MULTI

DIMENSIONAL

FLEXIBILITY

INTEGRATION

Foundational and Conceptual Challenges

Roland Saur

Multi-Dimensional Flexibility Integration

Foundational and Conceptual Challenges

Roland Saur

Multi-Dimensional Flexibility Integration: Foundational and Conceptual Challenges

Dissertation

for the purpose of obtaining the degree of doctor
at Delft University of Technology
by the authority of the Rector Magnificus,
Prof.dr.ir. H. Bijl,
chair of the Board for Doctorates
to be defended publicly on
Wednesday 14 October 2026 at 12.30

by

Roland Saur

This dissertation has been approved by the promotor.

Composition of the doctoral committee:

Rector Magnificus	chairperson
Prof.dr.ir. J.A. La Poutré	EWI-ESE-IEPG TU Delft, NL Centrum Wiskunde & Informatica, promotor
Dr. N. Yorke-Smith	EWI-ST-ALGOR TU Delft, NL, promotor

Independent members:

Prof.dr. F.M. Brazier	TBM-MAS-SK TU Delft, NL
Prof.dr.ir. G. Deconinck	KU Leuven, Belgium
Prof.dr. J.L. Hurink	U. of Twente, NL
Prof.dr. P. Palensky	EWI-ESE TU Delft, NL
Prof.dr. S.N. Singh	IIITM, India



The research in this dissertation was performed at the Intelligent & Autonomous Systems (IAS) group of Centrum Wiskunde & Informatica (CWI), the national research institute for mathematics and computer science in the Netherlands.



SIKS Dissertation Series No. 2026-51 - The research reported in this thesis has been carried out under the auspices of SIKS, the Dutch Research School for Information and Knowledge Systems.



This work is part of the research program Heat and Power Systems at Industrial Sites and Harbours (HaPSISH) with project number OND1363719, which is partly financed by the Dutch Research Council (NWO)

ISBN 978-94-6384-999-9

Copyright © 2026 Roland Saur

An electronic version of this dissertation is available at
http://repository.tudelft.nl/.

TU Delft

Summary

This work considers challenges around the idea of multi-energy flexibility as a cost effective facilitator of further renewable energy integration. Massive investments in renewable energy sources to combat global warming have put a strain on an ageing electrical infrastructure. While a further increase in renewable energy sources is still necessary, this cannot be achieved without an accompanying expansion of flexibility in the grid. Multi-energy flexibility has emerged in recent years as a potentially cheap source of flexibility: making use of different energy carriers gives such systems an additional dimension of flexibility.

More specifically this work considers both foundational and conceptual issues related to multi-energy systems and the integration of flexibility. These issues are addressed through four specific research questions: two foundational questions related to the practical operation of multi-energy systems and two broader questions exploring conceptual challenges related to flexibility.

Among the foundational questions, this work first explores the question of how to incentivize truthful behavior in settings with multiple energy carriers. For this, a two sided multi-energy unit-demand market is considered and it is shown that trade reduction can ensure a budget surplus as well as honest participation by removing the marginal trades from the allocation.

Tighter integration of multiple energy carriers also has ramifications for the bidding strategies of individual assets – addressed as the second research question. Access to several energy carriers allows agents participating in single carrier markets to settle some of their obligations outside of the market. We explore this topic by considering a flexible-in-time electricity consumer with access to a heating system. We show that the ability to shed excess electrical energy in a district heating network spreads the optimal bidding strategy over several time steps.

Besides foundational aspects of multi-energy flexibility, this work also takes a look at broader conceptual challenges regarding fairness and epistemic uncertainty in relation to multi-dimensional markets.

A sense of fairness is crucial for the adoption of new market mechanisms. Given that flexibility integration will require novel market designs, the question arises of how a sense of fairness can be established in new and not yet common market mechanisms. This work considers the question of what fairness means in the context of a two sided market mechanism with flexible-in-

time unit consumers and producers. We show that the addition of flexibility introduces some ambiguity to the more standard notion of non-discrimination. To address this ambiguity, we extend the concept of non-discrimination to a novel concept of consistent pricing over equivalent outcomes. Furthermore, we show that this new concept aligns with human notions of fairness from the psychology literature.

The second conceptual challenge relates to uncertainty. Since flexibility is a product designed to address uncertain future energy production, any market trading will require accurate predictions to function properly. While prediction models can describe some of the uncertainty, there is also uncertainty about the model itself. The difference between the uncertainty inherent in the system – the aleatoric uncertainty – and the uncertainty described by models is called the epistemic uncertainty. Ideally for a flexibility market, the epistemic uncertainty should be as small as possible, i.e. the prediction model is as close to the true uncertainty as possible. While providing accurate predictions is a well researched topic in isolated cases, for a market setting the task is further complicated by the competitive nature of markets: why would a fluctuating renewable energy source provide accurate predictions to a market? To address this question, we introduce the concept of best prediction incentivizing and provide solutions for classes of simplified distributions. The incentive works by tying the quality of a prediction to the payment. We then extend the usefulness of our solution by showing that our solution can be fine tuned to effectively approximate other – more common – distributions.

In conclusion, this work addresses both foundational as well as conceptual challenges related to multi-dimensional flexibility markets. While we provide some insights into the questions we pose, it would be remiss of us to not also mention the further question this work opens up. In particular the conceptual challenges surrounding fairness and epistemic uncertainty were only explored in simplified examples to establish a solid conceptual foundation. It would be interesting to see how these concepts unfold when applied in real world case studies or to explore the wider implications of epistemic uncertainty on mechanism design.

Samenvatting

Dit werk behandelt de uitdagingen rond het idee van Multienergie-flexibiliteit als een kosteneffectieve facilitator voor verdere integratie van hernieuwbare energie. Massale investeringen in hernieuwbare energiebronnen om de opwarming van de aarde tegen te gaan, hebben de verouderde elektrische infrastructuur onder druk gezet. Hoewel een verdere toename van hernieuwbare energiebronnen nog steeds noodzakelijk is, kan dit niet worden bereikt zonder een gelijktijdige uitbreiding van flexibiliteit in het net. Multienergie-flexibiliteit is de afgelopen jaren naar voren gekomen als een potentieel goedkope bron van flexibiliteit: door gebruik te maken van verschillende energiedragers krijgen dergelijke systemen een extra dimensie van flexibiliteit.

In het bijzonder behandelt dit werk zowel fundamentele als conceptuele vraagstukken gerelateerd aan Multienergie-systemen en de integratie van flexibiliteit. Deze kwesties worden aangepakt via vier specifieke onderzoeksvragen: twee fundamentele vragen die betrekking hebben op de praktische werking van Multienergie-systemen en twee bredere vragen die de conceptuele uitdagingen rondom flexibiliteit verkennen. Onder de fundamentele vragen onderzoekt dit werk eerst hoe billijk gedrag gestimuleerd kan worden in omgevingen met meerdere energiedragers. Hiervoor wordt een tweezijdige Multienergie unit-demand markt beschouwd en wordt aangetoond dat handelsreductie zowel een budgetoverschot als billijk deelname kan waarborgen door de marginale transacties uit de toewijzing te verwijderen. Een strakkere integratie van meerdere energiedragers heeft ook gevolgen voor de biedstrategieën van individuele activa – behandeld als de tweede onderzoeksvraag. Toegang tot verschillende energiedragers stelt agenten die deelnemen aan enkelvoudige drager markten in staat om een deel van hun verplichtingen buiten de markt af te wikkelen. We verkennen dit onderwerp door een elektriciteitsconsument die flexibel is in de tijd te beschouwen, die toegang heeft tot een verwarmingssysteem. We tonen aan dat de mogelijkheid om overtollige elektrische energie af te voeren in een wijkverwarmingsnetwerk de optimale biedstrategie over meerdere tijdstappen spreidt.

Naast de fundamentele aspecten van Multienergie-flexibiliteit behandelt dit werk ook bredere conceptuele uitdagingen met betrekking tot billijkheid en epistemische onzekerheid in multidimensionale markten.

Een gevoel van billijkheid is cruciaal voor de acceptatie van nieuwe marktmechanismen. Omdat de integratie van flexibiliteit nieuwe marktontwerpen

vereist, rijst de vraag hoe een gevoel van billijkheid tot stand kan worden gebracht in nieuwe en nog niet gangbare marktmechanismen. Dit werk onderzoekt wat billijkheid betekent in de context van een tweezijdig marktmechanisme met flexibele unit-consumenten en producenten. We laten zien dat de toevoeging van flexibiliteit enige ambiguïteit introduceert in de meer gangbare notie van non-discriminatie. Om deze ambiguïteit te adresseren, breiden we het concept van non-discriminatie uit tot een nieuw concept van consistente prijsstelling over equivalente uitkomsten. Voorts tonen we aan dat dit nieuwe concept in lijn is met menselijke opvattingen van billijkheid uit de psychologie.

De tweede conceptuele uitdaging betreft onzekerheid. Omdat flexibiliteit een product is dat is ontworpen om onzekere toekomstige energieproductie aan te pakken, vereist elke markttransactie nauwkeurige voorspellingen om goed te functioneren. Hoewel voorspellingsmodellen een deel van de onzekerheid kunnen beschrijven, bestaat er ook onzekerheid over het model zelf. Het verschil tussen de inherente onzekerheid in het systeem – de aleatorische onzekerheid – en de onzekerheid die door modellen wordt beschreven, wordt epistemische onzekerheid genoemd. Idealiter moet de epistemische onzekerheid voor een flexibiliteitsmarkt zo klein mogelijk zijn, d.w.z. dat het voorspellingsmodel zo dicht mogelijk bij de ware onzekerheid ligt. Hoewel het leveren van nauwkeurige voorspellingen een goed onderzocht onderwerp is in geïsoleerde gevallen, wordt de taak in een marktcontext verder gecompliceerd door de competitieve aard van markten: waarom zou een fluctuerende hernieuwbare energiebron nauwkeurige voorspellingen aan een markt verstrekken? Om deze vraag te beantwoorden introduceren we het concept van "best prediction incentivizing" en bieden oplossingen voor klassen van vereenvoudigde distributies. De prikkel werkt door de kwaliteit van een voorspelling te koppelen aan de betaling. Vervolgens breiden we de bruikbaarheid van onze oplossing uit door aan te tonen dat deze kan worden afgestemd om nauwkeuriger andere – meer voorkomende – distributies te benaderen.

Kortom, dit werk behandelt zowel fundamentele als conceptuele uitdagingen gerelateerd aan multidimensionale flexibiliteitsmarkten. Hoewel we enkele inzichten bieden in de gestelde vragen, zou het nalatig zijn om de verdere vragen die dit werk oproept niet te vermelden. Met name de conceptuele uitdagingen rondom billijkheid en epistemische onzekerheid werden alleen in vereenvoudigde voorbeelden onderzocht om een solide conceptueel fundament te leggen. Het zou interessant zijn om te zien hoe deze concepten zich ontwikkelen wanneer ze worden toegepast in praktijkgevallen, of om de bredere implicaties van epistemische onzekerheid op mechanisme-ontwerp te onderzoeken.

Acknowledgements

First and foremost I would like to thank my mother, the internet and BAFOG in this order. To my mother: without you I would have never ventured out. Thanks for inspiring me to be open to a wider world. I am grateful to the internet for giving me access to the basic information that makes higher education possible: where can I study? what programs are available? how much does it cost? To BAFOG, and by extension the German welfare state, I am grateful for providing the financial means for me to go to university. While, these things seem very basic, I think we should not take a loving mother, free internet and functioning welfare state for granted.

Beyond the basics, I am also thankful for the guidance provided by my thesis advisors Neil Yorke-Smith and Han La Poutré. Not only have your wise words helped me figure out things for myself, your compassion has also helped me deal with the many rejection letters and revisions along the way.

With regard to the thesis, I would also like to thank the committee for their insightful and detailed feedback. With your help this thesis has become a more polished and well rounded piece of work.

I would also like to thank everybody from the Intelligent Autonomous Systems group. Special thanks go to Brinn Hekkelman and Georgios Methenitis for being amazing office mates. You have helped me think more deeply about my research, indulged in silly ideas and made dealing with deadlines and rejection letters a bit lighter. But most of all you have made the office fun.

To Angelina, thank you for all the love and support you have given me. I wish you all the best for the future.

To Prof. Parspours, thank you for writing my letter of recommendation on such short notice. This small token of understanding has ended up enabling a lot in my life.

To Marcel, you are probably the person who I have stayed in contact with most consistently over the past 10 years. I have told you some of the strangest things on my mind. In general, I am grateful that you have stuck around and am happy to have you in my life.

