

Mechanism design for public goods, exclusive private goods and subcontracting

Citation for published version (APA):

Csapó, G. (2015). *Mechanism design for public goods, exclusive private goods and subcontracting*. [Doctoral Thesis, Maastricht University]. Datawyse / Universitaire Pers Maastricht. <https://doi.org/10.26481/dis.20150605gc>

Document status and date:

Published: 01/01/2015

DOI:

[10.26481/dis.20150605gc](https://doi.org/10.26481/dis.20150605gc)

Document Version:

Publisher's PDF, also known as Version of record

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.

Copyright © Gergely Csapó, Maastricht 2015

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form, or by any means, electronic, mechanical, photocopying, recording or otherwise, without the prior permission in writing from the author.

Cover illustration: author's own design prepared in the software of tagul.com.

This book was typeset by the author using \LaTeX and the ClassicThesis package.

Published by Universitaire Pers Maastricht

ISBN: 978-94-6159-443-3

Printed in the Netherlands by Datawyse

MECHANISM DESIGN FOR PUBLIC GOODS, EXCLUSIVE
PRIVATE GOODS AND SUBCONTRACTING

Dissertation

to obtain the degree of Doctor at Maastricht University,
on the authority of the Rector Magnificus
Prof. dr. Luc Soete,
in accordance with the decision of the Board of Deans,
to be defended in public
on Friday, 5th of June, 2015 at 10:00 o'clock

by

Gergely Csapó



Supervisor:

Prof. dr. Rudolf Müller

Co-Supervisor:

Dr. André Berger

Assessment Committee:

Prof. dr. A.J. Vermeulen (chairman)

Dr. T. Harks

Dr. K. Mierendorff (University College London)

Prof. dr. M. Uetz (University of Twente)

This research has been financially supported by the Graduate School of Business and Economics (GSBE).

To my family.

*“The mighty masterpiece at last is done;
The wheel revolves; the master in his throne
Doth rest.”*

— Imre Madách, *The Tragedy of Man*¹

ACKNOWLEDGEMENTS

This thesis is the embodiment of my four year adventure in the world of academia. Doing a PhD was special work- and lifestyle, hence a complex experience that constituted an amazing part of my life. Of course, it would not have been as remarkable or even possible without the help of many people. I thank all of you who contributed to this journey in one way or another. In the following, I would like to express my gratitude to some of you personally.

First and foremost, I thank my supervisor, Rudolf Müller, for trusting me with the opportunity of becoming a PhD student. He guided me in the maze of research by explaining the big picture and coached me by being critical of details, providing me feedback on presentations and articles to review. I also enjoyed wide freedom in selecting topics and experimenting with different approaches. Though this came with the certain uncertainty, it contributed a lot to my understanding and personal development.

I thank André Berger for jumping on the wagon and taking the responsibility of being my co-supervisor. I thank Alexander Grigoriev for the fun we had with teaching and for being enthusiastic and positive regardless of the topic or the time of the day. I thank Dries Vermeulen for keeping his door always open for professional and not that professional discussions.

I thank the members of my assessment committee for the thorough reading and constructive feedbacks. I thank Hemant Bargava for having me at UC Davis for a week and for the eye-opening discussions on different perspectives of research. I thank Tjark Vredeveld and Marc Uetz for the introduction to operations research and for infecting me with their enthusiasm for that field. Certainly, nothing could have been accomplished without the

¹ Translation by William N. Loew, October, 1908

administrative and personal help of the most professional and friendly secretaries I have ever met: Karin van den Boorn, Haydeé Hallmanns and Yolanda Paulissen. Thank you for your seemingly infinite patience, friendly words, continuous work and for laughing on my jokes.

I thank all the inhabitants of the fourth floor for creating a caring and sociable working environment that enriched me with friendships, deep talks, gossips and cultural shocks. It was great to share the joys and sorrows over teaching or the PhD life during boardgame/movie nights or coffee breaks with such a nice audience. I would like to personally mention Elnaz Bajoori, Thomas Goetz, Angie Mounir, Hanno Reuvers, Salahnejhad Ghalehjooghi Ahmad and Anna Zseleva. Special thanks to Jan Lohmeyer and Marc Schroeder for proof-reading the thesis and for many other things. I am indebted to Marc also for translating my summary to Dutch.

I thank Aleksandar Andonov, Janos Flesch and Norbert Metiu for being nice flatmates and for their friendship. The common lunches, dinners and small talks gave me a sense of home and serenity. An office with a couch and a view to the park of the university is a promising start, but it was Martijn van Brink, Birol Yüceoğlu, Hossein Naeemi and Cyriel Rutten who filled it with warmth, fun and the thrill of science. I thank you guys for being a constructive distraction. I thank Artem Duplinskiy, Bas Marban and Ruben van der Zwaan for being supportive friends and for all the nourishment they provided for my mind, soul and stomach. I am further indebted to Bas for helping me with countless things including getting the PhD position and my new job.

Finally, I would like to express my heartfelt gratitude to my family for their continuous support and love. Kedves Anya, Apa, Zoli és Jucus, köszönök szépen mindent! Milá Evi, doma je tam, kde jsi ty.

CONTENTS

i	PRELIMINARIES	1
1	INTRODUCTION	3
1.1	Preliminaries	8
1.1.1	Incentive compatibility	10
1.1.2	Individual rationality	11
1.1.3	Other useful concepts	12
1.2	Outline of the thesis	14
ii	OPTIMAL MECHANISM DESIGN	17
2	OPTIMAL MECHANISM DESIGN FOR THE PRIVATE SUPPLY OF A PUBLIC GOOD	19
2.1	Introduction	19
2.2	Model and Preliminaries	22
2.3	Finding the optimal mechanism	25
2.3.1	Transformation of the objective function	26
2.3.2	Polynomial time algorithms	27
2.3.3	Independent valuations and signal distributions	29
2.4	Other solution concepts	31
2.4.1	Bayes-Nash implementation	32
2.4.2	Interim individual rationality	34
2.5	Excludable public goods	38
2.5.1	Integer linear program formulation	38
2.5.2	A formulation having totally unimodular constraint matrix	41
2.5.3	From public to private goods	42
2.6	Conclusion	43
iii	APPROXIMATIONS	51
3	LEAD MARKETING AUCTIONS WITH PREMIUM FOR EXCLUSIVE POSSESSION	53

3.1	Introduction	53
3.2	Model formulation	61
3.3	Upper bounds	63
3.3.1	Relaxation 1: Public exclusivity margin	64
3.3.2	Relaxation 2: Bound with representatives	66
3.3.3	Significance of Relaxations	69
3.4	Heuristic mechanisms	71
3.4.1	VCG mechanism	72
3.4.2	Simple mechanisms	74
3.4.3	A two-dimensional candidate mechanism	77
3.5	Computational analysis	79
3.6	Conclusion	83
4	SUBCONTRACTOR SCHEDULING	93
4.1	Introduction	93
4.2	Preliminaries	98
4.3	Approximating scheduling algorithms	99
4.3.1	Shortest Remaining Processing Time	101
4.3.2	Highest Weight First	101
4.3.3	Other schedules	104
4.4	Mechanism Design	105
4.4.1	Equal weights	106
4.4.2	Arbitrary weights	112
4.5	Conclusion	117
iv	APPENDIX	129
	BIBLIOGRAPHY	131
	INDEX	139
	NEDERLANDSE SAMENVATTING	141
	VALORISATION	145
	CURRICULUM VITAE	151

Part I

PRELIMINARIES

“Mechanism design is the art of designing the rules of the game so that the desirable outcome is reached despite the fact that each agent acts in his own self-interest.”

— Sandholm [68]

INTRODUCTION

A fundamental question in economics is how to allocate scarce resources among economic agents. The designer of an economic system can model this question as an optimisation problem: given the input parameters the goal is to optimise an economic measure subject to resource constraints. For example given a set of bidders and their preferences over several items, what is the allocation of items and selling prices that maximise welfare? Or given a set of projects and their processing times and opportunity costs, how should we allocate the time of a subcontractor in order to minimise the total cost?

Despite the fact that in many cases such optimisation problems are already complicated enough, these models lack important characteristics of real-life systems. For example, parameters that serve as input for optimisation models are usually dispersed and privately held information of strategically behaving agents, who strive to achieve their own goals in interaction with the system. Indeed, bidders can lie about their preferences in order to lower their payments, or jobs may inflate their opportunity costs to get more help from the subcontractor. Failing to take incentives into account can lead to waste of resources, lower revenue, significant project delays or even to system meltdowns such as the recent financial crisis. On the other hand, well-designed institutions can help society. A well-known example is regarding matching algorithms in kidney exchange programs that facilitate finding compatible donors in a way that morally questionable incentives are taken care of. In a similar fashion sponsored search advertising helps companies targeting customers better, while yielding considerable revenue for the search engine providers. For these reasons we can conclude that the original question should rather be: how to allocate scarce resources among *competing* economic agents *despite their strategic behaviour*?

Further complications emerge when the designer faces constraints related to time or computational power. These issues can be vital for practical implementations where decisions have to be made fast and the size of the input

is large. Consider an online platform that runs parallel auctions for a wide range of items. The winner of these items and the corresponding prices have to be determined quickly subject to the limited capacity of the servers that process information from large amounts of bidders. Other examples emerge from project management when it comes to coordination of various projects in a huge supply chain. In this case speedy decisions are fundamental to meet deadlines. To take these matters also into account, the core question of this thesis is the following. How can we allocate scarce resources among competing economic agents despite their strategic behaviour in case the *computational power of the system is limited*? In particular, we ask this question for three specific problems.

PUBLIC GOODS The first problem is about the supply of a public good by a monopolistic provider. Samuelson [67] defines public goods as goods “which all enjoy in common in the sense that each individual’s consumption of such a good leads to no subtractions from any other individual’s consumption of that good”. In short this means that public goods are non-rivalrous. Most of the time they are also non-excludable, that is, no agent can be excluded from consuming the good. Examples include national defence, public parks and bridges. Non-rivalrous and excludable goods, e. g., cinemas, electricity or internet subscriptions, are called club goods. Pioneering papers in mechanism design, such as [17] and [35], studied incentives that would motivate agents to reveal their true preferences to help the system decide on the efficient level of public good provision. The bulk of the work that has been done on this topic is concentrated on the question whether efficient implementation of public good provisions is possible under different settings. As opposed to this, we take the view of a monopolistic provider, whose basis for supply decisions is profit maximisation. This institutional setting has been considered so far by only a handful of papers for basic settings (see Chapter 2 for references), hence our work intends to complete the picture in a more general framework. Although one can construct examples that are analogous to profit-maximising public good producers (cities bidding for regional infrastructure developments), it is not realistic in general that prices can be set on an individual basis, therefore the topic of our study is more of theoretical interest. Our analysis pinpoints some salient features of the underlying op-

timisation problem that makes the pricing of a public good computationally "easy" as opposed to the pricing of a single private good that is deemed to be computationally "hard" in general settings (Papadimitriou and Pierrakos [57]). Moreover, our structural results enables us to reduce profit maximisation to welfare maximisation, and to establish a connection with a graph theoretical problem.

ONLINE AD AUCTIONS By the emergence of the Internet a vast amount of information has become available regarding products and services. Companies, like Google, Yahoo or Bing, provide search engines that make it possible for customers to find the appropriate information by searching for keywords. Based on the keywords entered, a list of links is retrieved ordered by relevance. Next to these results the search engines preserve a separate space for links that are *sponsored* by companies who wish to display their advertisement related to the searched keywords. Unlike the results, the positions of these sponsored links are auctioned off generating enormous profit for the companies behind the search engines. Such auctions are called sponsored search auctions or ad auctions. Other companies go one step further and try to actively match demand with supply. First, they identify customers that are interested in certain products or services based on their online searches. Then they summarise the credentials of the customer and the request details in a condensed form, called *lead*. Finally, the lead is conveyed to possible providers of the desired product or service for a charge. Some companies charge flat prices, but there are more and more businesses that run auctions to determine the winner(s) and the price of the lead. The marginal cost of displaying sponsored links or providing copies of leads is almost zero, thus they can be considered as non-rivalrous goods and their pricing cannot be cost-based. If the provider of these products is uncertain about the preferences of its customers, i. e., the value it is worth for them, then auctions are an appropriate solution. Both kinds of platforms usually have to carry out millions of auctions in a small amount of time, thus computational aspects play a central role in the implementation design. This means that in spite of the fact that the bidders might have complex preferences, the bidding language, the rules of the allocation and the pricing have to be kept simple. Most of these auctions are using a variant of the generalised second price auction (GSP) that takes

one bid as input. This implies that the bidder's complex preferences have to be expressed via only one number. There are some important considerations that bidders cannot convey this way. For example firms might care about the total number of ads displayed next to theirs, the place of their ad in the ordered list or the particular neighbouring advertisements. Consequently the choice menu offered by such naïve auctions cannot be tailored to the special needs of the bidders, which might lead to some loss in revenue. Motivated by online ad auctions, in Chapter 3 we study revenue-maximising mechanisms for settings where the exclusivity of receiving a product or service matters. Our results have both theoretical and practical merits. In particular, we devise mechanisms that are simple to implement and yield an expected revenue that is close to the optimal one. In order to derive such results without knowing the optimal mechanism, we derive non-trivial upper bounds on the optimal revenue and relate the expected revenue of the candidate mechanism to these bounds. We also carry out simulations to demonstrate the power of the bounds and the constructed mechanisms.

SUBCONTRACTING Subcontracting is an act where an organisation contracts a third party to perform a part of its operations. Subcontracting has to be distinguished from outsourcing as in the latter one the whole process is contracted out. Benefits from involving third parties include cost reduction, risk mitigation, quality assurance, resource augmentation and expertise sharing. These advantages are appreciated particularly in the service industry where the main expenses arise from fluctuating demand and locked in capacities. Another frequent application area is related to complex projects that consist of a large amount of interdependent subprojects. Naturally, the resources of the subcontractors are also limited, therefore a firm that runs multiple similar operations has to decide on the allocation of the subcontractor's time. This creates a competition among the different operations and raises incentive issues in addition to the resource optimisation problem. To illustrate the importance of this topic we recall the case of Boeing's Dreamliner supply chain, where the lack of attention to these matters resulted in overloaded schedules of subcontractors and long delays in the overall production (see Vairaktarakis [79] for further details and references). Additional complications occur when there is no central authority to actively gather

information and allocate resources. In such *decentralised* systems the rules have to be set in a way that the actions of the autonomous agents are coordinated towards the common goal. An important tool in mechanism design to align incentives is the usage of monetary transactions. Moreover, like in the previous topics, it is often the goal of the system to maximise the collected payments. However, when the target is to maximise welfare or another non-monetary objective, then it is desirable to minimise the usage of money. For example, one could aim for minimising the total sum of payments or if compensations are allowed, then to achieve that the payments cancel each other out. There is a line of research focusing on the extreme case, where monetary transactions are not allowed at all. Schummer and Vohra [73] provides an extensive survey on mechanism design without money. In Chapter 4 we address these previously mentioned issues in a stylised subcontracting model by analysing various allocation algorithms and mechanisms with different practical properties. In this model there are several jobs characterised by their processing times and weights (e. g., opportunity costs per time unit). We assume that there is a subcontractor that can speed up the processing of the targeted job by a factor of two, but can only work on one job at a time. The role of the mechanism is to gather the information on processing times and weights, and construct a schedule for the subcontractor. The measure with which we evaluate the performance of the mechanisms is the weighted sum of time savings achieved via subcontracting. Notably, we show for this setting that ordering jobs simply according to their weights guarantees $2/3$ of the optimal objective value. We also construct a pricing scheme that implements this allocation. Furthermore, for the case of equal weights we devise a mechanism where the sum of the payments is zero, and introduce an auction variant that implements the optimal time schedule.

APPROACH Throughout the thesis we work under the aegis of algorithmic mechanism design, an intersection of Economics, Game Theory and Computer Science. Game Theory helps with modelling the behaviour and strategic interactions of rational agents. The role of Computer Science here is to provide analytical tools for shedding light on the complexity of the computations needed to derive mechanisms. The field of algorithmic mechanism design is grounded in the seminal paper of Nisan and Ronen [55]. Nisan et al.

[56] contains a broad range of results and applications that emerged from this explosively growing field. For a comprehensive introduction to classical mechanism design see Mas-Collel et al. [51]. Vohra [83] offers a more computer scientist-friendly approach by tackling various problems from mechanism design with linear programming.

In the next section we introduce the basic notation and terms used throughout the thesis. We also give the general definition of the most important concepts in mechanism design. The topic specific notions and symbols are discussed in the related chapters.

1.1 PRELIMINARIES

The focus of mechanism design is on economic situations where several *agents* interact with each other. In an auction environment the agents are referred to as bidders, while in case of machine scheduling they are called jobs. We denote the set of agents participating in the mechanism by $N := \{1, \dots, n\}$. Dependent on the topic the interaction of the agents can lead to different *outcomes*. The set of possible outcomes is denoted by \mathcal{O} . In auctions \mathcal{O} usually comprises sets of agents that can be served at the same time, whilst for scheduling it contains functions that assign agents to different time intervals on a machine. If we allow for randomised rules, then \mathcal{O} can be the set of distributions over a ground set of possible outcomes.

We assume that agents have preferences over the outcomes, which are dependent on their *private information* and can be expressed by their *valuation function*. We furthermore assume that this private information can be represented by a vector of m numbers $t^i \in T^i \subseteq \mathbb{R}^m$, to which we refer as the *type* or *signal* of agent i . In terms of auctions such private information can be the reservation price a bidder would pay for an object. Dependent on m we distinguish between single and multi-dimensional (parameter) type spaces. An example for multi-dimensional typed space could come from machine scheduling where jobs are endowed with different processing times and weights.

To be able to talk about the vector of agents we define $T := T^1 \times \dots \times T^n$ and $T^{-i} := \times_{j \neq i} T^j$. Each element in T is an $m \times n$ matrix and is denoted by t . We assume that the distribution of t is public knowledge. The valuation function

of agent i is defined as $v^i : \mathcal{O} \times T \rightarrow \mathbb{R}$. For example $v^i(z, t)$ is the valuation of agent i for outcome z if the agents have type t . When v^i is not dependent on t^{-i} for any $i \in N$, then we say that the valuations are independent. In case valuations are allowed to be interdependent types are usually referred to as signals. Settings with interdependent valuations are also said to have informational externalities. We refer to the sum of all valuations as *social welfare* or welfare in short. An outcome that maximises social welfare for a given type profile is called *efficient*. We denote by $SW^N(z, t) := \sum_{i \in N} v^i(z, t)$ the social welfare of agents in N for outcome $z \in \mathcal{O}$ and type profile $t \in T$.

During the interactions agents perform actions to influence the outcome. The set of possible actions for agent i is denoted by A^i , while the action space is denoted by $A := A^1 \times \dots \times A^n$. A *mechanism* is characterised by its allocation rule $x : A \rightarrow \mathcal{O}$ and its payment scheme $\pi : A \rightarrow \mathbb{R}^n$. In each mechanism, the designer announces the allocation rule and the payment scheme and the agents respond by taking actions. Based on the actions of the agents and the announced rules, the designer selects an outcome and elicits payments.

We assume that agents are *rational*, i. e., they choose actions in order to maximise their utility. Throughout the thesis we furthermore assume that this utility is *quasi-linear*, meaning that utility of agent i for action profile a and type profile t is

$$u^i(x(a), \pi(a), t) = v^i(x(a), t) - \pi^i(a). \quad (1)$$

A particular class of mechanisms that is in the focus of most of the studies in mechanism design is called *direct mechanisms*. In these mechanisms the set of possible actions is restricted to their set of possible types, i. e., $A^i = T^i$. In most of the settings, including all subjects of this thesis, it is without loss of generality to focus only on direct mechanisms. This is allowed by the celebrated result of Myerson [54], known as the *revelation principle*. This does not mean that studying other kinds of mechanisms is waste of effort. In practice general mechanisms that lead to the same outcome may have desirable properties compared to the direct ones such as less communication cost, more natural implementation or privacy preservation. Indeed, in Chapter 4 we provide a scheduling mechanism that demonstrates these advantages.

1.1.1 Incentive compatibility

As the utility of an agent depends on his realised type, so does the choice of the action that maximises it. A function that maps types of the agent to possible actions is called a *strategy*. Under direct mechanisms, a strategy of agent i can be given as $\theta^i : T^i \rightarrow T^i$. A strategy profile $\theta := \{\theta_1, \dots, \theta_n\}$ is an equilibrium if no agent wants to unilaterally change its strategy. In this respect equilibria are “stable” situations that a mechanism may lead to. In mechanism design we would like to incentivise strategies that lead to a “good” equilibrium, e. g., ones that maximise revenue, social welfare or minimise some other objective function on the possible outcome space. The main source of incentives for the designer is the combination of monetary payment and the promise of a desirable allocation. In the rest of this section we focus on direct mechanisms and define equilibrium notions and some other related concepts used in this dissertation. First we give the intuition, then the formal definition.

A strategy is *dominant* if it is the best strategy regardless of the strategies chosen by other agents. A direct mechanism is *dominant strategy incentive compatible* (DSIC) if for each agent truth-telling is a dominant strategy given any reports of the other agents.

Definition 1.1 (DSIC). A direct mechanism (x, π) is DSIC if for every i , type (t^i, t^{-i}) and report \hat{t}^i, \hat{t}^{-i} it holds that

$$u^i \left(x(t^i, \hat{t}^{-i}), \pi(t^i, \hat{t}^{-i}), (t^i, t^{-i}) \right) \geq u^i \left(x(\hat{t}^i, \hat{t}^{-i}), \pi(\hat{t}^i, \hat{t}^{-i}), (t^i, t^{-i}) \right). \quad (2)$$

This concept is used mainly for settings where valuations are independent, thus it prevails in Chapters 3 and 4, but it appears also in Chapter 2. When valuations are interdependent (like in Chapters 2 and 3) DSIC is very limiting, because it entails truth-telling even when the other agents lie about their type. Therefore in such cases we work with a less demanding concept, which requires that truth-telling is a Nash equilibrium in the complete information game defined by each type profile.

Definition 1.2 (EPIC). A direct mechanism (x, π) is ex-post incentive compatible (EPIC), if for all i , t^{-i} , t^i and \hat{t}^i it holds that

$$u^i \left(x(t^i, t^{-i}), \pi(t^i, t^{-i}), (t^i, t^{-i}) \right) \geq u^i \left(x(\hat{t}^i, t^{-i}), \pi(\hat{t}^i, t^{-i}), (t^i, t^{-i}) \right). \quad (3)$$

It is clear that every DSIC mechanism is EPIC. The other direction generally does not hold, except when valuations are independent: then every EPIC mechanism is also DSIC. In Chapter 2 we discuss an even less demanding property, which assumes that agents make their decisions under uncertainty with respect to the types of the others according to a common prior distribution. *Bayes-Nash incentive compatibility* entails that truth-telling maximises expected utility, given that the other agents report truthfully.

Definition 1.3 (BNIC). A direct mechanism (x, π) is Bayes-Nash incentive compatible (BNIC) if for all i , t^i and \hat{t}^i we have that

$$\begin{aligned} \mathbb{E}_{t^{-i}|t^i} \left[u^i \left(x(t^i, t^{-i}), \pi(t^i, t^{-i}), (t^i, t^{-i}) \right) \right] \\ \geq \mathbb{E}_{t^{-i}|t^i} \left[u^i \left(x(\hat{t}^i, t^{-i}), \pi(\hat{t}^i, t^{-i}), (t^i, t^{-i}) \right) \right]. \end{aligned} \quad (4)$$

Note that $\mathbb{E}_{t^{-i}|t^i} [\dots]$ denotes the conditional expected value of the underlying expression with respect to all t_{-i} given t^i .

1.1.2 Individual rationality

We assume that there is an outside option for every agent that guarantees them zero (expected) utility. Hence we require mechanisms to be *individual rational* not to deter agents from participation.

Definition 1.4 (EPIR). A direct mechanism (x, π) is ex-post individual rational (EPIR) if for all i and t we have that

$$u^i(x(t), \pi(t), t) \geq 0. \quad (5)$$

Similarly to incentive compatibility, there is a Bayes-Nash variant of individual rationality as well.

Definition 1.5 (IIR). A direct mechanism (x, π) is interim individual rational (IIR) if for all i and t^i we have that

$$\mathbb{E}_{t^{-i} | t^i} \left[u^i \left(x(t^i, t^{-i}), \pi(t^i, t^{-i}), (t^i, t^{-i}) \right) \right] \geq 0. \quad (6)$$

It is worth mentioning here some basic observations. If there is a dummy type for which the valuation of each allocation is zero and the payment is equal to zero if that type is reported, then EPIC implies EPIR and BNIC implies IIR. Regarding the relations of the different incentive compatibility concepts it is evident that every EPIC is BNIC, but the converse is not true generally. For more discussion on that issue see Gershkov et al. [30].

1.1.3 Other useful concepts

MONOTONICITY Constraint (2) can be satisfied by an appropriate pair of allocation rules and payment schemes. To ease the construction of counterexamples and proofs we derive the so-called monotonicity condition that is imposed solely on allocation rules. Failing to meet this condition implies that there is no payment scheme for this allocation rule such that together they can satisfy (2). The condition is obtained by adding up two (2) constraints such that the first constraint is applied to agent i having type t^i and reporting \hat{t}^i , while the second one is expressed for the same agent having \hat{t}^i and reporting t^i , to yield

$$v^i \left(x(\hat{t}^i, t^{-i}), \hat{t}^i \right) - v^i \left(x(\hat{t}^i, t^{-i}), t^i \right) \geq v^i \left(x(t, \hat{t}^i) \right) - v^i \left(x(t, t^i) \right). \quad (\text{MON})$$

Allocation rules satisfying (MON) will be called monotone. In single-item auction setting this condition entails that if a bidder receives the item for a given bid, then he should receive it for higher bids as well. An example from machine scheduling is that increasing a job's reported weight cannot increase the job's finishing time.

Another useful insight about DSIC mechanisms concerns their pricing scheme. The following lemma states that for the same outcome one has to be charged the same amount regardless of the reported type, if the reports of the others are kept unchanged. This is also referred to as the *taxation principle*.

In case of for deterministic mechanisms DSIC and EPIR lead to an intuitive, but non-trivial upper bound for each payment.

Lemma 1.6. *Let (x, π) be a DSIC mechanism. For given i, t^{-i}, t^i and \hat{t}^i , having $x(t^i, t^{-i}) = x(\hat{t}^i, t^{-i})$ implies that $\pi^i(t^i, t^{-i}) = \pi^i(\hat{t}^i, t^{-i})$. Moreover, if (x, π) is also EPIR, then for given i, t^{-i}, t^i we have that*

$$\pi^i(t^i, t^{-i}) \leq \inf_{\hat{t}^i} \{v^i(x(\hat{t}^i, t^{-i}), \hat{t}^i) \mid x(\hat{t}^i, t^{-i}) = x(t^i, t^{-i})\}.$$

Proof Follows directly from (4) and ex-post individual rationality. \square

COMPLEXITY If a problem can be solved in time polynomial in the size of its input, regardless of the specific input, then we say that the problem is in P, or equivalently, it is polynomial-time solvable. Non-deterministic polynomial time (NP) problems are decision problems for which “yes” answers are verifiable in polynomial time. Problems that are “at least as hard” as the ones in NP are called NP-hard. Here “at least as hard” means that if there was a polynomial algorithm to solve an NP-hard problem, then any problem in NP would be polynomial-time solvable. Problems that are NP and NP-hard are called NP-complete. It is generally believed by the scientific community that problems that are NP-hard are not polynomial-time solvable. One reason is that a counter example would go against many fundamental results in other areas of mathematics and would have implications that are highly unlikely to be true.

Complexity is an issue for all parties involved in a mechanism. The designer has to calculate the optimal mechanism, while the agents have to decide on an action that maximises their utility. If we consider mechanisms or strategies that are successful in practice then we can observe that they are rather simple and their main rules do not change by the number of participants or other input parameters (e. g., second price auction and its variants). This might serve as evidence that mechanism designers prefer systems that are easy to implement and robust to changes in the input even for some sacrifice in the performance. In consonance with that agents appreciate rules that are understandable and make their optimal action clear. Candidates for such successful mechanisms can be found in the class of approximation mechanisms.

APPROXIMATION MECHANISMS For maximisation problems an algorithm or a mechanism is a *factor ϱ approximation* if it can be constructed and executed in time polynomial of the input size, moreover, it yields an objective value that is at least ϱ times the optimal objective value. Approximation algorithms are appreciated in cases when the underlying problem is NP-hard, or the optimal algorithm/mechanism has a complicated structure. We make use of such mechanisms in Chapter 3 and Chapter 4. For some problems it is possible to design a polynomial-time approximation scheme (PTAS), that takes a parameter ϵ and for any instance produces a solution with an objective value at least $(1 - \epsilon)$ times the optimal objective.

1.2 OUTLINE OF THE THESIS

The dissertation investigates three application areas in more detail. The first part of the thesis is devoted to optimal mechanism design. In particular, Chapter 2 studies the private supply of a public good, where we take a linear programming approach to search for the profit-maximising mechanism for a broad range of settings. In our model we allow for correlated type distributions as well as for informational externalities. As main result we show that there is a polynomial-time algorithm to find the optimal mechanism that is deterministic, EPIC and EPIR. We furthermore discuss the effects of randomisation and of changing the underlying incentive compatibility and individual rationality assumptions to BNIC and IIR respectively.

To escape the intrinsic hardness of multi-dimensional mechanism design and to find simple mechanisms for practitioners, the third part of the thesis departs from optimal mechanisms design and turns to the study of approximations. In this spirit Chapter 3 deals with online lead auctions, a topic inspired by the cooperation with Reply!, a sales lead generator company. The twist in this setting, in contrast with the previous part, is that the private information held by the agents is multi-dimensional: their valuation is different for being served exclusively or simultaneously with others. We offer different simple mechanisms that are easy to implement in practice and still achieve a constant portion of the optimal revenue. To be able to show such results, we upper bound the optimal revenue via different single-dimensional relaxations of the problem. The yielded bounds are utilised in theoretical proofs

and in numerical simulations to evaluate the performance of the constructed mechanisms.

Chapter 4 is devoted to subcontractor scheduling, a problem of both practical and theoretical interest. We consider a given set of projects, each carried out normally by their own team. In addition to these, there is a single subcontractor that can support one of the projects at a time, thereby helping to complete the project by the respective amount of work earlier. Unlike in the previous chapters, where the focus is on revenue, here, our goal is to maximise social welfare. That is, we are interested in schedules for the subcontractor that could be implemented in practice and yield total saving that are within a constant factor of the optimal schedule's total saving. To this end, we investigate different simple heuristics that provide an ordering for the jobs (e. g., SPT, HWF, SMITH'S RULE) and derive approximation factors for them. From a mechanism design point of view we devise mechanisms that have special properties such as budget-balancedness or admitting natural, auction-type implementation.

Part II

OPTIMAL MECHANISM DESIGN

“...an auction design problem must be treated like any problem of decision-making under uncertainty. No auction mechanism can guarantee to the seller the full realization of his object’s value under all circumstances. Thus, the seller must make his best assessment of the probabilities and choose the auction design which offers him the highest expected utility, on average. The usual “garbage-in, garbage-out” warning must apply here, as in all operations research, but careful use of models and sensitivity analysis should enable a seller to improve his average revenues with optimally designed auctions.”

— Roger B. Myerson [54]

OPTIMAL MECHANISM DESIGN FOR THE PRIVATE SUPPLY OF A PUBLIC GOOD

This chapter is based on the following paper:

G. Csapó and R. Müller, “Optimal Mechanism Design for the Private Supply of a Public Good” in *Games and Economic Behavior*, Volume 80, July 2013, Pages 229–242.

2.1 INTRODUCTION

We study the problem of finding a mechanism that maximises the expected profit of a monopolistic provider of a single, non-excludable public good. The problem was first solved by Güth and Hellwig [36] for the case of Bayes-Nash implementation when valuations are independent and signals are independently distributed with monotone hazard rate. Our goal is to study the problem for possibly interdependent valuations, for general joint distributions of signals, for ex-post and Bayes-Nash implementation, and for ex-post as well as interim individual rationality. It is too much to hope for closed form solutions in each case. Instead, the mechanism design problem is modelled as a combinatorial optimisation problem, more precisely, as an integer linear program (ILP). The ILP has as set of feasible solutions the set of all mechanisms satisfying particular incentive and individual rationality constraints, and as objective the expected profit. The parameters of the constraints and the objective are determined by the set of possible signal profiles and their distribution. One of the challenges is to represent the expected profit as a linear function in the decision variables, the other is to provide, if possible, an efficient algorithm to solve the integer linear program. We are interested in polynomial-time algorithms in the number of agents and the number of signal profiles of the agents. We are also interested in gaining insights in eco-

nomic properties of optimal mechanisms, and classes of instances where a “simple” mechanism is optimal.

OUR CONTRIBUTION AND RELATED WORK To keep our model as general as possible we allow for informational externalities in the valuations¹ as well as for correlation in the signal distributions. The only restriction for the valuations is that for each agent they have to be strictly increasing in the agent’s own signal. This condition together with the assumption of single-dimensional signals help to avoid impossibility results regarding implementability (Jehiel et al. [41]). There has been a lot of literature on the role of informational externalities since the seminal papers of Wilson [84] and Milgrom and Weber [53]. An overview of major findings is provided by Jehiel and Moldovanu [43]. Branco [9] studies profit-maximising multi-unit auctions in the presence of informational externalities and provides regularity conditions under which the optimal mechanism can be characterised. For the welfare-maximising public good problem Laffont and Martimort [47] assume positively correlated signal distributions and they search for mechanisms that are strategy-proof against colluding coalitions.

Our main result is that the profit-maximising, ex-post individual rational, ex-post incentive compatible mechanism can be computed in polynomial time. This holds for any number of agents, which is in sharp contrast to a recent result in Papadimitriou and Pierrakos [57] on profit-maximising single-item auctions for general distributions of signals. Papadimitriou and Pierrakos [57] reduce the optimal auction problem to finding the maximum weight independent set on a k -partite graph, where k is the number of bidders. This yields a polynomial-time algorithm for the two bidder case. They further prove that the problem is NP-hard for more than two bidders. The latter result is also provided independently by Dobzinski et al. [24]. Moreover, they develop a constant factor approximation for the case of an arbitrary number of bidders and an efficient algorithm to compute the optimal randomised mechanism that is truthful in expectation. We show that finding a profit-maximising mechanism for the non-excludable public good problem for general distributions can be modelled as an integer linear program with

¹ We thank an anonymous referee to point out that our results in a previous version of this paper could be extended in this generality.

a totally unimodular (TUM) constraint matrix. More specifically, the ILP is equivalent to a maximal closure problem on partially ordered sets, where weights of the elements are equal to the sum of conditional virtual values minus the cost and the ordering is induced by the monotonicity constraints. Contrary to single-item auctions, this yields a polynomial-time algorithm for any number of agents. Notably, the optimal mechanism might choose to provide the public good for signal realisations where the sum of conditional virtual values is strictly smaller than the cost of providing the good. For the case of independent valuations and independent signal distributions with monotone hazard rate, we get the optimal Bayes-Nash implementable allocation rule of Güth and Hellwig [36] as a byproduct.

