

Lemma 2.1.3 (Myerson [1981])**.** *An auction satisfies incentive compatibility and individual rationality if and only if the following two conditions hold* $\forall i \in N, \forall v_i \geq t_i \in T_i, \forall \mathbf{v}_{-i} \in T_{-i}$*:*

1. *The allocation rule is monotone:*

$$x_i(v_i, \mathbf{v}_{-i}) \geq x_i(t_i, \mathbf{v}_{-i})$$

2. *Payments are computed as follows:*

$$p_i(v_i, \mathbf{v}_{-i}) = p_i(\underline{v}_i, \mathbf{v}_{-i}) + v_i x_i(v_i, \mathbf{v}_{-i}) - \int_0^{v_i} x_i(z, \mathbf{v}_{-i}) \mathrm{d}z$$

We start with the proof of monotonicity. Consider the definition of incentive compatibility: $\forall i \in N$, $\forall \mathbf{v}_{-i} \in T_{-i}$, and for any two types $v_i, t_i \in T$

$$v_i x_i(v_i, \mathbf{v}_{-i}) - p_i(v_i, \mathbf{v}_{-i}) \geq v_i x_i(t_i, \mathbf{v}_{-i}) - p_i(t_i, \mathbf{v}_{-i})$$
$$t_i x_i(t_i, \mathbf{v}_{-i}) - p_i(t_i, \mathbf{v}_{-i}) \geq t_i x_i(v_i, \mathbf{v}_{-i}) - p_i(v_i, \mathbf{v}_{-i})$$

Rearranging the terms to isolate payments and combine the inequalities gives:

$$v_i(x_i(v_i, \mathbf{v}_{-i}) - x_i(t_i, \mathbf{v}_{-i})) \geq p_i(v_i, \mathbf{v}_{-i}) - p_i(t_i, \mathbf{v}_{-i}) \geq t_i(x_i(v_i, \mathbf{v}_{-i}) - x_i(t_i, \mathbf{v}_{-i}))$$

Without loss of generality assume $v_i \geq t_i$. For the inequality to hold $x_i(v_i, \mathbf{v}_{-i})$ must be at least as large as $x_i(t_i, \mathbf{v}_{-i})$, therefore x_i must be monotone. There is no dependence on x_i or v_i being non-negative and the proof holds. Myerson's payment rule is derived from the alternate definition of incentive compatibility:

$$\frac{\mathrm{d}p_i(z, \mathbf{v}_{-i})}{\mathrm{d}z} = z\frac{\mathrm{d}x_i(z, \mathbf{v}_{-i})}{\mathrm{d}z}$$

Re-express this using integration by parts:

$$\int_{\underline{v}_i}^{v_i} \left[\frac{\mathrm{d}p_i(z, \mathbf{v}_{-i})}{\mathrm{d}z}\right] \mathrm{d}z = p_i(v_i, \mathbf{v}_{-i}) - p_i(\underline{v}_i, \mathbf{v}_{-i})$$

$$= \int_{\underline{v}_i}^{v_i} z \frac{\mathrm{d}x_i(z, \mathbf{v}_{-i})}{\mathrm{d}z} \,\mathrm{d}z$$

$$= z x_i(z, \mathbf{v}_{-i})|_{\underline{v}_i}^{v_i} - \int_{\underline{v}_i}^{v_i} x_i(z, \mathbf{v}_{-i}) \,\mathrm{d}z$$

$$= v_i x_i(v_i, \mathbf{v}_{-i}) - \underline{v}_i x_i(\underline{v}_i, \mathbf{v}_{-i}) - \int_{\underline{v}_i}^{v_i} x_i(z, \mathbf{v}_{-i}) \,\mathrm{d}z$$

Giving:

$$p_i(v_i, \mathbf{v}_{-i}) = p_i(\underline{v}_i, \mathbf{v}_{-i}) + v_i x_i(v_i, \mathbf{v}_{-i}) - \underline{v}_i x_i(\underline{v}_i, \mathbf{v}_{-i}) - \int_{\underline{v}_i}^{v_i} x_i(z, \mathbf{v}_{-i}) \,\mathrm{d}z$$

Comparing to Myerson's original payments, Lemma 2.1.3, we have the same $v_i x_i(v_i, \mathbf{v}_{-i})$ term. The integral term is identical except that we allow for an arbitrary minimum type $\underline{v}_i$. The other two terms normally cancel to zero because $\underline{v}_i = 0$, $x_i(0, \mathbf{v}_{-i}) = 0$ and individual rationality forces $p_i(0, \mathbf{v}_{-i}) = 0$. In particular, note the free term $p_i(\underline{v}_i, \mathbf{v}_{-i})$, which now can be chosen to ensure individual rationality. No step in this derivation depended on the non-negativity of x_i or v_i while Myerson's original lemma still holds. There are many offset terms which guarantee IR, however there is only one which also maximizes revenue. Formally,

Problem 1. *Assuming $v_i \in V_i = [\underline{v}_i, \overline{v}_i]$, find $p_i(\underline{v}_i, \mathbf{v}_{-i})$ such that:*

$$\min$$

Although we can relax the negativity constraints on both v_i and x_i, from this point forward we'll only consider problems where $\underline{v}_i = 0$. We let the product of the allocation function and an agent's type to represent their valuation of the auction result. Then with a negative allocation function we can model negative valuations. We believe setting $\underline{v}_i = 0$ is a reasonable assumption as a bidder should only enter an auction where they want *some* good.

3.1.1 Myerson Linear Example

Assume $v_i \in [0, 1]$ for all $i \in N$, and consider the following monotone allocation function:

$$x_i(v_i, \mathbf{v}_{-i}) = 2v_i - 1$$

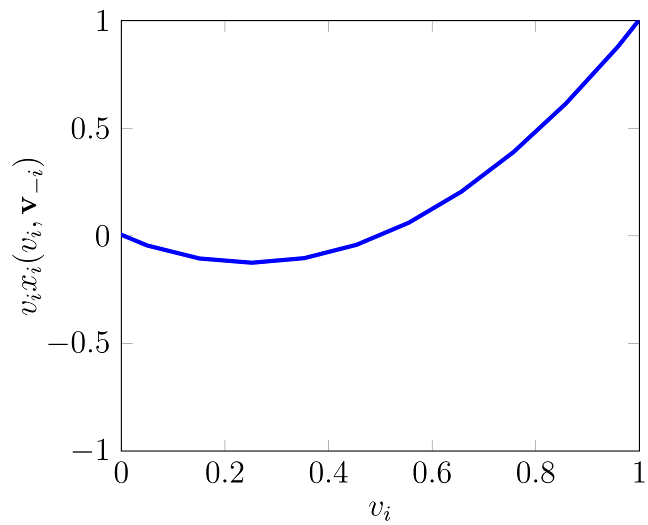

Figure 3.1: $v_i x_i(v_i, \mathbf{v}_{-i})$

Although this allocation function is monotone, it is not always non-negative. In particular, the allocation at the lowest type is -1. We now apply Myerson's lemma, (Lemma 2.1.3), to get a DSIC and IR payment formula. In particular, we show there are instances where an auctioneer has to pay agents to take goods.

Recall that bidder i's utility is given by $u_i(v_i, \mathbf{v}_{-i}) = v_i x_i(v_i, \mathbf{v}_{-i}) - p_i(v_i, \mathbf{v}_{-i})$. En route to plotting this utility, we first plot the first term, the product of bidder i's value and allocation. Observe in Figure 3.1 that this product is sometimes negative.

Applying Myerson's formula, assuming the lowest type is not allocated, and pays zero for said allocation (Lemma 2.1.3) yields:

$$\begin{aligned} p_i(v_i, \mathbf{v}_{-i}) &= v_i x_i(v_i, \mathbf{v}_{-i}) - \int_0^{v_i} x_i(z, \mathbf{v}_{-i}) \, \mathrm{d}z \\ &= v_i \left(2v_i - 1\right) - \int_0^{v_i} 2z - 1 \, \mathrm{d}z \\ &= 2v_i^2 - v_i - \left(z^2 - z \Big|_0^{v_i} \right) \\ &= v_i^2 \end{aligned}$$

We can now compute and plot the utility function (Figure 3.2). Note the

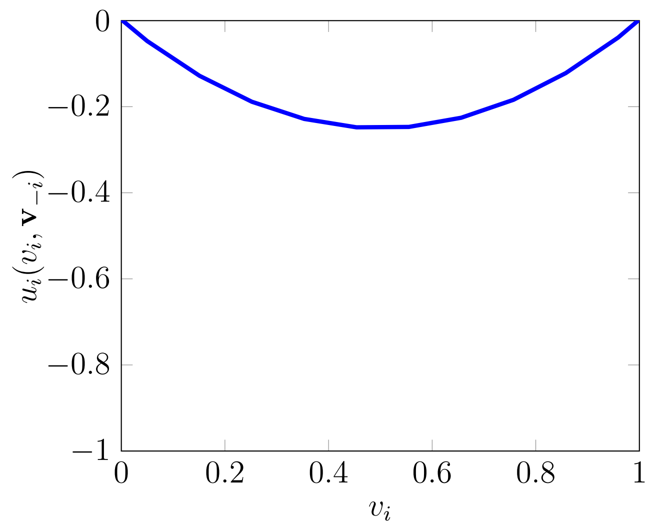

Figure 3.2: Utility Function

negative utilities:

$$\begin{aligned} u_i(v_i, \mathbf{v}_{-i}) &= v_i x_i(v_i, \mathbf{v}_{-i}) - p_i(v_i, \mathbf{v}_{-i}) \\ &= v_i(2v_i - 1) - v_i^2 \\ &= v_i^2 - v_i \end{aligned}$$