Special thanks to everybody at Albelli. Thursday drinks were amazing, more than just being open and compassionate, you are genuinely interesting people. You were a welcome reprieve from the thesis writing drudgery. I am grateful for having met you all.

However, this Thesis has taken much longer than it should have. For finishing it at all, I would like to thank Hazem and Vera.

To Hazem, thank you for being a great friend, the interesting conversations we have had over the years and for all your help along the way. In the end, you have kept me on track to finishing this thesis.

To my dear friend Vera. Thank you. You have brought so much joy into my life and helped me become a better person. Without you this thesis would have never been finished.

An meine Brüder, die weitere Familie und Familie Natzer. Ich bin dankbar dass ich euch alle habe. Durch Studium, Uni, und Arbeit bin ich irgendwie in Amsterdam gelandet, aber jedesmal wenn ich euch in der alten Heimat besucht habe, fühlte ich mich willkommen. Danke.

Vale, I love you. You are amazing fun. You make the good times feel brighter and the dark times more navigable. You make me a better person in so many ways. But ultimately words fail to describe how I feel. Instead, I hope I get to show you for many years to come.

Contents

Summary v

Samenvatting vii

Acknowledgements ix

1 Introduction 1

1.1 Multi-Dimensional Flexibility . . . 3

1.1.1 Forms of Multi-Dimensional Flexibility . . . 3

1.1.2 Trends in Energy Integration . . . 4

1.2 Markets and Mechanism Design . . . 5

1.2.1 What are Markets ? . . . 5

1.2.2 What is Mechanism Design ? . . . 7

1.2.3 Important Concepts in Mechanism Design . . . 7

1.2.4 Limitations in Mechanism Design . . . 8

1.3 Related Concepts . . . 10

1.3.1 Fairness from Psychology to Mathematics . . . 10

1.3.2 Uncertainty: Aleatoric and Epistemic . . . 10

1.4 Research Questions . . . 11

1.4.1 Multi-Energy Systems Integration: Incentives and Strategies . . . 12

1.4.2 Conceptual Challenges regarding Flexibility: Fairness and Uncertainty . 13

1.5 Thesis Outline . . . 14

1.6 List of Publications . . . 16

2 Combined Heat and Power Markets by Double-Sided Auction Mechanisms 17

2.1 Introduction . . . 17

2.2 The Model . . . 19

2.2.1 Properties Relevant to Mechanism Design . . . 20

2.3 Trade Reduction Mechanism . . . 21

2.3.1 General Properties . . . 22

2.3.2 Mechanism design relevant properties . . . 27
2.3.3 Social Welfare Loss . . . 29
2.4 Conclusion and Future Work . . . 30

3 Optimal Demand-Side Bidding in Parallel Auctions for Time-Shiftable Electrical Loads 31
3.1 Introduction . . . 31
3.2 Related Work . . . 33
3.3 Model . . . 33
3.4 Optimal Bidding Strategies . . . 35
3.4.1 Comparison to the Literature . . . 36
3.4.2 Identical Auctions . . . 36
3.4.3 Non-Identical Auctions . . . 39
3.5 Conclusion and Future Work . . . 43

4 Fair Pricing for Time-Flexible Smart Energy Markets 45
4.1 Introduction . . . 45
4.2 Market Setting . . . 48
4.3 Desired Properties of Pricing . . . 49
4.4 Price Discrimination in Time . . . 50
4.4.1 Consistent Pricing over Equivalent Outcomes . . . 50
4.4.2 Obtaining Consistent Pricing . . . 55
4.4.3 Consistent and Individually Rational Prices . . . 61
4.5 Pricing Mechanism . . . 62
4.6 Discussion: consistent pricing, envy-freeness and fairness . . . 63
4.6.1 Fairness: envy-free or consistent pricing . . . 63
4.6.2 Comparison to relaxations of envy-freeness . . . 64
4.7 Conclusion and Future Work . . . 65

5 Incentives for Accurate Energy Predictions: How to Reduce Epistemic Uncertainty 67
5.1 Introduction . . . 67
5.2 Problem Formulation . . . 70
5.2.1 Best Prediction Incentivizing . . . 71
5.2.2 Best Prediction Incentizing under epistemic uncertainty . . . 71
5.2.3 Comparison to Incentive Compatibility . . . 73
5.3 Applications of BPI contracts . . . 74
5.3.1 Solutions for uniform distribution . . . 74
5.3.2 Solutions for single term polynomial distributions . . . 75
5.3.3 Real World Applications . . . 77
5.4 Arbitrage: Reselling improved predictions . . . 88
5.4.1 Arbitrage Example: Middleman providing an improved Prediction . . . 89

5.5 Discussion 90
5.5.1 Comparison to Data Markets 91
5.5.2 Comparison to Prediction Markets 91
5.6 Conclusion 92

6 Conclusion **93**
6.1 Addressing Research Questions 93
6.1.1 Research Question 1 93
6.1.2 Research Question 2 94
6.1.3 Research Question 3 95
6.1.4 Research Question 4 96
6.1.5 Main Research Question 97
6.2 Further Research Directions 98
6.3 Concluding Remarks 99

A Appendix **101**
A.1 Regarding Chapter 3 101
A.1.1 Proofs 101
A.1.2 Solving $B_{-l,m}$ for a Pair of Bids in Lemma 3.3 107
A.1.3 Second Derivatives in Lemma 3.4 108
A.1.4 Considering case s=1 and |T|=2 109

Bibliography **113**

Curriculum Vitae **141**

1 Introduction

Global warming is one of the defining challenges facing humanity in the 21st century. The large amount of greenhouse gases, most notably CO_2 and methane, which have been put into the atmosphere, have already significantly warmed the planet since the beginning of the industrial revolution (Lee et al., 2023). The consequences of this warming, such as sea-level rise, ocean acidification and desertification, have had a wide range of negative impacts on global food supplies as well as the dying of species (Lee et al., 2023). The Intergovernmental Panel on Climate Change (IPCC) sets out a goal of total warming of at most $2°C$ by the end of the 21st century. This will require a massive and rapid reduction in greenhouse gas emissions.

As part of the required reduction of greenhouse gas emissions, electricity systems around the world will have to undergo substantial changes. With 13.6 Gt of CO_2 in 2023, electricity production is a major contributor to global climate change (IEA, 2024b). Therefore, decarbonization of this sector, which has historically been mostly coal and gas powered, is seen as a significant component in achieving the IPCC's $2°C$ goal.

Carbon-neutral renewable energies have slowly been replacing fossil fuels within the electricity sector. This development has not only been enabled by advancements in renewable energy technology but also through government-subsidies (Nicolini and Tavoni, 2017). As a result of these developments, so called variable, i.e. uncontrolled, renewables made up 15% of global electricity generation in 2024 and are projected to contribute 28% by 2030 (IEA, 2025b). With regards to europe, the IEA projects that the share of variable renewables will rise from 25% in 2024 to 40% in 2030 (IEA, 2025b). Going forward, most of the reduction in CO_2 emissions from the electricity system will be coming from variable energy source, as can be seen in Figure 1.1.

While the rapid increase in carbon free electricity production has been vital for the reduction of emissions, the scale of renewable energy integration poses a significant challenge to electricity systems. The weather-dependent nature of most renewable energy sources (RES) have resulted in uncontrolled spikes of electricity production. These spikes have led to strong price fluctuations on electricity markets (Auer and Haas, 2016; Pavlík et al., 2025) as well

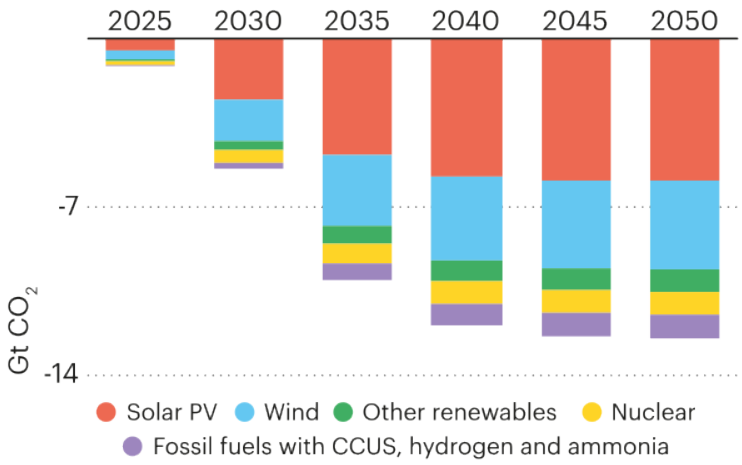


Figure 1.1: Global Power Sector CO_2 Emissions Reductions by 2050 (CC BY 4.0) (IEA, 2023b) (CCUS = Carbon Capture, Utilisation and Storage)

as network congestion sometimes going even so far as to cause negative wholesale electricity prices (Pavlík et al., 2025; Seel et al., 2021). However, Schöniger and Morawetz (2022) have pointed out that large RES integration does not lead to high price fluctuations if the grid contains a sufficient amount of flexibility, while Navia Simon and Diaz Anadon (2025) have argued that increased RES adoption reduces the sensitivity of electricity prices to changes in the price of natural gas.

Flexibility in the grid is broadly defined as the ability of a device to quickly adjust its level of electricity production or consumption. Such a device can be anything from dedicated generators over fleets of electric vehicles down to individual consumers.

However, the conventional approach of providing additional generator capacity to deal with power fluctuations will not be financially viable in the future. In the past, flexibility was provided in the form of spare capacity by conventional generators. To provide more flexibility, one had to make more spare capacity available, which come at a high investment cost. This means that the cost increases linearly with the need for flexibility, making further RES integration rather expensive.

To keep the cost of providing flexibility from growing, there has been a focus on making more intelligent use of existing infrastructure. This comes in many different forms: Demand Response (Jin et al., 2025; Meir et al.,

2017; Palensky and Dietrich, 2011; Wang et al., 2017), Time-shiftable loads (Graditi et al., 2015) as well as coupling of different sectors and energy carriers (Jin et al., 2025; Li et al., 2025, 2016; Luc et al., 2020; Wang et al., 2017), all integrate different dimensions of the overall energy system to provide flexibility. Rather than adding infrastructure to provide extra flexibility, the goal is to make use of the interconnected nature of modern energy systems.

This thesis aims to contribute to a better understanding of how this integration of different dimensions can be accomplished in existing and future energy markets.

1.1 Multi-Dimensional Flexibility

So far we provided a high-level understanding of the challenges posed by the renewable energy transformation and its need for cheap flexibility. Equipped with this overview, we will introduce the term *multi-dimensional flexibility* as a distinctive concept from more traditional forms of flexibility and show how it connects to current trends in energy system integration.

We will refer to multi-dimensional flexibility as the flexibility that arises from the interaction of individual uses of energy. By coupling two – hopefully independent – uses of energy, we obtain a fluctuation that is more easily managed than when considered separately.

1.1.1 Forms of Multi-Dimensional Flexibility

A generator can provide flexibility for a particular point in time by its ability to vary its energy output on short notice. Since providing flexibility to one point in time is functionally independent of flexibility in another point in time, we will refer to this kind of flexibility as one-dimensional flexibility.

Multi-dimensional flexibility, on the other hand, is about considering flexibility across multiple dimensions of energy use. We will consider two broad kinds of additional dimensions: *time* and *type* of energy carrier.

When considering time, electricity consumption at 3pm , 4pm and 5pm can be considered as one multi-dimensional entity. The clear benefit here is that an excess of production at 3pm could compensate for a lack of production at 5pm.

Another type of multi-dimensional flexibility is to combine several different energy carriers. For instance in residential housing blocks, the heat consumption and the electricity consumption can be viewed as one combined problem. Similarly to the time integration, excess demand in heat might be compensated by a lack of demand in electricity, resulting in an overall more predictable system.

While both kinds of multi-dimensional flexibility require certain infrastructure, we will see that this infrastructure often already exists.

1.1.2 Trends in Energy Integration

With this conceptual view of multi-dimensional flexibility, we will have a look at several trends in the field of energy integration that enable this type of flexbility.

Probably the most impactful trend in the energy sector relevant to flexibility is the tighter integration between electrical, gas and heating systems. Both gas and heating networks exhibit much higher inertia, making them ideal storage for excess energy from a much more dynamic electrical system. While electrical systems react very quickly to any kind of change in production or consumption mandating a strict power balance, heating and gas networks can withstand large changes in production/consumption without much immediate effect on their operation. Not only can these networks withstand sudden changes, they are also storing relatively large quantities of energy in already existing infrastructure.