For independent signal distributions and valuations without informational externalities, we illustrate, using recent results by Gershkov et al. [30], that allowing for interim individual rationality instead of only for ex-post individual rationality or Bayes-Nash implementation instead of only for ex-post implementation cannot increase the expected profit. However, by virtue of examples it is shown that for interdependent valuations or for correlated signals, Bayes-Nash implementation may yield more profit than ex-post implementation. Transferring techniques developed for auctions by Crémer and McLean [18] to the public good case, we illustrate that under some mild assumptions the profit-maximising interim individual rational ex-post implementable mechanism can be determined in polynomial time. Furthermore as for auctions, it yields full surplus extraction. The same holds for Bayes-Nash implementation in the setting without informational externalities. By virtue of an example it is shown that full surplus extraction breaks down if we require ex-post individual rationality. It remains an open problem whether the optimal deterministic Bayes-Nash incentive compatible, ex-post or interim individual rational mechanism can be efficiently computed for interdependent valuations and general signal distributions.

ORGANISATION In the next section we formally define our problem and introduce the related notation and main concepts. In Section 2.3 we show how to find the optimal solution for the problem. After that, Section 2.4 investigates the effects on the optimal solution of the changes in the underlying individual rationality and incentive compatibility assumptions. Section

2.5 adapts the developed techniques to the case of excludable public goods. Finally, Section 2.6 concludes this chapter by summarising our contributions and discussing some remaining open questions.

2.2 MODEL AND PRELIMINARIES

We are given a set of agents $N = \{1, \dots, n\}$ who hold private information, a signal $t^i \in T^i \subseteq \mathbb{R}$ about their values for the consumption of a public good. We denote the set of signal tuples of all agents by $T := T^1 \times \dots \times T^n$. All sets T^i are assumed to be finite. Signal tuples $t \in T$ occur according to the joint probability distribution φ and the cumulative distribution Φ . Furthermore, let φ_{t^i} be the probability that t^i occurs and Φ_{t^i} the corresponding cumulative distribution function.

We first define the model in such a way that it captures randomized mechanisms as well, therefore the set of possible outcomes is assumed to be $\mathcal{O} := [0, 1]$, that is, the probability that the public good is provided. The valuation of each $i \in N$ for outcome z is defined as

$$v^i(z, (t^i, t^{-i})) = v_t^i z.$$

It is assumed that $v_{t^i, t^{-i}}^i$ is strictly increasing in t^i for each t^{-i} . Such valuations are called interdependent as the value of agent i depends not only on his own signal but also on other agents' signals. Settings with interdependent valuations are also said to have informational externalities. Agents act via a direct mechanism (x, π) , which consists of an allocation rule $x : T \rightarrow \mathcal{O}$ and a payment scheme $\pi : T \rightarrow \mathbb{R}^n$. We allow for transfers. Agents are assumed to have quasi-linear utilities. In this setting the utility of agent i for signal t^i and report profile $(s^i, t^{-i}) \in T$ takes the form of $x(s^i, t^{-i}) v_{t^i, t^{-i}}^i - \pi^i(s^i, t^{-i})$.

Note that this particular choice of the utility function implicitly assumes that the agents' utility for not participating is zero independently of their signal. In settings with allocational externalities it may be assumed that the designer can set the outside option for the agents whose valuation for that option is dependent on their private information (e. g., Jehiel and Moldovanu [42]). The article of Figueroa and Skreta [28] points out that modelling outside options can have significant impact on the optimal mechanism.

The provision of the public good costs C . The profit $\pi(t)$ of the designer for a reported signal tuple t then is the sum of payments collected minus the cost incurred:

$$\pi(t) = \sum_{i=1}^n \pi^i(t) - x(t)C.$$

We are seeking for mechanisms (x, π) that maximise the expected profit, i. e., $\sum_t \varphi_t \pi(t)$. First, we search for the optimal mechanism among those that are deterministic, ex-post incentive compatible and ex-post individually rational. Then we discuss other solution concepts. By the revelation principle the restriction to direct mechanisms is without loss of generality.

For all $t \in T$ let $x_t \in \{0, 1\}$ be the decision variable that represents the allocation rule x , i. e., $x_t = 1 \Leftrightarrow x(t) = 1$. Let us also introduce decision variables for each payment, i. e., for all i and t we have that $\pi_t^i = \pi^i(t)$. Then our mechanism design problem can be formulated as the following integer linear program.

$$\max \sum_{t \in T} \varphi_t \left(\sum_i \pi_t^i - x_t C \right) \quad (\text{P1})$$

subject to

$$v_{t^i, t^{-i}}^i x_{t^i, t^{-i}} - \pi_{t^i, t^{-i}}^i \geq v_{s^i, t^{-i}}^i x_{s^i, t^{-i}} - \pi_{s^i, t^{-i}}^i \quad \forall i, t^{-i}, s^i, t^i, \quad (\text{EPIC})$$

$$v_{t^i}^i x_t - \pi_t^i \geq 0 \quad \forall i, t, \quad (\text{EPIR})$$

$$x_t \in \{0, 1\} \quad \forall t \quad (\text{BI})$$

Constraints (EPIC) ensure implementability in ex-post incentive compatible strategies, while constraints (EPIR) are responsible for ex-post individual rationality by definition of (3) and (5) respectively. As we go for deterministic mechanisms we also require the allocation variables to be binary (BI).

It is easy to verify that in this setting the monotonicity condition (MON) applied for type t^i and report s^i reduces to the following inequality:

$$(v_{s^i, t^{-i}}^i - v_{t^i, t^{-i}}^i)(x_{s^i, t^{-i}} - x_{t^i, t^{-i}}) \geq 0.$$

Therefore we refer to an allocation rule x as monotone if and only if for all i and t^{-i} , x is increasing in t^i , i. e., for all s^i, t^i if $s^i > t^i$, we have that $x_{s^i, t^{-i}} \geq x_{t^i, t^{-i}}$. It is a well-know result (see, e. g. Chung and Ely [16] or Vohra [83]) that monotonicity in such single-parameter domains is necessary and sufficient for implementability in ex-post incentive compatible strategies. Moreover, the maximal payment scheme is completely determined by the allocation rule. For economy of notation let us introduce

$$t_-^i := \begin{cases} \max\{s^i \in T^i | s^i < t^i\} & \text{if } t^i \neq \min\{s^i \in T^i\}, \\ t^i & \text{otherwise,} \end{cases}$$

and similarly

$$t_+^i := \begin{cases} \min\{s^i \in T^i | s^i > t^i\} & \text{if } t^i \neq \max\{s^i \in T^i\}, \\ t^i & \text{otherwise,} \end{cases}$$

that is, t_-^i is the next smaller type and t_+^i is the next larger type for t^i .

Proposition 2.1. *Let x be monotone. The maximal payments that can be charged in any ex-post incentive compatible and ex-post individual rational mechanism are given by:*

$$\pi_{t^i, t^{-i}}^i = \sum_{s^i \leq t^i} v_{s^i, t^{-i}}^i (x_{s^i, t^{-i}} - x_{s_-^i, t^{-i}}). \quad (7)$$

This is a general result for single-parameter domains (see, for example, Vohra [83]). Observe that for deterministic mechanisms (7) reduces to

$$\pi_{t^i, t^{-i}}^i = \begin{cases} \min\{v_{s^i, t^{-i}}^i | s^i \in T^i, x_{s^i, t^{-i}} = 1\} & \text{if } x_{t^i, t^{-i}} = 1, \\ 0 & \text{otherwise,} \end{cases}$$

which coincides with the upper bound from Lemma 1.6. Recall, that our objective is to maximise the expected total profit, which is the total payment collected minus the cost of allocating the public good. Observe, that in this case the payments should be maximal, hence, without loss of generality, we

can specify the payment scheme for a profit-maximising implementable allocation x as in (7).

Let us define the operator $>$ for vectors as follows: for every $s, t \in T$ we write $s > t$ if and only if $s^i \geq t^i$ for all $i \in N$. Note that $>$ induces a partial order on T .

Proposition 2.2. *Let x be an allocation rule and let $\{x_t\}_{t \in T}$ be the variables induced by x . Then x is monotone if and only if for all $s, t \in T$ such that $s > t$ it holds that $x_s \geq x_t$.*

Proposition 2.2 can be regarded as a special case of the PAD property (positive association of differences) that is implied by monotonicity (see Roberts [61]). A simple proof can be found in the appendix. A set $S \subseteq T$ is called upward closed if for all $s \in S, t \in T$, we have that $t > s$ implies $t \in S$.

Corollary 2.3. *An allocation function in the non-excludable public good setting is implementable in ex-post incentive compatible strategies if and only if the signal tuples for which we allocate the public good form an upward closed set on the partially ordered signal space $(T, >)$.*

2.3 FINDING THE OPTIMAL MECHANISM

In this section we provide an efficient algorithm to solve the integer linear program given by (P1), (EPIC), (EPIR), and (BI). First, we show how to eliminate variables π_t^i , second, we apply Corollary 2.3 to derive an integer linear program that can be solved in polynomial time. We conclude this section by discussing the special case of independent signal distributions and relating our results to the results by Güth and Hellwig [36].

2.3.1 Transformation of the objective function

In order to substitute payment variables with allocation variables, we reformulate the objective function such that we can plug-in (7).

$$\begin{aligned}
\sum_t \varphi_t \sum_i \pi_t^i &= \sum_i \sum_t \varphi_t \pi_t^i \\
&= \sum_i \sum_{t^{-i}} \sum_{t^i} \varphi_{t^i, t^{-i}} \pi_{t^i, t^{-i}}^i \\
&= \sum_i \sum_{t^{-i}} \sum_{t^i} \varphi_{t^i, t^{-i}} \sum_{s^i \leq t^i} v_{s^i, t^{-i}}^i (x_{s^i, t^{-i}} - x_{s^i, t^{-i}}) \\
&= \sum_i \sum_{t^{-i}} \sum_{t^i} x_{t^i, t^{-i}} \left(v_{t^i, t^{-i}}^i \sum_{s^i \geq t^i} \varphi_{s^i, t^{-i}} - v_{t^i, t^{-i}}^i \sum_{s^i \geq t^i_+} \varphi_{s^i, t^{-i}} \right) \\
&= \sum_i \sum_{t^{-i}} \sum_{t^i} x_{t^i, t^{-i}} \left(v_{t^i, t^{-i}}^i \varphi_{t^i, t^{-i}} - \left[v_{t^i_+, t^{-i}}^i - v_{t^i, t^{-i}}^i \right] \sum_{s^i \geq t^i_+} \varphi_{s^i, t^{-i}} \right)
\end{aligned}$$

The coefficient of each $x_{t^i, t^{-i}}$ can be interpreted as the added revenue from agent i if we allocate the public good at reported signal tuple (t^i, t^{-i}) , versus allocating only at (t^i_+, t^{-i}) . From now on we assume that $\varphi_t > 0$ for all t . This assumption facilitates an economic interpretation of the coefficients without making the notation too messy. All results with respect to the algorithmic complexity of our problem hold without this assumption. With the assumption we can further transform the objective function as follows.

$$\begin{aligned}
\sum_t \varphi_t \sum_i \pi_t^i &= \sum_t \varphi_t x_t \sum_i \left(v_{t^i, t^{-i}}^i - \left[v_{t^i_+, t^{-i}}^i - v_{t^i, t^{-i}}^i \right] \frac{\sum_{s^i \geq t^i_+} \varphi_{s^i, t^{-i}}}{\varphi_t} \right) \\
&= \sum_t \varphi_t x_t \sum_i \left(v_{t^i, t^{-i}}^i - \left[v_{t^i_+, t^{-i}}^i - v_{t^i, t^{-i}}^i \right] \frac{\varphi_{t^{-i}} \sum_{s^i \geq t^i_+} \varphi_{s^i | t^{-i}}}{\varphi_t} \right) \\
&= \sum_t \varphi_t x_t \sum_i \left(v_{t^i, t^{-i}}^i - \left[v_{t^i_+, t^{-i}}^i - v_{t^i, t^{-i}}^i \right] \left(\frac{1 - \Phi_{t^i | t^{-i}}}{\varphi_{t^i | t^{-i}}} \right) \right),
\end{aligned}$$

where $\varphi_{t^{-i}} = \sum_{r^i} \varphi_{r^i, t^{-i}}$ and $\varphi_{t^i | t^{-i}}$ is the conditional probability that s^i occurs assuming that the others have t^{-i} and $\Phi_{t^i | t^{-i}}$ is the corresponding con-

ditional cumulative distribution function.

The objective function of (P1) can then be expressed as

$$\begin{aligned} & \sum_t \varphi_t \left(\sum_i \pi_t^i - Cx_t \right) \\ &= \sum_t \varphi_t \left[\sum_i \left(v_{t^i, t^{-i}}^i - \left[v_{t_+, t^{-i}}^i - v_{t^i, t^{-i}}^i \right] \left(\frac{1 - \Phi_{t^i | t^{-i}}}{\varphi_{t^i | t^{-i}}} \right) \right) - C \right] x_t. \end{aligned}$$

Let

$$\mu_{t^i | t^{-i}} = v_{t^i, t^{-i}}^i - \left[v_{t_+, t^{-i}}^i - v_{t^i, t^{-i}}^i \right] \left(\frac{1 - \Phi_{t^i | t^{-i}}}{\varphi_{t^i | t^{-i}}} \right).$$

We call $\mu_{t^i | t^{-i}}$ the *conditional virtual valuation* of agent i having signal t^i while the rest of the agents have signal t^{-i} . In case of independent signal distributions this notion coincides with the virtual valuation introduced by Myerson [54]. Notice that if we interpret $\mu_{t^i | t^{-i}}$ as interdependent valuations of the agents then the objective function

$$\sum_t \varphi_t \left(\sum_i \mu_{t^i | t^{-i}} - C \right) x_t$$

can be regarded as expected welfare maximisation with respect to conditional virtual values. However, from what follows it can be seen that the optimal mechanism does not necessarily take the decision that maximises virtual welfare for each realisation t of signals. Indeed, the profit-maximising mechanism might set $x_t = 1$ even if its coefficient in the objective is negative.

2.3.2 Polynomial time algorithms

As the objective function has become linear, by introducing

$$b_t = \varphi_t \left(\sum_i \mu_{t^i | t^{-i}} - C \right)$$

we can reformulate our initial model (P1) to the following integer linear program to find the expected profit-maximising public good mechanism. Thereby, we make also use of Corollary 2.3.

$$\max \quad \sum_t b_t x_t \quad (\text{P2})$$

subject to

$$\begin{aligned} x_s - x_t &\geq 0 && \forall s, t \in T \text{ s.t. } s > t, \\ x_t &\in \{0, 1\} && \forall t \in T \end{aligned}$$

Notice that if $b_t \geq 0$ for some t implies that $b_s \geq 0$ for all $s > t$, then the solution is trivial: set $x_t = 1$ if and only if $b_t \geq 0$. In the general case, the following observation provides us with a polynomial time algorithm.

Proposition 2.4. *The incentive compatibility constraints of program (P2) form a totally unimodular matrix.*

Proof The coefficient matrix of the incentive compatibility constraints can be seen as the node-arc incidence matrix of the directed comparability graph of the partially ordered set $\{x_t\}_{t \in T}$. There is a directed arc between each comparable variable directed from the variable related to the larger signal vector to the one corresponding to the smaller signal vector. In the node-arc incidence matrix for each arc the start node is indicated by 1, whilst the end node is indicated by -1 , just like in the coefficient matrix of the constraints. The incidence matrix of any directed graph is totally unimodular (Schrijver [72]). \square

Corollary 2.5. *The linear program relaxation of (P2) has an integer optimal solution, hence the problem can be solved in time polynomial in $|N|$ and $|T|$. Moreover, it means that randomised algorithms cannot do better than deterministic ones.*

We can view (P2) also as a maximal closure problem, which is defined by Picard [59] as follows.

Definition 2.6 (Maximal closure problem). Given a directed graph $G = (V, A)$ where V is the set of nodes and A the set of arcs, a closure of G is

defined as a subset of vertices Y such that if a vertex belongs to Y then all its successors belong also to Y . If each vertex is associated with a weight (real number), then a maximal closure Y^* of G is defined as a closure of maximal sum of weights.

Picard [59] provides a combinatorial algorithm to find the optimal solution. It is done by reducing the problem to a network flow instance. First, he gives an equivalent quadratic formulation of the constraints. Then, it turns out that solving the Lagrangian dual of this construction is essentially the same task as finding the minimum cut of an associated network, which can be solved in polynomial time by maximising the flow on the related network.

Maximal closure problems are well studied by virtue of their connection to a wide range of practical problems from open pit mining, record segmentation in large shared data bases to portfolio selection under contingency constraints. For more examples see Faaland et al. [27] and Hochbaum and Chen [39]. The latter papers also provide an overview and comparison of network based algorithms and heuristics used in practice to deal with maximum closure problems.

As a summary of our discussion we get the following theorem.

Theorem 2.7. *Finding the profit-maximising, deterministic, ex-post individual rational and ex-post implementable mechanism for the provision of a single, non-excludable public good can be done in time polynomial in the number of agents and the number of possible signal tuples $(|N|, |T|)$.*

Güth and Hellwig [36] characterise the profit-maximising Bayes-Nash implementable mechanism for the case that signals are distributed independently and there are no informational externalities. The next subsection compares their result to ours.

2.3.3 Independent valuations and signal distributions

Suppose, that valuations are independent and the signals are independently distributed, i. e., for all $t \in T$ and for all $i \in N$ we have that $\varphi_t = \prod_i \varphi_{t^i}$,

and for all $i \in N$, for all $t^i \in T^i$ and for all $t^{-i}, s^{-i} \in T^{-i}$ we have that $v_{t^i, t^{-i}}^i = v_{t^i, s^{-i}}^i := t^i$. Readily, in this case for all t^{-i}

$$\mu_{t^i | t^{-i}} = \mu_{t^i} = t^i - \left[t_+^i - t^i \right] \left(\frac{1 - \Phi_{t^i}}{\varphi_{t^i}} \right).$$

This implies that we can express the coefficients of the objective function in terms of virtual valuations in the Myerson sense:

$$b_t = \varphi_t \left(\sum_i \mu_{t^i} - C \right).$$

Observe that b_t is non-negative if and only if the sum of virtual valuations at t covers the cost. Moreover, assuming signal distributions with monotone non-decreasing hazard rate $\frac{\varphi_{t^i}}{1 - \Phi_{t^i}}$ and valuations where $t_+^i = t^i + 1$, the virtual valuations along with the sum of virtual valuations become monotone non-decreasing in t . This implies that if there is a t such that $b_t \geq 0$, then for all $s \geq t$ we have that $b_s \geq 0$. Hence, the optimal solution can be characterised as follows: set $x_t = 1$ if and only if $\sum_i \mu_{t^i} \geq C$. This coincides with the findings of Güth and Hellwig [36], which implies that Bayes-Nash implementation leads to the same optimal mechanism in case of independent signal distributions. We elaborate on this observation in further details in Section 2.4.1.

In case virtual valuations are not monotone non-decreasing, we can use the *ironing* technique developed by Myerson [54]. It works as follows. Since valuations and signals are independent, we can treat the expected revenue yielded by each agent separately. For agent i we construct a *revenue curve* $R^i : T^i \rightarrow \mathbb{R}$, such that $R^i(t^i) = \sum_{s^i \geq t^i} \mu_{s^i}$. It is easy to see that the virtual valuations are the first derivatives of the revenue curve. This implies that when virtual valuations are increasing the corresponding part of the revenue curve is concave. By taking derivatives of the convex hull of the revenue curve, we arrive at “ironed” or monotone non-decreasing virtual valuations, hence the optimal solution can be characterised easily again. Note that the construction of the revenue curve and its convex hull takes time polynomial in the number of individual types, so not just in the number of type profiles.

2.4 OTHER SOLUTION CONCEPTS

In Section 2.3 we have shown how to find the optimal deterministic, ex-post individual rational, ex-post implementable mechanism. Corollary 2.5 tells us that allowing for randomised mechanisms does not increase the expected profit. In this section we investigate what happens if we relax ex-post incentive compatibility and ex-post individual rationality. First, we analyse Bayes-Nash implementation as a less demanding alternative for ex-post implementation, then we focus on the implications of interim individual rationality.

Definition 2.8. A direct mechanism, represented by its induced variables $\{x_t\}_{t \in T}$ and $\{\pi_t^i\}_{i \in N, t \in T}$, is Bayes-Nash incentive compatible if and only if for all $i \in N$ and $t^i, s^i \in T^i$ we have that

$$\sum_{t^{-i}} \varphi_{t^{-i}|t^i} \left(v_{t^i, t^{-i}}^i x_{t^i, t^{-i}} - \pi_{t^i, t^{-i}}^i \right) \geq \sum_{t^{-i}} \varphi_{t^{-i}|t^i} \left(v_{t^i, t^{-i}}^i x_{s^i, t^{-i}} - \pi_{s^i, t^{-i}}^i \right). \quad (\text{BNIC})$$

Definition 2.9. A direct mechanism, represented by its induced variables $\{x_t\}_{t \in T}$ and $\{\pi_t^i\}_{i \in N, t \in T}$, is interim individual rational if and only if for all $i \in N$ and $t^i \in T^i$ we have that

$$\sum_{t^{-i}} \varphi_{t^{-i}|t^i} \left(v_{t^i, t^{-i}}^i x_{t^i, t^{-i}} - \pi_{t^i, t^{-i}}^i \right) \geq 0. \quad (\text{IIR})$$

Definition 2.10. A direct mechanism, represented by its induced variables $\{x_t\}_{t \in T}$ and $\{\pi_t^i\}_{i \in N, t \in T}$, is dominant strategy incentive compatible if and only if for all $i \in N$, for all $t^i, s^i \in T^i$ and for all $t^{-i}, s^{-i} \in T^{-i}$ we have that

$$v_{t^i, t^{-i}}^i x_{t^i, s^{-i}} - \pi_{t^i, s^{-i}}^i \geq v_{t^i, t^{-i}}^i x_{s^i, s^{-i}} - \pi_{s^i, s^{-i}}^i. \quad (\text{DSIC})$$

Remember that when valuations are not interdependent, i. e., for all $i \in N$, for all $t^i \in T^i$ and for all $t^{-i}, s^{-i} \in T^{-i}$ we have that $v_{t^i, t^{-i}}^i = v_{t^i, s^{-i}}^i$, then (DSIC) coincides with (EPIC), hence ex-post implementation is equivalent to dominant strategy implementation.

2.4.1 *Bayes-Nash implementation*

Recent papers provide conditions under which Bayes-Nash and dominant strategy implementation are equivalent. They study environments without informational externalities where signals are one-dimensional and are distributed independently, and valuations for outcomes are linear functions of agents' signals. Gershkov et al. [31] show that in case of two outcomes for any Bayes-Nash incentive compatible mechanism there is a dominant strategy incentive compatible mechanism that generates the same interim probabilities with which each alternative is selected. The paper provides an example for a three outcome case where the equivalence breaks down. A positive result for environments with more than two outcomes is given in Manelli and Vincent [49], who prove the equivalence for not necessarily symmetric, single-item auction settings. If we define equivalence in a weaker sense, i.e. two implementations are equivalent if they generate the same interim utilities for each agent, then according to Gershkov et al. [30] Bayes-Nash and dominant strategy implementation remains equivalent in environments with more than two outcomes. All these results require that the signal space is single-dimensional.

Corollary 2.11. *In the profit-maximising non-excludable public good problem where signals are independently distributed and valuations are not interdependent, we can restrict ourselves to ex-post implementation without loss of generality.*

Proof Observe that this problem is a special case of the setting discussed in Gershkov et al. [31], which implies that for every Bayes-Nash incentive compatible mechanism there is a dominant strategy incentive compatible one that provides the same interim probabilities of providing the public good. In this case we can construct transfers for the dominant strategy implementable mechanism using the payments of the Bayes-Nash implementable mechanism in such a way that the expected profits of the two mechanisms are equal (see Gershkov et al. [30]). As in the absence of informational externalities ex-post implementation is equivalent to dominant strategy implementation, this completes the proof. \square

This result does not extend to the case of correlated signal distributions as we can see from the following example.

Example 2.12. Let $N = \{1, 2\}$, $T^1 = T^2 = \{1, 2\}$, $C = 1$, and $\varphi_{1,1} = 0.4$, $\varphi_{1,2} = 0.1$, $\varphi_{2,1} = 0.1$, $\varphi_{2,2} = 0.4$. Furthermore, let $v_{t^i}^i = t^i$ for all $i \in N$ and for all $t^i \in T^i$. Solving this problem with our approach yields the optimal solution given in the following table.

t	x_t	π_t^1	π_t^2
(1, 1)	0	0	0
(1, 2)	0	0	0
(2, 1)	0	0	0
(2, 2)	1	2	2

This mechanism yields an expected profit of 1.2, which is by definition the highest expected profit that any ex-post individual rational and ex-post implementable mechanism can attain in this setting. Next, consider the following mechanism.

t	x_t	π_t^1	π_t^2
(1, 1)	1	1	1
(1, 2)	0	0	0
(2, 1)	0	0	0
(2, 2)	1	1.75	1.75

The allocation rule of this mechanism is not monotone, hence the mechanism is not ex-post incentive compatible. On the other hand, it is Bayes-Nash incentive compatible, and has an expected profit of 1.4.

Thus, in case of correlated signal distributions there can be Bayes-Nash implementable mechanisms for the non-excludable public good problem that generate higher profit than any mechanism implementable in ex-post incentive compatible strategies.

Example 2.21 in the appendix demonstrates that the equivalence breaks down as well when the distributions of signals are independent, but the valuations exhibit informational externalities.

2.4.2 *Interim individual rationality*

Interim individual rationality constraints can be modelled as Bayes-Nash incentive compatibility constraints with dummy signals and the same can be done with ex-post individual rationality and ex-post incentive compatibility constraints (see Vohra [83] for details). Corollary 2.11 tells us that the optimal mechanisms yield the same profit for both notions of implementation when signals are distributed independently and there are no informational externalities. Hence we can conclude the following.

Corollary 2.13. *In the profit-maximising, non-excludable public good problem where signals are independently distributed and in the absence of informational externalities we can restrict ourselves to interim individual rationality without loss of generality.*

Example 2.12 shows that in case of correlated signal distributions the equivalence of Bayes-Nash and dominant strategy implementation does not hold, thus we can as well expect that interim individual rationality provides us more flexibility to elicit payments. This is indeed the case as the following example shows.

Example 2.14. Consider the following mechanism for the setting in Example 2.12.

t	x_t	π_t^1	π_t^2
(1,1)	1	2/3	2/3
(1,2)	1	$2\frac{1}{3}$	2/3
(2,1)	1	2/3	$2\frac{1}{3}$
(2,2)	1	$2\frac{1}{3}$	$2\frac{1}{3}$

Observe that the proposed mechanism is not ex-post individual rational, but ex-post incentive compatible, interim individual rational and yields an expected profit of 2. Therefore it extracts the whole social surplus, which is not possible when we require ex-post individual rationality. (Indeed, the second mechanism from Example 2.12 has been verified by a computer program to be optimal.)

A key result concerning the limitations of ex-post individual rationality compared to interim individual rationality is due to Crémer and McLean [18]. They deal with single-item auction settings with no informational externalities. Given a fixed signal space, they provide conditions on the joint distribution of these signals that allow full extraction of the agents' surplus by interim individual rational mechanisms which are either Bayes-Nash or dominant strategy implementable. When these conditions are met, full surplus extraction is achievable through a Vickrey auction with modified payments. McAfee and Reny [52] show that these results go well beyond the auction environment. To explain the conditions in our context, consider an agent i and let Γ^i denote the $T^i \times T^{-i}$ matrix whose elements are the conditional probabilities of the joint distribution function, that is, $\Gamma_{t^i, t^{-i}}^i = \varphi_{t^{-i}|t^i}$. For dominant strategy implementation the condition is that for each agent i the rows of Γ^i have to be linearly independent. For Bayes-Nash implementation the condition is milder: it requires only the convex independence of the rows of Γ^i (i. e., no row is contained in the convex hull of the other rows)². For a mechanism given by its induced variables $\{x_t\}_{t \in T}$ and $\{\pi_t^i\}_{t \in T, i \in N}$, let us define the expected social surplus as $\sum_t \varphi_t x_t (\sum_i v_t^i - C)$. The expected profit of the public good provider is always less than or equal to the expected social surplus due to (IIR). Therefore the objective value of a mechanism that maximises the expected surplus is an upper bound on the maximal expected profit that can be achieved by any IIR mechanism. If a mechanism yields profit equal to this upper bound, we say that the mechanism extracts the full social surplus.

An allocation of the public good is efficient if it allocates when the sum of valuations exceeds the cost of provision and retains the good otherwise. It is a folklore result that the expected social surplus is maximised by efficient mechanisms (see, for example, Nisan et al. [56]). Such mechanisms are monotone when there are no informational externalities, but if valuations are interdependent this is not generally true demonstrated by the following small example.

Example 2.15. Let $N = \{1, 2\}$, $T^1 = \{s^1, t^1\}$, $T^2 = \{s^2, t^2\}$, $t^i < s^i$ for all i and $C = 5.5$. Assume that the signals are distributed independently and

² Note that Crémer and McLean [18] require conic independence, but as the sum of the values in each row of Γ^i is equal to 1 the sum of the non-negative coefficients should be 1 as well, hence the condition can be strengthened to convex independence.

uniformly. Let the valuations be defined as in table 1a and 1b. Observe that

$T^1 \setminus T^2$	s^2	t^2
s^1	3	1
t^1	4	2

(a) Agent 1

$T^1 \setminus T^2$	s^2	t^2
s^1	3	4
t^1	1	2

(b) Agent 2

Table 1: Valuations

both valuations are increasing in their own signals, but decreasing in the other ones. The efficient allocation rule would be to allocate if and only if agents have signals (s^1, s^2) , but this wouldn't be monotone.

This phenomenon raises concern because in case of ex-post implementation the techniques of Crémer and McLean [18] only tells us how to extract the full social surplus, given the surplus maximising allocation rule is implementable. When valuations are interdependent we have to actually find the expected surplus maximising *monotone* allocation rule. For ex-post implementation we demonstrate below how to adapt our model to find this allocation rule. In case of Bayes-Nash implementation such problems do not occur as it is possible to construct payments via lotteries that enforce Bayes-Nash incentive compatibility for any allocation rule. Therefore two separate theorems are formulated in the following: one for ex-post and one for Bayes-Nash implementation.

Theorem 2.16. *Take an instance of the profit-maximising, non-excludable public good problem. Assume that the signal distribution satisfies the following: for each $i \in N$ there does not exist a vector of $\{\lambda_{t^i}\}_{t^i \in T^i}$, not all equal to zero, such that*

$$\sum_{t^i \in T^i} \lambda_{t^i} \varphi_{t^{-i}|t^i} = 0 \quad \forall t^{-i} \in T^{-i}. \quad (8)$$

Then the profit-maximising, ex-post implementable and interim individual rational mechanism can be constructed in time polynomial in $|T|$ and $|N|$. Moreover, the optimal mechanism is deterministic and extracts the full social surplus.

Theorem 2.17. *Take an instance of the profit-maximising, non-excludable public good problem. Assume that the signal distribution satisfies the following: for each $i \in N$ and each $t^i \in T^i$ there does not exist a vector of non-negative $\{\lambda_{s^i}\}_{s^i \in T^i \setminus t^i}$ such that*

$$\sum_{s^i \in T^i \setminus t^i} \lambda_{s^i} = 1 \quad (9)$$

and

$$\varphi_{t^{-i}|t^i} = \sum_{s^i \in T^i \setminus t^i} \lambda_{s^i} \varphi_{t^{-i}|s^i} \quad \forall t^{-i}. \quad (10)$$

Then the profit-maximising, Bayes-Nash implementable and interim individual rational mechanism can be constructed in time polynomial in $|T|$ and $|N|$. Moreover, the optimal mechanism is deterministic and extracts the full social surplus.

The proofs for both theorems are based on mild modifications of the construction given in Crémer and McLean [18] and the observation that these constructions can be done in polynomial time. We provide them in the appendix.

We emphasise that verifying whether a distribution satisfies the conditions of either Theorem 2.16 or Theorem 2.17 can be done in polynomial time. The linear independence of the conditional probability matrices can be verified by Gaussian elimination, whilst the question of convex independence can be reduced to a feasibility check of a series of linear programs (one program per row).

In terms of surplus extraction, the assumptions in Theorem 2.16 and Theorem 2.17 are not just sufficient, but also necessary in the following sense. Whenever a joint signal distribution fails condition (8) or (9) and (10) we can construct a T such that full surplus extraction is not possible in the corresponding settings (a proof can be given by modifying the proof for auctions by Crémer and McLean [18], but is omitted here). Independent distributions violate the assumptions of both theorems, which explains partly why we cannot achieve more profit with interim individual rationality than with ex-post individual rationality under such distributions. It remains an open problem

how to efficiently compute the optimal deterministic, interim individual rational mechanism for general dependent distributions.

One should note that the possibility of full value extraction in this setting is rather a technical result than an economic insight. The discussion in [18] points out that the ex-post utility of agents can be significantly negative for some signal tuples due to the large payments, which could discourage participants in a “real life” situation. Moreover, the construction of the lotteries is highly dependent on the assumption that every detail of the signal distribution is common knowledge. In practice agents can have different sets of information or they might not even agree on the interpretation of the same signal.

2.5 EXCLUDABLE PUBLIC GOODS

In this section we show how to extend our analysis to the excludable public good problem. We take the model of the non-excludable public good as starting point and highlight only the differences. The closest analysis to ours of this problem was performed by Schmitz [70] aiming for Bayes-Nash incentive compatibility and interim individual rationality assuming independent signal distributions that satisfy the increasing hazard rates condition. Their results were analogue to the ones of Güth and Hellwig [36] for the non-excludable public good. We will see that as in the case of the latter one, our model can reproduce the findings of Schmitz [70].

2.5.1 Integer linear program formulation

To keep our model capable of capturing general randomised mechanism as well, we introduce decision variables for all possible outcomes. Hence for all $t \in T$ and for all $S \subseteq N$ let x_t^S denote the probability that the group of agents S is provided with the public good in case of type report t . We assume that agents only care about their own probability of consuming the public good. Similarly to the non-excludable case, we also introduce decision variables for each payment, i. e., for all i and t we have that π_t^i denotes the quantity that agent i has to pay in case of signal report t . The production of the public

good costs C and further provision cost c^i occurs for each i for whom it is provided. We are going to make use of dummy variables $\{y_t\}_{t \in T}$ to express the production cost in a linear manner. Then our mechanism design problem can be formulated as the following integer linear program.

$$\max \sum_{t \in T} \varphi_t \left[\left(\sum_i \pi_t^i - \sum_{S \subseteq N, i \in S} x_t^S c^i \right) - y_t C \right] \quad (\text{P3})$$

subject to

$$v_{t^i, t^{-i}}^i \sum_{S \subseteq N, i \in S} x_{t^i, t^{-i}}^S - \pi_{t^i, t^{-i}}^i \geq v_{t^i, t^{-i}}^i \sum_{S \subseteq N, i \in S} x_{s^i, t^{-i}}^S - \pi_{s^i, t^{-i}}^i \quad \forall i, t^{-i}, s^i, t^i, \quad (\text{EPIC2})$$

$$v_t^i \sum_{S \subseteq N, i \in S} x_t^S - \pi_t^i \geq 0 \quad \forall i, t, \quad (\text{EPIR2})$$

$$y_t \geq \sum_{S \subseteq N} x_t^S \quad \forall t, \quad (\text{DU})$$

$$1 \geq \sum_{S \subseteq N} x_t^S \quad \forall t, \quad (\text{PROB})$$

$$x_t^S \in \{0, 1\} \quad \forall S, t. \quad (\text{BI2})$$

As before, constraints (EPIC2) ensure implementability in ex-post incentive compatible strategies, while constraints (EPIR2) are responsible for ex-post individual rationality, finally, (BI2) forces the allocation to be deterministic. The additional inequalities of (DU) make sure that the dummy variables are connected in the right way to variables x_t^S , while constraints (PROB) imply that the public good is provided with probability at most 1.