The utility function is negative at points in it's domain, and therefore does not satisfy individual rationality. At the beginning of this example we applied Myerson's payment formula assuming $p_i(\underline{v}_i, \mathbf{v}_{-i}) = 0$. However, for $v_i \in [0, .5]$, agents receive negative allocation, and deserve to be paid by the auctioneer for taking a good they negatively value. To satisfy individual rationality we need non-negative utility. Therefore we set $p_i(\underline{v}_i, \mathbf{v}_{-i})$ to be agent i's minimal utility when $p_i(\underline{v}_i, \mathbf{v}_{-i}) = 0$. In this example, utility is minimized at $v_i = \frac{1}{2}$, where $u_i(\frac{1}{2}, \mathbf{v}_{-i}) = -\frac{1}{4}$. Since $\underline{v}_i = 0$, we set $p_i(0, \mathbf{v}_{-i}) = -\frac{1}{4}$, which yields an individually rational payment rule:

$$\begin{aligned} p_i(v_i, \mathbf{v}_{-i}) &= v_i x_i(v_i, \mathbf{v}_{-i}) - \int_0^{v_i} x_i(z, \mathbf{v}_{-i})\, \mathrm{d}z + p_i(0, \mathbf{v}_{-i}) \\ &= v_i^2 - \frac{1}{4} \end{aligned}$$

We plot the adjusted payment rule and it's updated utility function in

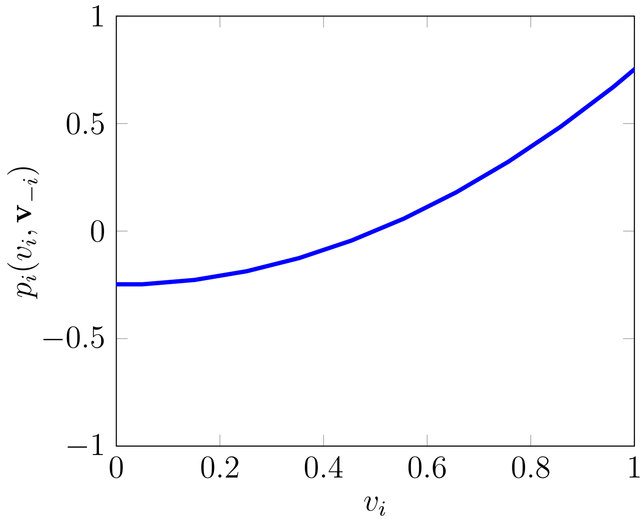

Figure 3.3: Shifted Payment Function

Figures 3.3 and 3.4 respectively. Note negative payments and non-negative utility.

Lastly, we confirm the optimal bidding strategy is to bid truthfully. We observe the derivative of agent i's utility function for bidding z given a true valuation of v_i and then set it to zero giving:

$$\begin{aligned}\frac{\mathrm{d}u_i(z, \mathbf{v}_{-i}; v_i)}{\mathrm{d}z} &= \frac{\mathrm{d}\left(v_i x_i(z, \mathbf{v}_{-i}) - p_i(z, \mathbf{v}_{-i})\right)}{\mathrm{d}z} \\ &= \frac{\mathrm{d}\left(v_i(2z-1) - z^2 + 1/4\right)}{\mathrm{d}z} \\ 0 &= 2v_i - 2z \\ z &= v_i\end{aligned}$$

Therefore agent i's utility is maximized by bidding $z = v_i$. There is no instance where agent utility is improved by lying about one's valuation. Therefore the payment rule is DSIC.

3.1.2 Visualizing Myerson Payments

Under standard assumptions (non-negative types and allocations), Myerson payments can be interpreted as the area to the left of the allocation function

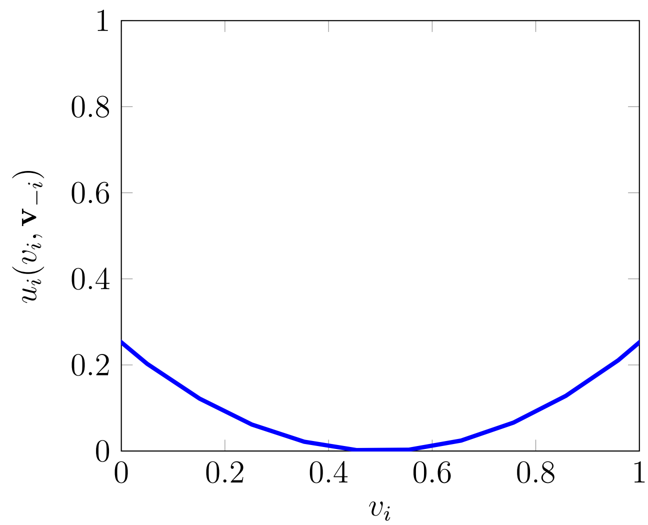

Figure 3.4: Shifted Utility Function

between $\underline{v}_i$ and v_i. We show even with negative allocation this remains true. We begin by drawing a box $v_i x_i(v_i, \mathbf{v}_{-i})$, as in Figure 3.5.

Now subtract the integral of the allocation curve, depicted in Figure 3.6.

Also note, if we observe an agents utility with Myerson payments and no offset applied, we get the same integral term. Therefore we can see in Figure 3.6 that utility will be negative if no offset is applied.

$$\begin{aligned} u_i(v_i, \mathbf{v}_{-i}) &= v_i x_i(v_i, \mathbf{v}_{-i}) - p_i(v_i, \mathbf{v}_{-i}) \\ &= v_i x_i(v_i, \mathbf{v}_{-i}) - (v_i x_i(v_i, \mathbf{v}_{-i}) - \int_0^{v_i} x_i(z, \mathbf{v}_{-i})\,\mathrm{d}z) \\ &= \int_0^{v_i} x_i(z, \mathbf{v}_{-i})\,\mathrm{d}z \end{aligned}$$

The remaining area, as in Figure 3.7, is the payment bidder i makes. Note the area is always strictly positive, even under negative allocation.

3.2 Negative Valuations Under VCG

We now repeat a similar exercise to understand negative valuations in the multi-parameter setting. Remember, a Vickrey-Clarke-Groves (VCG) mech-

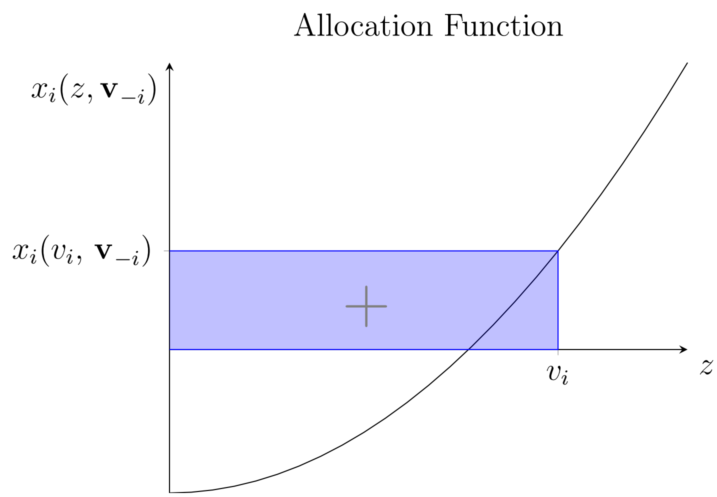

Figure 3.5: Area $v_i x_i(v_i, \mathbf{v}_{-i})$

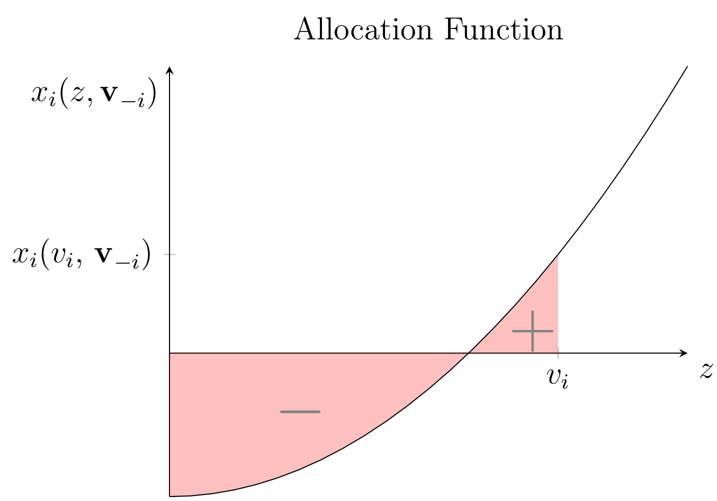

Figure 3.6: Area under the allocation curve, $\int_0^{v_i} x_i(z, \mathbf{v}_{-i})\,\mathrm{d}z$. Equivalent to non-offset Myerson utility.

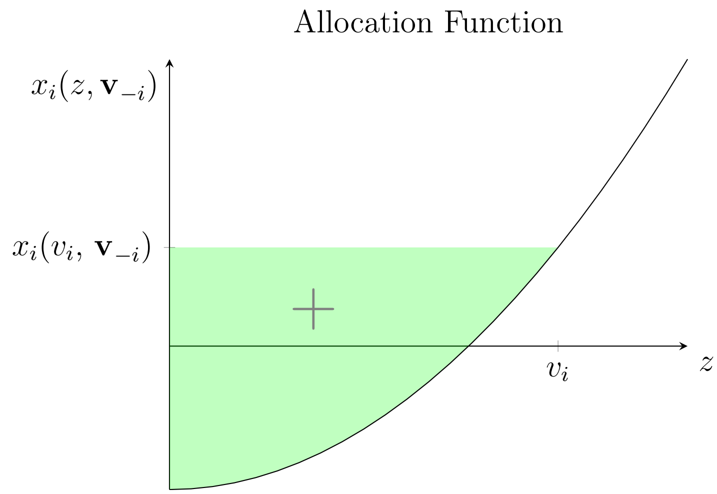

Figure 3.7: Payment, $v_i x_i(v_i, \mathbf{v}_{-i}) - \int_0^{v_i} x_i(z, \mathbf{v}_{-i})\, \mathrm{d}z$.

anism is defined as one where:

- $f(\mathbf{v}) = \arg\max_{a \in A} \sum_j v_j(a)$, i.e. social welfare is maximized. And
- $p_i = h_i(\mathbf{v}_{-i}) - \sum_{j \neq i} v_j(f(v_i, \mathbf{v}_{-i}))$, where $h_i(\mathbf{v}_{-i})$ is a function which does not depend on v_i. We call h_i a **pivot rule**.