The advent of power-to-gas technology has enabled bi-directional links between electrical and gas networks (Glenk and Reichelstein, 2022). Li et al. (2025) have considered the combined flexibility from district heating and hydrogen-enriched compressed natural gas networks. While the technology to turn electricity into hydrogen is not yet cost competitive (IEA, 2025a), the bridge from electrical to heating networks is much more established. Heat pumps are a mature technology and a range of thermal applications have already been considered as viable storage for uncontrolled renewable energies: from residential buildings (Foteinaki et al., 2018; Jin et al., 2025; Luc et al., 2020; Xu et al., 2020) over office buildings (Jin et al., 2025; Lu et al., 2022; Papachristou et al., 2021) all the way to greenhouses for growing crops (Roncancio et al., 2022).

While the flexibility stemming from combining different energy carriers is being enabled by adding connections between well established energy systems, the inclusion of different time steps of electricity is made possible by two developments: *demand response* and *storage*.

Generally speaking, demand response is the act of asking or incentivizng a consumer to change their consumption (Palensky and Dietrich, 2011) – often on short notice. This shift of consumption allows a system operator to react to fluctuations by simply asking a consumer to change when they consume electricity, without expensive infrastructure investment. While demand response is already well established for medium to large scale actors, recent development have shown that by combining many small consumers, an aggregator can provide demand response on a grid-relevant scale (Correa-Florez et al., 2019; González et al., 2020). On the individual household level, the electrification of more aspects of residential living is adding additional flexibility. In the past few decades, many home owners in Europe have added photovoltaic panels

to their rooftops (IEA, 2022), while cheaper electrical storage has made many consider installing batteries. Heat pumps are replacing gas or wood powered heating (Abbasi et al., 2021) and electric vehicles are becoming more commonplace (IEA, 2024a). Heat pumps provide storage by utilizing the thermal mass of buildings (Jin et al., 2025; Mueller et al., 2014) while electric vehicles can provide flexibility while plugged into the grid (Kumar et al., 2025; Sepetanc and Pandzic, 2021). These developments, in the context of aggregation, will make these homes more valuable to the grid as a potential source of flexibility.

Storage can also be viewed as providing a form of multi-dimensional flexibility by connecting energy usage at different time steps. Similar to how power-to-gas technology integrates different energy carriers, storage integrates energy usage across time. Advancements in battery technology over the past two decades have made grid-scale electric storage a viable option (Chen et al., 2020; Huang and Li, 2022; Jiang et al., 2025). For a current review of grid-scale battery technologies, we refer the reader to (Jiang et al., 2025).

As we can see, there is quite a potential for flexibility that emerges from the combination of already existing infrastructure.

1.2 Markets and Mechanism Design

So far we have discussed the need for and examples of energy flexibility. However, to integrate flexibility into modern energy systems, we need to be able to allocate flexibility optimaly, which requires an understanding of markets. Therefore, in this section, we will define markets as well as mechanism design, its most important concepts and limitations.

1.2.1 What are Markets ?

Markets are a setting in which two or more entities freely engage in a mutually beneficial trade. The freely part relates to the fact that neither of them are forced to buy or sell an item, if they do not want to, while mutually beneficial indicates that both get something out of the transaction.

A trade can be as simple as the exchange of a vase for €30. Curiously, both the buyer and seller will walk away from the exchange feeling somewhat superior. While the seller might say "I would have gone as low as €10", the buyer thinks the vase to be worth at least €50. Both buyer and seller conclude that they got the better deal, simply because they value the item differently.

But where does this discrepancy in value-perception come from? It stems from a difference in circumstance that the buyer and seller find themselves in. The seller might be a skilled craftsman with a good understanding of the vase making process allowing them to make vases at a low price. The buyer, on

the other hand, while lacking the technical expertise of the seller might have a particular use or possess knowledge that such vases sell for at least €50 in a nearby town. Ultimately, it is access to private skills and knowledge that makes buyer and seller value the vase differently.

But then what is the vase actually worth? To answer this question, we have to look at what happens when there is not just one trade happening, but many buyers and sellers, all with different kinds of private information. When there are more sellers the price of the vase will naturally go down as potential buyers can simply choose to trade with whoever sells at the lowest price. Increased competition between buyers will drive the price up as sellers can charge a higher price and still sell their wares. The price that most vases will be sold at is a result of these interactions. This means that there is no pre-determined value of the vase but that the value is an emergent property and its trading price an approximation thereof.

It has to be noted that the effects of this kind of trading go beyond value approximation through price discovery: It can guide economic behaviour on a large scale. The movements of prices can quickly and easily disseminate abstracted information and thereby guide the allocation and investment of both financial and human capital. Imagine for instance that the furnaces used for making the vases are coal powered and that due to a far away natural catastrophe the supply of cheap coal is severely restricted. This shortage causes the price of coal to increase significantly. Given the increase in price, a vase maker might shift some of their investment away from their marketing budget and towards the purchase or development of more energy efficient furnaces.[1] In the end, an event affecting coal production in one part of the world has led a vase maker in another part of the world to investigate energy efficient furnaces. In this sense, changes in price are signals that indicate whether a resource has become easier or harder to obtain. Without any need for centralized control, price signals are able to coordinate the economic behaviour of people across large distances. This effect is often referred to as the "invisible hand" of the market.

Looking back at our initial issue of renewable energy and flexibility integration, we hope it has become clear why we would like to leverage the "invisible hand of the market". Instead of trying to centrally plan a solution considering all aspects of a future energy system, markets can guide expertise and investment to a most cost efficient integration of fluctuating renewables.

[1] Alternatively, the vase maker might increase their marketing budget in the hope of convincing the buying public that more expensive vases are a more desirable good.

1.2.2 What is Mechanism Design ?

While markets refer to any kind of trading between individuals, a *mechanism* refers specifically to the rules that govern the trading. By designing a mechanism, i.e. the rules of trading, we can steer the market towards a desirable outcome. What is usually referred to as the invisible hand of the market takes on a more concrete shape through the study of mechanism design.

A mechanism is a way of aggregating individual preferences/valuations into a collective action. Given a set of agents N and a set of outcomes A, a mechanism aggregates valuations of the agents into a particular outcome $a \in A$ and prices p. The valuation of an agent $i \in N$ is expressed as a function from the space of outcomes A to the real numbers, i.e. $v_i : A \rightarrow \mathcal{R}$. V_i is the space of all such functions. The agent reports a particular valuation v_i to the mechanism. The mechanism consists of a social choice function and pricing functions. The social choice function f maps the combined valuation profile of the agents, $(V_1 \times V_2 \times ...V_N)$, to an outcome $a \in A$. Within the energy domain this collective action is often an allocation of resources, i.e. who gets what amount of energy. The pricing functions p_i maps the agents' valuation profile $(V_1 \times V_2 \times ...V_N)$ to the price paid by agent i.

Definition 1.1 (Direct Revelation Mechanism (Nisan et al., 2007))**.** A direct revelation mechanism consists of a social choice function $f : V_1 \times V_2 \times ...V_N \rightarrow A$ and a set of pricing functions $p_1, p_2, ..., p_N$.

The study of mechanism design is about coming up with social choice and pricing functions that result in a certain desirable outcome.

1.2.3 Important Concepts in Mechanism Design

With the use of mechanism design, the previously mostly unstructured interactions between individuals can be designed to fulfil certain properties. Most desirable properties in mechanism design relate to either the behaviour of the individual or the collective outcome.

An often stated purpose of a market is to maximize Social Welfare (SW), i.e. the overall "happiness". Social Welfare is the sum of valuations of the agents under a specific choice or outcome a. We refer to a mechanism as optimal if the outcome it provides maximizes social welfare.

Definition 1.2 (Social Welfare)**.** The social welfare of a choice a is the sum of the agents valuations:

$$SW(a) = \sum_{i \in N} v_i(a)$$

However, the goal of maximizing SW can only be achieved if the agents participate in the market honestly. If the agents stay away from the market,

the mechanism only optimizes for a subset of the relevant agents, while if the agents do participate but report their valuations dishonestly, the mechanism will optimize for an incorrect valuation profile not capturing the actual SW. The notions of participation and honesty are defined in mechanism design as Individual Rationality (IR) and Incentive Compatibility (IC).

To properly define these two terms, we need to introduce some additional notation: we denote the true valuation of agent i as v_i^* and the valuations reported by all agents except agent i as v_{-i}. We, further, denote the resulting action of the mechanism when agent i reports v_i as $a(v_{-i}, v_i)$ and the resulting price for agent i as $p_i(v_{-i}, v_i)$.

Individual Rationality (IR) states that no agent participating honestly will be left with a negative utility.

Definition 1.3 (Individual Rationality (Nisan et al., 2007))**.** A mechanism is Individually Rational if

$$v_i^*(a(v_{-i}, v_i^*)) - p_i((v_{-i}, v_i^*)) \geq 0 \qquad \forall v_{-i}$$

Incentive Compatibility (IC) states that the best strategy of an agent is to report their valuation honestly.

Definition 1.4 (Incentive Compatibility (Nisan et al., 2007))**.** A mechanism is Incentive Compatibility if

$$v_i^*(a(v_{-i}, v_i^*)) - p_i(v_{-i}, v_i^*) \geq v_i^*(a(v_{-i}, v_i)) - p_i(v_{-i}, v_i) \qquad \forall v_i, v_{-i}$$

Designing a mechanism that optimizes for the overall welfare of a system is a noble goal. However, this goal can only be realized if the market that implements it is feasible. In order for a market to be feasible it has to also be financially sustainable.

Financial sustainability can be expressed as budget balance – the payments made by and to the market cancel each other out.

Definition 1.5 (Budget Balance)**.** A mechanism is budget balanced if

$$\sum_{i \in N} p_i(v_{-i}, v_i) = 0 \qquad \forall v_{-i}, v_i$$

It should be noted that the weak budget balance version of this property allows the market to run a surplus (Dash et al., 2003), which might be desirable to cover operational costs.

1.2.4 Limitations in Mechanism Design

While mechanism design can be used to shape markets, there are also both theoretical as well as practical limitations to it.

One of the major theoretical limitations is the impossibility result regarding double sided auctions. To elaborate on this, we will first introduce the well known Vickrey–Clarke–Groves (VCG) mechanism.

Definition 1.6 (Vickrey–Clarke–Groves Mechanism (Nisan et al., 2007))**.** The Vickrey–Clarke–Groves mechanism picks the social welfare maximizing option:

$$a^* = argmax_a \left(SW(a)\right)$$

and agent i is charged an amount $h_i(v_i)$, which is independent of its reported valuation i, while being paid the sum of reported valuations of the other agents in the market, i.e.:

$$p_i = h_i(v_i) - \sum_{j \in N \setminus \{i\}} v_j(a^*)$$

The quantity h_i is often set to the sum of valuations that all other agents can obtain if i did not participate, i.e.:

$$h_i(v_i) = max_{a \in A} \sum_{j \in N \setminus \{i\}} v_j(a)$$

This way of setting h_i is referred to as the Clarke pivot rule.

In plain language, the VCG mechanism finds the optimal allocation while setting prices such that they reflect the positive or negative effect an agent has on the other agents in the market. This way of pricing does not only have a certain socially just ring to it but also makes the VCG mechanism Individually Rational as well as Incentive Compatible. However, it has been shown that the VCG mechanism always runs a budget deficit. Worse than that, it has been shown that no mechanism can be simultaneously social welfare optimizing, IR, IC as well as weakly budget balanced (Myerson and Satterthwaite, 1983). This means that it is impossible to design a double sided auction mechanism that has all the desired properties we outlined in the previous subsection.

In practical terms, we have the issue of actually computing the solutions, since we do not have infinite computational power. More specifically, with regard to the computability of the market, we require that the social choice and pricing functions are not only solvable but also exhibit low computational complexity. As the percentage of RES increases so will the demand for flexibility. This increase in demand will naturally drive more participants towards well functioning flexibility markets. The design of such markets has to take this potential increase in participation into account. A mechanism, therefore, has to scale well in computational terms. For flexibility markets to continue to

function in the long run, we would like the underlying mechanism to scale in polynomial time with the number of participants.

1.3 Related Concepts

So far we have surveyed different types of multi-dimensional flexibility, desired properties of energy markets as well as general notions and limitations of mechanism design. However, when it comes to real-world applications, concepts outside of the immediate technical domain also become relevant.

In the following we will discuss two such related concepts: fairness and epistemic uncertainty.