As agent i cares only about the total probability of providing the public good, we introduce $x_t^i = \sum_{S \subseteq N, i \in S} x_t^S$ to shorten the notation. By invoking the same reasoning as for non-excludable public goods, we can replace constraints (EPIC2) and (EPIR2) with the following monotonicity constraints:

$$x_{t^i, t^{-i}}^i \geq x_{s^i, t^{-i}}^i \quad \forall i, t^{-i}, t^i, s^i, t^i > s^i.$$

Furthermore, the payment variables can be expressed as

$$\pi_{t^i, t^{-i}}^i = \sum_{s^i \leq t^i} v_{s^i, t^{-i}}^i (x_{s^i, t^{-i}}^i - x_{s^i, t^{-i}}^i). \quad (11)$$

By taking these into account and repeating the analysis of Section 2.3 with minor modifications, we can rewrite the objective function in the following succinct form:

$$\sum_t \varphi_t \left[\sum_i x_t^i b_t^i - y_t C \right],$$

where

$$b_t^i = v_{t^i, t^{-i}}^i - \left[v_{t^i, t^{-i}}^i - v_{t^i, t^{-i}}^i \right] \left(\frac{1 - \Phi_{t^i | t^{-i}}}{\varphi_{t^i | t^{-i}}} \right) - c^i.$$

Consequently we can transform (P3) into the following integer program:

$$\max \quad \sum_t \varphi_t \left[\sum_i x_t^i b_t^i - y_t C \right] \quad (P4)$$

subject to

$$x_{s^i, t^{-i}}^i \geq x_{t^i, t^{-i}}^i \quad \forall i, t^{-i}, t^i, s^i, s^i > t^i$$

$$y_t \geq \sum_{S \subseteq N} x_t^S \quad \forall t, \quad (DU)$$

$$1 \geq \sum_{S \subseteq N, i \in S} x_t^S \quad \forall t, \quad (PROB)$$

$$x^i = \sum_{S \subseteq N, i \in S} x_t^S \quad \forall i, t$$

$$x_t^S \in \{0, 1\} \quad \forall S, t.$$

2.5.2 A formulation having totally unimodular constraint matrix

Now, we consider another integer program formulation that uses only $\{x_t^i\}_{t \in T}$ variables to describe the allocation rule.

$$\max \quad \sum_t \varphi_t \left[\sum_i x_t^i b_t^i - y_t C \right] \quad (\text{P5})$$

subject to

$$\begin{aligned} x_{s^i, t^{-i}}^i &\geq x_{t^i, t^{-i}}^i && \forall i, t^{-i}, t^i, s^i, s^i > t^i \\ y_t &\geq x_t^i && \forall i, t \\ x_t^i &\in \{0, 1\} && \forall i, t. \end{aligned} \quad (\text{DU2})$$

Note, that the only real difference is constraints (DU) and (DU2). It is apparent that any feasible solution for (P4) induces a feasible solution for (P5) that yields the same objective value. Conversely, any feasible solution for (P5) induces a feasible solution for (P4) that yields the same objective value. For us it is also consequential to establish the following relation between the linear relaxations of the two integer program.

Lemma 2.18. *Any fractional solution for the linear relaxation of (P4) induces a feasible fractional solution for the linear relaxation (P5).*

Proof We only have to check whether constraints (DU2) hold. Let $\{\hat{y}_t\}_{t \in T}$, $\{\hat{x}_t^S\}_{t \in T, S \subseteq N}$ and $\{\hat{x}_t^i\}_{i \in N, t \in T}$ represent a fractional solution for (P4). It follows from (DU) and the definition of \hat{x}_t^i that

$$\hat{y}_t \geq \sum_{S \subseteq N} \hat{x}_t^S \geq \max_i \{\hat{x}_t^i\},$$

which concludes that constraints (DU2) are satisfied as well. \square

Lemma 2.18 implies that the optimal value of the linear relaxation of (P5) is at least as much as the optimal value of the linear relaxation of (P4).

The advantage of formulation (P5) is that its constraint matrix can be considered as a node-arc incidence matrix of a comparability graph. It is because the monotonicity constraints define a partial ordering on the x variables,

while constraints (DU2) orders the y and x variables. Therefore the resulting constraint matrix is totally unimodular. This together with previous observations implies that there is an optimal integer solution for (P5), which in turn induces an optimal integer solution for (P4) that yields the same objective value. Consequently solving the linear relaxation of (P5) instead of optimising over (P3) is no loss of generality.

Furthermore, it is easy to see that there is no vital difference in the combinatorial structure of the excludable or the non-excludable public good problem. Therefore it can be shown that corollaries 2.5, 2.11 and theorems 2.16, 2.17 hold for the excludable public good problem as well.

2.5.3 From public to private goods

Observe that adding the simple constraints $\sum_i x_t^i \leq 1, \forall t \in T$ makes the excludable public good problem a single-item auction problem. Now, we leverage the model developed for excludable public goods to provide an independent proof of the result in Papadimitriou and Pierrakos [57], which states that the optimal single-item auction problem for two agents is solvable in polynomial time. Therefore consider (P5) with the additional constraint of having two agents and by assuming that $c^1 = c^2 = C = 0$.

$$\begin{aligned} \max \quad & \sum_t \varphi_t \left[x_t^1 b_t^1 + x_t^2 b_t^2 \right] && \text{(AUC)} \\ \text{subject to} \quad & && \\ & x_{s^1, t^2}^1 \geq x_{t^1, t^2}^1 && \forall t^1, t^2, s^1, s^1 > t^1 \\ & x_{t^1, s^2}^2 \geq x_{t^1, t^2}^2 && \forall t^1, t^2, s^2, s^2 > t^2 \\ & 1 \geq x_{t^1, s^2}^2 + x_{t^1, s^2}^1 && \forall t^1, t^2 \\ & x_t^1, x_t^2 \in \{0, 1\} && \forall t. \end{aligned}$$

Theorem 2.19 (Ghouila-Houri [33]). *Matrix A is totally unimodular if and only if any subset of its columns can be partitioned to J_1 and J_2 such that for each row i we have $|\sum_{j \in J_1} A_{i,j} - \sum_{j \in J_2} A_{i,j}| \leq 1$.*

Theorem 2.20 (Papadimitriou and Pierrakos [57]). *Given a discrete joint valuation distribution φ for two bidders, the optimal EPIC and EPIR deterministic auction can be computed in time polynomial in $|T|$.*

Proof The constraint matrix of (AUC) is totally unimodular and it consists of a polynomial number of constraints. To see this let J be a subset of the columns of the constraint matrix. Then we can define a partition satisfying Theorem (2.19) by letting J_1 and J_2 contain the columns related to agent 1 and agent 2 respectively. \square

For the case of three or more agents, the matrix is not totally unimodular anymore, which is consistent with the finding of Papadimitriou and Pierrakos [57] that the optimal auction problem is NP-hard.

2.6 CONCLUSION

We have studied the problem of finding the optimal, deterministic mechanism for providing a non-excludable public good when we allow for correlation in the signal distribution as well as for informational externalities. We considered ex-post as well as Bayes-Nash implementation, and ex-post as well as interim individual rational mechanisms. The results and open questions can be summarised as follows:

1. The optimal ex-post incentive compatible, ex-post individual rational mechanism can be computed in polynomial time.
2. The optimal ex-post incentive compatible, interim individual rational mechanism can be computed in polynomial time if the conditions of Theorem 2.16 hold. Example 2.14 shows that it may yield strictly higher profit than the optimal mechanism for requirements as in case (1). Also, the optimal Bayes-Nash incentive compatible, interim individual rational mechanism can be computed in polynomial time if the conditions of Theorem 2.17 hold.
3. The computation of the optimal deterministic Bayes-Nash incentive compatible, ex-post individual rational mechanism remains an open problem. Example 2.12 and 2.21 shows that it may yield strictly higher

profit for correlated signal distributions or for interdependent valuations.

4. The optimal deterministic Bayes-Nash incentive compatible, interim individual rational mechanism can be computed in polynomial time if the conditions of Theorem 2.17 hold. It follows from Crémer and McLean [18] that the optimal mechanism can yield strictly higher profit for correlated signal distributions. It remains an open problem how to efficiently compute the optimal deterministic mechanism for general distributions.

In each of the cases where we are able to show polynomial-time solvability, randomised mechanisms cannot yield more profit than deterministic mechanisms.

Finally, we have shown that our results can be extended to the excludable public good problem.

APPENDIX

Proof of Proposition 2.2. The “if” part follows from the fact that the condition of the theorem is a generalisation of the monotonicity constraints.

To prove the “only if” part, take any $t, s \in T$ such that $s > t$. Let I be the index set for dimensions i where $s^i > t^i$. Let us assume, without loss of generality, that s and t differs in the first m dimensions, so $I = \{1, \dots, m\}$. Consider the series of signals $(t_{(0)}, \dots, t_{(m)})$, where $t_{(i)} = (s^1, s^2, \dots, s^i, t^{i+1}, \dots, t^m, \dots)$. Clearly, $t_{(0)} = t$ and $t_{(m)} = s$. For all i by assumption and by the construction it holds that $t_{(i)}^j = t_{(i-1)}^j$ for $j \neq i$ and $t_{(i)}^i > t_{(i-1)}^i$. Monotonicity implies that $x_{t_{(i)}} \geq x_{t_{(i-1)}}$, hence

$$x_s = x_{t_{(m)}} \geq x_{t_{(m-1)}} \geq \dots \geq x_{t_{(0)}} = x_t.$$

□

Proof of Theorem 2.16. First, we have to find the expected surplus maximising allocation rule that is ex-post incentive compatible, i. e., monotone. The objective of expected surplus maximisation can be written as

$$\sum_t \varphi_t x_t \left(\sum_i v_t^i - C \right) = \sum_t b_t x_t,$$

where $b_t = \varphi_t (\sum_i v_t^i - C)$. Observe that the objective function in this form coupled with monotonicity and binary constraints of the $\{x_t\}_{t \in T}$ variables constitutes exactly the integer linear program P2, which can be solved in polynomial time.

According to Theorem 2.1, a payment scheme p defined as in (7) implements any monotone allocation rule and satisfies ex-post individual rationality. Now, we how to modify these payments $\pi_{t^i, t-i}^i$ by adding terms $g_{t^i}^i$ such that the resulting mechanism is interim individual rational and extracts full surplus in expectation. Such a modification does not harm ex-post incentive compatibility as the terms do not depend on t^i .

Introduce for all $i \in N$ and for all $t^i \in T^i$

$$h_{t^i}^i = \sum_{t^{-i}} \varphi_{t^{-i}|t^i} \left(v_{t^i, t^{-i}}^i x_{t^i, t^{-i}} - \pi_{t^i, t^{-i}}^i \right)$$

to measure the expected utility of agent i when he possesses signal t^i . Condition (8) ensures that Γ^i has full row rank, hence there exist vectors $(g_{t^{-i}}^i)_{t^{-i} \in T^{-i}}$ such that for all $t^i \in T^i$ we have that

$$\sum_{t^{-i} \in T^{-i}} \varphi_{t^{-i}|t^i} g_{t^{-i}}^i = h_{t^i}^i.$$

Now, modify payments p as follows: $\hat{p}_{t^i, t^{-i}}^i = \pi_{t^i, t^{-i}}^i + g_{t^{-i}}^i$. By definition

$$\sum_{t^{-i}} \varphi_{t^{-i}|t^i} \left(v_{t^i, t^{-i}}^i x_{t^i, t^{-i}} - \hat{p}_{t^i, t^{-i}}^i \right) = 0,$$

that is, the expected utility of each agent is 0. Thus the expected profit of the mechanism designer equals social surplus. By that it attains the upper bound on the maximal profit, hence the mechanism is optimal. Observe that

the allocation could be constructed in polynomial time due to the previous arguments and the payment rule can be determined in polynomial time as well using Gaussian elimination for instance. \square

Proof of Theorem 2.17 Let $x_t = 1$ if and only if $\sum_{i \in N} v_t^i \geq C$, i. e., take the allocation rule that maximises the expected surplus. Next, we define a payment scheme that enforces Bayes-Nash incentive compatibility under this allocation rule.

By the assumption, none of the rows of the matrices Γ^i is contained in the convex hull of the other rows of the same matrix. Since the sum of each row is 1, a row that is not in the convex hull of the other rows, is not contained in their cone either. Therefore due to the Separating Hyperplane Theorem, for each row there is a hyperplane, which strictly separates that row from the convex hull of the other rows, moreover the norm vector of that separating hyperplane is perpendicular to the row vector. Formally, for each $i \in N$ and $t^i \in T^i$ there exists a vector $(g_{t^i, t^i}^i)_{t^i \in T^i}$ such that:

$$\sum_{t^i \in T^i} \varphi_{t^i | t^i} g_{t^i, t^i}^i = 0$$

and

$$\sum_{t^i \in T^i} \varphi_{t^i | s^i} g_{t^i, t^i}^i > 0, \quad \forall s^i \neq t^i.$$

Now, define payments as

$$\pi_t^i = \sum_{t^i} \varphi_{t^i | t^i} v_{t^i, t^i}^i x_{t^i, t^i} + M g_t^i$$

for some large enough M . The construction of g_t^i implies that

$$\sum_{t^i} \varphi_{t^i | t^i} (v_{t^i, t^i}^i x_{s^i, t^i} - \pi_{s^i, t^i}^i) = 0$$

if $s^i = t^i$. Otherwise it is arbitrary negative, which ensures that the interim individual rationality constraints (IIR) are binding and the Bayes-Nash incentive compatibility constraints (BNIC) are satisfied. We get also that the mechanism extracts full social surplus, hence it is optimal.

We next discuss the computational complexity of constructing the optimal mechanism. The allocation function is given, thus it takes constant time to determine it. To calculate the payments, first, we have to find the collection $\{g_t^i\}_{t \in T}$, then set an appropriate M . The latter can be done in linear time, e. g., take

$$M = \max_{i, t^i} \left\{ \sum_{t^{-i}} \varphi_{t^{-i}|t^i} v_{t^i, t^{-i}}^i x_{t^i, t^{-i}} \right\} / \min_{i, t^i} \left\{ \sum_{t^{-i} \in T^{-i}} \varphi_{t^{-i}|t^i} g_{t^{-i}}^i \right\}.$$

For finding appropriate $\{g_t^i\}_{t \in T}$ for each agent i and each $t^i \in T^i$ in polynomial time one can use the following linear program.

$$\begin{aligned} \min \quad & \sum_{t^{-i} \in T^{-i}} \varphi_{t^{-i}|t^i} g_{t^i, t^{-i}}^i \\ \text{subject to} \quad & \sum_{t^{-i} \in T^{-i}} \varphi_{t^{-i}|s^i} g_{t^i, t^{-i}}^i > 0, \quad \forall s^i \neq t^i. \end{aligned}$$

This completes the proof. □

Example 2.21. Let $N = \{1, 2\}$, $T^i = \{a^i, b^i, c^i, d^i\}$, $a^i < b^i < c^i < d^i$ for all $i \in N$ and $C = 18.5$. We assume independent signal distributions, where $\varphi_{a^i} = 0.4$ and $\varphi_{b^i} = \varphi_{c^i} = \varphi_{d^i} = 0.2$ for all $i \in N$. The valuations are defined according to table 2. Solving this problem with our approach yields the optimal solution given in table 3a. The corresponding payments (tables 4a and 5a) have an expected value of 13.08, which results in an expected profit of 3.46.

Now, consider the allocation given in table 3b. It violates monotonicity, hence it is not EPIC, but it can be verified that together with the ex-post individual rational payments defined as in tables 4b and 5b it is Bayes-Nash implementable. This allocation and the corresponding payment rule achieve an expected revenue of 16.112 and an expected profit of 3.532, which is higher than the profit earned from the best ex-post individual rational and EPIC mechanism.

$T^1 \setminus T^2$	a^2	b^2	c^2	d^2	$T^1 \setminus T^2$	a^2	b^2	c^2	d^2
a^1	0.0	0.0	0.0	13.0	a^1	10.0	0.0	0.0	12.5
b^1	0.0	12.0	12.0	12.0	b^1	0.0	12.0	12.0	12.0
c^1	0.0	13.0	13.0	13.0	c^1	0.0	12.0	12.0	13.0
d^1	13.0	13.0	13.0	13.0	d^1	13.0	13.0	13.0	13.0

(a) EPIC (b) BNIC

Table 5: Payment rules for agent 2

Part III

APPROXIMATIONS

“Everything should be made as simple as possible, but not simpler.”

— Albert Einstein

LEAD MARKETING AUCTIONS WITH PREMIUM FOR EXCLUSIVE POSSESSION

This chapter is based on the following paper:

Hemant K. Bhargava, Gergely Csapó and Rudolf Müller, “Lead Marketing Auctions with Premium for Exclusive Possession”.

3.1 INTRODUCTION

When marginal costs of producing and distributing a good are negligible or zero, as for digital content and other non-rivalrous goods, owners can usually increase profits by providing the good to multiple buyers rather than just one. However, in many such cases, potential buyers place a premium on exclusive possession of the good. Higher value for exclusive purchase can be fuelled by the threat of competition, by a sense of privilege, or by special customer preferences (e. g., luxury goods). For example, a prospector would derive higher value from exclusive possession of information regarding a natural resource repository. A retailer may perceive greater value from exclusive right to sell a good because it avoids competition with other shops. A computing platform can entice more users when a *marqué* software application is solely available on that platform and not on competing ones. Similarly, publishers may acquire exclusivity rights for the distribution of books or movies. Owners of such goods must therefore decide whether to sell the good exclusively, at a high price, or to multiple buyers at lower prices. This paper analyses and proposes mechanisms for making such sales.

Our main motivating example comes from online marketing, specifically *sales lead* generation. A sales lead is condensed information about an individual who has exhibited interest in buying a product or service. Many firms specialise in generating such leads by identifying prospective customers *early* in

the customer's purchase process, and directing these customers to a suitable merchant. Because the referral is for a "potential" purchase than for a guaranteed sale, often a single prospect or lead can simultaneously be sold to multiple merchants who all find the prospect attractive. Internet-based lead generators therefore allocate the customer lead on the basis of an online auction, seeking to maximise their fee, either the fee from a single merchant who seeks exclusive access to the customer or the sum of fees from multiple merchants. Online lead generation is, in many industries, a cheaper means for acquiring a customer than traditional alternatives such as newspaper, television or billboard advertising. It is growing rapidly, with a total revenue of 1.75 billion USD earned in 2013.¹

Our understanding of lead marketing with exclusive and shared sales is based on detailed interactions with Reply!,² a lead generator company that deploys this business model. Reply! acquires "clicks" (web surfers) via several means including email marketing and search advertising, and then sells these leads to one or multiple merchants. Because of a large-scale and fast-moving dynamic environment (hundreds of leads sold per minute for distinct products), it is common to transact leads via online auctions. Reply! employs a particular first-price auction in which each advertiser submits two types of bids, one for getting the lead exclusively, and one for receiving it jointly with up to three other bidders. More generally, the fundamental challenge that we tackle in this paper is to design an auction protocol—a rule for allocation and payments—that will, in expectation, maximise the seller's profit, subject to surplus-maximising actions by participating buyers. The need for such auctions also occurs in other areas of online marketing including, notably, search advertising.

Sales lead generation is related to, but distinct from, other ways in which firms seek to acquire customers. For instance, trade shows help firms acquire new customers, especially in business marketing. However, the link from exposure to purchase is extremely long and hard to trace Gopalakrishna and Lilien [34], and trade show organisers usually are paid for space and time, rather than individual leads. Another category is referral infomediaries, who guide customers towards the purchase of specific products from other mer-

¹ http://www.iab.net/media/file/IAB_Internet_Advertising_Revenue_Report_FY_2013.pdf

² <http://www.reply.com>

chants; but such intermediaries are usually compensated only upon the customer's purchase, and moreover they usually contract with only a single merchant for each product category Chen et al. [14]. Search engines allocate paid search slots to multiple advertisers for every search query; these slots (which appear alongside organic search results) are treated as vertically differentiated goods and sold via variations of the Generalized Second Price auction (GSP, Edelman et al. [26], Varian [81]). But, recently, major online search engines like Google, Yahoo! and Bing have begun considering designs where advertisers can place two-dimensional bids for receiving the click exclusively or shared Sayedi [69], Jerath and Sayedi [45]. Yahoo! has patented such an auction, an extension of the GSP.

OUR CONTRIBUTION We study the problem of designing a mechanism for allocation and pricing of goods or services to unit-demand economic agents who might exhibit additional margin value over their basic valuation for situations when they are served exclusively. Our approach is to model this as a two-dimensional mechanism design problem, where the first parameter represents the base valuation for being served (called *shared* value), while the second one is the *exclusivity margin*. The total value for being served exclusively then is the sum of the shared value and the margin for exclusivity. The goal is to find the expected revenue maximising mechanism that is deterministic, implementable in dominant strategies and ex-post individual rational. Designing such optimal multi-dimensional (or multi-parameter) mechanisms is intrinsically hard and there hasn't been a general framework so far to treat these problems. To work around the boundaries some researchers choose to fold their problem into a single dimensional model e. g., Deng and Pekeć [21], such that the techniques of Myerson [54] can be applied. Others weaken the implementability condition and use linear programs to compute approximations e. g., Cai et al. [13].

We tackle the problem in the 2-dimensional setting to address the full complexity, but instead of going for optimality, our focus is on heuristics that provide *good enough* revenue relative to the optimal mechanism. In particular, we show that the revenue from naïve mechanisms, such as selling always exclusively or always non-exclusively, is a constant-factor approximation of the optimal revenue. Moreover, the same approximation factor holds for the

optimal Bayes-Nash implementable mechanism as well. We prove this by providing bounds on the optimal revenue for two separate single-dimensional relaxations of the problem. This technique can be of independent interest, especially the one for deriving the second bound. We modify the method used in Devanur et al. [22] in order to handle allocational externalities arising due to the exclusivity margin. This is done by decomposing the two-dimensional agents into single-parameter representatives. Under this decomposition, allocational externalities turn into interdependent valuations. We also devise a heuristic mechanism from the class of affine maximisers and confirm its good performance by means of simulation. Combining simulations with the upper bounds allows us a novel way of visualising how close are the different heuristics to the optimal mechanisms. Our results require mild conditions, that the valuations of each dimension are distributed independently and identically, and have monotone hazard rate. The distribution for shared values may differ from that for marginal values.

The choice of deterministic, ex-post individual rational and dominant strategy implementable mechanism might seem to be limiting. Indeed, it is folklore knowledge that in multi-dimensional settings randomisation improves revenue and that the class of Bayes-Nash implementable mechanisms allows more flexibility for the designer and admits different solving techniques e. g., Cai et al. [13]. However, the more restrictive setting is highly relevant for two reasons: the deep need in practice for speed and simplicity, and (as we demonstrate) little loss in revenue from doing so. Most of the applications, especially lead marketing and sponsored search, involve real-time auctions where participants exhibit varied levels of sophistication. Therefore, both the auctioneer and bidders desire fast, predictable and simple mechanisms that require minimal user interaction and that make the available strategies clear for the users. Moreover, the dominant strategy implementation is prior-free, hence it is more robust from a practical perspective. In contrast, Bayes-Nash implementation requires a lot more from the participants: all of them should have the same prior about the type distributions, and they have to be able to compute expected utilities based on their actions. The usage of lotteries might be acceptable in certain applications, but it is not yet desired in everyday business. Paying more than your bid is likely to alienate users from a system as they may regret their participation in the auction, therefore ex-post

individual rationality can be seen as a natural property of an applied mechanism. Moreover, our main theorem demonstrates that focusing on dominant strategy implementable, ex-post individual rational, simple therefore robust mechanisms costs only a small fraction of the optimal revenue. This means that the trade-off between optimality and simplicity is small.

RELATED WORK Our main motivating example stems from sales lead generation, an important section of online marketing. Companies engaging in sales lead generation act as infomediaries, gathering valuable information from prospective customers and conveying them to retailers. There is some literature on the impact of referral infomediaries Chen et al. [14], Ghose et al. [32], but we have not observed any literature on mechanism design for lead generators, specifically in this context where the lead generator has the option of selling the lead exclusively or jointly to multiple merchants, and when every merchant potentially has a pair of reservation prices, one for exclusive allocation and a second for shared allocation.

A potential simplification for addressing this exclusive-or-shared allocation problem is to reduce each merchant's pair of reservation prices down to a single type variable. For instance, one could assume that every merchant's value for exclusive allocation is a constant multiple of their value for shared allocation. In such single-dimensional settings, the framework introduced in the seminal Myerson [54] and Riley and Samuelson [60] papers prevails. His approach provides a closed form characterisation of the optimal mechanism under mild assumptions for most of the cases. In contrast, for settings with multi-dimensional types, the optimal mechanism design problem quickly becomes intractable. Even if one finds the optimal mechanism, it is generally too complicated to implement in practice. The optimal mechanism is usually tailored to the very details of the distribution on agent preferences. This precise knowledge can be considered a problematic requirement in some settings. These undesired properties are reflected in the works of Rochet and Choné [62], Armstrong [3] and Manelli and Vincent [50]. For more critics and thoughts on optimal mechanism design see Hartline [37].

There are different attempts to circumvent the inherent hardness in multi-dimensional mechanism design. One approach is to assume that only one parameter is private information and the the remaining ones are public, i. e.,

to stay in the single-dimensional domain. This path was chosen in Deng and Pekeč [21] with a focus on exclusivity contracts. In that paper an agent is allocated exclusively if nobody else from his direct neighbourhood receives it. The neighbourhood is defined on a publicly known network. Our setting covers the case when this neighbourhood graph is a clique. To get their main result they restrict themselves to clique neighbourhood graphs and single-dimensional valuations, where the value for the shared allocation is private information, while the exclusivity margin is public. The rest of the assumptions are similar to ours: private values are distributed according to a monotone hazard rate distribution and they search for deterministic mechanisms. Due to the single-dimensional assumption they can apply the techniques of Myerson and derive the optimal mechanism, which is the one that maximises the sum of virtual valuations for each type report. In contrast, our model keeps multi-dimensionality where the Myerson framework doesn't have a bite.

Along this line of research, Salek and Kempe [66] and Pei et al. [58] study the problem of selling a digital good with unlimited supply of copies to bidders whose value for the good is decreasing in the number of bidders obtaining it. The function according to which the valuation depreciates is public information, therefore their setting is single-dimensional. Salek and Kempe [66] provides the revenue maximising Bayes-Nash implementable mechanism based on the Myerson techniques, while Pei et al. [58] adapts the prior-free "single-sample" mechanism from Dhangwatnotai et al. [23] that yields a constant-factor approximation for that setting. Their theorems hinge on the assumption that types are single-dimensional and are independently distributed according to monotone hazard rate distributions. Moreover, besides that the "single-sample" mechanism is not deterministic, it does not even extend directly to our setting as applying reserve prices combined with the Vickrey-Clark-Groves mechanism (VCG, see Vickrey [82], Clarke [17] and Groves [35]) is not incentive compatible in general, demonstrated by Example 3.24 on page 79.

Jerath and Sayedi [45] studies an extension of GSP for sponsored search advertising patented by Yahoo!. This auction takes two bids as input: one for being displayed among other ads and a second bid for exclusive display. As truth-telling is not dominant strategy in that auction the authors aim to iden-

tify and analyse bidding strategies that lead to a Bayes-Nash equilibrium. To render their analysis tractable they restrict themselves to a highly stylised setting including only three agents: two having no exclusivity margin and one whose exclusivity margin is a fixed portion of the shared valuation. They find that allowing advertisers to bid for exclusivity usually increases the search engine's revenue (because of an increased competition effect) but that revenue may fall under certain conditions. Note that for such a single-dimensional setting the optimal auction format is characterised by Deng and Pekeč [21].

Cai et al. [13] and Cai et al. [12] weaken the requirement on implementability by looking for multi-dimensional mechanisms that are Bayes-Nash incentive compatible. They also allow for randomised mechanisms, which allows them to use linear programming to acquire a polynomial-time approximation scheme. Note, however, that their method is not directly applicable to our problem as it doesn't handle allocational externalities. Moreover, if we go for deterministic and dominant strategy incentive compatible mechanisms, then we have to face an integer linear program with exponential number of variables and constraints, which renders this approach impractical.

Another line of research focuses on heuristic mechanisms, which possess a succinct description and might exhibit a guaranty on its expected revenue. In relation to our problem the work of Devanur et al. [22] and Dhangwatnotai et al. [23] bear relevance. They deal with multi-parameter mechanism design involving unit-demand bidders, regular type distributions and matroid or downward closed feasibility constraints. They derive simple mechanisms that achieve constant-factor approximations of the optimal revenue. The key point of their proof is to employ the solution from a single-dimensional analogue as an upper bound for the optimal revenue. Then the revenue of their simple mechanism is compared to that upper bound. In spite of the similarities, these results cannot be utilised directly for our problem as their setting cannot accommodate allocational externalities. Even if one treats the different allocations as items, the corresponding feasibility constraints are not downward closed, which is a necessary assumption in their model. The way we derive the second upper bound is reminiscent to the technique developed in Devanur et al. [22] in the sense that we also split the multi-dimensional agents to single-dimensional representatives. The twist in our method is that

in order to handle allocational externalities we endow the representatives with interdependent valuations.

Allocational externalities are similar in nature to interdependent valuations and identity dependent externalities in the sense that they all try to catch the impact of the agents on each other. The implications of externalities have been studied in various settings, see, e. g., Jehiel et al. [44], Segal [74] and Figueroa and Skreta [29]. In Aseff and Chade [4] the optimal auction for selling two identical goods is considered, where bidders impose externalities on each other. Their model can be parametrised such that it coincides with a special case of Deng and Pekeč [21], but it is limited to selling two items, multiplicative externality and single-dimensional valuations.

The idea of using affine maximisers to improve on the revenue of the VCG mechanism has appeared mostly in connection with combinatorial auctions. Likhodedov and Sandholm [48] specifies a special class of affine maximisers and tries to fine-tune its parameters. Tang and Sandholm [78] considers the case of two bidders and two items and searches for the best parameters for a given affine maximiser. Jehiel et al. [40] study a special class of affine maximisers for multi-object auctions where agents have private and additive valuations for heterogeneous objects. They compute the partial derivatives of the affine maximisers' expected revenue with respect to their parameters and show that mixed bundling increases the expected revenue compared to pure bundling and to VCG. One of our heuristics is also an affine maximiser, but its parametrisation idea differs from that of the two previous papers. Moreover, with the help of simulation and the derived upper bound mechanisms, we can demonstrate in a novel way how small the gap is between the heuristic mechanisms and the optimal revenue.

ORGANISATION Next, we introduce the main concepts and notation for our model. In Section 3.3 we bound the revenue of the optimal mechanism by constructing two mechanisms that are single-dimensional variants of the original problem. These bounds are used in Section 3.4.2 to prove that simple mechanisms can approximate the optimal solution. We conclude that section by constructing a heuristic from the affine maximiser family of mechanisms. Finally, Section 3.5 presents a simulation study on the expected revenue of the discussed mechanisms for various settings.

3.2 MODEL FORMULATION

There are k copies of an item to sell and there are n unit-demand bidders who are interested in buying one of them.³ The set of bidders (or agents) is denoted by $N = \{1, \dots, n\}$. The type of agent i is characterised by the pair $t^i = (s^i, m^i)$, where s^i is the valuation for getting the item shared with some other agent, while m^i is the margin for exclusive possession, i. e., $s^i + m^i$ is the total valuation for receiving the item exclusively. The shared valuation s^i is drawn from the set S^i according with cumulative distribution F^i , while the exclusivity margin m^i is drawn from M^i with cumulative distribution G^i . The corresponding density functions are written in lower case. The set of possible types for agent i is then $T^i = S^i \times M^i$. To be able to talk about the vector of agents we define $T = T^1 \times \dots \times T^n$ and $T^{-i} = \times_{j \neq i} T^j$. Each element in T is a $2 \times n$ matrix and will be denoted by t . Throughout the chapter it is assumed that agents are identical, i. e., $T^i = T^j = [0, \tilde{s}] \times [0, \tilde{m}]$, $G^i = G^j = G$ and $F^i = F^j = F$ for all $i, j \in N$, where $\tilde{s}, \tilde{m} \in \mathbb{R}^+$. We define C , the distribution of the exclusive value, as the convolution of F and G . It is assumed that F and G are independent and strictly increasing. Furthermore, we let types be distributed independently among agents. This means that the distribution of type t^i is $F \times G$ and the distribution of type-tuple is $(F \times G)^n$.

Recall that for settings with single-dimensional type spaces, independent valuations and independent type distributions the virtual valuation of agent i for random variable t^i , that is distributed as D , is defined as

$$\phi_D(t^i) = t^i - \frac{1 - D(t^i)}{d(t^i)}. \quad (12)$$

In our context t^i could be s^i , m^i or $s^i + m^i$ for example. Distributions that have increasing virtual valuation are called regular. An interesting subclass of regular distributions are those having monotone hazard rate (MHR), otherwise known as log-concave probability distributions. In words, these are distributions with tail not fatter than that of the exponential distribution. This class

³ Alternatively, there is a single unit of a good which can be offered to k buyers simultaneously at no extra cost. The limit k is often set by policy. For instance, in automobile lead generation where customer contact information is sent to car dealers, k is often set at 3 or 4, due to concerns about customer's time and privacy.

is popular in many areas of economics as its properties usually yield special structure for the problem at hand and hence allows for insightful qualitative implications (see, e. g., Bagnoli and Bergstrom [6]). Main members of this class are the normal, exponential, some parametrisation of gamma, Pareto and the uniform distribution. Distributions D exhibiting MHR have a unique point where the corresponding virtual valuation hits zero. This value is called the reserve value r_D of D .

Definition 3.1 (Reserve Value). For a regular distribution D , the reserve value r_D is the unique value for which $\phi_D(r_D) = 0$.

Following the *revelation principle* Myerson [54] we restrict our attention to direct mechanisms. The set of possible allocations is $\mathcal{O} = \{0, 1\}^N$, i. e., the set of characteristic vectors related to the power set of agents. Each direct mechanism (x, π) can be characterised by its allocation rule $x : T \rightarrow \mathcal{O}$ and its payment scheme $p : T \rightarrow \mathbb{R}^N$. Let $a(t) = \{i : x^i(t) = 1\}$ be the set of agents who receive the item for report t (note that, here a does not denote an action profile as in Chapter 1). Given a mechanism (x, π) and report profile \hat{t} , the realised valuation of agent i having type $t^i = (s^i, m^i)$ is

$$v^i(x(\hat{t}), t^i) = \begin{cases} s^i & \text{if } i \in a(\hat{t}) \text{ and } |a(\hat{t})| \geq 2, \quad (\text{shared allocation}) \\ s^i + m^i & \text{if } a(\hat{t}) = \{i\}, \quad (\text{exclusive allocation}) \\ 0 & \text{otherwise,} \end{cases}$$

while his utility is $u^i(x(\hat{t}), \pi(\hat{t}), t^i) = v^i(x(\hat{t}), t^i) - \pi^i(\hat{t})$.