The general VCG mechanism is incentive-compatible. The most common choice of h_i is the Clarke pivot rule defined as follows:

$$h_i(\mathbf{v}_{-i}) = \max_{a \in A} \sum_{j \neq i} v_j(a) \tag{3.1}$$

The Clarke rule has the following two properties; payments are non-negative, and if valuations are non-negative the mechanism guarantees individual rationality. However, we consider negative valuations and want our mechanism to pay bidders under certain situations. VCG payments with the Clarke pivot rule are often interpreted as *maximum social welfare if i were absent - total social welfare of others when i is present*. But if we observe the pivot rule, it's actually maximizing the social welfare of the other agents over all possible $a \in A$. It still allows the option that agent i receives a good. Nisan [2007a] briefly mentions an adjusted Clarke pivot rule for negative valuations, where they consider simulate the auction as if agent i weren't participating (i.e.

not receiving a good). Let $A \setminus \{i\}$ be the set of allocations for the auction without agent i. Then the new Clarke pivot rule is,

$$h'_i(\mathbf{v}_{-i}) = \max_{a \in A \setminus \{i\}} \sum_{j \neq i} v_j(a) \tag{3.2}$$

Before proceeding we compare the two pivot rules and present the following Lemma:

Lemma 3.2.1. *For every bidder i, $h_i(\mathbf{v}_{-i}) \geq h'_i(\mathbf{v}_{-i})$*

Proof. Both pivot rules maximize the sum of agent's welfare excluding agent i. However, h_i maximizes over all possible allocations whereas h_i only maximizes over allocations where agent i is not allocated. Therefore h_i must be at least as large as h'_i. □

We can now ask what properties we get from this adjusted pivot rule. Observe agent i's utility:

$$\begin{aligned}
u_i(v_i, \mathbf{v}_{-i}) &= v_i(a) - p_i(v_i, \mathbf{v}_{-i}) \\
&= v_i(a) - \max_{a \in A \setminus \{i\}} \sum_{j \neq i} v_j(a) + \sum_{j \neq i} v_j(f(v_i, \mathbf{v}_{-i})) \\
&= \sum_{j} v_j(f(v_i, \mathbf{v}_{-i})) - \max_{a \in A \setminus \{i\}} \sum_{j \neq i} v_j(a) \\
&= \textit{Maximum social welfare with } i - \textit{Maximum social welfare without } i
\end{aligned}$$

Therefore the adjusted pivot rule guarantees individual rationality if and only if removing an agent from the auction doesn't increase total social welfare. For the majority of auctions this may be true. However, an auctioneer could add constraints where this might not be the case. Consider the following toy example:

3.2.1 Maximal Good Allocation

Imagine there are two goods a, b and two agents $1, 2$. Both agents have identical unit-demand valuations for each good:

$$v_i(a) = 1, v_i(b) = -1$$

The auctioneer enforces the added constraint that every participating agent must receive at least one good. Then the welfare maximizing allocation gives each agent one of the goods. Without loss of generality, let agent 1 receive a and agent 2 receive b. Using the generalized Clarke pivot rule, Equation (3.2), to compute payments gives:

$$p_1(a) = (v_2(\{a, b\})) - (v_2(b)) = 1 - (-1) = 2$$
$$p_2(b) = (v_1(\{a, b\})) - (v_1(a)) = 1 - 1 = 0$$

We see that with both agents total welfare is 0 and with one agent total welfare is 1. Now observe linear utilities:

$$u_1(a) = v_1(a) - p_1(a) = 1 - 2 = -1$$
$$u_2(b) = v_2(b) - p_1(b) = -1 - 0 = -1$$

Both agents have negative utility. If either agent left the auction, nobody is solely allocated b and the total social welfare would increase. Unfortunately, we've found an auction where the generalized Clarke pivot rule still gives negative utility. One might ask, if like Myerson's payment formula, we can add a constant term into our pivot rule which offsets the minimum possible utility in the auction. Let $\underline{u_i}(\mathbf{v}_{-i})$ represent the minimum utility agent i can achieve under the chosen pivot rule and given the other bids. Then the updated pivot rule would be:

$$h'_i(\mathbf{v}_{-i}) = \max_{a \in A \setminus \{i\}} \sum_{j \neq i} v_j(a) - \underline{u_i}(\mathbf{v}_{-i}) \tag{3.3}$$

However, there are two major roadblocks. First, if we allow unbounded negative valuations there is no minimal utility term. It too is unbounded. Assuming we enforce bounds on valuations, we still have to find the $\underline{u}(\mathbf{v}_{-i})$ term. Formally,

Problem 2. *Let agent valuations v_i be drawn from $V_1 \times V_2 \times \ldots \times V_n$ where V_i are bounded sets. Then find $\underline{u}(\mathbf{v}_{-i})$ such that:*

$$\min_{v \in V_1 \times \ldots \times V_n} \left(\max_{a \in A \setminus \{i\}} \sum_{j \neq i} v_j(a) - \sum_{j \neq i} v_j(f(v, \mathbf{v}_{-i})) \right) - \underline{u}(\mathbf{v}_{-i}) = 0$$

Note in the single-parameter setting this is equivalent to Problem 1. Now in multi-parameter settings it becomes increasingly difficult to explore the entire valuation space for the minimum utility allocation. In the next chapter, we provide a simplification of multi-good auctions to workaround this second issue.

Chapter 4

Reducing Multi-Good to Single Parameter

Myerson's Lemma limits our problem domain to those where valuations can be described by a single parameter. This is a natural constraint for single good auctions, as a bidder either receives the good or not. But it's hard to interpret negatively allocating a single good: is something taken from you? Instead, we broaden our problem domain. We consider multi-good auctions where agent's valuations can be reported with a single number. Agents and auctioneers agree to a multiplicative penalty function between goods. Goods with large penalties become associated with negative valuations.

In this chapter, we present a general technique for reducing multi-good auctions with special properties to a single parameter. We prove unadjusted Myerson payments are weakly greater than VCG with our modified Clarke pivot rule. We present a well-known theorem, that any two DSIC payment formulas only differ by an additive constant. Therefore, with the appropriate individual rationality guaranteeing shifts, the two mechanisms are identical. Lastly, we present an algorithm for computing Myerson payments, and the optimal offset value.

4.1 The Reduction

Now consider a subset of multi-good unit demand auctions in which a public penalty factor between goods is known. Therefore every agent can describe their complete valuations with a single parameter. For all the below examples, let G be the set of goods. Define the penalty function $d : G \times G \rightarrow [0, \infty)$. Each agent i publicly reports their preferred good g_i and their private valuation v_i. Because we're describing a multi-good auction with a single valued allocation function, we need a wrapper around the allocation function which matches goods to values. Let $V_i = [0, v_i]$ represent the space of agent i's valuations and $V = V_1 \times, \ldots, V_n$ the space of all agent valuations. Let $\mathcal{A} : V \rightarrow G$ be the allocation algorithm which determines which goods are given to each agent. To match standard notation, let $\mathcal{A}_i(v_i, \mathbf{v}_{-i}) = g_i'$ where g_i' is the good agent i is allocated under the given bids. Then the allocation function x_i is:

$$x_i(v_i, \mathbf{v}_{-i}) = 1 - d(g_i, \mathcal{A}_i(v_i, \mathbf{v}_{-i})) \tag{4.1}$$

We require two properties from the function; monotonicity, and if bidder i is given their preferred good then $x_i = 1$. Monotonicity implies if your valuation goes up you must be allocated a good with an equal to or lesser penalty term. Therefore if you fix any preferred good g_i, there must be a strict ordering of the remaining goods. And because this is a public penalty function, if two bidders share a preferred good, they must also have the same ordering over other goods. The second property implies $d(g_i, g_i) = 0$ for all g_i. Distance metrics satisfy these properties, which motivates applying this reduction to auctions where goods exist in physical spaces.

An important consequence of the allocation function is that bidder i's valuation for good g_i' is:

$$v_i(g_i') = v_i(1 - d(g_i, g_i')) \tag{4.2}$$

That is, the auctioneer can *infer* agents' valuations for non-preferred goods. Note, agents negatively value all goods with penalty terms larger than 1. We call problems which can be reduced in this manner problems with **inferred value assignments**.

Having reduced multi-good auctions to a single parameter we can now apply Myerson's Lemma (Lemma 2.1.3) to a larger problem space. Because the input domain is bounded and one dimensional, we can easily search the total space of agent utilities and apply our adjusted Clarke pivot rule (Equation 3.3). In the upcoming sections, we present proofs comparing the shifted and unshifted payments, multiple application domains, followed by a small example emphasizing the necessity of offsetting, regardless of the mechanism.

4.2 Bounding Payments

The Vickrey-Clarke-Groves Mechanism is normally seen as a generalization of the Myerson payment formula. With non-negative allocations/valuations both formula's produce Vickrey's second price payment formula. In our new setting with negative allocations and multiple distinct goods we don't have the same guarantee. Below we'll prove bounds on the varying payment rules.