1.3.1 Fairness from Psychology to Mathematics

Fairness, the feeling that one is neither being disadvantaged nor favoured by the world, is a hard to pin down yet important concept. What is or is not fair often lies in the eye of the beholder. Yet, a sense that one is being treated unfairly can easily lead someone to avoid a system. Given this subjectivity and importance, we need to consider both how humans perceive fairness as well as how fairness can be described mathematically.

Psychology provides us with conceptual lenses through which to understand how humans perceive fairness. While the psychology literature on fairness covers a wide spectrum, the relevant parts for our discussion are two contrasting concepts on how people generally derive a sense of fairness: the *distributive justice* approach states that people judge the fairness of a situation by comparing their personal outcome with that of people around them (Chapuis, 2012; Maxwell, 2002). In contrast, *procedural justices* states that humans can gain a perception of fairness of the outcome by judging the underlying mechanism itself (Chapuis, 2012; Maxwell, 2002).

However, insights into fairness perception are only useful to the design of markets if they can be expressed mathematically. In mechanism design, fairness is often codified as envy-freeness. Envy-freeness closely resembles the concept of distributive justice in psychology: a mechanism is considered envy-free if after the allocation and prices have been announced, no agent would prefer the assignment of another agent (Nisan et al., 2007). An alternative fairness property in mechanism design is non-discriminatory or anonymous pricing: the price that an agent has to pay does not depend on the identity of the agent (Conen and Sandholm, 2004; Sandholm and Suri, 2001).

1.3.2 Uncertainty: Aleatoric and Epistemic

In order to design markets that integrate flexibility, we need to have a proper understanding of uncertainty. Uncertainty, and the fluctuations that come with it, is what makes integrating renewable energies so difficult, while flexibility is

what allows us to react to unexpected fluctuations. However, using flexibility in response to fluctuations requires that we have a good understanding of the fluctuation.

A good understanding of future power fluctuations encompasses more than mathematical models. Through the use of probability distributions it is possible to describe power fluctuations and such mathematical models have in fact been used extensively to capture the uncertainty of renewable energy production. However, since none of us is omniscient, there is and always will be a difference between the true uncertainty, which philosophy refers to as aleatoric uncertainty (Der Kiureghian and Ditlevsen, 2009; Hüllermeier and Waegeman, 2021), and the uncertainty that is described by a model. This difference is referred to as the *epistemic uncertainty*: the uncertainty that exists about the mathematical model describing the uncertainty (Der Kiureghian and Ditlevsen, 2009; Hüllermeier and Waegeman, 2021).

Ultimately, to properly address the uncertainty that comes with most renewable energy sources, we need to be aware that there are limits to any model and that there will always exist a gap between modelled and aleatoric uncertainty. In short, we need to be aware that epistemic uncertainty exists.

1.4 Research Questions

So far we have outlined a potential solution to the challenges facing a clean energy transition: What we have seen is that multi-dimensional energy systems can provide additional flexibility via existing infrastructure (Section 1.1) while markets through price discovery can guide investments towards more cost efficient solutions (Section 1.2.1). This left us with a technical solution for providing cheap flexibility – multi-dimensional energy systems – and a process for integrating them efficiently – energy markets.

However, we have also seen that there are technical challenges and limitations to the design of markets itself (Section 1.2.2–1.2.4) and that the introduction of flexibility and uncertainty adds further conceptual challenges to the design process (Section 1.3).

These technical limitations and broader conceptual challenges lead us to the question of how multi-dimensional flexibility can generally be integrated into modern energy markets.

Main Research Question
What are ways to integrate multi-dimensional flexibility into current and future energy markets.

To address this overall research question we consider both general questions of multi-energy integration (Section 1.4.1) as well as conceptual challenges

of multi-dimensional flexibility (Section 1.4.2). With regard to integration we consider the problem of how to incentivize multi-energy agents to report honestly (Research Question 1) and explore the question of how access to multi-dimensional flexibility affects the bidding strategies of flexible-in-time energy consumers in conventional electricity markets (Research Question 2). Given the importance of markets to energy flexibility, we explore the conceptual challenges of fairness (Research Question 3) and epistemic uncertainty (Research Question 4) in energy flexibility itself.

1.4.1 Multi-Energy Systems Integration: Incentives and Strategies

We consider multi-energy integration both on a market and an individual level. For the market level, we address the question of incentives in multi-energy markets (Research Question 1), while on an individual level, we explore the use of multi-energy systems specifically for flexibility-providing assets. (Research Question 2).

Research Question 1
How can we incentivize honesty in multi-energy markets?

As mentioned earlier, incentivizing market-participants to be honest is vital in aligning the optimization problem that the market operator solves with the social goal that the market is supposed to address. However, as noted in Section 1.2.4, Incentive Compatibility cannot be achieved at the same time as optimality, Individual Rationality and Weak Budget Balance. While optimality is obviously a desired goal, from a market operator perspective Weak Budget Balance is important for sustaining the operation itself, while breaking Individual Rationality might cause agents to not participate in the market at all to avoid a negative return.

We explore this tension by providing a mechanism for a multi-energy market.

Research Question 2
How can access to a district heating network be used for energy procurement in current electricity markets?

For the foreseeable future, traditional energy markets will play a significant role in the integration of flexibility. While tailor-made flexibility markets would give flexible assets the best opportunity for selling their service, such markets have not yet been widely adopted. In the absence of dedicated markets, flexible assets have to make use of their abilities on traditional energy markets.

Simply having access to a district heating network allows for a wider range of bidding strategies in conventional energy markets. With only access to the electricity market, an agent is required to procure exactly the amount it consumes. With additional access to a district heating network, an agent

procuring above their needs is able to shed excess electricity in the form of heat.

How this access to a district heating network affects bidding strategies is the question we explore by providing procurement strategies. As a setting we consider a flexible-in-time consumer bidding on a day-ahead electricity market, while shedding excess electricity in a district heating network.

1.4.2 Conceptual Challenges regarding Flexibility: Fairness and Uncertainty

The previous research question considered a single agent utilizing multi-energy flexibility in conventional electricity markets. However, such a setup limits the flexibility benefit to one specific agent. Therefore, we also consider the trading of flexibility in the context of multi-agent systems. More specifically we explore some of the challenges that come with flexibility trading due to its peculiar nature.

Conceptually, flexibility differs from conventional energy both in what it does and what it is used for. Flexibility is often vaguely defined as the ability to adjust or shift energy output in response to an uncertain future. The two key points here are the adjustability of the output and the uncertainty of the future. The adjustability causes fairness to become a difficult to define concept (Research Question 3) while the uncertainty raises the question of how to obtain good predictions about future events (Research Question 4).

Research Question 3

What is fairness in flexible energy systems and how can it be achieved?

A sense of fairness is a pivotal contributor to acceptance and adoption of novel energy markets. Given that flexibility markets have been a more recent development and are not yet well established in current energy systems, it is important that we understand how fairness plays out in such markets.

Fairness in flexible energy systems is complicated because fairness is a vaguely defined human concept and flexibility introduces ambiguity about market outcomes. Given its vague nature, there exist different definitions in both the mechanism design and psychology literature. While there are some established approaches to fairness in the mechanism design literature such as envy-freeness and non-discrimination, these approaches become less clearly defined in the context of flexibility. Flexibility, the ability to adjust to changing circumstances, creates ambiguity about the optimal outcome by allowing for distinct outcomes to be the same in essence. For instance, if a market participant requires one unit of energy but is flexible about the exact time, then providing the energy at 3pm or 4pm are distinct outcomes that are essentially the same. However, these distinct outcomes might come with different prices giving rise to potentially unfair pricing.

We explore this very human concept by providing both conceptual and concrete solutions in the context of a flexible-in-time bilateral energy market.

Research Question 4
How to incentivize accurate predictions in flexibility markets?

Flexibility markets address the problem of uncertain fluctuating energy production. Flexibility in energy systems is broadly defined as the ability of an actor/device to quickly adjust its power consumption/production in response to power fluctuations. Unlike conventional energy production, the value of flexibility is not as much in the energy itself but rather in the ability to react to an uncertain future. For instance, when there is an unexpected drop in energy consumption, a flexible device adds value by being able to increase its consumption on short notice.

Providing the correct amount of flexibility requires accurate predictions. Given flexibility as the solution to unexpected energy fluctuations, the next question becomes "How much flexibility is needed?". To answer this question, we need to quantify the uncertainty of future energy production and consumption. This is usually done in the form of energy forecasts/predictions.

However, in flexibility markets obtaining accurate predictions becomes a question of incentives. In centralized systems, the system operator has access to all relevant information and forecasting becomes a purely technical issue. In a market system, the relevant information is provided by self-interested participants. Given the self-interest of agents, there needs to be a clear benefit for providing accurate predictions. This leads to the question of how to incentivize agents to provide accurate predictions. We explore this question by providing both conceptual and concrete solutions in the context of a peer-to-peer flexibility market.

1.5 Thesis Outline

Chapter 1 provided background information on the necessity and challenge of integrating flexibility into modern energy systems. After an overview of the relevant scientific fields and related concepts, four research question addressing both the foundational as well as broader conceptual challenges have been proposed.

These research question will be addressed in subsequent **Chapters 2–5**.

Chapter 2 addresses research question 1 in the context of a two sided auction setting. The main result of this chapter is that trade reduction mechanisms can provide a budget balanced, IC and IR mechanism for multi-energy systems while being close in terms of maximum social welfare.

Chapter 3 addresses research question 2 in the context of a flexible-in-time energy consumer with access to a district heating network participating on day ahead electricity markets. The main result of this chapter is that the ability to shed excess energy enables assets to make better use of their flexibility by allowing for more spread out bidding strategies.

Chapter 4 addresses research question 3 in the context of a two sided market with flexible-in-time unit-consumers and producers. The main result of this chapter is that flexibility in time complicates conventional fairness concepts such as non-discrimination. To address this complication, the concept of fairness by pricing consistency over equivalent outcomes is introduced.

Chapter 5 addresses research question 4 in the context of a peer-to-peer flexibility market. The main result of this chapter is that epistemic uncertainty complicates the idea of incentive compatibility as agents do not have perfect knowledge of the future. To address this complication, the concept of Best Prediction Incentivizing pricing functions is introduced.

In **Chapter 6**, answers to the research questions, which were provided in Chapter 1 and considered in detail in Chapters 2-5, will be discussed. From this discussion, conclusions will be drawn regarding the present state of the art and future research directions will be proposed.

1.6 List of Publications

- Roland Saur, Neil Yorke-Smith, and Han La Poutré. Combined heat and power markets by double-sided auction mechanisms. In *Proceedings of 2019 IEEE PES Innovative Smart Grid Technologies Europe, ISGT-Europe 2019*, pages 1–5, September 2019. doi: 10.1109/ISGTEurope.2019.8905714 **(Chapter 2)**
- Roland Saur, Han La Poutré, and Neil Yorke-Smith. Towards optimal demand-side bidding in parallel auctions for time-shiftable electrical loads. In *Proceedings of the IEEE International Conference on Industrial Informatics*, pages 340–347, June 2021. doi: 10.1109/INDIN45582.2020.9442077 **(Chapter 3)**
- Roland Saur, Han La Poutré, and Neil Yorke-Smith. Fair pricing for time-flexible smart energy markets. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '23, page 2703–2705, Richland, SC, May 2023. International Foundation for Autonomous Agents and Multiagent Systems. ISBN 9781450394321 **(Chapter 4)**
- Roland Saur, Han La Poutré, and Neil Yorke-Smith. Incentives for accurate energy predictions: How to reduce epistemic uncertainty. In *Proceedings of the 15th ACM International Conference on Future and Sustainable Energy Systems*, e-Energy '24, page 192–202, New York, NY, USA, June 2024. Association for Computing Machinery. ISBN 9798400704802. doi: 10.1145/3632775.3661956 **(Chapter 5)**

2 Combined Heat and Power Markets by Double-Sided Auction Mechanisms

Preface In order to reduce CO_2 emissions, energy systems using different energy carriers (e.g., heat and power) are becoming more intertwined and integrated. However, coordination between non-cooperative participants of these systems in the combined heat and power domain has been limited to single-sided auctions with one centralised seller. In this chapter, we present a double-sided auction mechanism in which prosumers as well as consumers and producers of heat and power can participate. By showing that our mechanism is Incentive Compatible and Individually Rational, we ensure that truthful bidding is the optimal strategy, simplifying the bidding process and thus accommodating agents with limited computational resources. Finally, we show that our mechanism is fiscally sustainable, i.e., Weakly Budget Balanced and close to optimal. [1]

2.1 Introduction

Global warming is one of the defining challenges of the 21st century. In order to reduce CO_2 emissions, significant investments in energy infrastructure and renewable energy sources (RES) have been made. These investments have had profound effects on the energy sector. Excess electricity production from fluctuating RES had to be curtailed at times (Böttger et al., 2014), while heating systems have been switched from gas-fired or electric boilers to combined heat and power plants (CHPs), which convert fuel much more efficiently by producing both heat and electricity.