The exclusivity auction problem is to find a DSIC and EPIR direct mechanism that maximises the expected revenue. This can be modelled as

$$\max \mathbb{E}_t \left[\sum_i \pi^i(t) \right] \quad (\text{Base})$$

subject to

$$u^i(x(t), \pi(t), t^i) \geq u^i(x(\hat{t}^i, t^{-i}), \pi(\hat{t}^i, t^{-i}), t^i) \forall i, t^{-i}, t^i, \hat{t}^i, \quad (\text{DSIC})$$

$$u^i(x(t), \pi(t), t^i) \geq 0 \quad \forall i, t, \quad (\text{EPIR})$$

$$\sum_i x^i(t) \leq k \quad \forall t, \quad (\text{Supply})$$

$$x^i(t) \in \{0, 1\} \quad \forall i, t. \quad (\text{BI})$$

Constraint (DSIC) is responsible for dominant strategy incentive compatibility, while (EPIR) enforces ex-post individual rationality. (Supply) sets the upper bound on the number of copies that can be sold, and (BI) ensures that the mechanism is deterministic. Note that with (Base) we refer to the whole mathematical program. The optimal solution (x, π) for (Base) will be called BASE^* .

3.3 UPPER BOUNDS

As the optimal mechanism is unknown we have to search for other benchmarks to measure the performance of any heuristic mechanism. One straightforward choice is the optimal expected welfare, as it is always larger or equal to the optimal revenue due to the EPIR assumption. The problem is that the gap between the two can be significant, therefore it is a too loose measure most of the time. See for example graph 6 on page 82 for justification, where the optimal revenue is less than 60% of the optimal welfare. In order to find tighter bounds we take the approach of relaxing the conditions of (Base) such that the resulting optimisation problem admits a closed form optimal solution. The plan is to relax the problem such that the resulting setting is single-dimensional, because in such environments the framework of Myerson [54] solves the revenue maximisation problem.

In this section we construct two such non-trivial upper bounds. For notational convenience the expected revenue achieved by any mechanism MECH will be symbolised by $\text{Rev}(\text{MECH})$.

3.3.1 Relaxation 1: Public exclusivity margin

The first upper bound is acquired by assuming that the exclusivity margin is public knowledge. The optimisation problem reflecting this change can be written as

$$\max \mathbb{E}_t \left[\sum_i \pi^i(t) \right] \quad (\text{UBM})$$

subject to

$$\begin{aligned} u^i(x(t), \pi(t), (s^i, m^i)) &\geq u^i(x((\hat{s}^i, m^i), t^{-i}), \pi((\hat{s}^i, m^i), t^{-i}), s^i, m^i) \\ &\quad \forall i, t^{-i}, s^i, m^i, \hat{s}^i, \quad (\text{DSIC}_2) \\ u^i(x(t), \pi(t), t^i) &\geq 0 \quad \forall i, t, \\ \sum_i x^i(t) &\leq k \quad \forall t, \\ x^i(t) &\in \{0, 1\} \quad \forall i, t. \end{aligned}$$

We will refer to the optimal mechanism for (UBM) as UBM^* . Note that the only difference between (Base) and (UBM) is that in (DSIC₂) agent i gets only incentives not to lie about valuation s^i . Therefore (UBM) is clearly a relaxation of (Base), which means that each feasible solution for (Base) is feasible for (UBM), moreover, $\text{Rev}(\text{UBM}^*) \geq \text{Rev}(\text{BASE}^*)$. In order to define UBM^* succinctly some additional notations are introduced. For given type profile $t = (s, m)$ define $a_s(t)$ and $a_e(t)$ to be the set of agents with the highest contribution to the expected revenue.

$$\begin{aligned} a_s(t) &= \arg \max_{\substack{a \subseteq N \\ |a| \geq 2 \\ |a| \leq k}} \left\{ \sum_{i \in a} \phi_F(s^i) \right\}, \\ a_e(t) &= \arg \max_i \left\{ \phi_F(s^i) + m^i \right\}. \end{aligned}$$

In words, $a_s(t)$ is the set of agents with k -highest shared virtual values, while $a_e(t)$ is the agent with highest sum of shared virtual value. Ties are broken ar-

bitrarily in case of multiple optimal arguments. Note that $\sum_{i \in a_s(t)} \phi_F(s^i) \geq 0$ as the empty set is also a solution. The k -highest virtual valuations (for shared possession) add up to $\sum_{i \in a_s(t)} \phi_F(s^i)$, while the highest sum of shared virtual value and exclusivity margin is $\phi_F(s^{a_e(t)}) + m^{a_e(t)}$. The next theorems demonstrate the advantage of this single-dimensional relaxation by showing that the optimal expected revenue can be expressed in terms of virtual valuations. Moreover, the set of winners in the optimal auction is equal to either $a_s(t)$ or $a_e(t)$, dependent on which one contributes the most to the expected revenue.

For cases when only shared valuations are private information and exclusivity margins are public constants Deng and Pekeć [21] characterises the optimal DSIC and EPIR mechanism by building on the framework of Myerson [54]. We rely on their result to determine the optimal mechanism for (UBM).

Theorem 3.2. *If F has MHR and the exclusivity margins are public information, then the allocation rule of UBM* given type profile $t = (s, m)$ can be characterised as*

$$x^i(t) = 1 \text{ for } \begin{cases} i \in a_s(t) & \text{when } \sum_{j \in a_s(t)} \phi_F(s^j) \geq \phi_F(s^{a_e(t)}) + m^{a_e(t)} \\ i = a_e(t) & \text{when } \sum_{j \in a_s(t)} \phi_F(s^j) < \phi_F(s^{a_e(t)}) + m^{a_e(t)} \end{cases} \quad (13)$$

and $x^i(t) = 0$ otherwise. The expected revenue under this allocation rule is equal to

$$\mathbb{E}_t \left[\sum_i \pi^i(t) \right] = \mathbb{E}_t \left[\sum_i \left(\phi_F(s^i) + m^i \prod_{j \neq i} (1 - x^j(s^i, m^i, t^{-i})) \right) x^i(s^i, m^i, t^{-i}) \right]. \quad (14)$$

The proof can be found in the appendix. In general, the optimal solution of (UBM) might yield strictly higher revenue than (Base). This is due to the fact that the optimal allocation in (UBM) is not always feasible for (Base) as the following example shows.

Example 3.3. Let $N = \{1, 2\}$, $s^i \sim U(0, 1)$, $m^i \sim U(0, 1)$ for all $i \in N$. Assume that $t^1 = (0.3, 0.3)$, $t^2 = (0.4, 0.1)$. Then we have that $\phi_{U(0,1)}(s^1) + m^1 =$

-0.1 , $\phi_{U(0,1)}(s^2) + m^2 = -0.1$, $\phi_{U(0,1)}(s^1) = -0.4$ and $\phi_{U(0,1)}(s^2) = -0.2$, therefore according to UBM* nobody gets anything. Now, change the valuations of agent 1 such that $\hat{t}^1 = (0.5, 0.05)$. As $\phi_{U(0,1)}(\hat{s}^1) + \hat{m}^1 = 0.05$ and $\phi_{U(0,1)}(\hat{s}^1) = 0$ agent 1 receives the item exclusively. This allocation rule violates monotonicity condition (MON), which means that there is no payment scheme such that the incentive compatibility constraints are satisfied. Therefore it cannot be part of a solution for (Base).

3.3.2 Relaxation 2: Bound with representatives

The derivation of the second upper bound mechanism is more involved. For every instance of (Base) we define an instance of an alternative setting by introducing two representatives for each agent: one for each dimension of his type. To distinguish this setting in the notation, we add a “bar” to each item, as in \bar{t} .

Definition 3.4 (Representative environment). For any instance of (Base) given by the set of agents N , the set of types T and distributions of types F, G , the representative environment is defined as follows:

- For each $i \in N$ introduce i_s and i_e , such that the set of agents is $\bar{N} = \bar{N}_s \cup \bar{N}_e$, where $\bar{N}_s = \{1_s, \dots, i_s, \dots, n_s\}$ and $\bar{N}_e = \{1_e, \dots, i_e, \dots, n_e\}$.
- $\bar{O} = \{a \subseteq \bar{N}_s \mid |a| \leq k\} \cup \{i_e\}_{i_e \in \bar{N}_e}$, the set of feasible allocations
- $\bar{T}^{i_s} = S^i$ for all $i_s \in \bar{N}_s$, $\bar{T}^{i_e} = M^i$ for all $i_e \in \bar{N}_e$, $\bar{T} = \times_{i \in \bar{N}} \bar{T}^i$, the set of type profiles
- Let the valuations for all $t \in \bar{T}$, for all $i \in \bar{N}$ and for allocation rule \bar{x} be:

$$\bar{v}^i(\bar{x}(t^i, t^{-i}), (t^i, t^{-i})) = \begin{cases} s^{i_s} & \forall i_s \in \bar{N}_s, \text{ if } \bar{x}^{i_s} = 1, \\ s^{i_s} + m^{i_e} & \forall i_e \in \bar{N}_e, \text{ if } \bar{x}^{i_e} = 1, \\ 0 & \forall i \in \bar{N} \text{ otherwise.} \end{cases}$$

- $\bar{u}^i(\bar{x}(t^i, t^{-i}), \bar{\pi}(t^i, t^{-i}), (t^i, t^{-i})) = \bar{v}^i(\bar{x}(t^i, t^{-i}), (t^i, t^{-i})) - \bar{\pi}^i(t^i, t^{-i})$, the utility function
- $t^i \sim F$ for all $i_s \in \bar{N}_s$, $t^i \sim G$ for all $i_e \in \bar{N}_e$

Note that agent i_e exhibits informational externality in his valuation. In such settings dominant strategy incentive compatibility is too demanding as it requires truthfulness for every possible report of the other agents, leaving small room for non-trivial mechanisms (see for example Roughgarden and Talgam-Cohen [64]). Therefore we relax the incentive compatibility condition to ex-post incentive compatibility (EPIC), which requires truthfulness only for every true type of the other agents.

Definition 3.5 (EPIC). A direct mechanism is ex-post incentive compatible (EPIC) if for each agent truth-telling is a Nash equilibrium in the complete information game defined by each type profile. Formally, a direct mechanism (x, π) is EPIC, if for every i , t^{-i} , t^i and \hat{t}^i it holds that

$$u^i(x(t), \pi(t), (t^i, t^{-i})) \geq u^i(x(\hat{t}^i, t^{-i}), \pi(\hat{t}^i, t^{-i}), (t^i, t^{-i})).$$

Another difference compared to the principles of (Base) is that we allow for allocations where only one shared agent receives the item. This relaxation increases the set of feasible allocations, hence the possibility to generate more revenue.

The revenue optimisation problem for the representative environment can be stated as

$$\max \mathbb{E}_t \left[\sum_i \bar{\pi}^i(t) \right] \quad (\text{UBR})$$

subject to

$$\bar{u}^i(\bar{x}(t), \bar{\pi}(t), (t^i, t^{-i})) \geq \bar{u}^i(\bar{x}(\hat{t}^i, t^{-i}), \bar{\pi}(\hat{t}^i, t^{-i}), (t^i, t^{-i})) \quad \forall i, t^i, \hat{t}^i, t^{-i} \quad (\text{EPIC})$$

$$\bar{u}^i(\bar{x}(t), \bar{\pi}(t), t^i) \geq 0 \quad \forall i, t \quad (\text{EPIR}_2)$$

$$(1 - \bar{x}^{j_e}(t))k \geq \sum_{i_s \in \bar{N}_s} \bar{x}^{i_s}(t) \quad \forall t, j_e \in \bar{N}_e \quad (\text{FeasA})$$

$$\sum_{i_e \in \bar{N}_e} \bar{x}^{i_e}(t) \leq 1 \quad \forall t \quad (\text{FeasB})$$

$$\bar{x}^i(t) \in \{0, 1\} \quad \forall i, \forall t.$$

(EPIC) is responsible for ex-post incentive compatibility, while (EPIR₂) for ex-post individual rationality. (FeasA) ensures that at most k items are allocated to agents in \bar{N}_s , and that there cannot be allocations where agents both from \bar{N}_s and \bar{N}_e receive an item. (FeasB) states that at most one agent from \bar{N}_e can be served. We will refer to the optimal mechanism for (UBR) as UBR*.

There are more ways of defining a single-dimensional variant of Base, but this formulation is particularly useful as the following proposition demonstrates.

Proposition 3.6. *For any feasible mechanism (x, π) for (Base), there is a feasible mechanism $(\bar{x}, \bar{\pi})$ for (UBR) such that $Rev(\bar{x}, \bar{\pi}) \geq Rev(x, \pi)$.*

Here, only the intuition behind the proof is given, for technical details the reader is referred to the appendix. For any feasible mechanism in (Base), we define a feasible mechanism in (UBR) as follows. An agent that is the representative of an exclusive valuation receives the item if and only if the original agent is allocated exclusively. The allocation rule for the representatives of the shared valuations is defined in a similar fashion. In (Base) EPIC constraints require an allocation rule that is monotone for each valuation dimension. We show that this monotonicity is preserved for each representative in (UBR) in the allocation rule defined earlier. We then exploit that as in single dimensional domains a monotone allocation rule is sufficient for the existence of a payment scheme that makes the whole mechanism EPIC. At the end we construct that particular payment scheme and show that its expected value is at least as large as the expected revenue of the original mechanism. The technical challenges in the proof are due to the fact that the desired structure of the allocation rule is lacking on a null set of types.

It is immediate from Proposition 3.6 that the optimal expected revenue of (UBR) serves as another non-trivial upper bound for $Rev(\text{BASE}^*)$. As (UBR) is

single-dimensional, it is possible to provide a closed form solution for UBR* in a similar way as before. For given type profile t let

$$\begin{aligned}\bar{a}_s(t) &= \arg \max_{\substack{a \subseteq \bar{N}_s \\ |a| \leq k}} \left\{ \sum_{i_s \in a} \phi_F(t^{i_s}) \right\} \\ \bar{a}_e(t) &= \arg \max_{i_e \in \bar{N}_e} \left\{ t^{i_s} + \phi_G(t^{i_e}) \right\}.\end{aligned}$$

Ties are broken arbitrarily in case of multiple optimal arguments. Note that $\sum_{i \in \bar{a}_s(t)} \phi_F(t^{i_s}) \geq 0$ as the empty set is also a solution.

Theorem 3.7. *When F and G have MHR, the allocation rule of UBR* represented by \bar{x} , for type profile t , is computed as follows,*

$$\bar{x}^i(t) = 1 \text{ for } \begin{cases} i \in \bar{a}_s(t) & \text{when } \sum_{j_s \in \bar{a}_s(t)} \phi_F(t^{j_s}) \geq \max_{j_e \in \bar{N}_e} \{t^{j_s} + \phi_G(t^{j_e})\} \\ i = \bar{a}_e(t) & \text{when } \sum_{j_s \in \bar{a}_s(t)} \phi_F(t^{j_s}) < \max_{j_e \in \bar{N}_e} \{t^{j_s} + \phi_G(t^{j_e})\} \end{cases}$$

and $\bar{x}^j(t) = 0$ otherwise. The optimal revenue under this allocation is

$$\text{Rev}(\text{UBR}^*) = \mathbb{E}_t \left[\sum_{i_s \in \bar{N}_s} \phi_F(t^{i_s}) \bar{x}^{i_s}(t) + \sum_{i_e \in \bar{N}_e} (t^{i_s} + \phi_G(t^{i_e})) \bar{x}^{i_e}(t) \right]$$

such that $\text{Rev}(\text{UBR}^*) \geq \text{Rev}(\text{BASE}^*)$.

3.3.3 Significance of Relaxations

The two relaxations provide us two alternative upper bounds. To get a sense for the quality of these bounds, we compare them against the welfare-maximising solution, which also provides an upper bound on optimal revenue. In order to do that we calculated the revenue of UBR* and UBM* along with the optimal welfare for different instances by means of simulation (more details on the technique in Section 3.5). The results are depicted in Figure 1 and 2. UBR denotes the upper bound with representatives, whilst UBM stands for the bound with public exclusivity margin. As the displayed revenue ratios are taken over the optimal welfare, it is apparent that both relaxation bounds are

much tighter than the welfare upper bound. In particular, in Figure 1 UBM is around half of the optimal welfare. Furthermore, dependent on the support of the two valuations, the difference between UBR and UBM might be significant. In general, when exclusivity margins are small relative to shared valuations, UBM* is tighter, while UBR* is closer to the optimal revenue when exclusivity margins dominate.

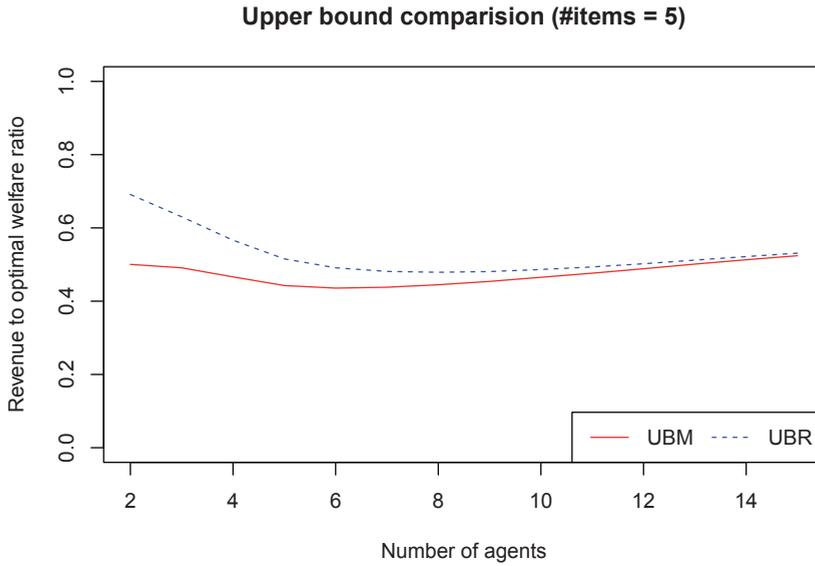


Figure 1: Revenue over optimal welfare of the two upper bounds ($s^i \sim Exp(1)$, $m^i \sim Exp(2)$).

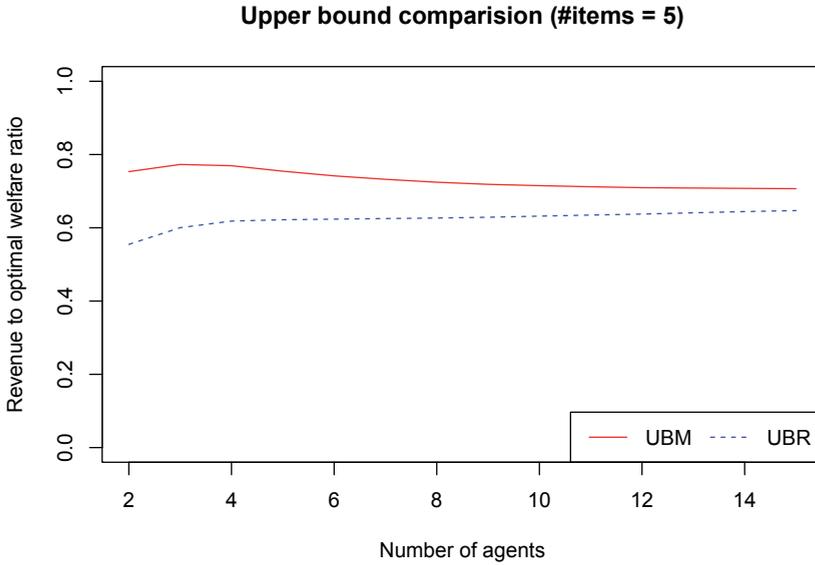


Figure 2: Revenue over optimal welfare of the two upper bounds ($s^i \sim \text{Exp}(1)$, $m^i \sim \text{Exp}(0.5)$).

3.4 HEURISTIC MECHANISMS

The previous section laid out two alternative formulations of the mechanism design problem, which provide upper bounds on the optimal revenue. Having established these, we now seek mechanisms that perform demonstrably well with respect to the optimal revenue. The first candidate will be the classic Vickrey-Clark-Groves mechanism. Then we turn our attention to simple mechanisms that allocate either always exclusively or always shared. It turns out that the combination of these two extreme mechanisms is a constant-factor approximation of the optimal revenue. We finish the section by trying to improve on the revenue of VCG with affine maximisers. The performance of the resulting mechanisms is evaluated in Section 3.5 via simulations.

3.4.1 VCG mechanism

In this subsection we introduce the well-known VCG mechanism, which will serve as a benchmark for our simulations in the next section. VCG is a natural choice for the role of a heuristic mechanism since it can be succinctly described and easily executed (at least in our case), moreover it is a good approximation of the optimal mechanism in some settings (see, e. g., Bulow and Klemperer [11] and Roughgarden and Sundararajan [63]). Furthermore, we prove a revenue guarantee for VCG in a single-item setting that is used in Section 3.4.2 to provide revenue guarantee for simple mechanisms for our setting.

Remember that $SW^N(a, t) = \sum_{i \in N} v^i(a, t^i)$ denotes the welfare of agents in N for allocation a and type profile t .

Definition 3.8 (VCG mechanism). Let O_N be the set of feasible allocations for agents in N and let $a(t) \in O_N$ be an allocation for each type profile t such that

$$a(t) \in \arg \max_{a \in O_N} SW^N(a, t).$$

Similarly, let $O_{N \setminus \{i\}}$ be the set of feasible allocations for agents in $N \setminus \{i\}$ and let $a(t^{-i}) \in O_{N \setminus \{i\}}$ be an allocation for each type profile t^{-i} such that

$$a(t^{-i}) \in \arg \max_{a \in O_{N \setminus \{i\}}} SW^{N \setminus \{i\}}(a, t^{-i}).$$

Then for each t VCG chooses $a(t)$ as allocation and elicits payment

$$\pi^i(t) = \sum_{j \neq i} v^j(a(t^{-i}), t^j) - \sum_{j \neq i} v^j(a(t), t^j). \quad (15)$$

VCG allocates efficiently, i. e., maximises the total welfare pointwise, moreover, it is EPIR. The following theorem helps computing its expected revenue.

Theorem 3.9. Let $N = \{1, \dots, n\}$ be the set of agents with i.i.d. types, and let $Q = \{1, \dots, n - 1\}$ denote the same set of agents with one less member. Let $T_N = \times_{i \in N} T^i$ denote the set of possible type profiles of agents in N and,

similarly let $T^Q = \times_{i \in Q} T^i$ be the set of possible type profiles of $n - 1$ agents. Then

$$\text{Rev}(\text{VCG}) = n \mathbb{E}_{t \in T^Q} [SW^Q(a(t), t)] - (n - 1) \mathbb{E}_{t \in T^N} [SW^N(a(t), t)] \quad (16)$$

Corollary 3.10. *If*

$$\frac{\mathbb{E}_{t \in T^Q} [SW^Q(a(t), t)]}{\mathbb{E}_{t \in T^N} [SW^N(a(t), t)]} \geq 1 - \varrho,$$

then the expected revenue of VCG is at least $1 - n\varrho$ times the expected optimal welfare.

The message of Corollary 3.10 is that if one wants to compare $\text{Rev}(\text{VCG})$ to the optimal expected welfare, then it is sufficient to know the added value of an extra agent to the welfare. Let $H_i = \sum_{j=1}^i 1/j$ represent the i^{th} Harmonic number and set $H_0 = 0$. The next lemma is useful for providing lower bounds on the revenue-welfare ratio.

Lemma 3.11 (Lemma 3 from Roughgarden and Sundararajan [63]). *Draw independently n times from an MHR distribution. Then the expected value of the l^{th} -largest value of n samples is at least $(H_n - H_{l-1})/(H_{n+j} - H_{l-1})$ times that of the l^{th} -largest value of $n + j$ samples.*

Theorem 3.12. *Consider the single-item auction problem with $n \geq 2$ agents who have unit demand and single-dimensional valuations i.i.d. according to an MHR distribution. Then VCG extracts at least $1 - 1/H_n$ fraction of the optimal welfare in terms of expected revenue.*

The proof is given in the appendix. According to Theorem 4 of Roughgarden and Sundararajan [63] the ratio of the VCG revenue to the optimal welfare is at least $1 - 1/n$ for monotone hazard rate distributions. Their result is apparently not precise as our bound is tight for exponential distributions and $1 - 1/H_n$ is generally lower than $1 - 1/n$.

3.4.2 Simple mechanisms

In this section we show that the optimal revenue can be approximated by two simple mechanisms. The first one is the optimal single item auction, i. e., the best out of those mechanisms that allocate to at most one agent. The second simple mechanism is the other extreme: it is the optimal mechanism among those which never allocate exclusively. Notice that these mechanisms are single-dimensional as in each auction only one valuation counts (shared or exclusive). Therefore we can utilise the framework of Myerson [54] that gives a characterisation of the optimal single-dimensional mechanism. For our purposes, we adapt a generalisation of the original result of Myerson [54] from Devanur et al. [22].

Theorem 3.13 (Myerson [54]). *Let NU be a set of unit-demand agents, and $O \subseteq 2^{NU}$ is the set of feasible allocations, i. e., the set of agents who can be served simultaneously. Assume that agent i 's type t^i is single-dimensional and drawn from set $T^i \subset \mathbb{R}$, with regular distribution D^i . Let $T = \times_i T^i$, and let $D = D^1 \times \dots \times D^n$ be the distribution over T . Then the single-parameter environment (NU, O, T, D) has the following properties.*

- A. *For every DSIC and EPIR mechanism (x, π) , where p is the maximal revenue given x and $x^i(t)$ is the probability that agent i is served, the expected value of total payments can be written as*

$$\mathbb{E}_t \left[\sum_i \pi^i(t) \right] = \mathbb{E}_t \left[\sum_i \phi_{D^i}(t^i) x^i(t) \right].$$

- B. *The revenue-maximising DSIC and EPIR allocation x can be characterised as $x^i(t) = 1$ if $i \in a(t)$ and $x^i(t) = 0$ otherwise, where*

$$a(t) = \arg \max_{a \in O} \left\{ \sum_{i \in a} \phi_{D^i}(t^i) x^i(t) \right\}.$$

With the help of Theorem 3.13 we can directly define both simple mechanisms.

Definition 3.14 (Optimal exclusive mechanism (OE)). Let

$$a_e(t) = \arg \max_{\substack{a \subseteq N \\ |a| \leq 1}} \left\{ \sum_{i \in a} \phi_C(s^i + m^i) \right\}.$$

Ties are broken arbitrarily in case of multiple optimal arguments. Then OE is given by its allocation and payment rule $(x_{\text{OE}}, p_{\text{OE}})$ as

$$x_{\text{OE}}^i(t) = \begin{cases} 1 & \text{if } i \in a_e(t) \\ 0 & \text{otherwise,} \end{cases} \quad (17)$$

$$\pi_{\text{OE}}^i = \begin{cases} \min\{\hat{s}^i + \hat{m}^i \mid x_{\text{OE}}^i((\hat{s}^i, \hat{m}^i), t^{-i}) = 1\} & \text{if } x_{\text{OE}}^i(t) = 1 \\ 0 & \text{otherwise.} \end{cases} \quad (18)$$

Lemma 3.15. *If F and G have MHR, then OE is feasible for (Base) and achieves the highest expected revenue among those mechanisms which allocate to at most one agent. Moreover, its expected revenue can be expressed as*

$$\text{Rev}(\text{OE}) = \mathbb{E}_t \left[\sum_i \phi_C(s^i + m^i) x_{\text{OE}}^i(t) \right]. \quad (19)$$

Proof It follows directly from Theorem 3.13 given that C has MHR as the set of MHR distributions is closed under convolution (see Barlow et al. [7]). \square

Definition 3.16 (Optimal shared mechanism (OS)). Let

$$a_s(t) = \arg \max_{\substack{a \subseteq N \\ |a| \leq k, |a| \neq 1}} \left\{ \sum_{i \in a} \phi_F(s^i) \right\}$$

Ties are broken arbitrarily in case of multiple optimal arguments. Then OS is given by its allocation and payment rule (x_{OS}, p_{OS}) as

$$x_{OS}^i(t) = \begin{cases} 1 & \text{if } i \in a_s(t) \\ 0 & \text{otherwise,} \end{cases} \quad (20)$$

$$\pi_{OS}^i = \begin{cases} \inf\{\hat{s}^i \mid x_{OS}^i((\hat{s}^i, m^i), t^{-i}) = 1\} & \text{if } x_{OS}^i(t) = 1 \\ 0 & \text{otherwise.} \end{cases} \quad (21)$$

Lemma 3.17. *OS is feasible for (Base) and achieves the highest expected revenue among those mechanisms which never allocate exclusively. Moreover, its expected revenue can be expressed as*

$$Rev(OS) = \mathbb{E}_t \left[\sum_i \phi_F(s^i) x_{OS}^i(t) \right]. \quad (22)$$

Proof It follows directly from Theorem 3.13. Note that Myerson requires only that for each agent the allocation is monotone non-decreasing in its own type given the others', therefore the fact that OS never allocates for singletons is not an issue for invoking Theorem 3.13. \square

Definition 3.18 (MAXSIMPLE mechanism). For any instance of the exclusivity auction problem MAXSIMPLE calculates the expected revenue of OS and OE, then commits the mechanism with the highest value.

Note that the choice is made upfront for every problem and not for every type realisation, therefore the actions of the agents don't influence which mechanism is executed. Despite the fact that MAXSIMPLE capitalise only on one type of bid and allocation, the following theorem reveals that the seller can still capture a constant fraction of the optimal revenue by implementing it. This is useful for practitioners because they can keep the bidding language and the rules of the auction simple for only a small sacrifice in the revenue. Furthermore, there is only one parameter (the reserve price) to change for different settings.

Theorem 3.19. *Take an instance of the exclusivity auction problem, where F and G are MHR distributions. Then expected revenue of MAXSIMPLE is at least $\left(\frac{1-1/H_n}{2-1/H_n}\right)$ the expected revenue of the optimal mechanism. The minimum of this approximation factor is $1/4$, attained in case of two bidders. Moreover, this approximation factor holds for the optimal Bayes-Nash implementable mechanism as well.*

Although MAXSIMPLE is intuitive and simple, it is a combination of two different mechanisms. Therefore if a seller has multiple similar products to offer, then changing the rules and the type of allocation product by product can create confusion among the buyers. Besides, MAXSIMPLE is only single-dimensional, thus it doesn't take advantage of the extra information and the wider offer menu it is possible to utilise by taking into account both types of valuations. The next natural step accordingly, is to devise mechanisms that consider two types of bids: one for shared and one for exclusive allocations.

3.4.3 A two-dimensional candidate mechanism

An obvious choice for a two-dimensional mechanism would be VCG. It is DSIC and achieves the highest revenue among those mechanism that maximise welfare. The problem is that its expected revenue can get quite bad compared to the optimal one if there is only a few bidders relative to the number of items offered. Our approach to improve on the revenue of VCG is to impose reserve prices. In order to recover incentive compatibility we do that by looking for other appropriate members of the affine maximiser family.

Definition 3.20 (Affine maximisers). Let O_N be the set of feasible allocations for agents in N . A mechanism is an affine maximiser if there are $\gamma^a, \forall a \in O$ and $\lambda^i > 0, \forall i \in N$ such that for each t its allocation is $a(t)$, where

$$a(t) \in \arg \max_{a \in O_N} \left[\gamma^a + \sum_{i \in N} \lambda^i v^i(a, t^i) \right]. \quad (23)$$

Similarly, define

$$a(t^{-i}) \in \arg \max_{a \in \mathcal{O}_{N \setminus i}} \left[\gamma^a + \sum_{j \in N \setminus i} \lambda^j v^j(a, t^j) \right].$$

Then the payment rule of the affine maximiser mechanism is equal to

$$\pi^i(t) = \frac{1}{\lambda^i} \left[\sum_{j \neq i} \lambda^j v^j(a(t^{-i}), t^j) + \gamma^{a(t^{-i})} - \sum_{j \neq i} \lambda^j v^j(a(t), t^j) - \gamma^{a(t)} \right] + h^i(t^{-i}), \quad (24)$$

where h^i is any real-valued function.

Fact 3.21. *Affine maximiser mechanisms are DSIC. Moreover, if h^i is the zero function for all i then they are EPIR as well (see, e. g., Roberts [61] or Nisan et al. [56]).*

Note that VCG is indeed an affine maximiser by taking $\lambda^i = 1, \forall i \in N$ and $\gamma^a = 0$ for all $a \in \mathcal{O}$. To motivate our choice for this class of mechanisms we show in the next proposition that under mild conditions the optimal mechanism is an affine maximiser for some classic problems.

Proposition 3.22. *In case of the k -unit auction where agents have unit demand and their valuation is i.i.d. according to an MHR distribution, the optimal mechanism is an affine maximiser.*

The role of the γ parameters can be interpreted as to introduce reserves for different allocations. This way we can withhold the item in case of low value bids and charge more from the higher valued bidders. Intuitively Base converges to a single-item auction as the exclusivity margins dominate the shared valuations, and to a k -unit auction in the other extreme. Therefore the following mechanism can be considered a reasonable candidate to boost the revenue compared to the VCG.

Definition 3.23 (Reserve mechanism (RM)). Let r_F and r_C denote the shared and exclusivity reserve values respectively (see Definition 3.1). Then RM is defined as an affine maximiser, where $\lambda^i = 1$ for all i , $\mathcal{O} = \{a \subseteq N \mid |a| \leq k\}$,

$\gamma^{(i)} = -r_C$ for all i , $\gamma^0 = 0$, $\gamma^a = -r_F|a|$ for all $a \subseteq N \mid 2 \leq |a| \leq k$ and h^i is the zero function for all i .

The reason for staying within the affine maximiser family is that just introducing reserves for the shared or exclusive allocation and allocating efficiently is not incentive compatible in general demonstrated by Example 3.24.

Example 3.24. Let $N = \{1, 2\}$, $k = 2$, $t^1 = (3, 3)$, $t^2 = (2, 1)$ and $r_F = 0$, $r_C = 5.5$. Then the efficient allocation that meets the reserve is that agent 1 receives the item exclusively. Now, set $\hat{t}^1 = (1, 4)$ and note that the efficient allocation that satisfies the reserve is to share the item between agent 1 and 2. This allocation violates monotonicity condition (MON).

Note that if $r_F = r_C = 0$, RM is equivalent to VCG. The idea behind the choice of these parameters is the following. The γ for the shared allocations are set such that they are mimicking the optimal parameters of an optimal k -item auction where types are distributed according to F . The difference is in the parameters for the exclusive allocations. For those γ is set such that it is optimal for single-item auctions given that types are distributed as C .

RM could be an appealing choice for practitioners because it is simple and fast enough to be implemented, but complex enough to be able to capitalise on the particularity of the multi-dimensional valuations. The reserve prices r_F and r_C are determined uniquely for each “product” based on knowledge about the distribution of reservation prices for that product. For instance, for Reply’s lead marketing auctions for automotive sales, this would mean computing these values for each “make” and “model” combination for groups of geographical locations which are similar in their distribution of reservation prices. Notably this requires no more information than that needed for designing a single-dimensional mechanism. Moreover, as the experimental evaluations in the next section highlight, its expected revenue is very close to the optimal one regardless of the number of agents.