We state a theorem first attributed to Myerson [1981] and directly quoted from Nisan [2007b].

Theorem 4.2.1. *Assume the domains of types V_i are connected sets in Euclidean space. Let $(f, p_1, p_2, \ldots, p_n)$ be an incentive compatible mechanism. The mechanism with modified payments is incentive compatible if and only if for some functions $h_i : V_{-i} \to R$ we have that $p'_i(v_1, \ldots, v_n) = p_i(v_1, \ldots, v_n) + h_i(\mathbf{v}_{-i})$ for all $v_1, \ldots, v_n$.*

Every function h_i depends solely on $\mathbf{v}_{-i}$. It is completely independent of bidder i's valuation v_i. Therefore if we observe a single bidder in two distinct IC mechanisms their payments will differ by an additive factor, regardless of the bidders type v_i. Note this theorem **does not** claim that the additive factors are the same for all bidders.

Every mechanism we've presented thus far has been DSIC and therefore, by Theorem 4.2.1 their payments only differ by an additive constant. On top of that I claim the three payments have a strict ordering.

Lemma 4.2.2. *Assume agent's valuations for goods are determined according to Equation (4.2). I.e. we're in a single parameter environment. Let P_M represent Myerson payments with $p_i(0, \mathbf{v}_{-i}) = 0$, P_{VCG} be VCG payments with the original Clarke pivot rule (Equation (3.1)), and P'_{VCG} be VCG payments with the Clarke pivot rule for negative valuations (Equation (3.2)). Then,*

$$P_M = P_{VCG} \geq P'_{VCG}$$

Proof. By Theorem 4.2.1 it suffices prove the Lemma for agents with type $v_i = 0$. By assumption Myerson always charges bidders with type zero, zero. I.e. $p_i(0, \mathbf{v}_{-i}) = 0$. Let Recall agent i's VCG payment with the original Clarke pivot rule is:

$$p_i(v_i, \mathbf{v}_{-i}) = \max_y \sum_{j \neq i} v_j y_j(v_j, \mathbf{v}_{-j}) - \sum_{j \neq i} v_j(v_i x_i(v_i, \mathbf{v}_{-i}))$$

Similarly, if we set bidder i's type to 0 we get:

$$p_i(0, \mathbf{v}_{-i}) = \max_y \sum_{j \neq i} v_j y_j(v_j, \mathbf{v}_{-j}) - \sum_{j \neq i} v_j(v_i x_i(v_i, \mathbf{v}_{-i}))$$

$v_i = 0$, regardless of agent i's allocation they don't contribute to total welfare. Therefore the second term is equivalent to the total welfare of the allocation (i.e. maximum welfare). The first term is the maximum possible welfare when we don't consider agent i. Agent i has type 0 and therefore is never considered. Therefore the first term also reduces to maximum total welfare. Therefore VCG with the original Clarke pivot rule sets $p_i(0, \mathbf{v}_{-i}) = 0$ and $P_M = P_{VCG}$.

Lastly, we reference Lemma 3.2.1 which gives $P_{VCG} \geq P'_{VCG}$ and the Lemma is complete. □

Therefore prior to shifting, Myerson payments are weakly greater than VCG. But when we adjust both payments to be the maximal IR (i.e. $\min_t p_i(t, \mathbf{v}_{-i}) = 0$, for all i), we claim they are the same payment formula:

Lemma 4.2.3. *Given two DSIC mechanisms,* $(f, p_1, p_2, \ldots, p_n)$ *and* $(f', p'_1, p'_2, \ldots, p'_n)$*. If* $\min_t p_i(t, \mathbf{v}_{-i}) = \min_t p'_i(t, \mathbf{v}_{-i})$*, for all* i*, then* $(f, p_1, p_2, \ldots, p_n) = (f', p'_1, p'_2, \ldots, p'_n)$*.*

Proof. By assumption both mechanisms are DSIC, therefore by Theorem 4.2.1 the payments only differ by an additive constant. Any two functions that differ by an additive constant always achieve their minimum at the same domain value t. Let $t_i = \arg\min p_i(t, \mathbf{v}_{-i}) = \arg\min p'_i(t, \mathbf{v}_{-i})$. By assumption, the two payment's minimum values are the same, therefore $p_i(t, \mathbf{v}_{-i}) - p'_i(t, \mathbf{v}_{-i}) = 0$. So the payments are identical and the Lemma holds. □

We shift both Myerson and VCG payments such that their minimum payment is always 0. By Lemma 4.2 this results in identical payments.

4.3 Computing Offset Myerson Payments

These multi-good allocation rules are more complex than canonical single-parameter ones. Without a closed form of the allocation rule, Myerson payments are difficult to compute. Remember when we generalized Myerson's payment formula to the multi-good setting, we defined our allocation function as follows (Equation (4.1)):

$$x_i(v_i, \mathbf{v}_{-i}) = 1 - d(g_i, \mathcal{A}_i(v_i, \mathbf{v}_{-i}))$$

We're allocating discrete goods monotonically, therefore every allocation function is a step function. Learning them reduces to finding the discontinuities in the corresponding step function (including the boundaries of the domain). This can be solved by performing a binary search over types. We can bound the number of discontinuities by the number of possible displacements an agent can force. Clearly an agent can only displace each agent/good exactly once. Therefore the number of discontinuities is bounded by the minimum of agents and goods. Call the function which finds these values `FindJumps`, an implementation can be found at Appendix A.3.

Given a sorted list of bidder types where the allocation function changes, our goal to find payments which adhere to Myerson's payment rule and preserve individual rationality. The psuedocode for this appears in Algorithm 1. We must guarantee bidders have non-negative utility across all possible types. To do this, we must find the minimum utility achieved under non-offset Myerson payments and shift the payment function accordingly (Problem 1). By Lemma 2.1.3 to preserve incentive compatibility, the $p_i(\underline{v}_i, \mathbf{v}_{-i})$ term must be independent of a bidders private reported type v_i. Therefore, when finding the minimal utility we must explore the *entire* type domain.

In the negative allocation setting the minimal utility occurs at the the last type t_i where the allocation function is negative. More formally

Lemma 4.3.1. *In the multi-good setting, utility $u_i(v_i, \mathbf{v}_{-i})$, is minimized by the type $v_i = t_i$ such that $x(t_i) < 0 \leq x(t_i + \epsilon)$ for all $\epsilon > 0$*

Proof. Observe bidder i's utility for bidding type t_i:

$$u_i(t_i, \mathbf{v}_{-i}) = t_i x_i(t, \mathbf{v}_{-i}) - p_i(t_i, \mathbf{v}_{-i})$$

Adding our generalized Myerson payments gives:

$$u_i(t_i, \mathbf{v}_{-i}) = t_i x_i(t_i, \mathbf{v}_{-i}) - (p_i(0, \mathbf{v}_{-i}) + t_i x_i(t_i, \mathbf{v}_{-i}) - \int_0^a x_i(z, \mathbf{v}_{-i}) \, \mathrm{d}z)$$

Canceling terms simplifies to:

$$u_i(t_i, \mathbf{v}_{-i}) = -p_i(0, \mathbf{v}_{-i}) + \int_0^{t_i} x_i(z, \mathbf{v}_{-i}) \, \mathrm{d}z$$

The first term is constant across all t for bidder i. Therefore utility is minimized by minimizing the integral. This is equivalent to taking the integral over the domain where x_i is negative. □

Because these are step functions, this t_i must occur at one of the discontinuities in x_i. With the sorted list of discontinuities and associated allocation values we can compute the integral of the allocation function, and therefore the non-offset utilities (lines 5-6 in Algorithm 1). By Lemma 4.3.1, the minimum of these values is guaranteed to be $\int_0^{t_i} x_i(z, \mathbf{v}_{-i}), \mathrm{d}z$. Lastly we compute the standard payment and shift according to the minimum utility value (lines 7-10 in Algorithm 1). When computing minimal utility we only

Algorithm 1 Get Shifted Payments

1: **procedure** GetShiftedPayments($x, v, \bar{v}$)
2:     **for** Agent i **do**
3:         *alloc, vals* $=$ `FindJumps`$(x_i(v, v_{-i}), 0, \bar{v})$
4:         $\mathit{util}(0) = 0$
5:         **for** j `in 1,2,...,length`(*alloc*) **do**
6:             $\mathit{util}(j) = \mathit{util}(j-1) + (\mathit{vals}(j) - \mathit{vals}(j-1)) \cdot \mathit{alloc}(j-1)$
7:         $j =$ Index of first jump that happens after v_i
8:         $\mathit{int} = \mathit{util}(j) - \mathit{alloc}(j-1) \cdot (\mathit{vals}(j) - v_i)$
9:         $\mathit{welf} = x_i(v_i, v_{-i}) \cdot v_i$
10:         $\mathit{pay}(i) = \mathit{welf} - \mathit{int} + \mathtt{min}(\mathit{util})$
11:     `Return` *pay*

adjust one bidders type. Allocations are monotone, and therefore an agent can only displace another agent once. This bounds the number of discontinuities in the allocation by $O(n)$. The discontinuities can be considered as the leaves of a binary search tree. Let A represent the time it takes to compute an allocation. Then finding all discontinuities takes $O(nA)$ time. Therefore finding all shifted payments according to this algorithm takes $O(n^2A)$ time.

Chapter 5

Real World Examples

Below we present two separate real world applications for our reduction. The first is the carpool assignment problem. There, the auctioneer's goal is maximize welfare while assigning riders to drivers. Every driver has a fixed pickup location and the penalty function is determined by the distances between driver locations. The second is the problem of allocating routes for a bike-share system. Every station has some number of bikes and empty spaces. Agents report a preferred source sink pair, and again penalties are determined according the distance from their ideal route.