Closer coordination between power and heat systems promises more efficient use of energy, meaning less curtailment of RES (Böttger et al., 2014; Mueller et al., 2014), better use of exhaust heat from industrial processes (Chae

[1] This chapter is based on publication (Saur et al., 2019)

et al., 2010) and a more economic operation of CHPs with District Heating Networks (Dimoulkas and Amelin, 2014; Javanshir et al., 2023). Attempts of better integration between heat and power have often focused on how to give individual CHPs the flexibility to react to fluctuating renewable electricity production, by considering heat storage – either dedicated (Chen et al., 2014; Dimoulkas and Amelin, 2014; Su et al., 2026) or indirectly in District Heating Networks (Li et al., 2016) – or by extending the operational range through the use of Power-to-Heat devices (Chen et al., 2014; Dimoulkas and Amelin, 2014; Gao et al., 2023). Zhong et al. (2018a) (with some extensions in (Zhong et al., 2018b)) introduced a multi-energy auction mechanism to coordinate between self-interested agents. However, this line of work was limited to single-sided auctions in which participants only buy electricity and heat from a single centralised seller. Since the publication (Saur et al., 2019) of the work underlying this chapter, the field of multi-energy market mechanisms has expanded into double sided auctions. Both (Li et al., 2023) and (Fu et al., 2025) provide two sided markets for combined heat and power. However, neither of them show that their approach is Incentive Compatible, i.e. guarantees honest behaviour by the participants.

To summarize, the engineering issue at hand is to design a market that allows participants to trade heat and electricity in an integrated way. This includes industrial processes that require both heat and power to operate, combined heat and power plants, whose heat and electricity output is coupled and heat pumps that can only sell heat if the necessary electricity is procured.

This work addresses an underlying foundational challenge of market design in a stylized setting that still captures the essential aspects of such integrated energy trading. We first develop a market model that allows both consumers and producers as well as power-to-heat devices to participate. We then design a double auction mechanism for allocating supply and demand in this market model. The mechanism design is based on the concept of trade reduction, which was first introduced in McAfee (1992). Gonen et al. (2007) provide a general framework for creating trade reduction mechanisms based on optimal dividing functions, but only provide a way of obtaining such dividing functions for procurement class domains, in which replacement sets are always of size one, i.e. an agent in a trade can only ever be replaced by one agent. However, in our setting, we consider agents that simultaneously buy heat and sell electricity. Such an agent can be replaced with two agents, one buying heat and another selling electricity. This work provides a trade reduction mechanism for such a setting. For a broader and more up to date overview of trade reduction mechanisms, we refer the reader to (Ehrlich et al., 2025). For our mechanism, we show that it incentivises agents to participate truthfully in the market, i.e., the mechanism is *Individually Rational* (IR) (Nisan et al., 2007) and *Incentive*

Compatible (IC) (Nisan et al., 2007). We thereby reduce the need for the agents to strategize. Lastly we prove that our market design does not require subsidies, i.e., the mechanism is *Weakly Budget Balanced* (WBB) (Dash et al., 2003) and show that our mechanism is close to optimal for large numbers of participants.

In this work, we contribute to the scientific literature by:

- Introducing a market model in which agents can simultaneously buy and sell heat and power, enabling power-to-heat devices to participate as service providers.
- Developing a double auction mechanism based on trade reduction, which we show to be not only WBB but also IR and IC for all market participants.

2.2 The Model

We introduce our market model and define desirable properties for its pricing mechanisms. We consider a unit demand/supply market, in which agents can both buy and sell bundles of heat and electricity. A unit in this market can be any amount of energy from Wh to MWh, as long as the unit stays consistent within a single run of the market. Every agent l in the set of agents N has a desired bundle (i_l, j_l), with $i_l, j_l \in \{0, 1, -1\}$, and associates a private valuation $v_l \in \mathbb{R}$ with this bundle. i indicates the amount of electricity while j is the amount of heat. Positive values in the bundle indicate consumption, while negative values are associated with production. For instance, a bundle $(1, -1)$ means an agent wants to consume one unit of electricity and produce one unit of heat. Every agent $l \in N$ submits one bid b_l consisting of its desired bundle and a reserve price $\rho_l \in \mathbb{R}$: $b_l = ((i_l, j_l), \rho_l)$. The reserve price is the maximum price an agent is willing to pay for its bundle, with negative values seen as payments to the agents. We will call a bid involving only one of the commodities a *simple bid*, e.g. $((1, 0), \rho_l), ((0, -1), \rho_l)$, while calling all other bids *complex bids*.

Agents posses a quasi-linear utility function $u_l(X, P)$, where X is the binary allocation vector, with $x_l = 1$ if agent l is allocated and $x_l = 0$ otherwise, and $P_{(i_l, j_l)}$ is the price set for bundle (i_l, j_l) by the mechanism. Not receiving any bundle leaves the agent with zero utility.

$$u_l(X, P) = v_l \cdot x_l - P_{(i_l, j_l)} \cdot x_l \tag{2.1}$$

An allocation is *feasible* if it results in material balance, i.e., $\sum_{l \in N} i_l \cdot x_l = 0 \wedge \sum_{l \in N} j_l \cdot x_l = 0$, ensuring that every unit of energy produced is also consumed within the system.

We assume a disinterested market operator concerned with maximising the social welfare of the system as a whole. Similarly to (Chu, 2009), we define social welfare as the sum of valuations of all of the participating agents.

Given this we can express our task as maximizing the following:

$$\begin{aligned} SW = & \sum_{l \in N} v_l x_l \\ s.t. \quad & \sum_{l \in N} i_l \cdot x_l = 0; \quad \sum_{l \in N} j_l \cdot x_l = 0 \\ & x_l \in \{0, 1\}, \end{aligned} \tag{2.2}$$

where the first constraint ensures strict material balance, while the second constraint prohibits partial allocation of bundles.

2.2.1 Properties Relevant to Mechanism Design

In this section, we explain some of the properties we want our market to have.

To coordinate between all the bidding agents, a mechanism takes in all bids and outputs an allocation of trading agents and prices. However, as the mechanism only receives reported bids (ρ_l) and not true valuations (v_l), it only maximizes social welfare, as described in equation 2.2, if the agents report their private valuations truthfully, i.e. $\rho_l = v_l \quad \forall l \in N$.

Mechanisms that are *Incentive Compatible* (IC) make truthful reporting the utility maximising strategy for every agent independent of the behaviour of other agents. The assumption of honesty is, therefore, fulfilled under the much weaker assumption of utility maximising agents if the mechanism is IC. Further, making truthful reporting the optimal strategy independent of the bids of other agents obviates the need for sophisticated bidding strategies. This enables small agents with limited computational power to efficiently operate in the market.

Additionally, to make participation appealing to agents, the mechanism should be *Individually Rational* (IR), i.e. reporting truthfully to the auctioneer never leads to negative utility for the reporting agent.

However, the seminal paper by Myerson and Satterthwaite (1983) showed that it is impossible to design a mechanism that is IC, IR, optimal in social welfare and *Weakly Budget Balanced* (WBB) at the same time, meaning any optimal, IC and IR mechanism necessarily runs a budget deficit, making it unsustainable. This impossibility result can be amended by trade reduction mechanisms, which artificially raise buying and lower selling prices by removing trading agents, in effect sacrificing some social welfare for WBB (McAfee, 1992).

The formal definitions of these concepts are as follows:

Definition 2.1 (IC)**.** (Nisan et al., 2007) A mechanism is IC if for all agents truthfully reporting their valuation maximises their utility independent of what others bid.

Definition 2.2 (IR)**.** (Nisan et al., 2007) A mechanism is IR, if for all agents truthful reporting never leads to negative utility.

Definition 2.3 (WBB)**.** (Dash et al., 2003) A mechanism is WBB, if the payments by the agents to the mechanism always outweigh the payments made by the mechanism.

We strive to develop a mechanism that is IC, IR and WBB.

2.3 Trade Reduction Mechanism

In this section, we introduce our trade reduction mechanism. We make the following two assumptions.

Assumption 2.1. No two bids have the exact same valuation.

Assumption 2.2. The market is deep: there is a sufficient number of bids per bundle allocated.

Our mechanism outputs both a final allocation and a set of prices, one for each bundle. As described in Alg. 1, it computes an optimal allocation X^* (line 1), removes the lowest two bids per bundle (lines 4 and 5), and sets the prices to the value reported by the removed bids (line 3).

Algorithm 1: Trade reduction mechanism returns the final allocation and prices.

Data: Set of Bids B
Result: Final allocation $X^{\square}$ and prices P

1 $X^* \leftarrow argmax_X SW(X, B)$ **for** $i \in \{0, 1, -1\}$ **do**
2 **for** $j \in \{0, 1, -1\}$ **do**
3 $P_{(i,j)} \leftarrow \rho_{s(i,j)}$
4 $x^*_{k(i,j)} \leftarrow 0$
5 $x^*_{s(i,j)} \leftarrow 0$
6 $P \leftarrow \bigcup_{i,j \in \{0,1,-1\}} P_{(i,j)}$
7 $X^{\square} \leftarrow X^*$
8 **return** $(X^{\square}, P)$

To describe the process in more detail, we introduce $k(i,j)$ as the allocated agent with the lowest bid for bundle (i,j).

$$k(i,j) = argmin_{m \in N}\{\rho_m | (i_m, j_m) = (i,j) \wedge x_m^* = 1\}$$

Next to $k(i,j)$, we define the next lowest bid for every bundle as $s(i,j)$.

$$s(i,j) = argmin_{m \in N}\{\rho_m | \rho_m > \rho_{k(i,j)}, (i_m, j_m) = (i,j)\} \tag{2.3}$$

The price for bundle (i,j) is defined by agent $s(i,j)$.

$$P_{(i,j)} = \rho_{s(i,j)}, \quad i,j \in \{0,1,-1\} \tag{2.4}$$

2.3.1 General Properties

Before we discuss the desired properties of our trade reduction mechanism (i.e., IR, IC, WBB), we establish some properties that are required for the proofs in Section 2.3.2. Notably, we establish that all agents in the final allocation reported a value above the price defining agent (see Lemma 2.2), remain allocated as long as they report a value above the value of the price defining agent (see Lemma 2.7) and cannot influence the price while allocated (see Lemma 2.8).

First we introduce some notation. Let N^* be the set of all agents allocated under allocation X^*: $N^* = \{l \in N | x_l^* = 1\}$.The set of all agents interested in bundle (i,j) is denoted as $N(i,j) = \{l \in N | (i_l, j_l) = (i,j)\}$.

The two agents, $k(i,j)$ and $s(i,j)$, which are removed by the trade reduction mechanism, will be denoted as $N^{\Delta}(i,j)$. This set is clearly defined because of Assumptions 2.1 and 2.2.

$$\begin{aligned} & N^{\Delta}(i,j) \subset N^*(i,j) \\ s.t. \quad & |N^{\Delta}(i,j)| = 2 \wedge (\rho_m > \rho_l) \\ \forall l \in & N^{\Delta}(i,j), \forall m \in N^*(i,j) \setminus N^{\Delta}(i,j) \end{aligned} \tag{2.5}$$

Agent $s(i,j)$ is therefore the highest valued bid in $N^{\Delta}(i,j)$.

$$s(i,j) = argmax_l\{\rho_l | l \in N^{\Delta}(i,j)\} \tag{2.6}$$

The set of all removed bids is defined as $N^{\Delta} = \bigcup_{(i,j)} N^{\Delta}(i,j)$, which allows us to describe the final allocation $N^{\square}$ as follows:

$$N^{\square} = N^* \setminus N^{\Delta}. \tag{2.7}$$

We consider an agent $l \in N$ interested in bundle (i,j) and consider the effects of only agent l changing its bid. For this, we introduce the bidding profile $b_{-l} = (b_1, b_2, ...b_{l-1}, b_{l+1}, ...b_n)$ as the set of bids from all agents except for agent l. We define $N^{1*}(((i,j), \rho_l^1), b_{-l})$ and $N^{2*}(((i,j), \rho_l^2), b_{-l})$

as an optimal allocations resulting from agent l bidding ρ_l^1 and ρ_l^2 under bidding profile b_{-l}, respectively. For brevity, we will write N^{1*} and N^{2*} instead of $N^{1*}(((i,j),\rho_l^1),b_{-l})$ and $N^{2*}(((i,j),\rho_l^2),b_{-l})$.