3.5 COMPUTATIONAL ANALYSIS

In this section we illustrate the performance of the discussed mechanisms in case of various type distributions. The expected revenues are acquired via

simulations in the following way. We sample from the type distribution many times and calculate the payment for each type report. The expected revenue is then estimated by the average of the simulated payments. To show the robustness of the results we also derive the 99% confidence intervals for each point estimate. The upper (lower) limit of the confidence interval is calculated as

$$\beta + (-)t_{\alpha} \frac{\sigma}{\sqrt{n^*}}, \quad (25)$$

where n^* is the number of simulations, β is the average and σ is the standard deviation of the simulated values, while t_{α} is Student's t -distribution value for the given critical level α and degrees of freedom $n^* - 1$. For our experiments we set $\alpha = 99\%$ and $n^* = 20000$, therefore the corresponding t -value is $t_{99\%} \approx 2.6$. \bar{a} and σ are estimated from the samples for each simulation. For each setting we included the upper bound on the optimal revenue (UB), which is calculated as the minimum of $Rev(UBM^*)$ and $Rev(UBR^*)$.

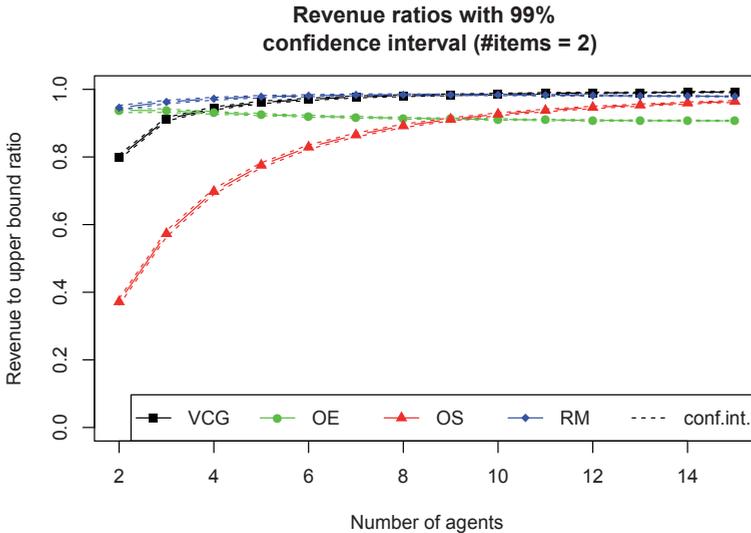


Figure 3: Revenue to upper bound ratios ($s^i \sim U(0, 1)$, $m^i \sim U(0, 1)$).

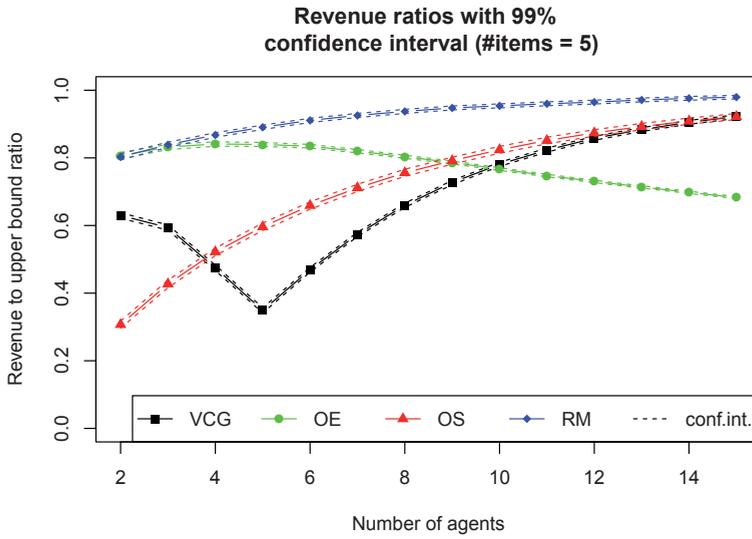


Figure 4: Revenue to upper bound ratios ($s^i \sim \text{Exp}(0.5)$, $m^i \sim \text{Exp}(0.5)$).

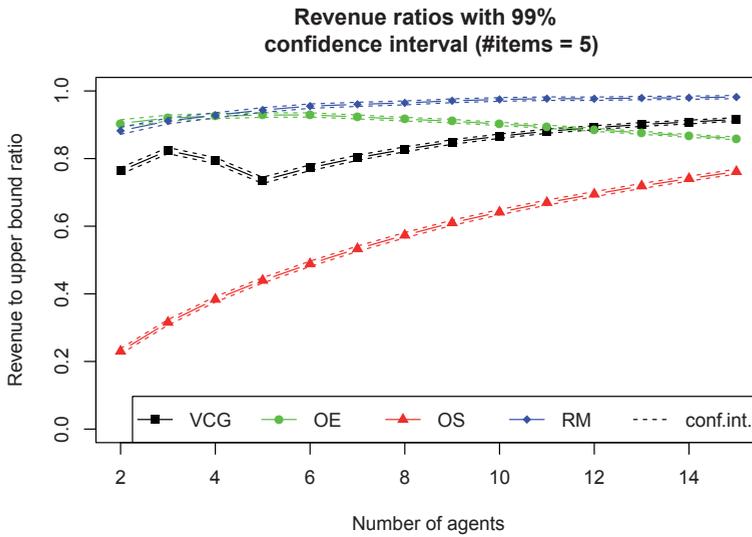


Figure 5: Revenue to upper bound ratios ($s^i \sim \text{Exp}(0.5)$, $m^i \sim \text{Exp}(0.25)$).

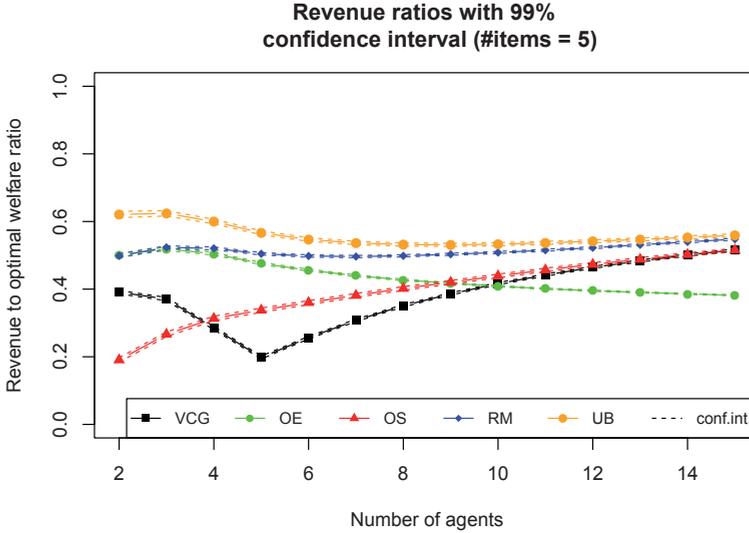


Figure 6: Revenue to optimal welfare ratios ($s^i \sim \text{Exp}(0.5)$, $m^i \sim \text{Exp}(0.5)$).

Graphs 3, 4 and 5 depict the revenues of different mechanisms relative to the upper bound. Graph 6 shows the revenue extraction from the optimal welfare. It is apparent from the graphs that RM is close to OE when the number of agents is small and close to OS when the number of agents is large. In the middle cases RM clearly outperforms the other two. The experiments indicate that RM exhibit a robust performance that is close to BASE^* regardless of the number of agents or items. This is another advantage of implementing RM, as it ensures that the seller doesn't have to worry about the number of agents showing up for the auction as opposed to the cases of VCG, OE or OS.

The intuition behind the simulation results is the following. When the exclusivity margin is high relative to the shared valuation and the number of agents is small then the exclusive allocation prevails and the more we allocate exclusively the closer OE is to the optimal auction. If there are a lot of bidders and the shared valuation dominates, then the shared allocation can bring more value therefore OS is converging to the optimal mechanism. As RM is parametrised such that it tries to catch the essence of these two mechanisms, it is close to them in the extreme cases and outperforms both in the

average scenarios. This phenomena seems to be independent of the underlying distributions therefore in this paper we present only a couple of results for the uniform and the exponential distribution.

3.6 CONCLUSION

We address the problem of designing a selling mechanism for non-rivalrous goods when exclusive purchase confers an additional value to buyers beyond their basic shared valuation for the good, so that the seller must tradeoff between a single sale at high price and multiple sales at lower prices. Online lead generators face this problem hundreds of times every second, and must make allocation and pricing decisions with a speedy, simple, and efficient mechanism. The problem also occurs in other applications with large scale (e. g., search engines) as well as others that exclusive alignments arise due to strategic considerations (e. g., games for video gaming consoles, applications for mobile operating systems, etc.).

While finding a complete, optimal multi-dimensional mechanism is intractable and the solution would be extremely complex in practice, we are able to design heuristic mechanisms that are simple, practical, and imply only a modest sacrifice on revenue. Moreover, we demonstrate how small the gap is between the heuristic mechanisms and the optimal revenue in a novel way via a combination of provable upper bounds on the optimal revenue and extensive numerical simulations. The contribution of this paper hence is manifold. We are able to derive two non-trivial upper bounds on the optimal revenue of a multi-dimensional problem by defining and solving single-dimensional relaxations of it. These bounds are shown to be significantly tighter than the traditional indicator, the optimal welfare. In addition, they also make possible to prove that a simple mechanism that either always sells exclusively or shared is a constant-factor approximation. We furthermore construct a heuristic mechanism that takes into account the multidimensionality of the valuations, but fits practical implementations: it is simple; easy to implement and its parameters can be easily tailored to the case at hand. Moreover, numerical experiments show that it yields a revenue that is very close to the upper bounds, demonstrating the power of our mechanism and devised bounds.

APPENDIX

Proof of Theorem 3.2. (UBM) can be decomposed by treating all possible exclusivity margin profiles separately, because none of the constraints involve variables that are related to different exclusivity margins, furthermore the probabilities are independent among agents and among the dimensions of their type. Therefore it is sufficient to solve the subproblems separately for each fixed m and acquire the expected revenue as the expectation over the optimal objective values of the subproblems. From Proposition 6 in Deng and Pekeč [21] it follows that for fixed m the allocation rule x of the optimal DSIC and EPIR mechanism given type profile $t = (s, m)$, is

$$x^i(t) = 1 \text{ for } \begin{cases} i \in a_s(t) & \text{when } \sum_{j \in a_s(t)} \phi_F(s^j) \geq \phi_F(s^{a_e(t)}) + m^{a_e(t)} \\ i = a_e(t) & \text{when } \sum_{j \in a_s(t)} \phi_F(s^j) < \phi_F(s^{a_e(t)}) + m^{a_e(t)} \end{cases} \quad (26)$$

and $x^i(t) = 0$ otherwise. The expected revenue under this allocation rule is equal to

$$\mathbb{E}_s \left[\sum_i \pi^i(s, m) \right] = \mathbb{E}_s \left[\sum_i \left(\phi_F(s^i) + m^i \prod_{j \neq i} (1 - x^j(s, m)) \right) x^i(s, m) \right]. \quad (27)$$

Since 26 holds for any fixed m , it concludes 13. Finally, 14 follows from taking expectation over 27 with respect to the exclusivity margins.

Note that Deng and Pekeč [21] originally studies ex-post incentive compatible (EPIC) mechanisms, because their exclusivity margin is assumed to be a linear combination of the shared valuations of the other agents. In our case exclusive margins do not depend on the valuation of the other agents, hence the notions DSIC and EPIC coincide (for formal definition of EPIC see Definition 3.5). \square

Proof of Proposition 3.6. Before proceeding to the main proof it is useful to get some insights on the allocation rule of any DSIC mechanism for BASE. The following lemma states that the higher your valuation is, the more chance you have to receive the item. The statements are grouped based on

which dimension of the type is changed. Fix agent i , the type of the others t^{-i} and the allocation rule x . Let $R_e(x, t^{-i}) \subseteq T^i$, $R_s(x, t^{-i}) \subseteq T^i$ be the set of types such that the allocation rule assigns the item to agent i exclusively and shared respectively. Similarly, $R_0(x, t^{-i}) \subseteq T^i$ is the set of types such that agent i gets nothing.

Lemma 3.25. *Let x be a monotone allocation rule and fix agent i and the report of the others t^{-i} . Then the following two statements hold:*

- A. *For all s^i, m^i, \hat{m}^i such that $\hat{m}^i > m^i$, we have that if $(s^i, m^i) \in R_e(x, t^{-i})$, then $(s^i, \hat{m}^i) \in R_e(x, t^{-i})$.*
- B. *There is an $m^* \in M^i \cup \{\infty\}$ such that for all s^i, m^i, \hat{s}^i , where $\hat{s}^i > s^i$ and $m^i \neq m^*$, we have that if $(s^i, m^i) \in R_s(x, t^{-i})$, then $(\hat{s}^i, m^i) \in R_s(x, t^{-i})$.*

Proof Assume that the first statement is not true, that is, $(s^i, \hat{m}^i) \in R_0(x, t^{-i})$ or $(s^i, \hat{m}^i) \in R_s(x, t^{-i})$. Applying (MON) for both cases with $t^i = (s^i, m^i)$ and $\hat{t}^i = (s^i, \hat{m}^i)$ results in

$$0 \geq \hat{m}^i - m^i.$$

This is a contradiction as $\hat{m}^i > m^i$ by assumption.

For the second statement define m^* as the infimum of exclusivity margins such that agent i is allocated exclusively, i. e.,

$$m^* = \inf\{\tilde{m}^i \mid (\tilde{s}^i, \tilde{m}^i) \in R_e(x, t^{-i})\}. \quad (28)$$

Now, assume the contrary of the second statement, that is, there are s^i, m^i, \hat{s}^i such that $\hat{s}^i > s^i$, $m^i \neq m^*$, $(s^i, m^i) \in R_s(x, t^{-i})$ and we have that $(\hat{s}^i, m^i) \in R_0(x, t^{-i})$ or $(\hat{s}^i, m^i) \in R_e(x, t^{-i})$. First tackle the case when $(\hat{s}^i, m^i) \in R_0(x, t^{-i})$. Applying (MON) with $t^i = (s^i, m^i)$ and $\hat{t}^i = (\hat{s}^i, m^i)$ results in

$$0 \geq \hat{s}^i - s^i.$$

This is a contradiction as $\hat{s}^i > s^i$ by assumption. Finally, assume that $(\hat{s}^i, m^i) \in R_e(x, t^{-i})$. Note that due to the definition of m^* this assumption implies that $m^i > m^*$ and that there exists a $(\tilde{s}^i, \tilde{m}^i)$ such that $\tilde{m}^i < m^i$ and $(\tilde{s}^i, \tilde{m}^i) \in R_e(x, t^{-i})$. Applying (MON) with $t^i = (s^i, m^i)$ and $\hat{t}^i = (\tilde{s}^i, \tilde{m}^i)$ results in

$$\tilde{s}^i + \tilde{m}^i - s^i - m^i \geq \hat{s}^i - s^i.$$

This is a contradiction as $m^i > \tilde{m}^i$ by assumption. \square

Now, we are ready to prove directly Proposition 3.6. By construction we have that for every $i \in N$ there is a pair $(i_s, i_e) \in \bar{N}$, and for every $t \in T$ there is a $\bar{t} \in \bar{T}$ such that $t^i = (s^i, m^i) = (\bar{t}^{i_s}, \bar{t}^{i_e})$. The steps of the proof are the following. First, for every feasible mechanism (x, π) for (Base), we define an allocation rule \bar{x} for (UBR) that satisfies (FeasA) and (FeasB). Then we construct a $\bar{\pi}$ such that $(\bar{x}, \bar{\pi})$ satisfies (EPIC) and (EPIR2). Finally, we show that $\mathbb{E}_{\bar{t} \in \bar{T}} \left[\sum_{i \in \bar{N}} \bar{\pi}^i(\bar{t}) \right] \geq \mathbb{E}_{t \in T} \left[\sum_{i \in N} \pi^i(t) \right]$.

Fix agent i and the type of the others t^{-i} . Let $t = (t^i, t^{-i})$, where $t^i = (s^i, m^i)$. Define m^* as in (28). Now, set $\bar{x}^{i_s}(\bar{t}) = 1$ if and only if $m^i \neq m^*$ and $t^i \in R_e(x, t^{-i})$, and set $\bar{x}^{i_e}(\bar{t}) = 1$ if and only if $t^i \in R_e(x, t^{-i})$. In words, whenever x allocates shared to agent i , then \bar{x} gives the item to representative i_s provided that $m^i \neq m^*$. Furthermore, if x allocates exclusively to agent i , then \bar{x} gives the item to representative i_e . It is easy to see that \bar{x} satisfies (FeasA) and (FeasB).

Next, let us define the payment rule $\bar{\pi}$ for agent $i_s \in \bar{N}_s$ as

$$\bar{\pi}^{i_s}(\bar{t}^{i_s}, \bar{t}^{-i_s}) = \begin{cases} \inf\{\hat{t}^{i_s} \mid \bar{x}^{i_s}(\hat{t}^{i_s}, \bar{t}^{-i_s}) = 1\} & \text{if } \bar{x}^{i_s}(\bar{t}^{i_s}, \bar{t}^{-i_s}) = 1 \\ 0 & \text{otherwise.} \end{cases} \quad (29)$$

and for agent $i_e \in \bar{N}_e$ as

$$\begin{aligned} & \bar{\pi}^{i_e}(\bar{t}^{i_e}, \bar{t}^{i_s}, \bar{t}^{-(i_s, i_e)}) \\ &= \begin{cases} \inf\{\hat{t}^{i_e} \mid \bar{x}^{i_e}(\hat{t}^{i_e}, \bar{t}^{i_s}, \bar{t}^{-(i_s, i_e)}) = 1\} + \bar{t}^{i_s} & \text{if } \bar{x}^{i_e}(\bar{t}^{i_e}, \bar{t}^{-i_e}) = 1 \\ 0 & \text{otherwise.} \end{cases} \end{aligned} \quad (30)$$

To see incentive compatibility, consider agent i_e and fix $(\bar{t}^{i_s}, \bar{t}^{-(i_s, i_e)})$. If there is no \bar{t}^{i_e} such that $\bar{x}^{i_e}(\bar{t}^{i_s}, \bar{t}^{i_e}, \bar{t}^{-(i_s, i_e)}) = 1$, then (EPIC) and (EPIR2) trivially hold. Otherwise let $\underline{m} = \inf\{\hat{t}^{i_e} \mid \bar{x}^{i_e}(\hat{t}^{i_e}, \bar{t}^{i_s}, \bar{t}^{-(i_s, i_e)}) = 1\}$. Note that Lemma 3.25 ensures that $\bar{x}^{i_e}(\bar{t}^{i_s}, \bar{t}^{i_e}, \bar{t}^{-(i_s, i_e)})$ is non-decreasing in \bar{t}^{i_e} for all \bar{t}^{-i_e} . This implies that $\bar{x}^{i_e}(\bar{t}^{i_s}, \bar{t}^{i_e}, \bar{t}^{-(i_s, i_e)}) = 1$ for all $\bar{t}^{i_e} > \underline{m}$, and $\bar{x}^{i_e}(\bar{t}^{i_s}, \bar{t}^{i_e}, \bar{t}^{-(i_s, i_e)}) = 0$ for all $\bar{t}^{i_e} < \underline{m}$. It is easy to see that this allocation rule together with pay-

ment $\bar{\pi}^{ie}(\bar{t}^{ie}, \bar{t}^{is}, \bar{t}^{-(is, ie)}) = \underline{m} + \bar{t}^{is}$ satisfy (EPIC) and (EPIR2). With respect to the comparison of the payments we can observe the following. When $x(t)$ allocates exclusively to agent i , then

$$\begin{aligned} \pi^i(s^i, m^i, t^{-i}) &\leq \inf\{v^i(x(\hat{s}^i, \hat{m}^i, t^{-i}), (\hat{s}^i, \hat{m}^i)) \mid x(\hat{s}^i, \hat{m}^i, t^{-i}) = x(s^i, m^i, t^{-i})\} \\ &\leq \inf\{\bar{t}^{is} + \hat{t}^{ie} \mid \bar{x}^{is}(\hat{t}^{is}, \bar{t}^{ie}, \bar{t}^{-(is, ie)}) = 1, \bar{t}^{is} = s^i\} \\ &= \bar{\pi}^{ie}(\bar{t}^{is}, \bar{t}^{ie}, \bar{t}^{-(is, ie)}). \end{aligned}$$

The inequalities follow from Lemma 1.6 and from the fact that the second infimum is taken over a smaller set. This means that in case of exclusive allocation $\bar{\pi}^{ie}(\bar{t}) + \bar{\pi}^{is}(\bar{t}) \geq \pi^i(t)$.

Finally, consider agent i_s and fix $(\bar{t}^{ie}, \bar{t}^{-(is, ie)})$. If $\bar{t}^{ie} = m^*$ or there is no \bar{t}^{is} such that $\bar{x}^{is}(\bar{t}^{is}, \bar{t}^{ie}, \bar{t}^{-(is, ie)}) = 1$, then $\bar{x}^{is}(\bar{t}^{is}, \bar{t}^{ie}, \bar{t}^{-(is, ie)}) = 0$ for all \bar{t}^{is} , therefore (EPIC) and (EPIR2) trivially hold. Otherwise assume that $\bar{t}^{ie} \neq m^*$ and let $\underline{s} = \inf\{\hat{t}^{is} \mid \bar{x}^{is}(\hat{t}^{is}, \bar{t}^{ie}, \bar{t}^{-(is, ie)}) = 1\}$. Note that Lemma 3.25 ensures that $\bar{x}^{is}(\bar{t}^{is}, \bar{t}^{ie}, \bar{t}^{-(is, ie)})$ is non-decreasing in \bar{t}^{is} for all $(\bar{t}^{ie}, \bar{t}^{-(is, ie)})$. As a result we have that $\bar{x}^{is}(\bar{t}^{is}, \bar{t}^{ie}, \bar{t}^{-(is, ie)}) = 1$ for all $\bar{t}^{is} > \underline{s}$, furthermore $\bar{x}^{is}(\bar{t}^{is}, \bar{t}^{ie}, \bar{t}^{-(is, ie)}) = 0$ for all $\bar{t}^{is} < \underline{s}$. It is easy to see that this allocation rule together with payment $\bar{\pi}^{is}(\bar{t}^{is}, \bar{t}^{ie}, \bar{t}^{-(is, ie)}) = \underline{s}$ satisfy (EPIC) and (EPIR2). With respect to the comparison of the payments we can observe the following. When $x(t)$ allocates shared to agent i and $\bar{t}^{ie} \neq m^*$, then

$$\begin{aligned} \pi^i(s^i, m^i, t^{-i}) &\leq \inf\{v^i(x(\hat{s}^i, \hat{m}^i, t^{-i}), (\hat{s}^i, \hat{m}^i)) \mid x(\hat{s}^i, \hat{m}^i, t^{-i}) = x(s^i, m^i, t^{-i})\} \\ &\leq \inf\{\hat{t}^{is} \mid \bar{x}^{is}(\hat{t}^{is}, \bar{t}^{ie}, \bar{t}^{-(is, ie)}) = 1, \bar{t}^{ie} = m^i\} \\ &= \bar{\pi}^{is}(\bar{t}^{is}, \bar{t}^{ie}, \bar{t}^{-(is, ie)}). \end{aligned}$$

The inequalities follow from Lemma 1.6 and from the fact that the second infimum is taken over a smaller set. This means that in case of shared allocation $\bar{\pi}^{ie}(\bar{t}) + \bar{\pi}^{is}(\bar{t}) \geq \pi^i(t)$.

Combining the results of all subcases implies that for any feasible mechanism (x, π) for (Base) there is a feasible $(\bar{x}, \bar{\pi})$ for (UBR) such that

$$\begin{aligned}
\mathbb{E}_t \left[\sum_{i \in N} \pi^i(t) \right] &= \sum_i \mathbb{E}_{t^{-i}} \left[\mathbb{E}_{(s^i, m^i)} \left[\pi^i(s^i, m^i, t^{-i}) \right] \right] \\
&= \sum_{i \in N} \mathbb{E}_{t^{-i}} \left[\mathbb{E}_{(s^i, m^i), m^i \neq m^*} \left[\pi^i(s^i, m^i, t^{-i}) \right] \right] \\
&\leq \sum_{i \in N} \mathbb{E}_{\bar{t}^{-(i_s, i_e)}} \left[\mathbb{E}_{t^{i_s, t^{i_e}}, t^{i_e} \neq m^*} \left[\bar{\pi}^{i_s}(t^{i_s}, t^{i_e}, \bar{t}^{-(i_s, i_e)}) + \bar{\pi}^{i_e}(t^{i_s}, t^{i_e}, \bar{t}^{-(i_s, i_e)}) \right] \right] \\
&= \sum_{i \in N} \mathbb{E}_{\bar{t}^{-(i_s, i_e)}} \left[\mathbb{E}_{t^{i_s, t^{i_e}}} \left[\bar{\pi}^{i_s}(t^{i_s}, t^{i_e}, \bar{t}^{-(i_s, i_e)}) + \bar{\pi}^{i_e}(t^{i_s}, t^{i_e}, \bar{t}^{-(i_s, i_e)}) \right] \right] \\
&= \mathbb{E}_{\bar{t}} \left[\sum_{i \in \bar{N}} \bar{\pi}^i(\bar{t}) \right].
\end{aligned}$$

The second and third equalities hold because type distributions are continuous, hence the measure of the type set where $t^{i_e} = m^*$ is zero. The inequality follows from the results of the previous subcases. This concludes the proof of Proposition 3.6. \square

Proof of Theorem 3.7. Roughgarden and Talgam-Cohen [64] extends the results of Myerson [54] to settings with informational externalities and correlated type distributions for mechanism that are ex-post incentive compatible and ex-post individual rational. In particular, it is shown that the expected revenue of any EPIC and EPIR mechanism equals to expected sum of virtual valuations provided that the payments are maximal, i. e., the utility of an agent with zero type is zero. Moreover, the revenue maximising mechanism allocates such that the sum of virtual valuations is maximal for each type profile, given that the resulting allocation is monotone non-decreasing for each agent in their own type. Note that the virtual valuation of agent i_e having type \bar{t}^{i_e} is $\bar{t}^{i_s} + \phi_G(\bar{t}^{i_e})$ due to the informational externality. Having all these in mind the only thing left to show is that all agent's probability of receiving the item is monotone non-decreasing in their reported types.

As F has monotone hazard rate, we have that $\phi_F(\bar{t}^{i_s})$ is monotone non-decreasing in \bar{t}^{i_s} for all $i_s \in \bar{N}_s$. Therefore if $i_s \in \bar{a}_s(\bar{t}^{i_s}, \bar{t}^{-i_s})$, then $i_s \in$

$\bar{a}_s(\hat{t}^{i_s}, \bar{t}^{-i_s})$ for any $\hat{t}^{i_s} > \bar{t}^{i_s}$. Similarly, as G has monotone hazard rate, we have that $\phi_G(\bar{t}^{i_e})$ is also non-decreasing in \bar{t}^{i_e} for all $i_e \in \bar{N}_e$. This ensures that once $i_e \in \bar{a}_e(\bar{t}^{i_e}, \bar{t}^{-i_e})$, then $i_e \in \bar{a}_e(\hat{t}^{i_e}, \bar{t}^{-i_e})$ also for any $\hat{t}^{i_e} > \bar{t}^{i_e}$. The last critical point is that when any agent $i_s \in \bar{N}_s$ increases \bar{t}^{i_s} , because then both $\phi_F(\bar{t}^{i_s})$ and $\bar{t}^{i_s} + \phi_G(\bar{t}^{i_e})$ increase at the same time. To see why it is not a problem observe that the MHR assumption ensures that $\phi_F(\bar{t}^{i_s})$ increases at least as fast as $\bar{t}^{i_s} + \phi_G(\bar{t}^{i_e})$ in \bar{t}^{i_s} . \square

Proof of Theorem 3.9. According to (15) the expected value of the VCG payments can be written as

$$\begin{aligned}
& Rev(\text{VCG}) \\
&= \mathbb{E}_{t \in T_N} \left[\sum_{i \in N} \left(\sum_{j \in N \setminus \{i\}} v^j(a(t^{-i}), t^j) - \sum_{j \in N \setminus \{i\}} v^j(a(t), t^j) \right) \right] \\
&= \mathbb{E}_{t \in T_N} \left[\sum_{i \in N} \sum_{j \in N \setminus \{i\}} v^j(a(t^{-i}), t^j) \right] - \mathbb{E}_{t \in T_N} \left[\sum_{i \in N} \sum_{j \in N \setminus \{i\}} v^j(a(t), t^j) \right] \\
&= \sum_{i \in N} \mathbb{E}_{t^{i \in T^i}} \left[\mathbb{E}_{t^{-i \in T_N^{-i}}} \left[\sum_{j \in N \setminus \{i\}} v^j(a(t^{-i}), t^j) \right] \right] - (n-1) \mathbb{E}_{t \in T_N} [SW^N(a(t), t)] \\
&= \sum_{i \in N} \mathbb{E}_{t^{i \in T^i}} \left[\mathbb{E}_{t \in T_Q} \left[\sum_{j \in Q} v^j(a(t^{-i}), t^j) \right] \right] - (n-1) \mathbb{E}_{t \in T_N} [SW^N(a(t), t)] \\
&= n \mathbb{E}_{t \in T_Q} \left[\sum_{j \in Q} v^j(a(t), t^j) \right] - (n-1) \mathbb{E}_{t \in T_N} [SW^N(a(t), t)] \\
&= n \mathbb{E}_{t \in T_Q} [SW^Q(a(t), t)] - (n-1) \mathbb{E}_{t \in T_N} [SW^N(a(t), t)].
\end{aligned}$$

The equalities are direct consequences of the assumption that types are i.i.d. \square

Proof of Theorem 3.12. Let $N = \{1, \dots, n\}$ be the set of agents with i.i.d. types, and let $Q = \{1, \dots, n-1\}$ denote the same set of agents with one less

member. Furthermore, denote the l^{th} largest value from n samples by $v_{[l:n]}$. Lemma 3.11 implies that

$$\mathbb{E} \left[v_{[l:n-1]} \right] \geq (H_{n-1} - H_{l-1}) / (H_n - H_{l-1}) \mathbb{E} \left[v_{[l:n]} \right].$$

Therefore we have that

$$\begin{aligned} \frac{\mathbb{E}_{t \in T_Q} \left[SW^Q(x(t), t) \right]}{\mathbb{E}_{t \in T_N} \left[SW^N(x(t), t) \right]} &= \frac{\mathbb{E} \left[v_{[1:n-1]} \right]}{\mathbb{E} \left[v_{[1:n]} \right]} \\ &\geq \frac{(H_{n-1}/H_n) \mathbb{E} \left[v_{[l:n]} \right]}{\mathbb{E} \left[v_{[l:n]} \right]} \\ &= 1 - \frac{1}{nH_n}. \end{aligned}$$

The proof is concluded by invoking Corollary 3.10 and letting $\varrho = \frac{1}{nH_n}$. \square

Proof of Theorem 3.19. According to Theorem 3.2 $Rev(\text{UBM}^*)$ is an upper bound on $Rev(\text{BASE}^*)$, therefore it is sufficient to show that $Rev(\text{MAXSIMPLE})$ approximates $Rev(\text{UBM}^*)$. Let x_{UBM^*} represent the allocation rule of UBM^* . Define $T_{\text{UBM}^*}^e = \{t \in T \mid \exists i : x_{\text{UBM}^*}^i(t) = 1, \forall j \neq i : x_{\text{UBM}^*}^j(t) = 0\}$ and $T_{\text{UBM}^*}^s = \{t \in T \mid \exists i, j \neq i : x_{\text{UBM}^*}^i(t) = x_{\text{UBM}^*}^j(t) = 1\}$. Then according to Theorem 3.2 the revenue of UBM^* can be split such that $Rev(\text{UBM}^*) = Rev(\text{UBM}^*)_s + Rev(\text{UBM}^*)_e$, where

$$Rev(\text{UBM}^*)_s = \mathbb{E}_{t \in T_{\text{UBM}^*}^s} \left[\sum_i \phi_F(s^i) x_{\text{UBM}^*}^i \right] \quad (31)$$

and

$$Rev(\text{UBM}^*)_e = \mathbb{E}_{t \in T_{\text{UBM}^*}^e} \left[\sum_i (\phi_F(s^i) + m^i) x_{\text{UBM}^*}^i \right]. \quad (32)$$

The idea of the proof is to bound the two terms separately.

We start with the term $Rev(\text{UBM}^*)_e$. As $\phi_F(s^i) + m^i \leq s^i + m^i$ for all i , we have that $Rev(\text{UBM}^*)_e$ is less than or equal to the optimal welfare of a single-item auction. Note that OE achieves at least as much expected rev-

enue as VCG does for the single-item auction. Therefore due to Theorem 3.12 $Rev(OE)$ is at least $1 - 1/H_n$ times the optimal welfare of a single-item auction. This leads to the conclusion that $Rev(OE) \geq (1 - 1/H_n) Rev(UBM^*)_e$.

To bound $Rev(UBM^*)_s$ observe that whenever UBM^* allocates shared for type report t , then $x_{UBM^*}(t) = x_{OS}(t)$. This is due to the fact $a_s(t)$ is defined the same way for both mechanism. Using this observation together with Lemma 3.17 we can write

$$\begin{aligned}
 Rev(OS) &= \mathbb{E}_t \left[\sum_i \phi_F(s^i) x_{OS}^i(t) \right] \\
 &= \mathbb{E}_{t \in T_{UBM^*}^s} \left[\sum_i \phi_F(s^i) x_{OS}^i(t) \right] + \mathbb{E}_{t \notin T_{UBM^*}^s} \left[\sum_i \phi_F(s^i) x_{OS}^i(t) \right] \\
 &= Rev(UBM^*)_s + \mathbb{E}_{t \notin T_{UBM^*}^s} \left[\sum_i \phi_F(s^i) x_{OS}^i(t) \right] \\
 &\geq Rev(UBM^*)_s.
 \end{aligned}$$

The last inequality holds because under OS allocation only occur when the sum of virtual valuations is non-negative. Putting together the two bounds results in

$$\begin{aligned}
 \frac{1}{1 - 1/H_n} Rev(OE) + Rev(OS) &\geq Rev(UBM^*)_s + Rev(UBM^*)_e \\
 &= Rev(UBM^*) \\
 &\geq Rev(BASE^*).
 \end{aligned}$$

To finish the proof note that $Rev(MAXSIMPLE) = \max\{Rev(OS), Rev(OE)\}$, hence

$$\left(\frac{1}{1 - 1/H_n} + 1 \right) Rev(MAXSIMPLE) \geq Rev(BASE^*).$$

Finally, note that we can replace DSIC to Bayes-Nash incentive compatibility in (Base) and relax the resulting mathematical program by letting the exclusivity margin public information as in (UBM). As we arrive at a single-dimensional setting it is folklore knowledge that the optimal Bayes-Nash mechanism is DSIC. This means that UBM^* is optimal even among Bayes-

Nash implementable mechanisms, therefore it is also an upper bound for the Bayesian relaxation of (Base). \square

Proof of Proposition 3.22. According to Myerson [54] the optimal mechanism gives the item to the k -highest value bidder who meets the reserve price and charges recipients the maximum of the $k + 1$ th highest value and the reserve price. Now, let \mathcal{O} be the set of possible allocations that consists of sets of agents that can be served simultaneously, i. e., $\mathcal{O} = \{a \subseteq N \mid |a| \leq k\}$, and let r be the reserve price. By taking $\lambda^i = 1, \forall i \in N$ and $\gamma^a = -r|a|, \forall a \in \mathcal{O}$ it is easy to see that (23) allocates the item only to those agents who meet the reserve price and gives away at most k items. Moreover, the payment of (24) is also equivalent to the maximum of the $k + 1$ th highest value and the reserve price. \square

SUBCONTRACTOR SCHEDULING

This chapter is based on the following paper:

Martijn van Brink, Gergely Csapó and Ruben van der Zwaan, “Sub-contractor Scheduling”.