We work through examples in both domains and show how some agents are negatively allocated in the welfare-maximizing solutions. We also compute unadjusted payments and utilities under each mechanism. Earlier we proved how under standard non-negativity assumptions, the two unshifted payments formulas are identical. In our new multi-good domain they often differ. We see different cases where each mechanism finds the optimal payments, and others where shifting is necessary for individual rationality.

5.1 Carpool

Imagine a company wanted to organize carpools for it's employees. Drivers have volunteered, and now the company wants to assign employees to cars such that welfare is maximized.

This problem is known as fixed maximum carpool matching. It was first explored by Hartman et al. [2014]. Let $P = \{p_1, p_2, \ldots, p_n\}$ be the set of passengers and $S = \{s_1, s_2, \ldots, s_m\}$ the set of drivers. Represent them as vertices of the graph $G = (P \cup S, A)$. An arc (p_i, d_j) exists if passenger p_i is willing to ride with driver d_j. There's a capacity function $c : D \rightarrow N$ which indicates how many passengers each driver can drive. There is a weight function $w : A \rightarrow R$ which defines the riders valuation $v(p_i, s_j)$, that passenger i gets for riding with driver j. The goal is to assign members of P to D, such that capacity constraints hold and passengers welfare is maximized. Note we require that the maximum number of passengers ride in cars. Everyone needs to go to work.

This problem is more general than our framework. It allows for arbitrary valuation functions. We specifically consider the problem where $A = P \times S$ and valuations are defined by the multiplicative allocation function using a distance metric over pick-up locations as the penalty function. Each agent i has some *public* preferred carpool d_i and some *private* valuation v_i for receiving good s_i. Let D be the maximum distance some agent is willing to defer from their ideal carpool before they have negative valuation. Then agent i's penalty for good s'_i is:

$$d(s_i, s'_i) = \frac{d_G(s_i, s'_i)}{D} \tag{5.1}$$

The penalty function is public information, therefore the auctioneer can compute each agents full space of valuations. The problem of finding the maximum assignment of riders to drivers is in P, by Hartman et al. [2014]. Therefore, we can maximize welfare for all instances. In addition, we proved the following theorem:

Theorem 5.1.1. *The constraint matrix for the carpool assignment problem is totally unimodular, and therefore the linear program has an integral solution.*

The proof can be found in Appendix A.1. As a result of the theorem, we're guaranteed the linear program is integral [Heller and Tompkins, 1956] and implementing an algorithm is simplified.

5.1.1 Example

We present a small example where unshifted Myerson and VCG payments differ. Note for different agents, each mechanism finds a better payment rule (either IR or higher revenue) prior to shifting. There are four agents and three cars, a, b and c. Each car only has one empty seat. Bidders report their type

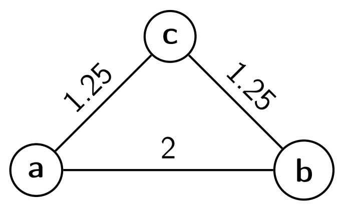

Figure 5.1: Distances between goods a, b and c

v_i and their preferred car g_i. Remember there's forced participation, i.e. the maximum number of seats must be filled. Agents report their bids as follows:

$$\{(v_1, a), (v_2, a), (v_3, a), (v_3, b)\}$$

We choose publicly agreed upon penalty function based off distances shown in Figure 5.1. The allocation function follows from Equation 4.1 and the carpool penalty function (Equation 5.1). It follows that,

$$x_1(b) = x_2(b) = x_3(b) = -1$$

and

$$x_1(c) = x_2(c) = x_3(c) = -.25$$

Without loss of generality, assume $v_i \geq v_2 \geq v_3$. Then the welfare maximizing allocation gives agent one good a, agent three good c and agent four good b.

Agent One - Possibly Negative

First we compute agent one's Myerson payments where the offset term is zero. Observing the allocation function:

$$x_1(t, \mathbf{v}_{-1}) = \begin{cases} -.25 & t < v_3 \\ 0 & v_3 \leq t < v_2 \\ 1 & \text{otherwise} \end{cases}$$

Then according to Myerson's payment formula (Lemma 2.1.3)

$$\begin{aligned} p_1(v_1, \mathbf{v}_{-1}) &= v_1 - \int_0^{v_1} x_1(z, \mathbf{v}_{-1}) dz \\ &= v_1 - ((v_1 - v_2) - .25 v_3) \end{aligned}$$

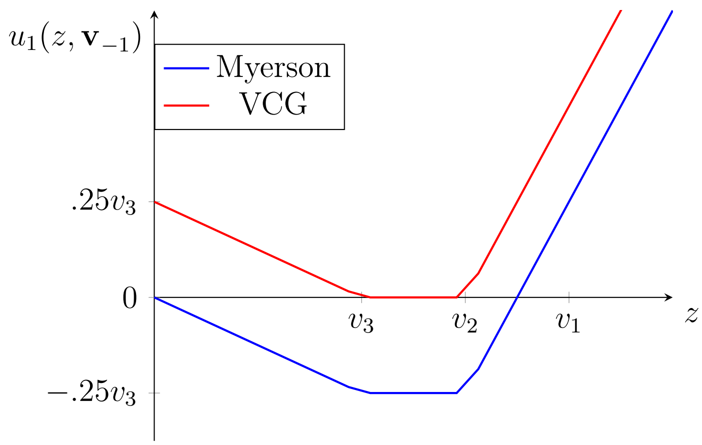

Figure 5.2: Agent one's utilities before adjusting for individual rationality.

$$= v_2 + .25v_3$$

The allocation and payment are constant when $v_1 \geq v_2$. We repeat the computation with $v_3 \leq v_1 < v_2$ and find agent one pays $.25v_3$ to receive nothing. Lastly, when $v_1 < v_3$, they pay nothing to get good c. We plot the agents utility in Figure 5.2 along with the utility from the VCG mechanism.

Before computing VCG payments, remember that $f : V_1 \times \ldots \times V_n \to A$ is the social welfare maximizing allocation function. It takes a vector of bids and returns an allocation $a \in A$. For agent i, a_i represents agent i's allocation and $v_i(a)$ their valuation of the goods in a_i. Then VCG payments with the Clarke pivot rule for negative valuations (Equation 3.2) is:

$$\begin{aligned} p_1(v_1, \mathbf{v}_{-1}) &= \max_{a \in A \setminus \{1\}} \sum_{j \neq 1} v_j(a) - \sum_{j \neq 1} v_j(f(v_1, \mathbf{v}_{-1})) \\ &= (v_2 + v_4 - .25v_3) - (v_4 - .25v_3) \\ &= v_2 \end{aligned}$$

Exactly as Myerson, this payment is constant when $v_1 \geq v_2$. When $v_3 \leq v_1 < v_2$ shows agent one pays nothing to get nothing. Lastly, when $v_1 < v_3$ they *get paid* $.25v_3$ to take good c. The utility function under the VCG mechanism is also plotted in Figure 5.2. Note how the utility functions of each mechanism only differ by an additive constant. The Clarke Pivot Rule for negative valuations finds the optimal IR payments, but Myerson payments needs to be offset by $-.25v_3$.

Agent Two - No Allocation

We'll now repeat the exercise with agents two and three. They prefer the same good as agent one, and therefore have very similar allocation/utility functions. Agent two's allocation function:

$$x_2(t, \mathbf{v}_{-1}) = \begin{cases} -.25 & t < v_3 \\ 0 & v_3 \leq t < v_1 \\ 1 & \text{otherwise} \end{cases}$$

Then payments according to Lemma 2.1.3 follow:

$$\begin{aligned} p_2(v_2, \mathbf{v}_{-2}) &= v_2 - \int_0^{v_2} x_2(z, \mathbf{v}_{-1}) dz \\ &= v_2 - v_3 \end{aligned}$$

According to Myerson, agent two should pay to receive nothing. He should have to pay to avoid receiving a negative good. Observe in Figure 5.3 that agent two has negative utility under Myerson. Now computing VCG payments with the updated Clarke rule:

$$\begin{aligned} p_2(v_2, \mathbf{v}_{-2}) &= \max_{a \in A \setminus \{2\}} \sum_{j \neq 2} v_j(a) - \sum_{j \neq 2} v_j(f(v_2, \mathbf{v}_{-2})) \\ &= (v_1 - .25v_3 + v_4) - (v_1 - .25v_3 + v_4) \\ &= 0 \end{aligned}$$

Here VCG acknowledges that agent two has no impact on the auction. In particular, it understands nobody should pay to not receive a good. Observe Figure 5.3 to see agent two's utilities under both payment rules. Note how the shape is nearly identical to agent one. And it also needs to be shifted by $-.25v_3$.

Agent Three - Negative Allocation

Agent three is almost exactly like agent's one and two. The allocation function:

$$x_3(t, \mathbf{v}_{-3}) = \begin{cases} -.25 & t < v_2 \\ 0 & v_2 \leq t < v_1 \\ 1 & \text{otherwise} \end{cases}$$

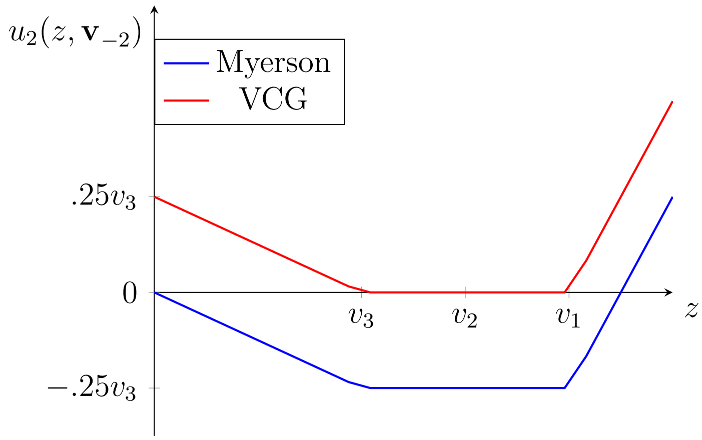

Figure 5.3: Agent two's utilities before adjusting for individual rationality.