We show that an allocated agent cannot affect an optimal allocation by increasing its bid (see Lemma 2.1 and Corollary 2.1) and describe all allocated agents (see Lemma 2.2).

Lemma 2.1. Given a bidding profile b_{-l} and an agent l interested in bundle (i,j) the following statement holds:

$$\left(\rho_l^2 > \rho_l^1 \wedge l \in N^{1*}\right) \implies N^{1*} = N^{2*}.$$

Proof. Assume by contradiction that $\rho_l^2 > \rho_l^1 \wedge l \in N^{1*}$ and $N^{1*} \neq N^{2*}$. We define $T^{\Delta} = N^{1*} \setminus N^{2*}$ and $T^{-\Delta} = N^{2*} \setminus N^{1*}$, with which we can express N^{2*} as follows:

$$N^{2*} = \left(N^{1*} \setminus T^{\Delta}\right) \cup T^{-\Delta}. \tag{2.8}$$

As N^{2*} is an optimal allocation given bids $B^2 = (((i,j),\rho_l^2),b_{-l})$, we know that $\sum_{m\in N^{2*}} \rho_m^2 \geq \sum_{m\in N^{1*}} \rho_m^2$. Applying equation (2.8) gets us to the following expression:

$$\begin{aligned}\sum_{m\in N^{1*}} \rho_m^2 - \sum_{m\in T^{\Delta}} \rho_m^2 + \sum_{m\in T^{-\Delta}} \rho_m^2 &\geq \sum_{m\in N^{1*}} \rho_m^2 \\ \sum_{m\in T^{-\Delta}} \rho_m^2 - \sum_{m\in T^{\Delta}} \rho_m^2 &\geq 0.\end{aligned} \tag{2.9}$$

Given the fixed bidding profile b_{-l} and $\rho_l^2 > \rho_l^1$, we know that $\rho_m^2 \geq \rho_m^1 \quad \forall m \in N$ and therefore $\sum_{m\in T^{\Delta}} \rho_m^2 \geq \sum_{m\in T^{\Delta}} \rho_m^1$. Furthermore, $\rho_m^1 = \rho_m^2$ holds for all $m \neq l$ from which follows that $\sum_{m\in T^{-\Delta}} \rho_m^2 = \sum_{m\in T^{-\Delta}} \rho_m^1$ as $l \notin T^{-\Delta}$. Applying these insights to equation (2.9) results in:

$$\begin{aligned}\sum_{m\in T^{-\Delta}} \rho_m^1 - \sum_{m\in T^{\Delta}} \rho_m^1 \geq \sum_{m\in T^{-\Delta}} \rho_m^2 - \sum_{m\in T^{\Delta}} \rho_m^2 &\geq 0 \\ \sum_{m\in T^{-\Delta}} \rho_m^1 - \sum_{m\in T^{\Delta}} \rho_m^1 &\geq 0.\end{aligned} \tag{2.10}$$

There exist two cases for equation (2.10) – strictly larger and equality. When $\sum_{m\in T^{-\Delta}} \rho_m^1 - \sum_{m\in T^{\Delta}} \rho_m^1 > 0$ holds, equation (2.10) contradicts our assumption that N^{1*} is optimal, as $N^{1*} \setminus T^{\Delta} \cup T^{-\Delta}$ results in higher social welfare. When equality holds, either $T^{-\Delta} = T^{\Delta} = \varnothing$ in which case $N^{1*} = N^{2*}$ or there exist multiple optimal allocations, one of which is N^{2*}; consistent tie breaking will lead to $N^{1*} = N^{2*}$. □

Corollary 2.1. Given a bidding profile b_{-l} and an agent l the following holds: $\left(l \in N^{1*} \wedge l \in N^{2*}\right) \implies N^{1*} = N^{2*}$.

Lemma 2.2 (Allocated Agents)**.** $\rho_l > \rho_{s(i,j)} \quad \forall l \in N^{\square}(i,j)$

Proof. $s(i,j) \in N^{\Delta}(i,j)$. The result follows from the definitions of $N^{\Delta}(i,j)$ and $N^{\square}(i,j)$ in (2.5) and (2.7). □

We next introduce the central concept of a trade (see Definition 2.4) as the building block of allocations (see Lemma 2.3) and show that trades can at most contain 2 bids on the same bundle (see Lemmas 2.4 and 2.5).

Definition 2.4 (Trade)**.** A trade T is a nonempty set of agents m s.t. $(\sum_{m\in T} i_m = 0 \wedge \sum_{m\in T} j_m = 0)$ while $\nexists \tilde{T} \subsetneq T \;\; s.t. \;\; \tilde{T} \neq \varnothing \wedge \sum_{m\in \tilde{T}} i_m = 0 \wedge \sum_{m\in \tilde{T}} j_m = 0$

Lemma 2.3. Every feasible allocation set N^X can be partitioned into a set of trades: $N^X = \cup_m T_m$

Proof. Take any feasible allocation set N^X. Since N^X is feasible we know that $(\sum_{m\in N^X} i_m = 0 \wedge \sum_{m\in N^X} j_m = 0)$.
One of two cases applies. N^X can be partitioned into two feasible non-empty allocation sets, which can be partitioned again. N^X cannot be partitioned into two feasible allocation sets. Therefore there does not exist a non-empty strict subset of N^X that is feasible and N^X is a trade by Def. 2.4. □

Lemma 2.4. No trade T can contain more than one complex bid for the same bundle (i,j) $i,j \neq 0$.

Proof. The proof goes to show that assuming the opposite always leads to a trade within a trade, which contradicts our definition of a trade. We start by assuming that there is a trade T and two agents $l, m \in T$ bidding on bundle (i,j) $i,j \neq 0$. Let $T^- = T \setminus \{l,m\}$ be the residual trade. For T to be materially balanced $\sum_{r\in T^-} i_r = -2i \wedge \sum_{r\in T^-} j_r = -2j$ has to hold. Therefore, T^- needs to contain at least two agents bidding on bundles $\{(-i,0),(-i,-j),(-i,j)\}$ and at least two agents bidding on bundles $\{(0,-j),(-i,-j),(i,-j)\}$.

All of these options will lead to a contradiction. If there exists an agent $r \in T^-$ bidding on bundle $(-i,-j)$, then $\{r,l\} \subsetneq T$ is a trade. If there exists an agent $r \in T^-$ bidding on bundle $(-i,j)$, then for the trade to be materially balanced the following has to hold: $\sum_{r\in T^-\setminus r} j_r = -3j$. This means that T has to contain three agents bidding on bundles $\{(0,-j),(i,-j),(-i,-j)\}$. However, there cannot be an agent $r_2 \in T$ bidding on bundle $(i,-j)$, because

$\{r, r_2\}$ constitutes a trade. As mentioned before T cannot contain an agent bidding on bundle $(-i, -j)$. Therefore, T has to contain three agents r_2, r_3, r_4 bidding on $(0, -j)$. However, this leads to the following trade within a trade: $\{l, r, r_2, r_3\} \subsetneq T$. The same reasoning holds for bundle $(i, -j)$. Therefore, for T to be materially balanced, T^- contains agents r_1, r_2 bidding on $(-i, 0)$ and agents r_3, r_4 bidding on $(0, -j)$. However, then $\{l, r_1, r_3\} \subsetneq T$ is a balanced trade.

□

Lemma 2.5. No trade can contain more than two bids for the same simple bundle $(i, j) \in \{(1, 0), (0, 1), (-1, 0), (0, -1)\}$

Proof. We assume, contrary to the lemma, that there exists a trade T with bids $l, m, o \in T$ on a bundle $(i, 0)$, $i \in \{1, -1\}$. Let $T^- = T \setminus \{l, m, o\}$. For T to be feasible $\sum_{r \in T^-} i_r = -3i$ and therefore there have to be three bids on bundles $\{(-i, 0), (-i, 1), (-i, -1)\}$. However, there cannot be a bid $r \in T$ on bundle $(-i, 0)$. Otherwise $\{r, l\} \subsetneq T$ and $T \setminus \{r, l\}$ are feasible trades, contradicting our definition of a trade. Therefore, there have to be 3 bids on bundles in $\{(-i, 1), (-i, -1)\}$, which means that at least two bids are on the same complex bundle which violates Lemma 2.4. This contradiction holds similarly for 3 bids on a bundle $(0, j)$, $j \in \{1, -1\}$ □

With this knowledge about the composition of trades, we establish the value of the price defining agent as a lower bound on what an allocated agent could have bid and still be allocated (see Lemmas 2.6 and 2.7).

Lemma 2.6. Given a bidding profile b_{-l}, an agent l interested in bundle (i, j) and its bid $((i, j), \rho_l^1)$ such that $l \in N^{1*}$:

$$\rho_l^2 > \rho_{s^1(i,j)} \implies l \in N^{2*}.$$

Proof. Given an agent l with bid $((i, j), \rho_l^1)$ such that $l \in N^{1*}$, assume contrary to the lemma that $\rho_l^2 > \rho_{s^1(i,j)}$ and $l \notin N^{2*}$. There are two cases to be considered.

Case 1: $l \in \{s^1(i, j), k^1(i, j)\}$: If so, we know that $\rho_l^2 > \rho_{s^1(i,j)} \geq \rho_l^1$ and thus $l \in N^{2*}$ by Lemma 2.1. This contradicts the initial assumption.

Case 2: $l \notin \{s^1(i, j), k^1(i, j)\}$: Since, by the initial assumption, $l \notin N^{2*}$ and $\rho_l^2 > \rho_{s^1(i,j)} > \rho_{k^1(i,j)}$, it follows that $s^1(i, j), k^1(i, j) \notin N^{2*}$. Otherwise N^{2*} could be improved by swapping $k^1(i, j)$ or $s^1(i, j)$ for l.

We will arrive at a contradiction by showing that N^{1*} is not optimal when l bids ρ_l^1. Define $T^\Delta = N^{1*} \setminus N^{2*}$ and $T^{-\Delta} = N^{2*} \setminus N^{1*}$. Since $l, s^1(i, j), k^1(i, j) \in N^{1*}$ and $l, s^1(i, j), k^1(i, j) \notin N^{2*}$ we know that $l, s^1(i, j), k^1(i, j) \in T^\Delta$. Consider improving N^{1*} by identifying a balanced

exchange of agents occurring in the process of going from N^{1*} to N^{2*} that is independent of agent l.

First, we express the removal of T^{Δ} and addition of $T^{-\Delta}$ as a new hypothetical allocation. To do so, we introduce M^{Δ} as a mirror image of T^{Δ} as follows. For every agent $m \in T^{\Delta}$ define one corresponding agent $m^M \in M^{\Delta}$ with bid $b_{m^M} = ((-i_m, -j_m), -\rho_m)$. Let k^M, s^M, $l^M \in M^{\Delta}$ be the corresponding agents for agents $k^1(i,j), s^1(i,j), l \in T^{\Delta}$, respectively. $(M^{\Delta} \cup T^{-\Delta})$ is a feasible allocation, disjoint from N^{1*}, that represents the exchange of agents when going from N^{1*} to N^{2*}.

Second, to show that N^{1*} can be improved upon, we identify a trade in $(M^{\Delta} \cup T^{-\Delta})$ that is independent of l as follows. Let $T_k \subseteq (M^{\Delta} \cup T^{-\Delta})$ be a maximum value trade such that $k^M \in T_k$. We know that $l^M \notin T_k$. Otherwise it would follow from Lemmas 2.4 and 2.5, and $l^M, k^M \in T_k$, that $s^M \notin T_k$. Then, since $-\rho_{s^1(i,j)} > -\rho_l$, $T_k \setminus \{l^M\} \cup \{s^M\}$ would have a higher value, contradicting our definition of T_k as a maximum value trade.