4.1 INTRODUCTION

By offering external resources and expertise subcontractors help companies to reduce the completion time of their projects and hence to save cost. Subcontracting is a common practice in industries that face volatile demand and hard capacity constraints. When the time of the subcontractor has to be allocated to multiple projects, special attention has to be paid to scheduling. The Semiconductor Product Analysis and Design Enhancement (SPADE) centre of Hong Kong University of Science and Technology serves as a practical example for subcontractor scheduling. SPADE provides consulting services to semiconductor companies who can book their service time via an online reservation system. The customers are served in a first-come, first-book manner and the price of different time slots may vary.

Further complications arise from the fact that larger projects are usually run by different divisions or teams whose interest is to maximise their own savings. Therefore companies face a global optimisation problem in which coordination and incentives have to be taken into account. To illustrate the importance of this issue we recall the case of Boeing’s Dreamliner supply chain, where the mismanagement of subcontractors resulted in overloaded schedules and long delays in the overall production. See Vairaktarakis [79] for further details and references.

We model this problem as follows. There are multiple jobs endowed with a dedicated machine, where each job is characterised by a release date, a

processing time, and a weight that represents the opportunity costs of being finished later. We assume that there is only a single subcontractor that assists jobs in reducing their completion time by speeding up their operation by a factor of two. The subcontractor’s capacity is limited, hence it can only handle one job at a time. A schedule for the subcontractor therefore has to indicate for each time point which job to assist. We allow for preemption, i. e., it is possible to switch to more urgent tasks any time. This assumption is natural when the subcontractor’s resources are only used for accelerating processes, e. g., in case of consulting services. Given this setting, our goal is to find a schedule for the subcontractor that maximises the weighted sum of savings. To facilitate the understanding of the problem, a small instance is visualised in Figure 7.

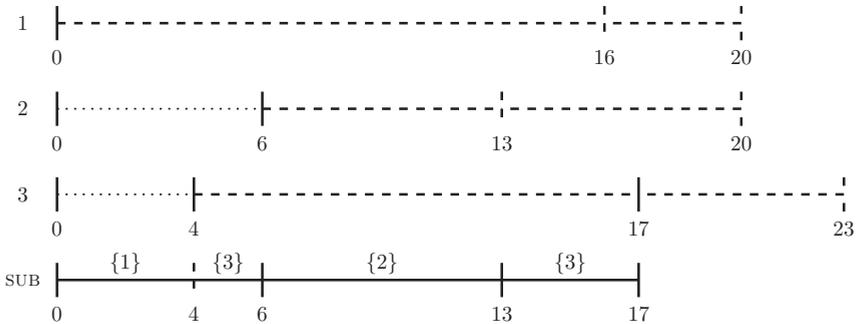


Figure 7: A small example of a feasible schedule for three jobs, where 1-3 are dedicated machines, and SUB is the subcontractor machine. The release dates are (0,6,4), the processing times are (20,14,19), and the weights are (5,10,8), respectively. Dashed lines indicate the additional time required on the dedicated machine, if the subcontractor is not employed. The time savings obtained by the jobs are (4,7,6), respectively. The weighted sum of savings is 138.

Note that although there are multiple machines, only the schedule of the subcontractor is a decision variable. Therefore in essence we face a single machine scheduling problem, in which the gain of each allocated time unit equals the weight of the job it is spent on. Furthermore, jobs are being processed on their own machine independently and the subcontractor only provides a twofold speed up for the scheduled job.

We can also interpret it as a problem of scheduling malleable and preemptible jobs on parallel machines, where each job has its own fixed processor and there is only one extra machine with which load can be shared. The original problem of malleable job scheduling is \mathcal{NP} -hard, but due to the restrictions on the processors' usage, it is not apparent whether this complexity is preserved.

In our context, minimising the weighted sum of completion times leads to the same solution as maximising the weighted sum of savings. The reason for working with the latter objective is the economic interpretation. Companies are interested in the monetary value of the time they save by subcontracting. As the optimal solution is the same for both objectives, we can relate this problem to the vast amount of research done on the minimisation of the weighted sum of completion times in different settings. We only have to keep in mind that approximation ratios derived for one objective are not valid for the other one in general. For example, take the worst scheduling policy possible, i. e., not using the subcontractor at all just leaving it idle. The total savings yielded by this protocol is zero, which is infinitely worse than a schedule that makes any use of the subcontractor. Nevertheless, the total weighted completion time is at least half of the optimal total weighted completion time. This is because jobs are continuously processed on their own machine irrespective of the subcontractor's schedule, hence finished within their original processing time.

We study this problem both from an algorithmic and a mechanism design viewpoint. For the former, we are interested in developing and evaluating algorithms that construct feasible, or even optimal solutions, taking all inputs as given. In terms of mechanism design, our main interest is to find simple mechanisms that are speedy, robust to changes in the instances and simple to implement for a practitioner. We also would like to give guarantees for the trade-off between these desirable features and the objective value. To this end we deploy schedules from the algorithmic part of the chapter that are good approximations and we search for practical mechanisms that implement these schedules.

RELATED WORK This chapter considers the problem where jobs are not only processed on a dedicated machine, but can also be outsourced to a third

party. Hence, jobs can be processed in parallel on two machines. This type of jobs resemble infinitely divisible job-shops (see, e. g., Anderson [2]). Practical examples of such production types include parallel and distributed computer systems (see, e. g., Bharadwaj et al. [8], Drozdowski [25]). In these applications computations are performed by private computers, but additional computation power can be acquired via shared servers.

In the pioneering work of Smith [76] it is shown that in a single-stage production, the total completion time is minimised by processing the jobs in non-decreasing order of their processing times, known as the shortest processing time first (SPT) algorithm. Furthermore, if every job i has a positive weight that represents its relative importance, the total weighted completion time is minimised by ordering the jobs in non-increasing order of weight over processing time, known as SMITH'S RULE. Unfortunately, by introducing release dates the problem becomes \mathcal{NP} -hard. When jobs can be processed on parallel machines, then minimising the total weighted completion time is already \mathcal{NP} -hard without release dates even if we allow for preemption (Labetoulle et al. [46]. Bruno et al. [10]). Afrati et al. [1] contains several polynomial time approximation schemes for different variants of scheduling problems, in which the aim is to minimise the weighted sum of completion times.

There is a rich history for the classic problem of minimising the total weighted completion time of jobs in parallel identical machines. Due to the substantial differences between the original and our setting (e. g., decreasing processing times and default machines for each job) we only mention two papers that are related to ours. The article of Sitters [75] analyses the performance of the shortest remaining processing time algorithm (SRPT) and shows that its competitive ratio is $5/4$ with respect to the total completion time objective. For minimising total weighted completion time for malleable jobs it is shown in [65] that there is always an optimal schedule in which all parts of a job is finished at the same time. From mechanism design point of view the paper of [20] bears importance. In order to enforce truth-telling in a single-server system with online arrivals, the authors devise a game where jobs can boost the priority of some of their parts using tokens. The resulting behaviour in this game is that jobs schedule themselves according to an approximation of the shortest remaining processing time rule and reveal their processing time.

There is a strong link between subcontractor scheduling and scheduling jobs whose processing time is decreasing over time. One can relate the two problems by setting the decay rate for each job to one. Bachman et al. [5] prove that if the decay rate can differ per job, then minimising the weighted sum of completion times is \mathcal{NP} -hard. For a concise review of the related literature see Cheng et al. [15].

A restricted version of our problem, in which all release dates and all weights are equal, is considered in multiple papers. Vairaktarakis and Aydinliyim [80], Vairaktarakis [79] compare the performance of an optimal centralised schedule with the result of a decentralised subcontracting game under different protocols. Our problem coincides with the “overlapping” setting from that paper. The naming refers to the fact that jobs can be processed simultaneously on their own and on the subcontractor’s machine. In particular, by assuming that jobs can only be scheduled for one time interval on the subcontractor’s machine it is shown in Vairaktarakis and Aydinliyim [80] shows that the schedule which maximises the total saving orders jobs according to the SPT rule and do not preempt. We extend this result to the case of different release dates by showing that the optimal schedule always works on the job that has the shortest remaining processing time.

The previous line of research was picked up by Hezarkhani and Kubiak [38]. Building on the results of Vairaktarakis and Aydinliyim [80] they approach the problem under the aegis of mechanism design and try to devise pricing schemes and decentralised mechanisms to steer selfish jobs towards the optimal outcome. They show that even if we allow jobs to be scheduled for different time intervals on the subcontractor’s machine, the SPT algorithm remains optimal. For the case where processing times are public information, the authors construct a pricing scheme that coordinates the jobs to buy the right time intervals on the subcontractor machine. For the private information case they show that under the Vickrey-Clark-Groves mechanism (see Vickrey [82], Clarke [17], Groves [35]) jobs are strictly better off by reporting their true processing time. Moreover, their payments can be expressed in a closed-form expression. By means of an example it is demonstrated in Hezarkhani and Kubiak [38] that if the weights of jobs may differ, then there is no coordinating pricing scheme in general. Finding the optimal schedule for the weighted case is left as an open question by the authors.

OUR CONTRIBUTION We consider several scheduling algorithms and analyse their performance under different assumptions regarding the jobs' characteristics. Notably, for the most general case we show that scheduling in non-decreasing order of weights is a factor $2/3$ approximation, that is, the total weighted saving of that schedule is at least $2/3$ of the optimal value. In contrast, applying Smith's rule surprisingly is at best a factor $1/2$ approximation. For practical implementations we devise three different mechanisms. One is an auction implementation of the optimal schedule that operates under the assumption that all weights and all release dates are equal. This is a step towards the decentralisation of the scheduling process because there is no central authority that directly acquires the information and constructs the schedules, but jobs allocate themselves via the auction indirectly. For settings, in which the designer does not want to invest or gain money, we propose a mechanism where the total sum of payments is zero. The existence of such a mechanism is rare in the scheduling domain demonstrated by [77], where it is shown that for single machine sequencing there is no incentive compatible and individual rational mechanism that is budget-balanced. Finally, for the case where weights can vary, we construct a mechanism that is a factor $2/3$ approximation.

ORGANISATION In the next section we proceed with defining the model formally. Then in Section 4.3 we introduce and investigate some heuristic schedules from an algorithmic point of view. Building on the insights gained there, in Section 4.4 we implement the best schedules by constructing payment schemes to incentivise jobs and align their behaviour. Finally, Section 4.5 concludes this chapter by summarising our contributions and discussing some remaining open questions.

4.2 PRELIMINARIES

The subcontractor problem is defined as follows. We are given a set of n jobs $N = \{1, \dots, n\}$, each characterised by a triplet of private information (r^i, p^i, w^i) : a release date $r^i \in R^i$, a processing time $p^i \in P^i$, and a weight $w^i \in W^i$. Let $R = R^1 \times \dots \times R^n$, $P = P^1 \times \dots \times P^n$ and $W = W^1 \times \dots \times W^n$ denote the set of possible release date, processing time and weight tuples

respectively. For each job a dedicated machine is available, complemented with a subcontractor machine, hereafter called SUB, which can work on at most one job at a time.

A schedule $\delta = \{\delta^i\}_{i \in N}$ is a set of measurable functions $\{\delta^i : \mathbb{R}_+ \rightarrow \{0, 1\}\}_{i \in N}$ that assigns jobs to time points on SUB. The set of possible outcomes is the set of possible schedules that is denoted by \mathcal{O} . Given schedule δ , job i is completed at time $C^i(\delta, r^i, p^i)$ if the scheduled time between r^i and $C^i(\delta, r^i, p^i)$ on the dedicated machine and on SUB is equal to p^i , i. e., $C^i(\delta, r^i, p^i)$ is the unique solution to the equation

$$C^i(\delta, r^i, p^i) + \int_{r^i}^{C^i(\delta, r^i, p^i)} \delta^i(t) dt = r^i + p^i.$$

The saving that job i obtains from schedule δ is denoted by $s^i(\delta, r^i, p^i)$. Note that this is exactly the time that job i spends on SUB before $C^i(\delta, r^i, p^i)$, hence $s^i(\delta, r^i, p^i) = \int_{r^i}^{C^i(\delta, r^i, p^i)} \delta^i(t) dt$. The value of job i for schedule δ given input (r, p, w) is $v^i(\delta, r, p, w) = w^i s^i(\delta, r^i, p^i)$. The social welfare equals the sum of values over all jobs, i. e., $SW^N(\delta, r, p, w) = \sum_{i \in N} v^i(\delta, r, p, w)$. As mentioned before, preemption is allowed, that is, the session of any job on SUB can end before the completion time and can be split into any number of intervals. Whenever the input is clear from the context, we write $v^i(\delta)$, $SW(\delta)$, $s^i(\delta)$ and $C^i(\delta)$ instead of $v^i(\delta, r, p, w)$, $SW^N(\delta, r, p, w)$, $s^i(\delta, r^i, p^i)$ and $C^i(\delta, r^i, p^i)$.

Our first aim is to find schedules that maximise the total weighted saving or at least provide a good approximate solution for any given (r, w, p) . To this end, in the next section we restrict our attention to schedules and disregard incentives by assuming that all information is public. The analysis of incentives and mechanisms is postponed to Section 4.4.

4.3 APPROXIMATING SCHEDULING ALGORITHMS

In this section we compare several schedules that operate according to the same principle. At each release date or completion time the next job that is scheduled on SUB is selected from the available jobs. A job is available if it is released but not yet completed. This implies that if a job is available, then SUB cannot be idle in the optimal solution. Note that jobs can be preempted,

but only at release dates of other jobs. For any triplet (r, w, p) the goal is to find a feasible schedule that maximises the social welfare:

$$\max \sum_{i \in N} v^i(\delta, r, p, w) \text{ subject to } \sum_i \delta^i(t) \leq 1, \quad \forall t \in \mathbb{R}.$$

A schedule that maximises social welfare is called efficient. The following two observations are implied immediately by the definitions of C^i and s^i .

Observation 4.1. *For any schedule δ and any job i having type (r^i, w^i, p^i) , we have that $s^i(\delta, r^i, p^i) \leq p^i/2$.*

Observation 4.2. *For any two schedules $\delta, \hat{\delta}$, and for any job $i \in N$ we have that $s^i(\delta) \leq s^i(\hat{\delta})$ if and only if $C^i(\delta) \geq C^i(\hat{\delta})$.*

In this chapter we make repeatedly use of a basic result combined from Vairaktarakis and Aydinliyim [80] and Hezarkhani and Kubiak [38], therefore it is formally cited here for later reference.

Theorem 4.3. *Vairaktarakis and Aydinliyim [80] and Hezarkhani and Kubiak [38] Assume that $r^i = 0$ and $w^i = 1$ for all jobs $i \in N$ and index jobs in non-decreasing order of their processing times. Then the optimal schedule is non-preemptive and it schedules jobs according to the SPT algorithm. Moreover, the saving of job i is equal to*

$$s^i(\text{SPT}, p^i) = \frac{p^i}{2} - \sum_{j=1}^{i-1} \frac{s^j(\text{SPT}, p^j)}{2} = \frac{p^i}{2} - \sum_{j=1}^{i-1} \frac{p^j}{2^{i+1-j}},$$

while the sum of savings can be expressed as

$$SW^N(\text{SPT}) = \sum_{i=1}^n s^i(\text{SPT}, p^i) = \sum_{i=1}^n \frac{p^i}{2^{n+1-i}}.$$

In the following subsections we investigate schedules that simply order jobs according to some measure such as processing time, weight, weight over processing time or weight times remaining processing time.

4.3.1 Shortest Remaining Processing Time

The shortest remaining processing time (SRPT) schedule selects at any time an available job that has the lowest remaining processing time. In case of a tie, the job with the highest weight is selected. If there is still a tie, the job with the lowest index is chosen. First, consider the case where all weights are equal, i. e., $w^i = 1$ for all $i \in N$. Theorem 4.3 states that scheduling according to SRPT is optimal in the absence of release dates, i. e., if $r^i = 0$ for all jobs i . We extend this result by showing that SRPT remains optimal in the presence of release dates, if preemption is allowed.

Theorem 4.4. *If preemption is allowed, and all weights are equal, then SRPT is optimal.*

The proof of Theorem 4.4 can be found in the appendix. Note that Schrage [71] shows the optimality of SRPT under the assumption that processing times are independent of the schedule, which does not hold for our problem. Unfortunately, Theorem 4.4 does not extend to the case of arbitrary weights as the following theorem demonstrates.

Theorem 4.5. *SRPT is a factor $\frac{1}{2^{n-1}}$ approximation in case of arbitrary weights, and this ratio is tight.*

The proof of Theorem 4.5 can be found in the appendix.

4.3.2 Highest Weight First

Under the highest weight first (HWF) algorithm, at any time an available job with the highest weight is processed on SUB . In case of a tie, the job with the lowest remaining processing time is selected. If there is still a tie, the job with the lowest index is chosen.

Theorem 4.6. *HWF is a factor $\frac{2}{3}$ approximation.*

Before proceeding to the proof, first, we derive some supporting lemmas. Let OPT denote the optimal schedule. Define N^+ as the set of jobs that are scheduled on SUB in HWF for at least as much time as in OPT and let N^- be

its complement, i. e., $N^+ := \{i \in N \mid s^i(\text{HWF}) \geq s^i(\text{OPT})\}$ and $N^- := N \setminus N^+$. Let \mathcal{T} denote all time points t such that OPT schedules some job $i \in N^+$ on SUB, while HWF has completed job i before t , i. e., $\mathcal{T} := \{t \mid \exists i \in N^+ : \text{OPT}^i(t) = 1 \text{ and } t > C^i(\text{HWF})\}$. Finally, define $\overline{\text{OPTRES}}$ as OPT restricted to time points in \mathcal{T} , and its complement $\overline{\overline{\text{OPTRES}}}$,

$$\text{OPTRES}^i(t) := \begin{cases} \text{OPT}^i(t) & \text{if } t \in \mathcal{T} \\ 0 & \text{else,} \end{cases} \quad \overline{\overline{\text{OPTRES}}}^i(t) := \begin{cases} 0 & \text{if } t \in \mathcal{T} \\ \text{OPT}^i(t) & \text{else.} \end{cases}$$

Lemma 4.7. *For any $i \in N^+$, $s^i(\text{OPTRES}) \leq \frac{1}{2}s^i(\text{HWF})$.*

Proof Consider time point $t = C^i(\text{HWF})$. For job i the remaining processing time in OPTRES at time t is at most $s^i(\text{HWF})$. Of this remaining time at most half can be scheduled on SUB, therefore $s^i(\text{OPTRES}) \leq \frac{1}{2}s^i(\text{HWF})$. \square

Lemma 4.8. *For any time point $t \notin \mathcal{T}$, if HWF is idle, then OPT is idle as well.*

Proof Take an arbitrary time point $t \notin \mathcal{T}$ where HWF is idle. Assume to the contrary that OPT schedules job i at time t . Since HWF is idle, i is already finished in schedule HWF, thus $C^i(\text{OPT}) \geq t > C^i(\text{HWF})$. This implies by Observation 4.2 that $s^i(\text{HWF}) > s^i(\text{OPT})$. On the other hand, as $t \notin \mathcal{T}$ and $t > C^i(\text{HWF})$, we have that $i \in N^-$, which leads to the contradiction that $s^i(\text{HWF}) < s^i(\text{OPT})$. \square

Lemma 4.9. *For any time point $t \notin \mathcal{T}$, if HWF processes job i and OPT processes job j , then $w^i \geq w^j$.*

Proof Take an arbitrary time point $t \notin \mathcal{T}$ where HWF and OPT are working on jobs i and j respectively. Now, assume to the contrary that $w^i < w^j$. As HWF schedules the job with the largest weight, we have that job j is already finished before t in schedule HWF. Hence $C^j(\text{HWF}) < C^j(\text{OPT})$, which means that $s^j(\text{HWF}) > s^j(\text{OPT})$ and thus $j \in N^+$. This together with our observation that $t > C^j(\text{HWF})$ leads to the contradiction that $t \in \mathcal{T}$. \square

Lemma 4.10. $SW^N(\text{HWF}) \geq SW^N(\overline{\overline{\text{OPTRES}}})$.

Proof From Lemma 4.8 and the definition of $\overline{\text{OPTRES}}$, it follows that if HWF is idle at time t , then $\overline{\text{OPTRES}}$ is also idle at time t . Lemma 4.9 ensures that whenever HWF is working on a job at time t , then $\overline{\text{OPTRES}}$ is working on a job with at most the same weight. Hence, at each time point, HWF obtains at least as much weighted time saving as $\overline{\text{OPTRES}}$. \square

Now, equipped with these lemmas, we are ready to prove Theorem 4.6.

Proof of Theorem 4.6 Note that $\sum_{i \in N^+} w^i s^i(\text{OPTRES}) = \sum_{i \in N} w^i s^i(\text{OPTRES})$ as OPTRES does not work on jobs in N^- by definition. Consequently, we can write

$$\begin{aligned} \frac{3}{2} SW^N(\text{HWF}) &= \frac{3}{2} \sum_{i \in N} w^i s^i(\text{HWF}) \\ &\geq \frac{1}{2} \sum_{i \in N^+} w^i s^i(\text{HWF}) + \sum_{i \in N} w^i s^i(\text{HWF}) \\ &\geq \sum_{i \in N} w^i s^i(\text{OPTRES}) + \sum_{i \in N} w^i s^i(\overline{\text{OPTRES}}) \\ &= \sum_{i \in N} w^i s^i(\text{OPT}) = SW^N(\text{OPT}), \end{aligned}$$

where the second inequality follows from Lemmas 4.7 and 4.10. This proves that HWF is a factor $2/3$ approximation. \square

Lemma 4.11. *The approximation factor for Theorem 4.6 is tight.*

Proof Consider the instance of n jobs consisting of $n - 1$ “light” jobs with processing time and weight 1, and one “heavy” job with processing time 2 and weight $1 + \epsilon$. All release dates are zero. Without loss of generality assume that jobs are indexed in SRPT order. First, consider HWF, which starts with the heavy job. It is easy to see that $C^n(\text{HWF}) = s^n(\text{HWF}) = \frac{1}{2} p^n = 1$ and $s^i(\text{HWF}) = 0$ for all $i \neq n$ as at time 1 all other jobs are finished. This means that $SW(\text{HWF}) = w^n s^n(\text{HWF}) = 1 + \epsilon$.

Consider now the solution in which jobs are processed in SRPT order, that is, all light jobs are processed first and the heavy job comes last. We can apply Theorem 4.3 to see that the total time spent on SUB for the light jobs is

$$\sum_{i=1}^{n-1} s^i(\text{SRPT}) = \sum_{i=1}^{n-1} \frac{1}{2^{n-i}} = 1 - \frac{1}{2^{n-1}}, \quad (33)$$

while the time spent on the heavy one is

$$s^n(\text{SRPT}) = \frac{p^n}{2} - \sum_{i=1}^{n-1} \frac{s^i(\text{OPT})}{2} = 1 - \frac{1}{2} \left(1 - \frac{1}{2^{n-1}}\right) = \frac{1}{2} + \frac{1}{2^n}. \quad (34)$$

Combining equations (33) and (34) we get that

$$\begin{aligned} SW(\text{SRPT}) &= \left(1 - \frac{1}{2^{n-1}}\right) + \left(\frac{1}{2} + \frac{1}{2^n}\right) (1 + \epsilon) \\ &= \frac{3}{2} - \frac{1}{2^n} + \left(\frac{1}{2} + \frac{1}{2^n}\right) \epsilon. \end{aligned}$$

From these observations we get that

$$\frac{SW(\text{HWF})}{SW(\text{OPT})} \leq \frac{SW(\text{HWF})}{SW(\text{SRPT})} = \frac{1 + \epsilon}{\frac{3}{2} - \frac{1}{2^n} + \left(\frac{1}{2} + \frac{1}{2^n}\right) \epsilon}.$$

By letting $\epsilon \rightarrow 0$ and $n \rightarrow \infty$, we can conclude that the approximation ratio of $2/3$ is indeed tight. \square

4.3.3 Other schedules

Our investigation included some other classic and intuitive ordering algorithms, but all of them were outperformed by HWF. For the sake of completeness we present two such ordering algorithms. The first one is the well-known SMITH'S RULE, which at any time selects the job that maximises the ratio between weight and remaining processing time. In case of a tie, the job with the highest weight is selected. If there is still a tie, the job with the lowest index is selected. Although SMITH'S RULE takes also the processing time into account, it performs worse than HWF.

Theorem 4.12. SMITH'S RULE is at most a factor $\frac{1}{2}$ approximation.

Our greedy schedule, called highest profit first (HPF), at any time processes the job that maximises the product of the remaining processing time and the weight. In case of a tie, the job with the highest weight is selected. If there is still a tie, the job with the lowest index is chosen.

Theorem 4.13. *HPF is a factor $\frac{1}{n}$ approximation and this ratio is tight.*

The proofs of Theorems 4.13 and 4.12 can be found in the appendix.

4.4 MECHANISM DESIGN

From now on we assume that jobs are selfish agents that act strategically and part of the job characteristics is their private information. To align the behaviour of jobs we construct mechanisms to provide proper incentives. We first focus on direct mechanisms. This choice is without loss of generality due to the revelation principle. A direct mechanism (x, π) consists of an allocation rule $x : R \times W \times P \rightarrow \mathcal{O}$ and a payment scheme $\pi : R \times W \times P \rightarrow \mathbb{R}^n$. For each report the allocation rule assigns a schedule on SUB , while the payment scheme determines the related monetary transfers. The utility of every job is quasi-linear, so given type profiles (r, p, w) , report profile $(\hat{r}, \hat{w}, \hat{p})$, schedule $\delta = x(\hat{r}, \hat{w}, \hat{p})$ and payment $\pi(\hat{r}, \hat{w}, \hat{p})$ we have that

$$u^i(\delta, \pi(\hat{r}, \hat{w}, \hat{p}), r^i, w^i, p^i) = w^i s^i(\delta, r^i, p^i) - \pi^i(\hat{r}, \hat{w}, \hat{p}).$$

Under any direct mechanism we require truth-telling to be dominant strategy for each job, i. e., reporting the true type maximises the job's utility. Mechanisms that adhere to this property are *dominant strategy incentive compatible* (DSIC). Formally, a mechanism (x, π) is DSIC if and only if for any job i , report profile (r^{-i}, w^{-i}, p^{-i}) , type (r^i, w^i, p^i) and alternative report $(\hat{r}^i, \hat{w}^i, \hat{p}^i)$ we have that

$$u^i(x(r, w, p), \pi(r, w, p), r^i, w^i, p^i) \geq u^i(x(\hat{r}, \hat{w}, \hat{p}), \pi(\hat{r}, \hat{w}, \hat{p}), r^i, w^i, p^i), \quad (\text{DSIC})$$

where $[\hat{r}, \hat{w}, \hat{p}] = [(\hat{r}^i, r^{-i}), (\hat{w}^i, w^{-i}), (\hat{p}^i, p^{-i})]$. Furthermore, we require that participation in the proposed mechanism does not inflict a loss on any jobs, i. e., truth-telling guarantees a utility of at least zero. Mechanisms that satisfy

this property are *ex-post individual rational* (EPIR). Formally, a mechanism (x, π) is EPIR if and only if for any job i and type profile (r, w, p) we have that

$$u^i(x(r, w, p), \pi(r, w, p), r^i, w^i, p^i) \geq 0. \quad (\text{EPIR})$$

In the following we devise mechanisms for various settings that implement schedules from Section 4.3.

4.4.1 Equal weights

In this subsection we investigate the special case in which all weights are equal and jobs arrive at the same time. We assume that processing time is private information and may vary among jobs. The objective then reduces to maximising the total amount of time spent on SUB. We assume that jobs are indexed in non-decreasing order of their processing time, i. e., $p^i \leq p^{i+1}$ for all $i < n$. In this case the optimal schedule follows the SPT rule by Theorem 4.3. Since the optimal allocation rule is known, we can use the VCG mechanism (see Subsection 3.4.1 from Chapter 3) to implement the schedule that maximises the total saving. Remember, that VCG is DSIC and EPIR. In the current setting VCG allocates according to the SPT algorithm and the amount that job i is charged for report profile p is

$$\pi_{\text{VCG}}^i(p) = \sum_{j \neq i} s^j(x_{\text{SPT}}^j(p^{-i})) - \sum_{j \neq i} s^j(x_{\text{SPT}}^j(p)), \quad (35)$$

where $s^j(x_j(p^{-i}))$ denotes the saving of job j in the absence of job i and report profile p^{-i} . Intuitively, the payment equals the total net impact of job i on the savings of the other jobs. It turns out that the VCG payments can be expressed in a more concise closed-form expression that is easier to work with:

Lemma 4.14 (Lemma 4 in Hezarkhani and Kubiak [38]). *For a given job i and report profile p we have that*

$$\pi_{\text{VCG}}^i(p) = s^i(x_{\text{SPT}}^i(p)) \left(1 - \frac{1}{2^{n-i}}\right) \quad (36)$$

and

$$u^i(x_{\text{VCG}}^i(p), \pi_{\text{VCG}}^i(p), p) = \frac{s^i(x_{\text{SPT}}^i(p))}{2^{n-i}}. \quad (37)$$

4.4.1.1 An ascending auction implementation

From a theoretical viewpoint it is true that working with a direct mechanism is without loss of generality, but from a practical perspective it is desirable to find more natural implementations of the corresponding allocation rule and payment scheme. The advantages of such mechanisms can stem from lower administration cost, more privacy preservation, or less complex and more intuitive algorithms. To this end, we present an auction that implements the VCG mechanism. Before formal definition of the auction, we introduce the notation necessary for describing the model. In contrast to direct mechanisms, the action space of jobs is not anymore their type space. We denote by A^i the set of possible actions for job i , and by $A := A^1 \times \dots \times A^n$ the action space of all jobs. A strategy of job i is defined by $\theta^i : P^i \rightarrow A^i$, i. e., by a mapping from types to actions. The set of possible strategies for job i is denoted by Θ^i , and the strategy space of all jobs by $\Theta := \Theta^1 \times \dots \times \Theta^n$. The allocation rule $x : A \rightarrow \mathcal{O}$ and payment scheme of $\pi : A \rightarrow \mathbb{R}^n$ do not operate directly on type space P , but on action space A .

The idea behind the auction is that each position in the schedule of `sub` is auctioned separately, one after the other in n steps. In each step jobs bid the time they want to spend on `sub`. The smallest bid wins and the payment is a function of their bid and the step. We call this auction `SUBA` and formally define it in Algorithm 1. In this context the action space of job i is $A^i := \{f : \mathbb{R}_0^+ \rightarrow \mathbb{R}_0^+\}$, that is, the set of functions that map time points to non-negative real numbers. Let $\theta_t^i(p^i)$ be the bid that job i quotes given type p^i at time t and let $\theta(p) := (\theta^1(p^1), \dots, \theta^n(p^n))$ the action profile given type profile p . By a little stretch of notation we denote by $x^i(\theta(p)) = [k, l]$ that for action profile $\theta(p)$ time interval $[k, l]$ is allocated to job i on `sub`.

Special attention is paid to strategy profiles that lead to a stable outcome in a sense that nobody has an incentive to deviate even after all types have been realized.

Algorithm 1: SUBCONTRACTOR'S AUCTION (SUBA)

Input: $N = \{1, \dots, n\}$, $\theta(p) = (\theta^1(p^1), \dots, \theta^n(p^n))$
Output: $x^i(\theta(p))$, $\pi^i(\theta(p)) \forall i \in N$
1 Initialisation: $S_1 = N$, $t_1 = 0$, $x^i(\theta(p)) = [0, 0]$, $\forall i \in N$
2 **for** $j \in \{1, \dots, n\}$ **do**
3 $\alpha = \arg \min_{i \in S_j} \{\theta_{t_j}^i(p^i)\}$, ties are broken in favour of lower index
4 $t_{j+1} = t_j + \theta_{t_j}^\alpha(p^\alpha)$
5 $x_{\text{SUBA}}^\alpha(\theta(p)) = [t_j, t_{j+1}]$
6 $\pi_{\text{SUBA}}^\alpha(\theta(p)) = [t_j, t_{j+1}] \left(1 - \frac{1}{2^{n-j}}\right)$
7 $S_{j+1} = S_j \setminus \alpha$

Definition 4.15. Under mechanism (x, π) and type profile p , a strategy profile $\theta(p) = (\theta^1(p^1), \dots, \theta^n(p^n))$ is an ex-post equilibrium, if for all jobs i and alternative strategy $\hat{\theta}^i(p^i)$ we have that

$$u^i(x(\theta(p)), \pi(\theta(p)), p) \geq u^i(x(\hat{\theta}(p)), \pi(\hat{\theta}(p)), p), \quad (38)$$

where $\hat{\theta}(p) = (\theta^1(p^1), \dots, \hat{\theta}^i(p^i), \dots, \theta^n(p^n))$.

We would like to show that there is a strategy profile that is an ex-post equilibrium and leads to the efficient allocation under SUBA.

Theorem 4.16. *SUBCONTRACTOR'S AUCTION implements the allocation that maximises the total saving. In particular, the strategy profile $\theta(p)$ in which $\theta_t^i(p^i) = \max\{(p^i - t)/2, 0\}$ is an ex-post equilibrium and $x_{\text{SUBA}}(\theta(p)) = x_{\text{SPT}}(p)$.*

Proof Assume that $p^i \leq p^{i+1}$ for all $i < n$ and $\theta_t^i(p^i) = \max\{(p^i - t)/2, 0\}$. As $(p^i - t)/2 \leq (p^j - t)/2 \iff p^i \leq p^j$ and the bid of job i at each point of time equals the maximum time it can spend on SUB, we can conclude that SUBA allocates the same way for strategy profile $\theta(p)$ as the VCG mechanism for report p . Moreover, the payment rule of SUBA in this case coincides with that of the VCG payment, hence

$$u^i(x_{\text{SUBA}}(\theta(p)), \pi_{\text{SUBA}}(\theta(p)), p) = \frac{s^i(x_{\text{SPT}}^i(p))}{2^{n-i}}.$$

Let $\hat{\theta}^i(p^i)$ be an alternative strategy for job i . The alternative strategy profile is denoted by $\hat{\theta}(p) := (\theta^1(p^1), \dots, \hat{\theta}^i(p^i), \dots, \theta^n(p^n))$. As in Algorithm 1, t_k denotes the starting point of the time interval sold in step k . In the following we show that (38) holds for $\theta^i(p^i)$ against any $\hat{\theta}^i(p^i)$ for all job i . To this end, consider job i for the rest of the proof. Assume that with alternative strategy $\hat{\theta}^i(p^i)$, job i can obtain position k .