Then payments according to Lemma 2.1.3 follow:

$$\begin{aligned} p_3(v_3, \mathbf{v}_{-3}) &= v_3 - \int_0^{v_3} x_3(z, \mathbf{v}_{-3})dz \\ &= v_3 - v_3 \\ &= 0 \end{aligned}$$

Without shifting, Myerson doesn't pay agent three to take a good they don't want. Note the negative utility in Figure 5.4 Now computing VCG payments with the updated Clarke rule:

$$\begin{aligned} p_3(v_3, \mathbf{v}_{-3}) &= \max_{a \in A \setminus \{3\}} \sum_{j \neq 3} v_j(a) - \sum_{j \neq 3} v_j(f(v_3, \mathbf{v}_{-3})) \\ &= (v_1 - .25v_2 + v_4) - (v_1 + v_4) \\ &= -.25v_2 \end{aligned}$$

VCG pays agent three because it understands their value of preventing agent two from entering the auction. Figure 5.4

Agent Four - Always Allocated

We'll now repeat the exercise with agent four. Note how they're the only agent who *wants* good b. We'll see VCG values this, whereas Myerson will

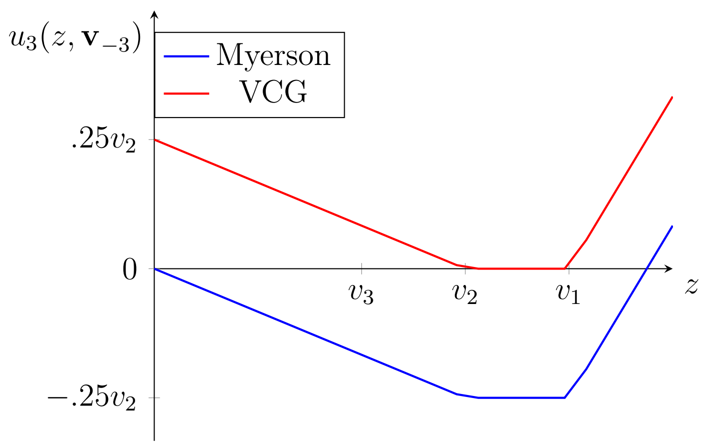

Figure 5.4: Agent three's utilities before adjusting for individual rationality.

implicitly abuse this fact. In computing Myerson payments we first note agent four is always allocated b, i.e.

$$x_4(z, \mathbf{v}_{-4}) = 1 \quad \forall z, 0 \leq z \leq v_4$$

Then payments according to Lemma 2.1.3 follow:

$$\begin{aligned} p_4(v_4, \mathbf{v}_{-1}) &= v_4 - \int_0^{v_4} x_1(z, \mathbf{v}_{-1}) dz \\ &= v_1 - v_4 \\ &= 0 \end{aligned}$$

No matter how much agent four wants b they always get it for free according to Myerson. This gives the linear utility function seen in Figure 5.5. When computing VCG payments with our updated Clarke pivot rule, we note if agent four were not participating in the auction, agents two and three would get goods c and b, respectively. Agent four's payment follows:

$$\begin{aligned} p_4(v_4, \mathbf{v}_{-1}) &= \max_{a \in A \setminus \{4\}} \sum_{j \neq 4} v_j(a) - \sum_{j \neq 4} v_j(f(v_4, \mathbf{v}_{-4})) \\ &= (v_1 - .25v_2 - v_3) - (v_1 - .25v_3) \\ &= -.25v_2 - .75v_3 \end{aligned}$$

Here VCG *pays* agent four because there are no adequate replacements and therefore has a high externality. But if we look at agent four's utility we can

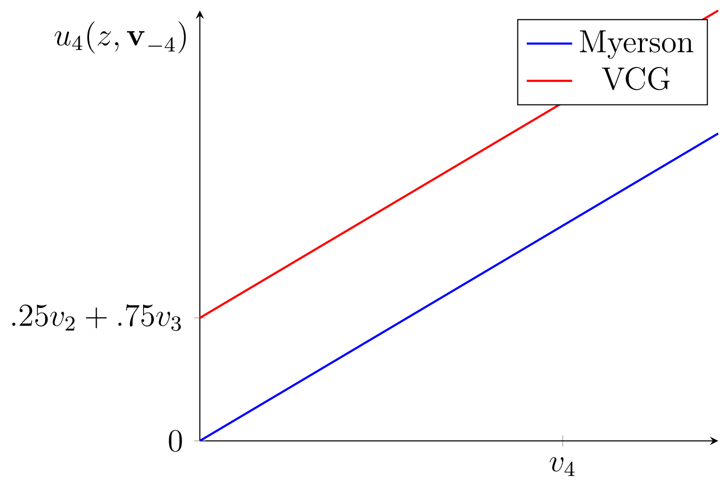

Figure 5.5: Agent four's utilities prior to adjusting for individual rationality.

see they don't care, they're just happy to get good b. Comparing agent four's utilities under Myerson and VCG payments (Figure 5.5), we see Myerson is already the optimal IR payment formula. But, adjusting VCG to Myerson would increase total revenue.

In one small example we've seen how both mechanisms only sometimes find the optimal payments. In both cases, finding revenue maximizing individually rational payments requires exploring the allocation space and finding minimal agent utilities. However, we can also note that differences occur when agents report unique preferences. VCG had the optimal payment rule for all three agents that wanted good a, whereas Myerson was more effective for agent four.

5.2 Bike-Sharing

As it currently stands, bike-sharing models allow users to pick up and drop off bikes from any station. However, this leads to system imbalance. Bike-sharing companies have to pay "rebalancers" to move bikes around the city [Grabar, 2017]. We consider an alternate model where biker riders announce a public preferred source and sink (s_i, t_i) and a bid v_i, then the auctioneer allocates routes to bidders.

Formally, there are $S = \{1, 2, \ldots, s\}$ stations, and a distance function $d : S \times S \mapsto [0, \infty]$ indicating the distance between pair of stations s and

t (the distance might come from an underlying graph, but we make no use of this graph at the moment). There are a fixed number of identical bikes $b_s \in N$, and a number of empty spaces $e_t \in N$ available at each station.

There are $N = \{1, 2, \ldots, n\}$ agents. In general, each agent $i \in A$ is endowed with a valuation function $v_i : S \times S \mapsto R$ that associates a real value with every pair of source and target stations.

The goal is to design a truthful mechanism that maximizes welfare. We call this problem the *bike-share problem*, and a solution the *bike-share mechanism*. In the mechanism, all agents report their preferences to the center, who allocates a (possibly null) source-target pair (s'_i, t'_i) to each agent i, and charges a cost c_i.

In our single-parameter setting, we assume each agent i is characterized by a *public* preferred source station s_i and a preferred target station t_i, as well as a single *private* value $v_i \in R^+$ associated with pair (s_i, t_i). Let D be the globally defined maximum distance allowed, then for allocation (s'_i, t'_i) define the penalty function as:

$$d(s'_i, t'_i) = \frac{d(s_i, s'_i) + d(t_i, t'_i)}{D} \tag{5.2}$$

This penalty function can be substituted into Equation (4.1) to complete the reduction. See Appendix A.2 for proof of NP-hardness, even under the single-parameter reduction.

5.2.1 Example

Below we present a bike-sharing instance where both mechanisms require shifting to find the revenue maximizing IR payments. Additionally, the welfare maximizing allocation *pays* an agent to take a good they value negatively. We have two bidders and three bike stations X, Y and Z. Bids take the form of their valuation v_i and their preferred source sink pair (s_i, t_i). The two agents report as follows:

$$\text{Agent One: } \{4, (X, Z)\}$$
$$\text{Agent Two: } \{2, (Y, Z)\}$$

Note the maximal agent type is 10, i.e. $\bar{v}_i = 10$. Bike availability, station capacity, and station distances are shown in Figure 5.6. Specifically, $b_X = b_Y = b_Z = 1$, $c_X = c_Y = 2$ and $c_Z = 1$. Agent valuations are defined by the

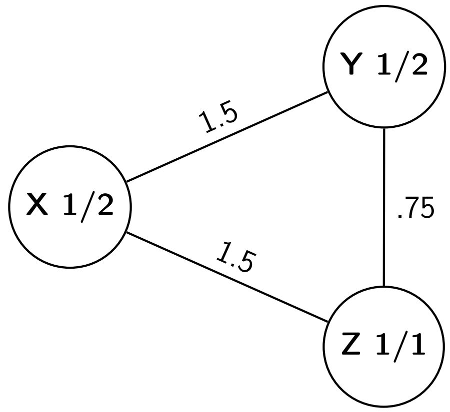

Figure 5.6: Distances between goods X, Y and Z

penalty function Equation (5.2) with $D = 1$. Applying Equation 4.1 agent one's allocation function for receiving the source sink pair (s, t) is:

$$x_1(v_1, \mathbf{v}_{-1}) = 1 - (d(s_i, X) + d(t_i, Z)) \tag{5.3}$$

Therefore when agent one receives (X, Y), $x_1 = 1-(0+.75) = .25$. Repeating for agent two, their allocation function for receiving (s, t) is:

$$x_2(v_2, \mathbf{v}_{-2}) = 1 - (d(s, Y) + d(t, Z)) \tag{5.4}$$

It follows that when agent two receives (Z, Y), $x_2 = 1 - (.75 + .75) = -.5$. The welfare maximizing allocation routes agent one from X to Z and agent two from Z to Y. Note how agent two opens a space at station Z allowing agent one to ride there. The total social welfare of this allocation is 3.75.