Third, we now use T_k to improve upon N^{1*}, as follows. Let $T_k^{-\Delta} = T_k \cap T^{-\Delta}$ and let $T_k^{\Delta} \subseteq T^{\Delta} (\subset N^{1*})$ be the set of agents that correspond to the agents in the mirror image $T_k \cap M^{\Delta}$. As N^{2*} is optimal, $N^{1*} \setminus T_k^{\Delta} \cup T_k^{-\Delta}$ has a higher social welfare than N^{1*}, when l bids ρ_l^2. However, since $l^M \notin M^{\Delta}$ and therefore $l \notin T_k^{\Delta}$, $N^{1*} \setminus T_k^{\Delta} \cup T_k^{-\Delta}$ has a higher social welfare than N^{1*} independent of what l bids, contradicting our assumption that N^{1*} is optimal when l bids ρ_l^1. □

Lemma 2.7. Given a bidding profile b_{-l}, an agent l interested in bundle (i,j) and its bid $((i,j), \rho_l^1)$ such that $l \in N^{1\square}$:

$$\rho_l^2 > \rho_{s^1(i,j)} \implies l \in N^{2\square}.$$

Proof. As $l \in N^{1\square}$ and $N^{1\square} \subset N^{1*}$, we know that $l \in N^{1*}$. Given $l \in N^{1*}$ and $\rho_l^2 > \rho_{s^1(i,j)}$ it follows by Lemma 2.6 that $l \in N^{2*}$ and thus $N^{1*} = N^{2*}$, see Corollary 2.1. By our initial assumption, we know that $\rho_l^2 > \rho_{s^1(i,j)} > \rho_{k^1(i,j)}$ and by $N^{1*} = N^{2*}$ we know that $s^1(i,j), k^1(i,j) \in N^{2*}$. This implies that l is not one of the lowest two bids on bundle (i,j) in N^{2*}, i.e., $l \notin N^{2\Delta}$. From $l \notin N^{2\Delta} \wedge l \in N^{2*}$ it follows that $l \in N^{2\square}$. □

Further, we show that agents in the final allocation cannot at the same time affect the price they pay for their bundle and remain in the final allocation (see Lemma 2.8).

Lemma 2.8. Given a bidding profile b_{-l} and an agent l interested in bundle (i,j), the following holds:

$$l \in N^{1\square} \wedge l \in N^{2\square} \implies P^1_{(i,j)} = P^2_{(i,j)}.$$

Proof. Recall that $s(i,j)$ is the price defining agent for bundle (i,j) (see Equation (2.4)) and therefore it suffices to show that $\rho_{s^1(i,j)} = \rho_{s^2(i,j)}$. Since $N^{\square} \subset N^*$, it follows from our initial assumption that $l \in N^{1*} \wedge l \in N^{2*}$ and therefore $N^{1*} = N^{2*}$, by Corollary 2.1. As $l \in N^{1\square} \wedge l \in N^{2\square}$, we know that $l \notin N^{1\Delta} \wedge l \notin N^{2\Delta}$. Since all bids of agents other than l are fixed by b_{-l}, we know that $N^{1\Delta} = N^{2\Delta}$ and by extension $\rho_{s^1(i,j)} = \rho_{s^2(i,j)}$. □

2.3.2 Mechanism design relevant properties

With the above general properties in place, we can now show that our trade reduction mechanism (Alg. 1) satisfies all three desired properties: IR, IC and WBB (see section 2.2).

Individual Rationality (IR)

To ensure IR, we need to show that agents truthfully reporting their valuation can never obtain negative utility.

Theorem 2.1 (Individual Rationality)**.** *The mechanism is Individually Rational, i.e.,* $u_l(X(\rho_l = v_l), P(\rho_l = v_l)) \geq 0$

Proof. Assume that trading agent l has a bid $b = ((i,j), \rho_l)$ and true valuation v_l. If agent l is not allocated $l \notin N^{\square}$, then its utility is 0. If agent l is allocated $l \in N^{\square}$, its utility is $u_l = v_l - P_{(i,j)} = v_l - \rho_{s(i,j)} = \rho_l - \rho_{s(i,j)} \geq 0$. The inequality holds because of Lemma 2.2. For the auctioneer agent IR holds due to WBB, see Theorem 2.3. □

Incentive Compatible (IC)

To show that the mechanism satisfies IC, we need to show that no agent can strictly increase its utility by bidding anything but its true valuation.

Theorem 2.2 (Incentive Compatibility)**.** *The mechanism is Incentive Compatible, i.e.* $u_l(\rho_l = v_l) \geq u_l(\rho_l \neq v_l) \quad \forall l \in N$.

Proof. To show that no agent can improve its utility by deviating from bidding its true valuation, we separately examine agents that would and would not be allocated under truthful bidding. Let agent l have private valuation v_l and let $\rho_l^1 = v_l$. Case $l \in N^{1\square}$ (allocated under truthful bidding):
There are two cases to be considered for the bid of agent l. Either it is allocated or not when bidding ρ_l^2.

- $l \in N^{2\square}$: Through Lemma 2.8, we know that $l \in N^{1\square} \wedge l \in N^{2\square}$ implies $P^1_{(i,j)} = P^2_{(i,j)}$ and the utility of the agent can be written as $u_l(\rho_l^2) = v_l - P^2_{(i,j)} = v_l - P^1_{(i,j)} = u_l(\rho_l^1 = v_l)$.

- $l \notin N^{2\square}$: Given $l \notin N^{2\square}$, agent l receives zero utility and its utility is: $u_l(\rho_l^2) = 0 \leq u_l(\rho_l^1 = v_l)$. The inequality holds because the mechanism is Individually Rational.

Case $l \notin N^{1\square}$ (not allocated under truthful bidding):
There are two cases to be considered for the bid of agent l. Either it is allocated or not when bidding ρ_l^2.

- $l \in N^{2\square}$: Recalling that $s(i,j)$ is the price defining agent, the utility of agent l can be written as $u_l(\rho_l^2) = v_l - \rho_{s^2(i,j)} \leq 0 \leq u(\rho_l^1 = v_l)$. Indeed, we know that $v_l \leq \rho_{s^2(i,j)}$ as otherwise $\rho_l^1 = v_l > \rho_{s^2(i,j)}$ leads to a contradiction: By Lemma 2.7 (index reversed), $\rho_l^1 > \rho_{s^2(i,j)}$ and $l \in N^{2\square}$ implies that $l \in N^{1\square}$ – a contradiction. The second inequality holds because the mechanism is Individually Rational.
- $l \notin N^{2\square}$: As both $l \notin N^{1\square}$ and $l \notin N^{2\square}$, it follows that the utility of agent l is zero either way.

□

Weak Budget Balance (WBB)

To ensure WBB, we need to show that payments made to the mechanism are larger than payments made by the mechanism.

Theorem 2.3 (Weak Budget Balance)**.** *The mechanism is weakly budget balanced.*

Proof. Given Lemma 2.3, we can split any allocation $N^{\square}$ into trades and therefore only have to show that no trade can run a budget deficit. The budget of a trade $T \subset N^{\square}$ is the sum of payments by the agents in the trade and can be expressed as the prices associated with the bundles of the agents in the trade. Recall that the price for a bundle (i,j) is set by the corresponding agent $s(i,j) \in N^*$, see equation (2.4). Alternatively, we can sum over all possible bundles (i,j) and multiply each price by the number of times the bundle occurs in T. For this, let $|T_{(i,j)}|$ be the number of agents in T that bid on bundle (i,j). Then:

$$\begin{aligned}\text{Budget}(T) &= \sum_{m \in T} P_{(i_m,j_m)} = \sum_{m \in T} \rho_{s(i_m,j_m)} = \sum_{(i,j)} \rho_{s(i,j)} |T_{(i,j)}| \\ &\geq \sum_{(i,j)} \begin{cases} \rho_{s(i,j)} + \rho_{k(i,j)} & \text{if } |T_{(i,j)}| = 2 \\ \rho_{s(i,j)} & \text{if } |T_{(i,j)}| = 1 \\ 0 & \text{if } |T_{(i,j)}| = 0 \end{cases} \geq 0.\end{aligned}$$

By Lemmas 2.4 and 2.5, we know that $|T_{(i,j)}|$ can only be 0, 1 or 2. Therefore the inequality holds if it holds for $|T_{(i,j)}| \in \{0, 1, 2\}$. In the case that $|T_{(i,j)}| = 2$, we know that $2 \cdot \rho_{s(i,j)} \geq \rho_{s(i,j)} + \rho_{k(i,j)}$, because we defined $\rho_{s(i,j)} \geq \rho_{k(i,j)}$, see (2.3). For $|T_{(i,j)}| = 1$, the inequality holds trivially as $\rho_{s(i,j)} \geq \rho_{s(i,j)}$. For $|T_{(i,j)}| = 0$, the inequality holds trivially as $0 \geq 0$.

The second inequality holds, since we assumed N^* to be optimal in Social Welfare, and therefore the sum of the valuations of any trade in N^* has to yield a non-negative value. Otherwise, there exists a trade with negative sum of valuations, which means that N^* can be improved by removing said trade, contradicting our assumption that N^* is optimal.

□

2.3.3 Social Welfare Loss

The mechanism presented here is financially sustainable, i.e. does not lose money, and incentivises truthful participation. Given truthful bidding, we can express the social welfare of any feasible allocation N^* as $\sum_{l \in N^*} \rho_l$.

In order to gauge the loss in social welfare due to trade reduction we introduce the efficiency ratio as the ratio of social welfare achieved by our mechanism ($SW^{\square}$) divided by the optimal social welfare (SW^*). Our mechanism removes N^{Δ} from the optimal allocation N^* and therefore loses the value of agents $s(i,j)$ and $k(i,j)$ per bundle. The efficiency ratio can be expressed as follows:

$$\begin{aligned} \text{Efficiency-Ratio} &= \frac{SW^{\square}}{SW^*} = \frac{SW^* - \sum_{l \in N^{\Delta}} \rho_l}{SW^*} \\ &= 1 - \frac{SW^{\Delta}}{SW^*} \end{aligned} \tag{2.11}$$

The number of agents in the optimal allocation grows as participation in the market grows, while the number of agents in N^{Δ} stays steadily at 2 agents per bundle. Furthermore, N^{Δ} consists of the two lowest bidding agents in N^*. The ratio of social welfare of removed agents (SW^{Δ}) to social welfare of agents in the optimal allocation (SW^*) therefore tends towards zero and the efficiency ratio towards one as more agents participate.

This makes our mechanism especially well suited for so called deep markets, where many agents engage in trading and therefore the loss to social welfare is minimal.

2.4 Conclusion and Future Work

This chapter addressed combined heat and power markets. We developed a market setting in which not only consumers and producers, but also agents that want to buy one commodity and simultaneously sell another, e.g. power-to-heat devices, can participate. We introduced a trade reduction mechanism which does not require subsidies and which incentivises agents to participate truthfully, creating an efficient marketplace in which agents can participate without the need for sophisticated bidding strategies. By removing the computational burden from individual agents, we enable a future energy system to coordinate between small self-interested energy devices. Allowing bids on arbitrary bundles and accommodating more than two energy carriers are future extensions of this work.

3 Optimal Demand-Side Bidding in Parallel Auctions for Time-Shiftable Electrical Loads

Preface Increasing electricity production from renewable energy sources has, by its fluctuating nature, created the need for more flexible demand side management. How to integrate flexible demand in the electricity system is an open research question. We consider the energy-procurement problem of a time-shiftable load with access to a District Heating Network. For the setting of simultaneous second price auctions, we derive a required condition for optimal bidding strategies. We then show the following results regarding bidding strategies under different market assumptions. For identical uniform auctions and multiple units of demand, we show that the global optimal strategy is to bid uniformly across all auctions. For non-identical auctions and multiple units, we provide a way to find solutions through a recursive approach and a non-linear solver. We show that our approach outperforms the literature under higher uncertainty conditions. [1]

3.1 Introduction

Climate change is a defining challenge of the 21st century. To reduce CO_2 emissions, massive investments in renewable energies have been made and will still be required. Since renewable energy source are often not controlled generators, electricity production as well as electricity prices are subject to greater fluctuation than before.

Given the increased fluctuation in electricity production, both time shiftable loads (Graditi et al., 2015; Mohsenian-Rad, 2014; Palensky and Dietrich,

[1] This chapter is based on publication (Saur et al., 2021). Supporting code can be found under doi.org/10.4121/a6db0891-31ee-491f-b5c1-cf5a0feca4f3

2011; Strbac, 2008) and coupling between electrical and heat systems have emerged as potential remedies. The coupling of electrical and heat systems often comes in the form of District Heating Networks (DHN) (Dimoulkas et al., 2017; Foteinaki et al., 2018), while time-shiftable and other flexible loads can be found in various domains, such as: data centres (Ghamkhari and Mohsenian-Rad, 2013; Liu et al., 2024), water distribution systems (Oikonomou and Parvania, 2018), and household energy consumption (Agnetis et al., 2013). Particular attention should be given to industrial processes, which by their high energy intensity can achieve a significant impact on the electricity system (Barth et al., 2018).