First, consider the case when $k \leq i$. The rules of the auction dictate that in each step k the lowest bid wins, therefore all bidders who haven't won in the previous steps has at least as high bid as the winner of the current step, that is, $\theta_{t_k}^j(p^j) \geq \theta_{t_k}^k(p^k)$ for all $j \in S_k$. Therefore $\hat{\theta}_{t_k}^i(p^i) \leq \theta_{t_k}^k(p^k) = s^k(x_{\text{SPT}}^k(p))$. The last equality holds because SUBA allocates the same way as SPT. As $p^k \leq p^i$ it is feasible for job i to be processed for time $s^k(x_{\text{SPT}}^k(p))$ on SUB. Furthermore, the utility of job i is linearly increasing in $\hat{\theta}_{t_k}^i(p^i)$ due to the payment scheme, hence

$$u^i(x_{\text{SUBA}}(\hat{\theta}(p)), \pi_{\text{SUBA}}(\hat{\theta}(p)), p) \leq \frac{s^k(x_{\text{SPT}}^k(p))}{2^{n-k}}.$$

Lemma 2 in Hezarkhani and Kubiak [38] states that for all $j < n$ we have that $s^j(x_{\text{SPT}}^j(p)) \leq 2s^{j+1}(x_{\text{SPT}}^{j+1}(p))$. Consequently, $s^k(x_{\text{SPT}}^k(p)) \leq 2^{i-k}s^i(x_{\text{SPT}}^i(p))$. This leads to the fact that

$$\begin{aligned} u^i(x_{\text{SUBA}}(\theta(p)), \pi_{\text{SUBA}}(\theta(p)), p) - u^i(x_{\text{SUBA}}(\hat{\theta}(p)), \pi_{\text{SUBA}}(\hat{\theta}(p)), p) \\ &\geq \frac{s^i(x_{\text{SPT}}^i(p))}{2^{n-i}} - \frac{s^k(x_{\text{SPT}}^k(p))}{2^{n-k}} \\ &\geq \frac{s^i(x_{\text{SPT}}^i(p))}{2^{n-i}} - \frac{s^i(x_{\text{SPT}}^i(p))}{2^{n-i}} = 0. \end{aligned}$$

Finally, assume that $k > i$, that is, in the new order jobs $\{i+1, \dots, k\}$ precede job i . Due to the definition of $\hat{\theta}(p)$ and the assumption that $p^{j+1} \geq p^j$ for all $j < n$, the total saving of the first $k-1$ jobs cannot be smaller in the new order than in the original order induced by $\theta(p)$. Hence, job i has to wait at least $\sum_{j=1}^{k-1} s^j(x_{\text{SPT}}^j(p))$ amount of time before being processed. This means that the saving of job i is at most $\max\{0, [p^i - \sum_{j=1}^{k-1} s^j(x_{\text{SPT}}^j(p))]/2\}$.

As $u^i(x(\theta(p)), \pi(\theta(p)), p) \geq 0$, we only have to check the situation when $p^i > \sum_{j=1}^{k-1} s^j(x_{\text{SPT}}^j(p))$. This implies that

$$\begin{aligned}
& u^i(x(\theta(p)), \pi(\theta(p)), p) - u^i(x(\hat{\theta}(p)), \pi(\hat{\theta}(p)), p) \\
& \stackrel{(37)}{\geq} \frac{s^i(x_{\text{SPT}}^i(p))}{2^{n-i}} - \left(\frac{p^i}{2^{n-k+1}} - \sum_{j=1}^{k-1} \frac{s^j(x_{\text{SPT}}^j(p))}{2^{n-k+1}} \right) \\
& \stackrel{\text{Th. 4.3}}{=} \frac{p^i}{2^{n-i+1}} - \sum_{j=1}^{i-1} \frac{p^j}{2^{n-j+1}} - \left(\frac{p^i}{2^{n-k+1}} - \sum_{j=1}^{k-1} \frac{p^j}{2^{n-j+1}} \right) \\
& = p^i \left(\frac{1}{2^{n-i+1}} - \frac{1}{2^{n-k+1}} \right) + \sum_{j=i}^{k-1} \frac{p^j}{2^{n-i+1}} \\
& \geq p^i \left(\frac{1}{2^{n-i+1}} - \frac{1}{2^{n-k+1}} + \sum_{j=i}^{k-1} \frac{1}{2^{n-i+1}} \right) = 0.
\end{aligned}$$

The first equality follows from Theorem 4.3, while the last inequality is due to the fact $p^i \leq p^{i+1}$ for all $i < n$. \square

The advantage of such an implementation is that by setting up an auction the designer is freed from the task of acquiring information and deciding on the schedule. This can be interpreted as a step towards decentralisation, as centralised, direct communication is replaced by an auction that gathers information and makes the assignments indirectly without involving the mechanism designer. Moreover, during the auction jobs do not have to disclose their private information directly.

4.4.1.2 A mechanism with strongly budget-balanced payments

Definition 4.17. A payment scheme π is strongly budget-balanced if and only if for all reports p we have that $\sum_i \pi^i(p) = 0$. A mechanism is strongly budget-balanced if its payment scheme is strongly budget-balanced.

This property is useful in cases where the designer neither wants to invest money into the mechanism nor wants to profit from it. In case payments are necessary to ensure truth-telling, strong budget balancedness can only

be achieved if jobs compensate each other for using the common resource, thus we have to allow for negative payments. To this end, we alter the VCG payments such that the resulting mechanism is not just ex-post individual rational, but also strongly budget-balanced. We define our Budget-Balanced (BB) mechanism as follows. BB allocates according to the SPT algorithm and charges job i for report profile p

$$\pi_{\text{BB}}^i(p) = s^i(x_{\text{SPT}}^i(p)) - \frac{p^i}{2^{n+1-i}}. \quad (39)$$

Theorem 4.18. *When only the processing time of jobs may differ, the BB mechanism is efficient, dominant strategy incentive compatible, individual rational and strongly budget-balanced.*

Proof BB is an efficient mechanism as it follows the SPT algorithm. From equation (39) it is clear that the payment never exceeds the time saving, implying ex-post individual rationality. To see dominant strategy incentive compatibility, rewrite (39) as

$$\begin{aligned} \pi_{\text{BB}}^i(p) &= s^i(x_{\text{SPT}}^i(p)) - \frac{p^i}{2^{n+1-i}} \\ &\stackrel{\text{Th. 4.3}}{=} s^i(x_{\text{SPT}}^i(p)) - \left(\sum_j s^j(x_{\text{SPT}}^j(p)) - \sum_{j \neq i} \frac{p^j}{2^{n+1-j}} \right) \\ &= \sum_{j \neq i} \left(\frac{p^j}{2^{n+1-j}} - s^j(x_{\text{SPT}}^j(p)) \right). \end{aligned} \quad (40)$$

The second equality follows from Theorem 4.3. Note that adding any function $h^i : P^{-i} \rightarrow \mathbb{R}$ to (36) does not affect dominant strategy incentive compatibility as it does not depend on the reported p^i . By taking

$$h^i(p^{-i}) = \sum_{j \neq i} \left(\frac{p^j}{2^{n+1-j}} - s^j(x_{\text{SPT}}^j(p^{-i})) \right),$$

the payment of (40) can be derived from (35) as follows:

$$\begin{aligned} \pi_{\text{VCG}}^i(p) + h^i(p^{-i}) &= \sum_{j \neq i} \left(s^j(x_{\text{SPT}}^j(p^{-i})) - s^j(x_{\text{SPT}}^j(p)) \right) + \sum_{j \neq i} \left(\frac{p^j}{2^{n+1-j}} - s^j(x_{\text{SPT}}^j(p^{-i})) \right) \\ &= \sum_{j \neq i} \left(\frac{p^j}{2^{n+1-j}} - s^j(x_{\text{SPT}}^j(p)) \right) = \pi_{\text{BB}}^i(p). \end{aligned}$$

This implies that BB is also DSIC. Finally, by equation (39) and Theorem 4.3 we get that

$$\begin{aligned} \sum_i \pi_{\text{BB}}^i(p) &\stackrel{(39)}{=} \sum_i \left(s^i(x_{\text{SPT}}^i(p)) - \frac{p^i}{2^{n+1-i}} \right) \\ &= \sum_i s^i(x_{\text{SPT}}^i(p)) - \sum_i \frac{p^i}{2^{n+1-i}} \stackrel{\text{Th.4.3}}{=} 0, \end{aligned}$$

which proves strong budget balancedness. □

4.4.2 Arbitrary weights

In this subsection we take a step forward in generalising the model and assume that jobs may have different weights. To our current knowledge, the solution of welfare maximisation problem in this setting is not yet known, not even its computational complexity. Therefore our aim in this subsection is to find a simple mechanism that is a close approximation of the optimal solution. Theorem 4.6 states that the HWF algorithm is a factor 2/3 approximation, therefore we choose it as a base for our approximation mechanism.

Let us define the HWF mechanism $(x_{\text{HWF}}, \pi_{\text{HWF}})$ as the one that schedules jobs on SUB in non-increasing order of their reported weight. For jobs with identical reported weight, ties are broken in favour of lower index. As in this subsection all schedules are represented by orderings, to ease the exposition we write $s^i(j, p^i)$ to denote the saving of job i if it is scheduled at position

j and has processing time p^i . To take into account that job i might report a processing time $\hat{p}^i \neq p^i$, we define

$$s^i(j, \hat{p}^i, p^i) = \begin{cases} s^i(j, p^i) & \text{if } \hat{p}^i \geq p^i, \\ s^i(j, \hat{p}^i) & \text{otherwise,} \end{cases} \quad (41)$$

for the time saving of job i at position j , having processing time p^i , but reporting \hat{p}^i . In the rest of the section we focus only on job i , and fix the reports (w^{-i}, p^{-i}) of the other jobs. To reduce the size of the notation, we leave out the reference to the HWF allocation rule and to the reports of the others (w^{-i}, p^{-i}) when they are clear from the context.

Let $\sigma^i(w^i, p^i)$ denote the position of job i in the schedule when it reports (w^i, p^i) , and let $\sigma_{-1}^i(w^i, p^i)$ denote the job at position i under the same report. Additionally, assume that jobs are indexed in non-increasing order of their weight, i. e., $w^i \geq w^{i+1}$ for all $i < n$. The payment scheme of mechanism HWF is defined as

$$\pi_{\text{HWF}}^i(w^i, p^i) = \sum_{j=\sigma^i(w^i, p^i)}^{n-1} w^{\sigma_{-1}^{j+1}(w^i, p^i)} [s^i(j, p^i, p^i) - s^i(j+1, p^i, p^i)].$$

Note that we cannot use the payment scheme of the VCG mechanism because HWF generally does not allocate the same way, that is, efficiently. Therefore, it is a challenge to find an appropriate payment scheme that, in combination with the HWF ordering, constitutes a DSIC and EPIR mechanism. Moreover, there is no guarantee that such payment scheme even exists. The idea behind our proposed payment scheme is the following. For fixed processing times the time saving is monotone non-increasing in the position defined by the ordering, i. e., the earlier the better. For each incremental gain in position, the payment equals the additional time spent on SUB times the minimum weight with which the new position could be obtained.

Theorem 4.19. *The HWF mechanism is DSIC, EPIR and a factor $\frac{2}{3}$ approximation of the efficient mechanism.*

Proof Theorem 4.6 states that the HWF mechanism is a factor 2/3 approximation. To see individual rationality observe that

$$\begin{aligned} \pi_{\text{HWF}}^i(w^i, p^i) &= \sum_{j=\sigma^i(w^i, p^i)}^{n-1} w^{\sigma_{-1}^{j+1}(w^i, p^i)} \left[s^i(j, p^i, p^i) - s^i(j+1, p^i, p^i) \right] \\ &\leq \sum_{j=\sigma^i(w^i, p^i)}^{n-1} w^j \left[s^i(j, p^i, p^i) - s^i(j+1, p^i, p^i) \right] \\ &= w^i \left(s^i(i, p^i, p^i) - s^i(n, p^i, p^i) \right) \leq w^i s^i(i, p^i, p^i). \end{aligned}$$

For the rest of the proof we concentrate on incentive compatibility. Let (w^i, p^i) be the true type of job i . Now, assume that the following inequalities hold.

$$u^i(\hat{w}^i, \hat{p}^i, w^i, p^i) \leq u^i(\hat{w}^i, p^i, w^i, p^i) \quad \forall \hat{w}^i, \forall \hat{p}^i \geq p^i \quad (42)$$

$$u^i(\hat{w}^i, \hat{p}^i, w^i, p^i) \leq u^i(w^i, \hat{p}^i, w^i, p^i) \quad \forall \hat{w}^i, \forall \hat{p}^i \leq p^i \quad (43)$$

$$u^i(w^i, \hat{p}^i, w^i, p^i) \leq u^i(w^i, p^i, w^i, p^i) \quad \forall \hat{w}^i, \forall \hat{p}^i \leq p^i. \quad (44)$$

Then for all $\hat{p}^i > p^i$ we have that

$$u^i(\hat{w}^i, \hat{p}^i, w^i, p^i) \stackrel{(42)}{\leq} u^i(\hat{w}^i, p^i, w^i, p^i) \stackrel{(43)}{\leq} u^i(w^i, p^i, w^i, p^i),$$

as the first inequality is implied by (42), and the second one by (43). Similarly, for all $\hat{p}^i \leq p^i$ we have that

$$u^i(\hat{w}^i, \hat{p}^i, w^i, p^i) \stackrel{(43)}{\leq} u^i(w^i, \hat{p}^i, w^i, p^i) \stackrel{(44)}{\leq} u^i(w^i, p^i, w^i, p^i),$$

as the first inequality is implied by (43), and the second one by (44). Combining these two results implies that $u^i(\hat{w}^i, \hat{p}^i, w^i, p^i) \leq u^i(w^i, p^i, w^i, p^i)$, which is equal to requirement (DSIC) and concludes dominant strategy incentive compatibility.

What remains to be shown is that the inequalities (42), (43), (44) hold. Consider first inequality (42). Note, that by inflating the processing time, job i can buy more time on SUB, but cannot increase its saving. Furthermore, reporting a different p^i does not influence the order. Thus, we have that

$\sigma^i(\hat{w}^i, \hat{p}^i) = \sigma^i(\hat{w}^i, p^i)$ and $s^i(j, p^i, p^i) = s^i(j, \hat{p}^i, p^i)$ for all $\hat{p}^i \geq p^i$ and j . When we set $\sigma^i(\hat{w}^i, \hat{p}^i) = k$ these previous observations imply that

$$\begin{aligned}
& u^i(\hat{w}^i, p^i, w^i, p^i) - u^i(\hat{w}^i, \hat{p}^i, w^i, p^i) = \pi^i(\hat{w}^i, p^i) - \pi^i(\hat{w}^i, \hat{p}^i) \\
&= \sum_{j=k}^{n-1} w^{\sigma_{-1}^{j+1}(\hat{w}^i, \hat{p}^i)} \left[s^i(j, \hat{p}^i, \hat{p}^i) - s^i(j+1, \hat{p}^i, \hat{p}^i) \right] \\
&\quad - \sum_{j=k}^{n-1} w^{\sigma_{-1}^{j+1}(\hat{w}^i, p^i)} \left[s^i(j, p^i, p^i) - s^i(j+1, p^i, p^i) \right] \\
&= \sum_{j=\sigma^i(\hat{w}^i, p^i)}^{n-1} w^{\sigma_{-1}^{j+1}(\hat{w}^i, p^i)} \left[s^i(j, \hat{p}^i, \hat{p}^i) - s^i(j+1, \hat{p}^i, \hat{p}^i) \right. \\
&\quad \left. + s^i(j+1, p^i, p^i) - s^i(j, p^i, p^i) \right].
\end{aligned}$$

To prove that the last expression is non-negative it is sufficient to show that

$$s^i(j, \hat{p}^i, \hat{p}^i) - s^i(j, p^i, p^i) \geq s^i(j+1, \hat{p}^i, \hat{p}^i) - s^i(j+1, p^i, p^i),$$

for any j . Let T_j^i denote the starting time of job i on SUB if it is at position j in the schedule. By using the definition of savings, we can express the last inequality as

$$\max\{0, \hat{p}^i - T_j^i\} - \max\{0, p^i - T_j^i\} \geq \max\{0, \hat{p}^i - T_{j+1}^i\} - \max\{0, p^i - T_{j+1}^i\}.$$

It is easy to see that this holds by noting that $\hat{p}^i \geq p^i$ and $T_j^i \leq T_{j+1}^i$. Now, we turn our attention to inequality (43). In this case we have that $\hat{p} \leq p$ and

$\sigma^i(w^i, \hat{p}^i) = i$. Let $\sigma^i(\hat{w}^i, \hat{p}^i) = k$ and assume that $k \geq i$. Using the fact that jobs are indexed in non-decreasing order of their weights we can write

$$\begin{aligned}
& u^i(w^i, \hat{p}^i, w^i, p^i) - u^i(\hat{w}^i, \hat{p}^i, w^i, p^i) \\
&= w^i \left(s^i(i, \hat{p}^i, p^i) - s^i(k, \hat{p}^i, p^i) \right) - \sum_{j=i}^{k-1} w^{j+1} \left[s^i(j, \hat{p}^i, p^i) - s^i(j+1, \hat{p}^i, p^i) \right] \\
&\geq w^i \left(s^i(i, \hat{p}^i, p^i) - s^i(k, \hat{p}^i, p^i) \right) - \sum_{j=i}^{k-1} w^i \left[s^i(j, \hat{p}^i, \hat{p}^i) - s^i(j+1, \hat{p}^i, p^i) \right] \\
&= w^i \left(s^i(i, \hat{p}^i, p^i) - s^i(k, \hat{p}^i, p^i) \right) - w^i \left(s^i(i, \hat{p}^i, p^i) - s^i(k, \hat{p}^i, p^i) \right) = 0.
\end{aligned}$$

Similarly, if $k < i$, we have that

$$\begin{aligned}
& u^i(w^i, \hat{p}^i, w^i, p^i) - u^i(\hat{w}^i, \hat{p}^i, w^i, p^i) \\
&= w^i \left(s^i(i, \hat{p}^i, p^i) - s^i(k, \hat{p}^i, p^i) \right) + \sum_{j=k}^{i-1} w^j \left[s^i(j, \hat{p}^i, \hat{p}^i) - s^i(j+1, \hat{p}^i, p^i) \right] \\
&\geq w^i \left(s^i(i, \hat{p}^i, p^i) - s^i(k, \hat{p}^i, p^i) \right) + \sum_{j=k}^{i-1} w^i \left[s^i(j, \hat{p}^i, p^i) - s^i(j+1, \hat{p}^i, p^i) \right] \\
&= w^i \left(s^i(i, \hat{p}^i, p^i) - s^i(k, \hat{p}^i, p^i) \right) + w^i \left(s^i(k, \hat{p}^i, p^i) - s^i(i, \hat{p}^i, p^i) \right) = 0.
\end{aligned}$$

This proves that inequality (43) holds. Finally, consider inequality (44). We know that $\hat{p}^i \leq p^i$ and that $\sigma^i(w^i, \hat{p}^i) = \sigma^i(w^i, p^i) = i$, so we can write

$$\begin{aligned}
& u^i(w^i, p^i, w^i, p^i) - u^i(w^i, \hat{p}^i, w^i, p^i) \\
&= w^i \left(s^i(i, p^i, p^i) - s^i(i, \hat{p}^i, p^i) \right) \\
&\quad - \sum_{j=i}^{n-1} w^{j+1} \left[s^i(j, p^i, p^i) - s^i(j+1, p^i, p^i) + s^i(j+1, \hat{p}^i, p^i) - s^i(j, \hat{p}^i, p^i) \right].
\end{aligned}$$

If $s^i(j, p^i, p^i) - s^i(j, \hat{p}^i, p^i) \geq s^i(j+1, p^i, p^i) - s^i(j+1, \hat{p}^i, p^i)$, then we can bound the last expression from below by

$$\begin{aligned} & w^i \left(s^i(i, p^i, p^i) - s^i(i, \hat{p}^i, p^i) \right) \\ & - \sum_{j=i}^{n-1} w^i \left[s^i(j, p^i, p^i) - s^i(j+1, p^i, p^i) + s^i(j+1, \hat{p}^i, p^i) - s^i(j, \hat{p}^i, p^i) \right] \\ & = w^i \left(s^i(n, p^i, p^i) - s^i(n, \hat{p}^i, p^i) \right) \stackrel{41}{\geq} 0. \end{aligned}$$

The last inequality is a consequence of the fact that reporting shorter processing time can only shorten the allocated time on SUB. What remains to be shown is that

$$s^i(j, p^i, p^i) - s^i(j, \hat{p}^i, p^i) \geq s^i(j+1, p^i, p^i) - s^i(j+1, \hat{p}^i, p^i).$$

Using the same formulation as in the proof of inequality (42) leads to

$$\max\{0, p^i - T_j^i\} - \max\{0, \hat{p}^i - T_j^i\} \geq \max\{0, p^i - T_{j+1}^i\} - \max\{0, \hat{p}^i - T_{j+1}^i\}.$$

It is easy to see that this holds by noting that $\hat{p}^i \leq p^i$ and $T_j^i \leq T_{j+1}^i$. This concludes the proof. \square

4.5 CONCLUSION

We study the problem of subcontractor scheduling from both an algorithmic and a mechanism design perspective. First, we look at simple schedules that can serve as a base for practical mechanisms. We show that scheduling jobs greedily in non-increasing order of weight is a $2/3$ approximation, while other intuitive algorithms such as SMITH'S RULE perform worse. For practical implementations we devise three mechanisms. For the case of equal weights and release dates we present an auction implementation of the optimal schedule that decentralises the processes of information acquisition and scheduling. For settings in which the designer does not want to invest nor to gain money, we propose a mechanism in which the total sum of payments is zero. Finally, for the general case where weights and processing times are arbitrary, we construct a mechanism that is a $2/3$ -approximation of the op-

timal one. The complexity of finding the schedule that maximises the total weighted saving for the case of arbitrary weights remains an intriguing open question.

APPENDIX

Proof of Theorem 4.4 Assume to the contrary that SPT is not optimal. Let OPT denote the optimal schedule in which the first time point T that differs from SPT is maximal, i. e., $\text{SPT}^i(t) = \text{OPT}^i(t)$ for any job $i \in N$ and time point $t < T$. This means that only the remaining instance after time T is of interest, and all choices up until T can be “ignored”. Therefore, without loss of generality, we consider the instance remaining after time T , and assume that the time point of the first difference is equal to zero.

It is trivial that neither SPT nor OPT is idle at time zero, hence SPT and OPT are respectively working on some jobs i and j , with $i \neq j$ and $p^i \leq p^j$. Consider the feasible schedule Feas , that is identical to OPT for all jobs except for i and j . The difference is that at time zero, Feas starts working on i , but only at time points where OPT is working on i or j . If Feas finishes job i , then it allocates time for j , but again only at time points where OPT is working on i or j . Hence, we have that

$$\text{Feas}^i(t) = \begin{cases} \text{OPT}^i(t) + \text{OPT}^j(t) & t \leq C^i(\text{Feas}) \\ 0 & \text{else,} \end{cases}$$

$$\text{Feas}^j(t) = \begin{cases} 0 & t \leq C^i(\text{Feas}) \\ \text{OPT}^i(t) + \text{OPT}^j(t) & \text{else.} \end{cases}$$

Observation 4.20. $C^j(\text{Feas}) \geq C^i(\text{OPT}) \geq C^i(\text{Feas})$ and $C^j(\text{Feas}) \geq C^j(\text{OPT}) \geq C^i(\text{Feas})$.

This follows from the definition of Feas .

Claim 4.21. *Feas is an optimal schedule.*

Proof As already observed before, Feas and OPT only differ in jobs i and j , so we are only interested in the savings of these two jobs:

$$\begin{aligned}
 s^i(\text{OPT}) + s^j(\text{OPT}) &= \int_0^{C^i(\text{OPT})} \text{OPT}^i(t) dt + \int_0^{C^j(\text{OPT})} \text{OPT}^j(t) dt \\
 &= \int_0^{C^j(\text{Feas})} \text{OPT}^i(t) + \text{OPT}^j(t) dt \\
 &= \int_0^{C^j(\text{Feas})} \text{Feas}^i(t) + \text{Feas}^j(t) dt \\
 &= \int_0^{C^i(\text{Feas})} \text{Feas}^i(t) dt + \int_0^{C^j(\text{Feas})} \text{Feas}^j(t) dt \\
 &= s^i(\text{Feas}) + s^j(\text{Feas}).
 \end{aligned}$$

The second equality follows from Observation 4.20, while the third and fourth equalities from the definition of Feas. \square

This implies that Feas is an optimal schedule that works on job i at time zero. This contradicts our initial assumption that OPT is the optimal schedule where the time point of the first difference is maximal. Therefore, SPT is an optimal schedule. \square

Proof of Theorem 4.5 Let OPT denote the optimal schedule. Consider the SRPT schedule and let $p_T^i = p^i - (T - r^i) - \int_{r^i}^T \text{SPT}^i(t) dt$ denote the remaining processing time of job i at time $T \in [r^i, C^i(\text{SPT})]$. Fix job i for the rest of the proof.

Claim 4.22. *Take any time point $T \in [r^i, C^i(\text{SPT})]$ and let C denote the smallest completion time that is larger than T . Then, $C - T \leq \frac{1}{2}p_T^i$, i. e., the time between T and the next completion time C is at most half the remaining processing time of job i at time T .*

Proof Under SRPT, at any time point the job with the shortest remaining processing time is processed. Hence, during an interval of length ϵ in which no job is released, the lowest remaining processing time decreases with 2ϵ , while at release dates the lowest remaining processing time can only decrease. Combining these observations we have that after at most half of the lowest

remaining processing time at T , a job is completed. Clearly, p_T^i is an upper bound for this. \square

Let $k \in \{0, \dots, n-1\}$ denote the number of completion times between r^i and $C^i(\text{SPT})$. Furthermore, without loss of generality assume that $r^i < C^1 < \dots < C^k < C^i(\text{SPT})$. We also define $C^0 = r^i$ and $C^{k+1} = C^i(\text{SPT})$. By Claim 4.22 we have for $j \in \{1, \dots, k+1\}$ that

$$C^j - C^{j-1} \leq \frac{1}{2} p_{C^{j-1}}^i \leq \frac{1}{2} (p^i - C^{j-1} + r^i) \Rightarrow C^j \leq \frac{1}{2} (p^i + C^{j-1} + r^i).$$

Corollary 4.23. $C^j \leq \left(1 - \frac{1}{2^j}\right) p^i + r^i, \quad \forall j \in \{1, \dots, k+1\}$.

Proof We prove the corollary by induction. Thereto, first consider the case where j equals 1. We get

$$C^1 \leq \frac{1}{2} (p^i + C^0 + r^i) = \frac{1}{2} p^i + r^i = \left(1 - \frac{1}{2}\right) p^i + r^i.$$

Assume that the corollary holds for $1, \dots, j$. For $j+1$ we have that

$$\begin{aligned} C^{j+1} &\leq \frac{1}{2} (p^i + C^j + r^i) \leq \frac{1}{2} p^i + \frac{1}{2} \left(1 - \frac{1}{2^j}\right) p^i + \frac{1}{2} r^i + \frac{1}{2} r^i \\ &= \left(1 - \frac{1}{2^{j+1}}\right) p^i + r^i, \end{aligned}$$

which shows that the corollary holds for $j+1$ as well. \square

Lemma 4.24. For any $i \in N$, $s^i(\text{SPT}) \geq \frac{1}{2^n} p^i$.

Proof Combining the definition of $C^i(\text{SPT})$, the observation that $C^{k+1} = C^i(\text{SPT})$, and Corollary 4.23 results in

$$p^i + r^i - s^i(\text{SPT}) = C^i(\text{SPT}) = C^{k+1} \leq \left(1 - \frac{1}{2^{k+1}}\right) p^i + r^i \Rightarrow s^i(\text{SPT}) \geq \frac{1}{2^n} p^i,$$

where the last inequality follows from the fact that $k \in \{0, \dots, n-1\}$. \square

From Observation 4.1 and Lemma 4.24 it follows that

$$2^{n-1} SW(\text{SPT}) = 2^{n-1} \sum_{i \in N} w^i s^i(\text{SPT}) \geq \sum_{i \in N} \frac{1}{2} w^i p^i \geq \sum_{i \in N} w^i s^i(\text{OPT}) = SW(\text{OPT}),$$

which shows that SPT is indeed a factor $1/2^{n-1}$ approximation.

To see that this ratio is tight, consider the following instance. There are n jobs, split into $n - 1$ “light” jobs with processing time 2^n and weight ϵ , and one “heavy” job with processing time $2^n + \epsilon$ and weight 1. All release dates are equal to zero. Without loss of generality, we assume that jobs are indexed in SPT order.

First, consider the solution obtained by SPT. Since all release dates are equal to zero, we can apply Theorem 4.3 to obtain the time spend on SUB for the light jobs and the heavy job.

$$\sum_{i=1}^{n-1} s^i(\text{SPT}) = \sum_{i=1}^{n-1} 2^i = 2^n - 2, \quad (45)$$

$$s^n(\text{SPT}) = \frac{p^n}{2} - \sum_{i=1}^{n-1} \frac{s^i(\text{SPT})}{2} = 2^{n-1} + \frac{1}{2}\epsilon - \frac{1}{2}(2^n - 2) = 1 + \frac{1}{2}\epsilon. \quad (46)$$

Combining equations (45) and (46) we get that

$$SW(\text{SPT}) = \sum_{i \in N} w^i s^i(\text{SPT}) = (2^n - 2)\epsilon + \left(1 + \frac{1}{2}\epsilon\right) = 1 + \left(2^n - \frac{3}{2}\right)\epsilon. \quad (47)$$

Second, consider the feasible schedule Feas, where at time zero SUB starts working on the heavy job, until it is finished. We get that

$$SW(\text{Feas}) = w^n s^n(\text{Feas}) = 2^{n-1} + \frac{1}{2}\epsilon \Rightarrow SW(\text{OPT}) \geq 2^{n-1} + \frac{1}{2}\epsilon. \quad (48)$$

Equations (47) and (48) imply that the ratio of the total saving SPT and OPT is

$$\frac{SW(\text{SPT})}{SW(\text{OPT})} \leq \frac{1 + \left(2^n - \frac{3}{2}\right)\epsilon}{2^{n-1} + \frac{1}{2}\epsilon} \rightarrow \frac{1}{2^{n-1}} \text{ as } \epsilon \rightarrow 0,$$

which proves that the approximation ratio is indeed tight. \square

Proof of Theorem 4.12 Let SMITH and OPT denote the SMITH’S RULE schedule and the optimal schedule, respectively. Consider the following instance

consisting of $n = 2k$ jobs, with $k \geq 2$. The processing time for jobs $i = 1, \dots, k$ is $p^i = i + 1$ and $p^i = k + 2$ for jobs $k + 1, \dots, 2k$. The weight of each job is equal to the processing time plus ϵ , i. e., $w^i = p^i + \epsilon$, and the release date is equal to zero. Note that jobs are indexed according to the SPT rule.

First, consider SMITH schedule.

Claim 4.25. *In SMITH jobs are processed according to the SPT order, i. e., $\frac{w^1}{p^1} \geq \frac{w^2}{p^2} \geq \dots \geq \frac{w^n}{p^n}$.*

Proof The order of jobs does not change over time. To see that, consider two jobs i and j , with $i < j$, that have not been scheduled yet. Let t denote an arbitrary time point between zero and p^i . At time t we have

$$\frac{w^i}{p^i - t} = \frac{p^i + \epsilon}{p^i - t} = 1 + \frac{t + \epsilon}{p^i - t} \geq 1 + \frac{t + \epsilon}{p^j - t} = \frac{p^j + \epsilon}{p^j - t} = \frac{w^j}{p^j - t},$$

which shows that the order of jobs does not change over time. \square

Claim 4.26. $s^i(\text{SMITH}) = 1$ for $i = 1, \dots, k$ and $s^i(\text{SMITH}) = \frac{1}{2^{i-k-1}}$ for $i = k + 1, \dots, 2k$.

Proof We prove this by induction. If $i = 1$, then $s^1(\text{SMITH}) = \frac{p^1}{2} = 1$. Now, assume the claim holds for $j = 1, \dots, i$, and consider job $i + 1$. We have that

$$\begin{aligned} s^{i+1}(\text{SMITH}) &= \frac{p^{i+1}}{2} - \sum_{j=1}^i \frac{s^j(\text{SMITH})}{2} \\ &= \begin{cases} \frac{i+2}{2} - \frac{i}{2} = 1 & \text{if } i \leq k \\ \frac{k+2}{2} - \frac{k}{2} - \sum_{j=k+1}^i \frac{1}{2^{j-k}} = \frac{1}{2^{i-k}} & \text{if } i > k, \end{cases} \end{aligned}$$

which concludes the proof. \square

This claim implies that

$$\sum_{i=k+1}^{2k} s^i(\text{SMITH}) = \sum_{i=k+1}^{2k} \frac{1}{2^{i-k-1}} = 2 - \frac{1}{2^{k-1}}.$$

Combining these observations results in that the social welfare for SMITH is

$$\begin{aligned}
 SW^N(\text{SMITH}) &= \sum_{i \in N} w^i s^i(\text{SMITH}) \\
 &= \sum_{i=1}^k (i + 1 + \epsilon) + \left(2 - \frac{1}{2^{k-1}}\right) (k + 2 + \epsilon) \\
 &= \frac{k^2}{2} + \frac{7k}{2} - \frac{k+2}{2^{k-1}} + 4 + \left(k + 2 - \frac{1}{2^{k-1}}\right) \epsilon. \tag{49}
 \end{aligned}$$

Now, consider the optimal schedule that process only jobs $k + 1, \dots, 2k$. Observe that all these jobs have a processing time of $k + 2$, hence they are processed in non-decreasing order of their processing time. By Theorem 4.3 we get that

$$\sum_{i=1}^k s^{k+i}(\text{OPT}) = \sum_{i=1}^k \frac{k+2}{2^{k-i+1}} = (k + 2) \left(1 - \frac{1}{2^k}\right).$$

Note that for $k \geq 2$ we have that $(k + 2) \left(1 - \frac{1}{2^k}\right) > k + 1$, which implies that jobs $1, \dots, k$ have already finished by the time jobs $k + 1, \dots, 2k$ are processed. Therefore $s^i(\text{OPT}) = 0$ for $i = 1, \dots, k$. The social welfare for OPT is then

$$\begin{aligned}
 SW^N(\text{OPT}) &= \sum_{i \in N} w^i s^i(\text{OPT}) \\
 &= (k + 2) \left(1 - \frac{1}{2^k}\right) (k + 2 + \epsilon) \\
 &\geq (k^2 + 4k + 4) \left(1 - \frac{1}{2^k}\right). \tag{50}
 \end{aligned}$$

Equations (49) and (50) provide the ratio of

$$\frac{SW(\text{SMITH})}{SW(\text{OPT})} \leq \frac{\frac{k^2}{2} + \frac{7k}{2} - \frac{k+2}{2^{k-1}} + 4 + \left(k + 2 - \frac{1}{2^{k-1}}\right) \epsilon}{(k^2 + 4k + 4) \left(1 - \frac{1}{2^k}\right)} \rightarrow \frac{1}{2} \text{ as } \epsilon \rightarrow 0, k \rightarrow \infty,$$

which shows that SMITH's RULE is at best a factor 1/2 approximation. □

Proof of Theorem 4.13 Let OPT denote the optimal schedule. Let k denote the job for which the processing time multiplied with weight is maximised, i. e., $p^k w^k \geq p^i w^i$ for all $i \in N$. For OPT the total weighted saving equals

$$SW(\text{OPT}) = \sum_{i \in N} w^i s^i(\text{OPT}) \leq \sum_{i \in N} \frac{1}{2} p^i w^i \leq \sum_{i \in N} \frac{1}{2} p^k w^k = \frac{n}{2} p^k w^k, \quad (51)$$

where the first inequality follows from Observation 4.1. Consider HPF, and let p_t^i denote the remaining processing time of job i at time $t \in [r^i, C^i(\text{HPF})]$. Outside of this range, we define p_t^i to be equal to zero. Let $\mathcal{T} = \{r^i, C^i(\text{HPF}) \mid i \in N\}$ denote the set of all release dates and completion times. Let $SW_t(\text{HPF})$ denote the weighted sum of savings obtained after time t .