Agent One

First we compute agent one's Myerson payments where the offset term is zero. When agent one has type zero, they are allocated the route Z, Y for -1.25 times their type. With a very small type, agent one is now enlisted to move a bike for agent two. As their valuation increases they get allocated route X, Y because it's not worth having agent two move a bike from station Z. Lastly, once their valuation is high enough, we require agent two to move

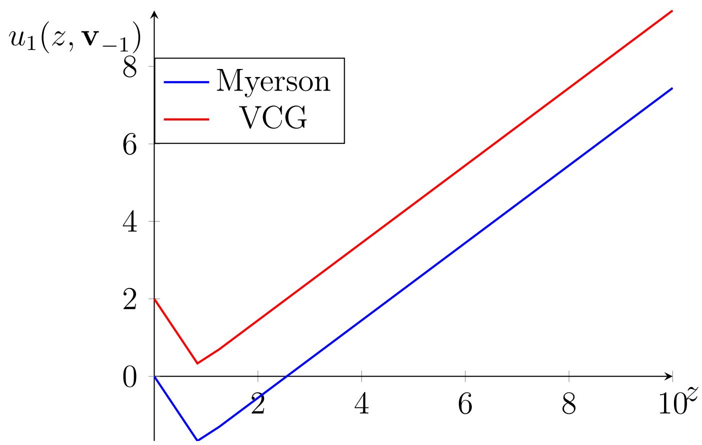

Figure 5.7: Agent one's utilities before shifting.

a bike for agent one. Therefore agent one's allocation function is as follows:

$$x_1(t, \mathbf{v}_{-1}) = \begin{cases} -2 & t < 8/9 \\ .25 & 8/9 \le t < 1 \\ 1 & \text{otherwise} \end{cases}$$

Then according to Myerson's payment formula (Lemma 2.1.3)

$$\begin{aligned} p_1(4, \mathbf{v}_{-1}) &= 4 - \int_0^4 x_1(z, \mathbf{v}_{-1}) dz \\ &= 4 - (-2(8/9) + .25(1/9) + (4 - 1)) \\ &= 2.75 \end{aligned}$$

We observe this payment is constant for all types larger than 1. With types between 8/9 and 1 they pay 1.778. And for types less than 8/9 they pay nothing. We plot agent one's utility in Figure 5.7 along with the utility from VCG computed below.

Before computing VCG payments, remember that $f : V_1 \times \ldots \times V_n \to A$ is the social welfare maximizing allocation function. It takes a vector of bids and returns an allocation $a \in A$. $v_j(a)$ represents agent j's valuation for the goods they receive in a. Then VCG payments with the Clarke pivot rule for negative valuations (Equation 3.2) is:

$$p_1(4, \mathbf{v}_{-1}) = \max_{a \in A \setminus \{1\}} \sum_{j \neq 1} v_j(a) - \sum_{j \neq 1} v_j(f(v_1, \mathbf{v}_{-1}))$$

$$
\begin{aligned}
&= (0) - (-.5 \cdot 2) \\
&= 1
\end{aligned}
$$

Exactly as Myerson, this payment is constant when $v_1 \geq 1$. When $8/9 \leq v_1 < 1$, they pay nothing. Lastly, when $v_1 < 8/9$ they're *paid* 2. Observing Figure 5.7, note how payment rules can be improved. Myerson payments need to be decreased by $-16/9$ to restore it's IR guarantee. VCG payments can actually be increased by $2/9$ while preserving IR. In this case, shifting VCG *increases* revenue.

Agent Two

We'll now repeat the exercise with agent two. For types $0 \leq v_2 < 6$ agent two is used to open up a space for agent one. For types $6 \leq v_2 < 9$ agent two is unallocated because it isn't worth using them to open the space. Then for types $v_2 \geq 9$ the equation flips and it's worth more to have agent one move a bike for agent two. This results in the following allocation function:

$$
x_2(t, \mathbf{v}_{-2}) = \begin{cases} -.5 & t < 6 \\ 0 & 6 \leq t < 9 \\ 1 & \text{otherwise} \end{cases}
$$

Then payments according to Lemma 2.1.3 follow:

$$
\begin{aligned}
p_2(2, \mathbf{v}_{-2}) &= -.5 \cdot 2 - \int_0^2 x_1(z, \mathbf{v}_{-1}) dz \\
&= -.5 \cdot 2 - (-.5 \cdot 2) \\
&= 0
\end{aligned}
$$

Unshifted Myerson is incapable of charging an agent for a negative good. On the other hand, when computing VCG payments with our updated Clarke pivot rule, we note if agent two weren't in the auction agent one would suffer.

$$
\begin{aligned}
p_2(2, \mathbf{v}_{-2}) &= \max_{a \in A \setminus \{2\}} \sum_{j \neq 2} v_j(a) - \sum_{j \neq 2} v_j(f(2, \mathbf{v}_{-2})) \\
&= .25 \cdot 4 - 4 \\
&= -3
\end{aligned}
$$

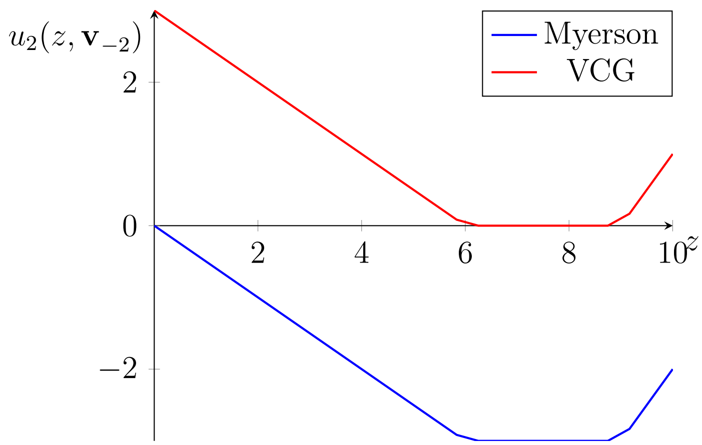

Figure 5.8: Agent two's utilities before shifting.

Here VCG *pays* agent two because there is no replacement and therefore has a high externality.

Observing Figure 5.8 we note for agent two VCG finds the revenue maximizing IR payment. On the other hand, Myerson payments need to be decreased by 3 to guarantee IR. In one small example we've seen how neither mechanisms are guaranteed to find the revenue maximizing individually rational payments. Throughout this paper we've shown neither payment rule is always IR without offsets. Additionally, we've seen exploring the allocation space and finding minimal agent utilities can increase revenue over unshifted VCG.

Chapter 6

Discussion

We confirmed that the Myerson and VCG mechanisms remain DSIC, even under negative valuations and allocations. Restricting our attention to negative allocations, we then showed how to set offset terms which ensure the two mechanisms are also IR. Then we presented an algorithm that searches for the revenue-maximizing offset in the single-parameter setting. We described two application domains, the carpool problem and bike-sharing. In both examples, agents have unit-demand valuations, so would be traditionally treated as multi-parameter problems. We showed how to reduce problems like these, in which goods are "comparable," to single-parameter settings. Finally, we present instances of these problems where the search for an offset is necessary, because negatively allocating agents benefits the greater good.

We make two crucial assumptions in this thesis, in solving inferred value assignment problems. The first is *public* agent preferences. We think this is a safe assumption in physical spaces where we have access to location data. The second assumption is that agents have identical penalty functions. Our model actually allows for multiple penalty functions, as long as bidder penalty types are *public*. Consider the bike-sharing problem. Some agents may be commuters while others are tourists. Commuters need to get to work and would be negatively value most other destinations. On the other hand, tourists might be more interested in just going for a bike ride, rather than ending up at a particular destination. They wouldn't penalize destinations nearly as much. This expansion would allow agents to declare more diverse preferences, while preserving our ability to use the machinery developed here.

Even in simple problem domains, finding optimal IR payments is expensive. It is necessary to explore allocations over the entire valuation domain.

This presents an interesting direction for research. When we fix all other agents' valuations, can we efficiently find one allocation function? Even more importantly, can we a find a utility-minimizing type? This research direction is not limited to the single-parameter setting, but added structure in valuations can only help. In our examples, we saw that when agents share preferences and therefore penalty functions, their allocation functions are very similar. In such cases, some of our computations were redundant. We expect improvements in learning agents' allocation functions to be domain specific, but interesting work nonetheless.

It is possible that in many instances, the Clarke pivot rule for negative valuations happens to guarantee individual rationality. In these cases, it is not necessary to look for offsets; but it still might be beneficial to do so, even with the added complexity, to increase revenue. That is, even when payments happen to be IR, searching for offsets can increase revenue.

As a subroutine, the VCG mechanism requires that we solve for optimal welfare-maximizing allocations. In NP-hard problem domains, such as bike-sharing, this might not be feasible. In problems with inferred value assignments, however, through our reduction, we can use Myerson instead of VCG. Then, if we happen to have greedy approximation algorithms, we can still design DSIC auctions with IR payments.

Bibliography

David J. Abraham, Avrim Blum, and Tuomas Sandholm. Clearing algorithms for barter exchange markets: Enabling nationwide kidney exchanges. In *Proceedings of the 8th ACM Conference on Electronic Commerce*, EC '07, pages 295–304, New York, NY, USA, 2007. ACM. ISBN 978-1-59593-653-0. doi: 10.1145/1250910.1250954. URL `http://doi.acm.org/10.1145/1250910.1250954`.

E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11(1):17–33, 1971.

B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1):242–259, 2007.

Henry Grabar. How new york's bike-share system pays riders to make it run better. *Slate*, 2017.