Lemma 4.27. *For any $t \in \mathcal{T}$, $SW_t(\text{HPF}) \geq \frac{1}{2} p_t^k w^k$.*

Proof Observe that by definition the lemma holds for the maximum time point in \mathcal{T} . This time point is the completion time of the last job, thus the remaining processing time of all jobs is zero. Assume to the contrary that the lemma does not hold for all time points in \mathcal{T} . Let t denote the maximum time point for which the lemma does not hold, and let τ denote the first time point in \mathcal{T} after t . By our previous observation, τ always exists. Let i and j denote the jobs for which the remaining processing time multiplied with the weight is maximised at times t and τ , respectively. Observe that between t and τ SUB is working on job i , as there are no other release dates or completion times between these two time points. Together with the definition of jobs i and j , we have that $p_\tau^j w^j \geq p_t^i w^i = (p_\tau^i - 2(\tau - t)) w^i$. This leads to

$$\begin{aligned} SW_t(\text{HPF}) &= (\tau - t) w^i + SW_\tau(\text{HPF}) \\ &\geq (\tau - t) w^i + \frac{1}{2} p_\tau^j w^j \\ &\geq (\tau - t) w^i + \frac{1}{2} p_\tau^i w^i \\ &\geq (\tau - t) w^i + \frac{1}{2} (p_\tau^i - 2(\tau - t)) w^i = \frac{1}{2} p_t^i w^i, \end{aligned}$$

which contradicts the assumption. \square

Consider time point r^k , i. e., the release date of the job for which processing time multiplied with weight is maximised. By definition $p_{r^k}^k = p^k$, thus

$$SW(\text{HPF}) \geq SW_{r^k}(\text{HPF}) \geq \frac{1}{2}p^k w^k. \quad (52)$$

Combining equations (51) and (52) results in $nSW(\text{HPF}) \geq \frac{n}{2}p^k w^k \geq SW(\text{OPT})$, which shows that HPF is indeed a factor $1/n$ approximation.

To see that this ratio is tight, consider the following instance. There are n jobs, with release date $r^i = 0$, processing time $p^i = k^i$, and weight $w^i = k^{n-i}$, except for job n , which has weight $w^n = 1 + \epsilon$. Let $k \geq 2$. Observe that the jobs are indexed in SPT order (see Section 4.3.1).

First, consider HPF, which starts with job n . In this case we have that

$$C^n(\text{HPF}) = s^n(\text{HPF}) = \frac{1}{2}p^n = \frac{1}{2}k^n.$$

Observe that at time $\frac{1}{2}k^n$ all other jobs are also finished, so $s^i(\text{HPF}) = 0$ for all $i \neq n$. Combining these observations we get that

$$SW(\text{HPF}) = w^n s^n(\text{HPF}) = \frac{1}{2}k^n + \frac{1}{2}k^n \epsilon. \quad (53)$$

I would like to thank the kind reader who made it that far in my thesis. Your enthusiasm is what kept me going on writing these details. Now, consider the optimal schedule, where jobs are processed in order of their index. As all release dates are equal to zero, we immediately get the time spend on SUB from Theorem 4.3, that is, $s^i(\text{OPT}) = \frac{k^i}{2} - \sum_{j=1}^{i-1} \frac{k^j}{2^{i-j+1}}$ for all jobs $i \in N$. Consequently the optimal social welfare is

$$\begin{aligned} SW(\text{OPT}) &= \sum_{i \in N} w^i s^i(\text{OPT}) \\ &= \sum_{i=1}^n \left(\frac{k^i}{2} - \sum_{j=1}^{i-1} \frac{k^j}{2^{i-j+1}} \right) k^{n-i} + \left(\frac{k^n}{2} - \sum_{j=1}^{n-1} \frac{k^j}{2^{n-j+1}} \right) \epsilon \\ &\geq \sum_{i=1}^n \frac{k^n}{2} - \sum_{i=1}^n \sum_{j=1}^{i-1} \frac{k^{n+j-i}}{2^{i-j+1}} = \frac{n}{2}k^n - \sum_{i=1}^{n-1} (n-i) \frac{k^{n-i}}{2^{i+1}}. \end{aligned} \quad (54)$$

Equations (53) and (54) imply that the social welfare ratio of the two schedule is

$$\frac{SW(\text{HPF})}{SW(\text{OPT})} \leq \frac{\frac{1}{2}k^n + \frac{1}{2}k^n \epsilon}{\frac{n}{2}k^n - \sum_{i=1}^{n-1} (n-i) \frac{k^{n-i}}{2^{i+1}}} \rightarrow \frac{1}{n} \text{ as } \epsilon \rightarrow 0, k \rightarrow \infty,$$

which shows that the approximation ratio is indeed tight. \square

Part IV

APPENDIX

BIBLIOGRAPHY

- [1] AFRATI, F., E. BAMPIS, C. CHEKURI, D. KARGER, C. KENYON, S. KHANNA, I. MILIS, M. QUEYRANNE, M. SKUTELLA, C. STEIN, AND M. SVIRIDENKO (1999): “Approximation Schemes for Minimizing Average Weighted Completion Time with Release Dates,” in *Proceedings of the 40th Annual Symposium on Foundations of Computer Science*, IEEE Computer Society, FOCS '99, 32–43.
- [2] ANDERSON, E. (1981): “A new continuous model for job-shop scheduling,” *International Journal of Systems Science*, 12, 1469–1475.
- [3] ARMSTRONG, M. (1996): “Multiproduct Nonlinear Pricing,” *Econometrica*, 64, pp. 51–75.
- [4] ASEFF, J. AND H. CHADE (2008): “An Optimal Auction with Identity-Dependent Externalities,” *The RAND Journal of Economics*, 39, pp. 731–746.
- [5] BACHMAN, A., T. CHENG, A. JANIAK, AND C. NG (2002): “Scheduling Start Time Dependent Jobs to Minimize the Total Weighted Completion Time,” *Journal of the Operational Research Society*, 53, 688–693.
- [6] BAGNOLI, M. AND T. BERGSTROM (2005): “Log-Concave Probability and Its Applications,” *Economic Theory*, 26, pp. 445–469.
- [7] BARLOW, R. E., A. W. MARSHALL, AND F. PROSCHAN (1963): “Properties of probability distributions with monotone hazard rate,” *The Annals of Mathematical Statistics*, 34, 375–389.
- [8] BHARADWAJ, V., D. GHOSE, V. MANI, AND T. ROBERTAZZI (1996): *Scheduling Divisible Loads in Parallel and Distributed Systems*, IEEE Computer Society Press.
- [9] BRANCO, F. (1996): “Multiple unit auctions of an indivisible good,” *Economic Theory*, 8, 77–101.

- [10] BRUNO, J., E. C. JR., AND R. SETHI (1974): "Scheduling Independent Tasks to Reduce Mean Finishing Time," *Communications of the ACM*, 17, 382–387.
- [11] BULOW, J. AND P. KLEMPERER (1996): "Auctions versus Negotiations," *American Economic Review*, 86, 180–94.
- [12] CAI, Y., C. DASKALAKIS, AND S. WEINBERG (2013): "Understanding Incentives: Mechanism Design Becomes Algorithm Design," in *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on*, 618–627.
- [13] CAI, Y., C. DASKALAKIS, AND S. M. WEINBERG (2011): "On Optimal Multi-Dimensional Mechanism Design," *SIGecom Exch.*, 10, 29–33.
- [14] CHEN, Y., G. IYER, AND V. PADMANABHAN (2002): "Referral Infomediaries," *Marketing Science*, 21, 412–434.
- [15] CHENG, T., Q. DING, AND B. LIN (2004): "A concise survey of scheduling with time-dependent processing times," *European Journal of Operational Research*, 152, 1–13.
- [16] CHUNG, K.-S. AND J. C. ELY (2002): "Ex-Post Incentive Compatible Mechanism Design," Discussion Papers 1339, Northwestern University, Center for Mathematical Studies in Economics and Management Science.
- [17] CLARKE, E. H. (1971): "Multipart pricing of public goods," *Public Choice*, 11, 17–33.
- [18] CRÉMER, J. AND R. P. MCLEAN (1988): "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions," *Econometrica*, 56, 1247–1257.
- [19] CSAPÓ, G. AND R. MÜLLER (2013): "Optimal mechanism design for the private supply of a public good," *Games and Economic Behavior*, 80, 229–242.
- [20] DEBO, L., P. ENDERS, A. GANDHI, V. GUPTA, M. HARCHOL-BALTER, AND A. SCHELLER-WOLF (2014): "Inducing Optimal Scheduling with Selfish Users," .

- [21] DENG, C. AND S. PEKEČ (2013): “Optimal Allocation of Exclusivity Contracts,” .
- [22] DEVANUR, N. R., J. D. HARTLINE, A. R. KARLIN, AND C. T. NGUYEN (2011): “Prior-Independent Multi-parameter Mechanism Design,” in *WINE*, 122–133.
- [23] DHANGWATNOTAI, P., T. ROUGHGARDEN, AND Q. YAN (2010): “Revenue Maximization with a Single Sample,” in *Proceedings of the 11th ACM Conference on Electronic Commerce*, New York, NY, USA: ACM, EC ’10, 129–138.
- [24] DOBZINSKI, S., H. FU, AND R. D. KLEINBERG (2011): “Optimal auctions with correlated bidders are easy,” in *STOC’11*, 129–138.
- [25] DROZDOWSKI, M. (2009): *Scheduling for Parallel Processing*, Springer-Verlag.
- [26] EDELMAN, B., M. OSTROVSKY, AND M. SCHWARZ (2007): “Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords,” *American Economic Review*, 97, 242–259.
- [27] FAALAND, B., K. KIM, AND T. SCHMITT (1990): “A New Algorithm for Computing the Maximal Closure of a Graph,” *Management Science*, 36, 315–331.
- [28] FIGUEROA, N. AND V. SKRETA (2009): “The role of optimal threats in auction design,” *Journal of Economic Theory*, 144, 884–897.
- [29] ——— (2011): “Optimal allocation mechanisms with single-dimensional private information,” *Review of Economic Design*, 15, 213–243.
- [30] GERSHKOV, A., J. K. GOEREE, A. I. KUSHNIR, B. MOLDOVANU, AND X. SHI (2013): “On the Equivalence of Bayesian and Dominant Strategy Implementation,” *Econometrica*, 81, 197–220.
- [31] GERSHKOV, A., B. MOLDOVANU, AND X. SHI (2011): “Bayesian and Dominant Strategy Implementation Revisited,” Working papers, University of Toronto, Department of Economics.

- [32] GHOSE, A., T. MUKHOPADHYAY, AND U. RAJAN (2003): "Strategic Benefits of Internet Referral Services," in *Proceedings of the 5th International Conference on Electronic Commerce*, New York, NY, USA: ACM, ICEC '03, 240–247.
- [33] GHOUILA-HOURI, A. (1962): "Caractérisation des matrices totalement unimodulaires," *Comptes Rendus Hebdomadaires des Séances de l'Académie des Sciences*, 254, pp. 1192–1194.
- [34] GOPALAKRISHNA, S. AND G. L. LILIEN (1995): "A Three-Stage Model of Industrial Trade Show Performance," *Marketing Science*, 14, 22–42.
- [35] GROVES, T. (1973): "Incentives in Teams," *Econometrica*, 41, pp. 617–631.
- [36] GÜTH, W. AND M. HELLWIG (1986): "The private supply of a public good," *Journal of Economics*, 5, 121–159.
- [37] HARTLINE, J. D. (2012): "Approximation in mechanism design," *The American Economic Review*, 102, 330–336.
- [38] HEZARKHANI, B. AND W. KUBIAK (forthcoming): "Decentralized Subcontractor Scheduling with Divisible Jobs," *Journal of Scheduling*.
- [39] HOCHBAUM, D. S. AND A. CHEN (2000): "Performance Analysis and Best Implementations of Old and New Algorithms for the Open-Pit Mining Problem," *Operations Research*, 48, pp. 894–914.
- [40] JEHIEL, P., M. MEYER-TER VEHN, AND B. MOLDOVANU (2007): "Mixed bundling auctions," *Journal of Economic Theory*, 134, 494–512.
- [41] JEHIEL, P., M. MEYER-TER VEHN, B. MOLDOVANU, AND W. R. ZAME (2006): "The Limits of Ex Post Implementation," *Econometrica*, 74, pp. 585–610.
- [42] JEHIEL, P. AND B. MOLDOVANU (2001): "Efficient Design with Interdependent Valuations," *Econometrica*, 69, pp. 1237–1259.
- [43] ——— (2006): *Allocative and Informational Externalities in Auctions and Related Mechanisms*, Discussion paper, Centre for Economic Policy Research (Great Britain).

- [44] JEHIEL, P., B. MOLDOVANU, AND E. STACCHETTI (1996): "How (Not) to Sell Nuclear Weapons," *The American Economic Review*, 86, pp. 814–829.
- [45] JERATH, K. AND A. SAYEDI (2012): "Exclusive Display in Sponsored Search Advertising," Working paper.
- [46] LABETOULLE, J., E. LAWLER, J. LENSTRA, AND A. R. KAN (1984): "Preemptive scheduling of uniform machines subject to release dates," .
- [47] LAFFONT, J.-J. AND D. MARTIMORT (2000): "Mechanism Design with Collusion and Correlation," *Econometrica*, 68, pp. 309–342.
- [48] LIKHODEDOV, A. AND T. SANDHOLM (2004): "Methods for Boosting Revenue in Combinatorial Auctions," in *Proceedings of the 19th National Conference on Artificial Intelligence*, AAAI Press, AAAI'04, 232–237.
- [49] MANELLI, A. AND D. VINCENT (2010): "Bayesian and Dominant-Strategy Implementation in the Independent Private-Values Model," *Econometrica*, 78, 1905–1938.
- [50] MANELLI, A. M. AND D. R. VINCENT (2007): "Multidimensional mechanism design: Revenue maximization and the multiple-good monopoly," *Journal of Economic Theory*, 137, 153–185.
- [51] MAS-COLLEL, A., M. WHINSTON, AND J. GREEN (1995): *Microeconomic Theory*, Oxford student edition, Oxford University Press.
- [52] MCAFEE, R. P. AND P. J. RENY (1992): "Correlated Information and Mechanism Design," *Econometrica*, 60, pp. 395–421.
- [53] MILGROM, P. R. AND R. J. WEBER (1982): "A Theory of Auctions and Competitive Bidding," *Econometrica*, 50, pp. 1089–1122.
- [54] MYERSON, R. (1981): "Optimal auction design," *Mathematics of Operations Research*, 6, 58–73.
- [55] NISAN, N. AND A. RONEN (2001): "Algorithmic Mechanism Design," *Games and Economic Behavior*, 35, 166–96.

- [56] NISAN, N., T. ROUGHGARDEN, E. TARDOS, AND V. V. VAZIRANI (2007): *Algorithmic Game Theory*, New York, NY, USA: Cambridge University Press.
- [57] PAPADIMITRIOU, C. H. AND G. PIERRAKOS (2011): “On Optimal Single-item Auctions,” in *Proceedings of the Forty-third Annual ACM Symposium on Theory of Computing*, New York, NY, USA: ACM, STOC '11, 119–128.
- [58] PEI, J., D. KLABJAN, AND W. XIE (2014): “Approximations to auctions of digital goods with share-averse bidders,” *Electronic Commerce Research and Applications*, 13, 128–138.
- [59] PICARD, J.-C. (1976): “Maximal Closure of a Graph and Applications to Combinatorial Problems,” *Management Science*, 22, pp. 1268–1272.
- [60] RILEY, J. G. AND W. F. SAMUELSON (1981): “Optimal Auctions,” *American Economic Review*, 71, 381–92.
- [61] ROBERTS, K. (1979): “The characterization of implementable choice rules,” in *Aggregation and Revelation of Preferences*, ed. by J.-J. Laffont, North-Holland, 321–349.
- [62] ROCHET, J.-C. AND P. CHONÉ (1998): “Ironing, Sweeping, and Multidimensional Screening,” *Econometrica*, 66, pp. 783–826.
- [63] ROUGHGARDEN, T. AND M. SUNDARARAJAN (2007): “Is efficiency expensive,” in *In Third Workshop on Sponsored Search Auctions*.
- [64] ROUGHGARDEN, T. AND I. TALGAM-COHEN (2013): “Optimal and Near-optimal Mechanism Design with Interdependent Values,” in *Proceedings of the Fourteenth ACM Conference on Electronic Commerce*, New York, NY, USA: ACM, EC '13, 767–784.
- [65] SADYKOV, R. (2012): “A Dominant Class of Schedules for Malleable Jobs in the Problem to Minimize the Total Weighted Completion Time,” *Computers and Operations Research*, 39, 1265–1270.

- [66] SALEK, M. AND D. KEMPE (2008): "Auctions for Share-Averse Bidders." in *WINE*, ed. by C. H. Papadimitriou and S. Zhang, Springer, vol. 5385 of *Lecture Notes in Computer Science*, 609–620.
- [67] SAMUELSON, P. A. (1954): "The Pure Theory of Public Expenditure," *The Review of Economics and Statistics*, 36, 387–389.
- [68] SANDHOLM, T. (2003): "Automated Mechanism Design: A New Application Area for Search Algorithms," in *Principles and Practice of Constraint Programming - CP 2003*, ed. by F. Rossi, Springer Berlin Heidelberg, vol. 2833 of *Lecture Notes in Computer Science*, 19–36.
- [69] SAYEDI, A. (2012): "Essays on Sponsored Search Advertising," Ph.D. thesis, Pittsburgh, PA, USA.
- [70] SCHMITZ, P. W. (1997): "Monopolistic Provision of Excludable Public Goods under Private Information," *Public Finance / Finances publiques*, 52, 89–101.
- [71] SCHRAGE, L. (1968): "A proof of the optimality of the shortest remaining processing time discipline," *Operations Research*, 16, 687–690.
- [72] SCHRIJVER, A. (2003): *Combinatorial optimization*, Springer.
- [73] SCHUMMER, J. AND R. V. VOHRA (2007): "Mechanism design without money," in *Algorithmic Game Theory*, ed. by N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, New York, NY, USA: Cambridge University Press.
- [74] SEGAL, I. (1999): "Contracting with Externalities," *The Quarterly Journal of Economics*, 114, pp. 337–388.
- [75] SITTERS, R. (2010): "Efficient Algorithms for Average Completion Time Scheduling," in *Integer Programming and Combinatorial Optimization*, ed. by F. Eisenbrand and F. Shepherd, Springer Berlin Heidelberg, vol. 6080 of *Lecture Notes in Computer Science*, 411–423.
- [76] SMITH, W. (1956): "Various optimizers for single-stage production," *Naval Research Logistics Quarterly*, 3, 59–66.

- [77] SUIJS, J. (1996): “On incentive compatibility and budget balancedness in public decision making,” *Economic design*, 2, 193–209.
- [78] TANG, P. AND T. SANDHOLM (2012): “Mixed-bundling Auctions with Reserve Prices,” in *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 2*, Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems, AAMAS '12, 729–736.
- [79] VAIRAKTARAKIS, G. (2013): “Noncooperative Games for Subcontracting Operations,” *Manufacturing & Service Operations Management*, 15, 148–158.
- [80] VAIRAKTARAKIS, G. AND T. AYDINLIYIM (2007): “Centralization vs. Competition in Subcontracting Operations,” Technical memorandum number 819, Case Western Reserve University.
- [81] VARIAN, H. R. (2007): “Position auctions,” *International Journal of Industrial Organization*, 25, 1163–1178.
- [82] VICKREY, W. (1961): “Counterspeculation, auctions, and competitive sealed tenders,” *The Journal of Finance*, 16, 8–37.
- [83] VOHRA, R. V. (2011): *Mechanism Design: A Linear Programming Approach*, Econometric Society Monographs, Cambridge University Press.
- [84] WILSON, R. B. (1967): “Competitive Bidding with Asymmetric Information,” *Management Science*, 13, pp. 816–820.

INDEX

affine maximiser, 62, 80
approximation, 14, 80, 107

Bayes-Nash incentive compatible,
12

dominant strategy incentive com-
patible, 11, 111

ex-post equilibrium, 114
ex-post incentive compatible, 11
ex-post individual rational, 12

interim individual rational, 12

monotone hazard rate (MHR), 19,
64

monotonicity, 13, 24

NP-hardness, 44, 102

social welfare, 9

taxation principle, 13
trongly budget-balanced, 117

NEDERLANDSE SAMENVATTING

Een fundamentele vraag binnen economie is de verdeling van schaarse goederen tussen personen. Een ontwerper van een economisch systeem kan deze vraag modelleren als een optimalisatieprobleem: gegeven de eigenschappen van het probleem, maximaliseer een doelfunctie met randvoorwaarden die de beperkingen op goederen weergeven. Bijvoorbeeld, gegeven een aantal bidders met bijbehorende voorkeuren voor de objecten, wie moet welk object en tegen welke prijs ontvangen, zodat de sociale welvaart wordt gemaximaliseerd. Of gegeven een aantal projecten met bijbehorende tijdsduur en kosten, hoe moet de tijd van een onderaannemer worden verdeeld zodanig dat de totale kosten worden geminimaliseerd.

Ondanks dat dergelijke optimalisatieproblemen in veel gevallen al lastig zijn, ontbreekt het deze modellen aan realistische kenmerken. Denk bijvoorbeeld aan personen die informatie privé houden, zodat deze op strategische wijze in eigenbelang kan worden gebruikt. Immers, een bidder kan liegen zodat de prijs voor het object lager uitvalt, of een projectleider kan de verwachte kosten van zijn project verhogen, zodat de onderaannemer meer tijd besteedt aan het project. Doordat veel modellen geen rekening houden met deze vorm van manipulatie worden grondstoffen verspild, opbrengsten verlaagd en projecten vertraagd. In het ergste geval kan dit zelfs leiden tot het instorten van een systeem, zoals tijdens de recente financiële crisis. Aan de andere kant, als een systeem goed ontworpen is, dan komt dit ten goede aan de gemeenschap. Een bekend voorbeeld is het Amerikaanse programma voor het verhogen van niertransplantaties. Het ontwerp is dusdanig dat er op een moreel verantwoorde manier donateurs en ontvangers worden zamen gebracht. Een ander voorbeeld zijn gesponsorde zoekresultaten op Internet die zorgen voor een effectievere marketing van bedrijven en tegelijkertijd omzet genereren voor zoekmachines. Vanwege deze redenen zouden we de oorspronkelijke vraag moeten vervangen door: hoe verdeel je schaarse goederen over *competitieve* personen ondanks hun *strategisch gedrag*?

Verdere complicaties treden op als de ontwerper te maken heeft met beperkte tijd en rekencapaciteit. Deze problemen zijn veelal praktisch van aard vanwege de grote hoeveelheid aan data en de snelheid van beslissingen. Denk aan een online veiling waarin meerdere objecten tegelijk worden verkocht. De winnaar van ieder object met bijbehorende prijs dient snel bepaald te worden, rekening houdend met de beperkte capaciteit van de servers. Andere voorbeelden zijn de coördinatie van projecten in lange productieketens. In dit geval zijn snelle beslissingen essentieel om deadlines te halen. Om ook deze zaken in acht te nemen, is de hoofdvraag van dit proefschrift: hoe verdeel je schaarse goederen over *competitieve* personen ondanks hun *strategisch gedrag* en een *beperkte rekencapaciteit*?

Dit proefschrift behoort tot de tak der Algoritmisch Mechanisme Ontwerp, een combinatie van economie, speltheorie en informatica. Speltheorie helpt om het gedrag en de strategische interactie van rationele agenten te modelleren. Informatica wordt gebruikt om de complexiteit van de berekeningen van de mechanismes te analyseren.

Dit proefschrift onderzoekt drie toepassingen in detail. Het eerste gedeelte heeft betrekking tot optimaal mechanisme ontwerp. In het bijzonder, Hoofdstuk 2 houdt zich bezig met het aanbod van een publiek goed. Hierbij wordt aan de hand van lineair programmeren gezocht naar het mechanisme dat de winst maximaliseert in allerlei situaties. In ons model zijn gecorreleerde types alsook externe informatie-effecten toegestaan. Als belangrijkste resultaat bewijzen we het bestaan van een algoritme met polynomiale rekentijd die het optimale mechanisme vindt dat deterministisch, ex-post incentive compatibel en ex-post individueel rationeel is. Daarnaast bespreken we het effect van randomisatie en de zwakkere condities Bayes-Nash incentive compatibiliteit en interim individuele rationaliteit.

Om te ontkomen aan de moeilijkheid van multidimensionaal mechanisme ontwerp en om eenvoudige mechanismen te vinden voor gebruikers, richt het tweede gedeelte van het proefschrift zich op het ontwerp van approximaties van het optimale mechanisme. Hoofdstuk 3 gaat over online veilingen, geïnspireerd door de samenwerking met Reply!, een online marktplaats voor verkoopinformatie. Het verschil met het vorige hoofdstuk is dat de persoonlijke informatie uit meerdere dimensies bestaat: de waardering voor een exclusieve service is anders dan voor een niet-exclusieve service. We presen-

teren eenvoudige mechanismen die makkelijk te implementeren zijn in de praktijk en een constante fractie garanderen van de optimale opbrengst. Om zulke resultaten aan te tonen, begrenzen we de optimale opbrengst aan de hand van relaxaties naar een één-dimensionaal probleem. Deze begrenzingen worden gebruikt in zowel de theoretische bewijzen alsook in de numerieke simulaties.

Hoofdstuk 4 is gewijd aan het maken van een planning voor een onderaannemer. Dit probleem is zowel theoretisch als praktisch interessant. Gegeven zijn een aantal projecten met ieder zijn eigen team. Daarnaast is er een onderaannemer die mee kan helpen aan één project tegelijk, zodat de tijdsduur van het project wordt verminderd. In tegenstelling tot de vorige twee hoofdstukken zijn niet de opbrengsten, maar het maximaliseren van de sociale welvaart het hoofddoel. In andere woorden, we zoeken naar een planning voor de onderaannemer die in de praktijk kan worden gebruikt en binnen een constante fractie zit van de optimale kostenbesparing. Hiervoor gebruiken we eenvoudige heuristieken die een ordening bepalen voor de projecten en berekenen de bijbehorende benaderingsfactor. Vanuit het oogpunt van een ontwerper ontwerpen we mechanismen met speciale eigenschappen, zoals een evenwichtige distributie van geldstromen.

“process of creating value from knowledge, by making knowledge suitable and/or available for social (and/or economic) use and by making knowledge suitable for translation into competitive products, services, processes and new commercial activities.”

— adapted definition based on the National Valorization Committee 2011:8

VALORISATION

In accordance with the principle of the National Valorization Committee cited above, this chapter discusses the social and economic benefits of my thesis. In my view, the underlying purpose of adding a valorisation chapter to a thesis is to make transparent the direct and indirect return for society on the public sector’s research investments. Therefore, in order to thoroughly assess that investment I am going to mention also the benefits of the process that leads to the final thesis, which I think bears at least as much social value.

BENEFITS OF A PHD STUDENT The set of diverse tasks associated with the PhD track includes active research, teaching, giving presentations at seminars and conferences, refereeing articles for journals, publishing papers and networking. These activities constitute the cornerstones of scientific research: the creation and sharing of knowledge. The benefits of supporting PhD students don’t stop at the defence of the thesis. Researchers pursuing a career outside of academia strengthen the cooperation between universities and the private sector. A good illustration for the importance of establishing stronger ties between academia and business is the Study Group Mathematics with Industry workshop organised every year. The idea of this event is to gather researchers from different fields and invite industrial companies to present a particular problem they don’t have the knowledge or capacity to deal with. In particular, I was a participant of that workshop in 2011 and worked on the problem offered by Statkraft, a German company that is involved in energy generation and trading. The task was to analyse from a game theoretical point of view the bidding behaviour of the participants in the auction that

sells capacity for the German electricity primary reserve. Although the reality usually doesn't fall perfectly into the theoretical framework, with careful and critical thinking our working team managed to give useful insights. An interview version of our findings can be found online ¹.

RELEVANCE OF THE FIELD The methodology applied throughout the thesis belongs to the field of algorithmic mechanism design (AMD), an intersection of economics, game theory and computer science. The additional insight of AMD is that it brings optimisation problems one step closer to reality by taking into account also the strategic interactions of the involved parties in a formalised manner. This is an important aspect because parameters that serve as input for optimisation models are usually dispersed and privately held information of strategically behaving agents, who strive to achieve their own goals in interaction with the system. Indeed, bidders can lie about their preferences in order to lower their payments, or jobs may inflate their opportunity costs to get more help from the subcontractor. Failing to take incentives into account can lead to waste of resources, lower revenue, significant project delays or even to system meltdowns such as the recent financial crisis. On the other hand, well-designed institutions can help society. A well-known example is regarding matching algorithms in kidney exchange programs that facilitate finding compatible donors in a way that morally questionable incentives are taken care of. In a similar fashion sponsored search advertising helps companies targeting customers better, while yielding considerable revenue for the search engine providers.

VALORISATION OF CHAPTER 2 In Chapter 2 we search for a mechanism that maximises the profit of a monopolistic provider of a public good. We approach this problem in a general setting since we allow for correlated type distributions as well as for informational externalities. As main result we show that the optimal solution can be found in polynomial time, moreover we offer concrete algorithms for that purpose by establishing connection to a well-studied optimisation problems.

For practitioners the main message of this chapter is that formulating optimization problems in an integer linear programming framework can lead

¹ <http://www.nieuwarchief.nl/serie5/pdf/naw5-2012-13-4-287.pdf>

to additional structural insights and may reveal relations to other problems for which solutions have been already developed.

The paper, on which it is based, has also been published in the journal of *Games and Economic Behavior* (Csapó and Müller [19]). Journal publication is one of the most important and rewarding form of knowledge sharing. This way the work can reach a wide audience and gets strong quality assurance via the peer review system.

VALORISATION OF CHAPTER 3 The scientific work underlying Chapter 3 is built on the cooperation with Reply!, a sales lead generator company. A sales lead contains condensed information about an individual who has already exhibited interest in buying a product or service. These leads are sold to providers of the related goods via online platforms. Due to the threat of competition exclusive ownership of such a sales lead have higher value for the buyers compared to outcomes in which it is sold to multiple parties. In many industries online lead generation is a cheaper means for acquiring a customer than traditional alternatives such as newspaper, television or billboard advertising. This business is growing rapidly, its total revenue reached 1.75 billion USD in 2013.² Reply! has currently a first price auction in place for selling leads, but constantly tries to come up with new innovations to boost its revenue. The outcomes of that project have not just theoretical but practical merits as well. In particular, the relevance of our results for this company is that we devise auction formats that are simple, require minimal user interaction and easy to implement. The advantage of our design is that we have a theoretical proof that the aforementioned useful properties come at a small loss of revenue compared to the optimal auction format. Such findings bear relevance as the construction of the optimal mechanism for such settings is in general very complicated, moreover, the optimal solution might not be suitable for practical implementation. Our proposed mechanisms also have some advantages over the first price auction that Reply! runs. For example, bidders are incentivised to bid their true valuation, while in the first price auction it is beneficial for them to lower their bids. The resulting bidding behaviour of the latter is hard to predict and might lead to diverse and unwanted outcomes. Our solution also introduces carefully set minimum bids

² http://www.iab.net/media/file/IAB_Internet_Advertising_Revenue_Report_FY_2013.pdf

that can help retain profit in situations when the auction has only a handful of participants. This feature is currently missing from the company's practice. Contacts from Reply! have already expressed interest in our insights and would like to apply them to improve their business. There are plans for further cooperation that would entail the implementation of a version of our proposed auction format for certain products.

The crux of the implementation is how to set the minimum bids below which the item is not sold. Too high thresholds can result in loss of sale, while setting them too low would drive the prices down. For a certain class of bid distributions we propose a way of finding the appropriate levels, but for general distributions further sensitivity analysis and simulations have to be carried out.

Our results can be appealing to other enterprises as well that are interested in selling copies of digital goods for which buyers would pay a higher price if they can own them exclusively. One potential area of application is, for example, sponsored search advertising. Recently, major online search engines like Google, Yahoo! and Bing have begun considering designs where advertisers can place two-dimensional bids for receiving the click exclusively or shared Sayedi [69], Jerath and Sayedi [45]. Yahoo! has even patented such an auction, an extension of the GSP.

VALORISATION OF CHAPTER 4 A subcontractor helps companies with their projects by offering them external resources and expertise. This way companies can reduce the completion time of their projects and save cost. Subcontracting is a common practice in industries that face volatile demand and hard capacity constraints. When multiple projects can benefit from the use of a subcontractor, it also becomes a question of how to schedule the subcontractor over the projects involved. Further complications arise from the fact that such projects are usually run by different divisions or teams whose interest is to maximise their own savings. Therefore the company is facing a global optimisation problem where special attention has to be paid to coordination and incentives.

The focus of this chapter is to construct intuitive and simple schedules that have good approximation guarantee, i. e., for any instance the total weighted saving is at least a fixed proportion of the optimal schedule. In particular, we

show that greedily working on the job with the highest weight at any point in time is a $2/3$ approximation and devise a payment scheme that implements this schedule. For companies the policy implication of that result is that the cost of not having a complicated scheduling software/scheme is bounded and for some businesses it can be considered reasonable. Therefore focusing on tasks with the highest opportunity cost could be already a good enough solution. As the maximum loss is quantified, if decision makers consider it as unacceptable, then the proven approximation ratio serves as indication how much money is worth spending on new research or more complicated solutions. Next to the theoretical guarantee, the power of such an approximation mechanism lies in its simplicity. Only one characteristic, the weight/cost is needed to be taken into account and no further considerations are needed. Hence, the administration cost of such a policy is low and it is very easy to implement. This schedule can be fine-tuned by breaking ties between identical jobs favouring the shorter one.

Our negative results can yield also positive value for decision makers. By means of some examples we show that other intuitive policies like scheduling based only on processing time, weight over processing time or weight times processing times can result in bad performance in certain cases. This serves as indication for companies to abandon such current practices or avoid them in their search for new designs.

The results of this chapter are applicable to companies that, for example, produce malleable goods, i. e., products with the property that any portion of them can be processed parallel. Further examples include also the execution of computational tasks on multiple processors or projects where consultant's time is the resource.

CURRICULUM VITAE

Gergely Csapó was born on May 17, 1983 in Veszprém, Hungary. In 2002, he passed the maturity exam with distinction at Lovassy László Gymnasium in Veszprém. Subsequently, he studied economics at Corvinus University of Budapest. As part of his higher education studies, he has completed the fall term of the academic year 2006/2007 at the University of Maastricht following operations research related courses. In 2008, he obtained his master's degree in economics with a major in actuarial science. His master thesis is titled "Constructing Life Table for Pension Funds".

During the last year of his studies, he joined ING Hungary and worked there as an actuary until 2010. In September 2010, he became a PhD student at the Quantitative Economics department of Maastricht University under the supervision of Prof. dr. Rudolf Müller. The results of his research are presented in this thesis. Gergely presented his work at various international conferences and parts of his thesis are published or under review in leading academic journals.