T. Groves. Incentives in teams. *Econometrica*, 41(4):617–631, 1973.

Irith Ben-Arroyo Hartman, Daniel Keren, Abed Abu Dbai, Elad Cohen, Luk Knapen, Davy Janssens, et al. Theory and practice in large carpooling problems. *Procedia Computer Science*, 32:339–347, 2014.

I. Heller and C. B. Gh Tompkins. Dantzig. In *Kuhn*, pages 247–254. Princeton University Press, H.W.; Tucker, A.W., Linear Inequalities and Related Systems, Annals of Mathematics Studies, 38, Princeton (NJ), 1956.

R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.

Noam Nisan. *Clarke Pivot Rule*, chapter 9.3.4, pages 219–220. Cambridge University Press, 2007a.

Noam Nisan. *Uniqueness of Prices*, chapter 9.5.5, pages 230–231. Cambridge University Press, 2007b.

W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1):8–37, 1961.

Appendix A

A.1 Carpool Totally Unimodular

First we restate Theorem 5.1.1:

Theorem 5.1.1. *The constraint matrix for the carpool assignment problem is totally unimodular and therefore the linear program has integral solution and the problem is in P.*

Proof. Observe the problem formulation: Let $P = \{p_1, p_2, \ldots, p_n\}$ be the set of passengers and $D = \{d_1, d_2, \ldots, d_m\}$ the set of drivers. There are a collection of pairs A, where if passenger p_i is willing to ride with driver d_j then $(p_i, d_j) \in A$. There's a capacity function $c : D \to N$ which indicates how many passengers each driver can drive. Represent the problem as a graph $G = (P \cup D, A)$. There is a weight function $w : A \to R$ which defines the riders valuation $v(p_i, d_j)$, that passenger i gets for riding with driver j. The goal is to assign members of P to D, such that capacity constraints hold and passengers welfare is maximized. Note we require that the maximum number of passengers ride in cars. Everyone needs to go to work.

First define the integer program. The goal of the problem is to maximize total weight. For every passenger p_i, driver d_j pair define a variable $x_{i,j} \in \{0, 1\}$ that is 1 if the passenger isn't in the car and 0 otherwise. Our goal is to maximize

$$\sum_{i,j} v(p_i, d_j) x_{i,j}$$

But we want to guarantee the maximum number of seats are filled, regardless of negative valuations. Therefore we shift every valuation by the same additive factor such that all $v(p_i, d_j)$ are positive. Therefore it's always the case that the maximum number of seats are filled.

Now there are two types of constraints. First that no passenger is in more than one car. I.e

$$\sum_j x_{i,j} \leq 1 \quad \forall i$$

And the second is that no car is overfilled:

$$\sum_i x_{i,j} \leq c_j \quad \forall j$$

Assuming $X = [x_{1,1}, x_{1,2}, \ldots, x_{1,m}, \ldots, x_{i,1}, x_{i,2}, \ldots, x_{i,m}, \ldots, x_{n,m}]$. Then we can write the two associated constraint matrices B and C, where:

$$B = \begin{bmatrix} 1 & 1 & \ldots 1 & 0 & 0 & \ldots 0 & 0 & 0 & \ldots 0 \\ 0 & 0 & \ldots 0 & 1 & 1 & \ldots 1 & 0 & 0 & \ldots 0 \\ \vdots & \ddots & & & & & & & \\ \vdots & & \ddots & & & & & & \\ 0 & 0 & \ldots 0 & 0 & 0 & \ldots 0 & 0 & 0 & \ldots 1 \end{bmatrix}$$

$$C = \begin{bmatrix} 1 & 0 & \ldots 0 & 1 & 0 & \ldots 0 & 1 & 0 & \ldots 0 \\ 0 & 1 & \ldots 0 & 0 & 1 & \ldots 0 & 0 & 1 & \ldots 0 \\ \vdots & \ddots & & & & & & & \\ \vdots & & \ddots & & & & & & \\ 0 & 0 & \ldots 1 & 0 & 0 & \ldots 1 & 0 & 0 & \ldots 1 \end{bmatrix}$$

Therefore our full contrainst matrix $A = \begin{bmatrix} B \\ C \end{bmatrix}$. Heller and Tompkins [1956] provides sufficient conditions for a matrix to be totally unimodular. If there exists a complete and disjoint partition of the rows of A into two B and C such that

- Every column of A contains at most two non-zero (i.e., $+1$ or -1) entries.
- Every entry in A is 0, $+1$, or -1
- If two non-zero entries in a column of A have the same sign, then the row of one is in B, and the other in C
- If two non-zero entries in a column of A have opposite signs, then the rows of both are in B, or both in C

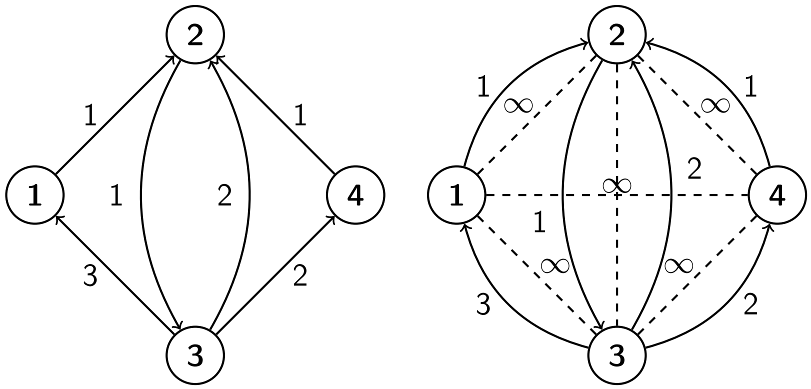

Figure A.1: Clearing Problem instance (left) and Bike Sharing reduction (right)

In this specific case, let B and C be the matrices associated with the two types of contraints defined above. Every $x_{i,j}$ value, a column, appears in exactly one constraint in each set. In this case, the values are always $+1$. Every other value is 0. Therefore conditions above are satisfied, A is totally unimodular, and there's always an integral solution to our carpool assignment problem. □

A.2 Bike-Sharing NP Hardness

Theorem A.2.1. *Finding the welfare-maximizing allocation for the bike sharing problem with a single valuation is NP-hard.*

Proof. We'll show the problem is NP-hard by a reduction of the Clearing Problem [Abraham et al., 2007], over a weighted directed graph (generalization of Kidney Exchange Problem). The clearing problem can be formulated as follows: Given a directed graph $G = (V, E)$ with edge weights w_e find the weight maximizing collection of disjoint cycles in G with maximum cycle length L. (Insert ref) showed the problem was NP-complete for $L \geq 3$.

Assume we're given an instance of the Clearing Problem with $L = \infty$, then the reduction procedes as follows:

Figure A.1 is an example of a clearing problem instance and it's respective bike sharing reduction. Create a metric space G' over the vertex space of G. I.e. for every $v \in V$ make a station in G'. For any two stations $i \neq j$ in

G' set $d(i, j) = \infty$. In figure 2, the dotted lines represent the new distance metric. Now consider the original edge set E. Let $e = (i, j) \in E$ be an arbitrary edge. Create an agent who's preferred source and sink are i and j, respectively. Their private valuation is w_e. Repeat this for every edge. In figure two, each directed edge represents one of these agents with their valuation represented as the edge weight. Lastly, let each station have exactly one bike $b_s = 1$ and no empty spaces $e_s = 0$. Solve this instance of the single parameter bike sharing welfare maximization problem and return the collection of agent journeys (source sink pairs) as the edges comprising the cycles for the original problem.

I claim these solutions are equivalent. First, there are $|V|$ stations in the new graph and $|E|$ agents. Therefore the reduction can be formulated in polynomial time. Note by construction if some agent is allocated a source-sink pair it must be their preferred one. Otherwise, the distance penalty in welfare would be $-\infty$ and they wouldn't be chosen in a maximal solution. By the supply constraints (1 bike at each station) only one departing agent can be chosen for each station. In the original problem this correlates to vertex independence. By the capacity constraints (0 empty spaces at each station) a station is a valid sink if and only if it's also a source for another agent. Therefore every edge belongs to a cycle. The single parameter integer program aims to maximize welfare.

Because agent valuations are equivalent to original edge weights, maximizing welfare in the reduction is equivalent to maximizing edge weight in the original problem. Therefore the reduction finds the maximum weight disjoint cycles in the original graph and the proof is complete. □

A.3 Find Jumps Implementation

The `FindJumps` algorithm, see Algorithm 2, computes the sorted list of discontinuities and their associated function values in a step function. Note this includes the domain boundary values. The algorithm maintains a queue of regions where we know the function changes values. Whenever the width of region at the top of the queue is smaller than a chosen tolerance we add the jump and function value to our list. Otherwise we check the allocation function at the middle value. If *tail*, *mid*, *head* are all distinct, there are at least two value changes in this region and we add [*mid*, *head*] to our queue. Then we continue with the lower region. If there are only two distinct values, then we continue our search in the smaller region bounded by different function values. Lastly we return the lists sorted according to the jump values.

Algorithm 2 Find Jumps Algorithm, which computes the discontinuities in a step function.

```
procedure FindJumps(f, v_min, v_max)
    Queue q = ∅
    List jumps = [v_min, v_max], List vals = [f(v_min), f(v_max)]
    tail = v_min,  head = v_max
    if f(tail) ≠ f(head) then
        q.add([tail, head])
    while q not empty do
        if head − tail < tol then
            res.add(tail, f(head)])
            q.pop
            [tail, head] = q.peek()
        else
            mid = (tail + head)/2
            if f(mid) ≠ f(tail) then
                if f(mid) ≠ f(head) then
                    q.add([mid,head])
                head = mid
            else
                tail = mid
    Return jumps, vals sorted according to jumps
```