In order for time shift-able loads to take advantage of fluctuating prices on day-ahead electricity markets, bidding strategies have been investigated. The approach of Mohsenian-Rad (2014) suggests to bid only in a single time slot – the cheapest one in expectation. However, their setting only considers time-shiftable electric assets by themselves, not an integration of time-shiftable loads and power-to-heat coupling.

The engineering setting we consider consists of a time-shiftable load such as electric vehicle charging or an industrial process coupled with a District Heating Network for one way energy balancing. The challenge is to design a bidding strategy for electricity procurement that takes into consideration both the ability to shift the electricity consumption in time and to balance excess electricity by converting it to heat.

We extend the literature on the underlying foundational challenge as follows. In the general case, we extend the optimality requirement derived in Gerding et al. (2008) from single unit to multi-unit demand (see Lemma 3.1 in Section 3.4). For the case of identical auctions, we show that uniform bidding, i.e. not just bidding in a single auction but participating in all auctions, is the optimal bidding strategy for uniform distributions (see Theorem 3.2). For the non-identical auction setting, in Section 3.4.3 we

1. provide a way to guide a non-linear solver to converge more often to a non-trivial solution.
2. provide a dynamic programming approach to make the problem computationally feasible.
3. show that our approach outperforms solutions found in literature (Mohsenian-Rad, 2014) under higher uncertainty conditions.

3.2 Related Work

Given our abstract procurement problem of a time shift-able load, we will consider both literature particular to demand side bidding and literature on bidding strategies more generally.

The problem of demand side bidding in electricity markets with fixed demand in a single time slot is considered in Liu and Guan (2003); Philpott and Pettersen (2006), while Herranz et al. (2012) considers multiple time slots and stochastic demand. Neither, however, consider any control over the demand they have to satisfy. In contrast, an electric vehicle aggregator, as considered in Bessa et al. (2011); Vagropoulos and Bakirtzis (2013); Zheng et al. (2023), can control or incentivize the charging rate of the electric vehicles in its fleet. However, this flexibility is in magnitude only and their bidding strategies, therefore, ignore the flexibility to shift demand in time. Demand side bidding for time shift-able loads is considered in Mohsenian-Rad (2014). However, we will show that their solution is no longer optimal under the assumption of free disposal.

The general problem of bidding in simultaneous auctions is considered in the Trading Agent Competition (Greenwald and Boyan, 2004; Wellman et al., 2003). The goal of their agent, however, was to construct bundles of non-identical items, which were strong complements. Our agent on the other hand values all electricity equally up to a particular demand. Bidding in simultaneous auctions with identical items has been considered in Rothkopf (1977). However, in their work winning any auction yields some valuation independently of the other auctions and what connects the auctions is a shared budget that should not be exceeded. A shift-able load, as considered by our work, on the other hand values electricity only up to its demand and has zero valuation beyond that. Closest to our work, Gerding et al. (2008) develops optimal bidding strategies for simultaneous auctions of identical items with a unit demand agent. We extend their work from single unit to multi unit demand.

3.3 Model

We consider the procurement problem of a deferrable electric load connected to a large District Heating Network. The deferrable load has to run for a total of $s \geq 1$ possibly non-consecutive time units within a bounded discretized time horizon. We assume that this deferrable load requires the same amount of energy for every time step in which it is switched on. As the amount of energy stored in a DHN is usually quite large, we assume that any excess electrical energy procured can be balanced via the DHN.

For the given time horizon, there exists a set of electricity auctions T, which are held in parallel and ahead of time. For readability reasons, we

will use the terms unit(s) and auction(s) interchangeably, where clear. To win the necessary s units/auctions to run the load, the agent submits a bid vector $B = (b_1, b_2, ..., b_{|T|})$ over auctions T. $f_t(\tau)$, supported on the interval $[0, \lambda_{el}]$, describes the clearing price distribution of auction $t \in T$, while $F_t(\tau)$ is its corresponding cumulative distribution, with $F_t(0) = 0$ and $F_t(\lambda_{el}) = 1$ [2]. The agent wins auction $t \in T$ if the submitted bid b_t is at least the clearing price of auction t.

The expected cost the agent incurs when submitting bid B given a demand of s units and a set of auctions T is composed of a market cost and a backup cost, see (3.1).

$$Cost(B|T, s) = Cost_M(B|T, s) + Cost_B(B|T, s) \tag{3.1}$$

The expected market cost is the sum of payments to auctions T. We assume that the agent is a price taker.

Assumption 3.1 (Price Taker). The agent has no effect on the clearing price.

The payment per auction $t \in T$ is the clearing price of that auction, conditionalized on the agent winning auction t.

$$Cost_M(B|T, s) = \sum_{t \in T} \int_0^{b_t} \tau f_t(\tau) d\tau \tag{3.2}$$

The expected backup cost is the payment made to a backup generator in case of shortfall. If fewer than the desired s units are acquired, the agent incurs the maximum price of λ_{el}. We are making the assumption of free disposal.

Assumption 3.2 (Free Disposal). We assume free disposal: any amount of electricity acquired beyond the agent's need can be disposed of at zero cost.

This assumption is justified by the connection of the deferrable load to the district heating network. Any excess energy that the deferrable load aqcuires can be turned into heat in the district heating network. We are assuming that the electrical load is small relative to the district heating network, which is, therefore able to absorb any excess electrical power.

The backup cost is zero if the agent wins more than s units. For any number of won auctions $j \in [0, s-1]$ there is a set of subsets $w \subset T$ such that $|w| = j$. For every subset of auctions w, the probability that the agent wins the subset and loses all other auctions is

$$\prod_{t \in w} F_t(b_t) \prod_{t \in T \setminus w} [1 - F_t(b_t)]$$

[2] We assume that f is non-zero as well as continuous and differentiable on the interval $[0, \lambda_{el}]$.

By summing over all possible sets $w \subset T \quad s.t. \quad |w| = j$, we obtain the probability of the agent winning j of the T auctions and can therefore calculate the expected backup cost as follows

$$Cost_B(B|T,s) = \underbrace{\sum_{j<s} \sum_{\substack{w \subset T \\ |w|=j}} \prod_{t \in w} F_t(b_t) \prod_{t \in T \setminus w} [1 - F_t(b_t)]}_{Prob(x=j|T,B)} (s-j)\lambda_{el}. \quad (3.3)$$

Here $Prob(x = j|T, B)$ represents the probability of winning j auctions out of the set of T when submitting bid B.

3.4 Optimal Bidding Strategies

The most straight-forward bidding strategy is to bid λ_{el} in s auctions and zero in all other auctions. We define this as a special strategy (Definition 3.1).

Definition 3.1. Let Λ_s be the set of bid vectors which consist of s bids of λ_{el} and $(n - s)$ bids of zero value, where $n = |T|$.

Among the Λ_s-bidding strategies there exists one that submits λ_{el} to the auctions with the lowest clearing price in expectation. Whenever we compare any strategy to Λ_s-bidding strategies, we mean this lowest price version of it. We will show that any strategy that meets the requirement set out in Lemma 3.1 and wins or loses an auction with certainty, i.e., $F_t(b_t) = 1$ or $F_t(b_t) = 0$, is a Λ_s-bidding strategy (Theorem 3.1).

We begin our theoretical solution by introducing a requirement that any optimal bid has to satisfy, see Lemma 3.1, and provide a strategy that always meets this requirement, see Theorem 3.1.

Lemma 3.1 introduces an optimality condition. We define B_{-k} as the bid vector B but with the bid for auction k removed. x is the random variable indicating the number of auctions the agent won. Intuitively speaking (3.4) states that the bid submitted to any particular auction $k \in T$ is proportional to the probability of falling short in the remaining auctions.

Lemma 3.1. Any optimal bid for s out of T auctions has to satisfy the requirement in (3.4).

$$b_k = Prob(x < s|T \setminus \{k\}, B_{-k})\lambda_{el} \quad \forall k \in T. \quad (3.4)$$

Proofs for Lemmas and Theorems are given in Appendix A.1.1.

Theorem 3.1. *Any bid vector B that satisfies condition* (3.4) *and contains a bid of value λ_{el} or 0 is an element of Λ_s.*

This means that an optimal bidding strategy is either a Λ_s-bidding strategy or participates in all auctions with some non-zero bid.

3.4.1 Comparison to the Literature

Having established our model, condition for optimality and first bidding strategy, we connect our work to that of (Mohsenian-Rad, 2014). Mohsenian-Rad addresses the energy acquisition problem of a time-shiftable load that can be run within a single time step. His approach considers submitting energy-price bid pairs to a set of day ahead auctions with a secondary intra-day market for recourse. Our bids on the other hand only consider the price and we consider a backup generator instead of a second stage intra-day market.

However, Theorem 3 in (Mohsenian-Rad, 2014) states that the optimal bidding strategy on the day ahead market is to acquire the entire energy need from a single auction, i.e. submitting a single non-zero bid for the entire energy need with a price equal to the expected price on the intra-day market. In our setting the intra-day market is replaced with a backup generator of deterministic cost λ_{el}. The expected price on the intra-day market in (Mohsenian-Rad, 2014) can be viewed as our backup generator cost λ_{el}. Translated to our setting, Mohsenian-Rad suggests bidding λ_{el} in a single auction and 0 in all others. We will therefore associate the strategy suggested by (Mohsenian-Rad, 2014) with Λ_s-bidding as defined in Definition 3.1.

3.4.2 Identical Auctions

As a first step, we consider an identical auctions setting. We will show that uniform bidding satisfies (3.4) and provide an algorithm that can quickly find the correct uniform bid value.

Assumption 3.3 (Identical Auctions). All auctions are identical: $f(b) = f_t(b) \quad \forall t \in T$.

We define the uniform bidding strategy as a bid B that submits the same value to all auctions. For this setting, uniform bidding satisfies (3.4).

Definition 3.2 (Uniform Bid)**.** A uniform bid vector B is a bid vector such that $b_k = b_l \quad \forall k, l \in T$

Lemma 3.2. Under the assumption 3.3 , there exists a unique uniform bid $B_u = (b_u, b_u, ..., b_u)$ s.t. condition (3.4) holds.

This uniform bid value, b_u, can be approximated via interval halving, as the right side of (3.4) strictly increases while the left side decreases with increasing b_u.

Identical Uniform Auctions

Next, we consider the case where every auction has an identical uniform price distribution.

Assumption 3.4 (Identical Uniform Auctions). The price of every auction is uniformly distributed on the interval $[0, \lambda_{el}]$: $f_t(b) = \frac{1}{\lambda_{el}} \quad \forall t \in T$.

We will establish uniform bidding as described in Lemma 3.2 as the optimal bidding strategy for settings of identical uniform distributions (Theorem 3.2). To do so, we first show that uniform bidding is the only non-Λ_s bidding strategy that satisfies (3.4) (Lemma 3.3) and that the uniform bid is a local minimum(Lemma 3.4). Finally, we compare our results to a strategy suggested by Mohsenian-Rad (2014) and discuss the difference in modelling assumptions we make.

Lemma 3.3. Under Assumption 3.4 and $|T| > 2$, any bid B that satisfies condition (3.4) is either a uniform bid across all auctions or $B \in \Lambda_s$.

Lemma 3.4. Under Assumption 3.4 and $|T| > 2$, the uniform bid B_u as defined in Lemma 3.2 is a local minimum of the cost function or is in Λ_s.

Regarding a discussion of the edge case where $n = |T| = 2$, see Appendix A.1.4.

Theorem 3.2. *Under the assumption of identical uniform distributions, uniform bidding is a global optimum, i.e. for any number of auctions $n = |T| \geq 2$ and demand $s \in [1, n-1]$, the following holds*

$$Cost_u(B^u_{n,s}|n,s) \leq Cost_u(B|n,s) \quad \forall B \in [0, \lambda_{el}]^n,$$

where $B^u_{n,s}$ is the unique uniform Bid as described in Lemma 3.2 and $Cost_u(B|n,s)$ is the cost of submitting bid B in n auctions with uniform price distributions when requiring s units.

In Theorem 3.2, we only establish optimality of uniform bidding, but say nothing yet about how the demanded quantity s or the number of available auctions affect the size of the cost difference..

Figure 3.1 shows numerical results of uniform bidding for loads of different demand s. The graph displays the cost difference $Cost_u(B_\Lambda|n,s) - Cost_u(B^u_{n,s}|n,s)$ between Λ_s-bidding and optimal uniform bidding $B^u_{n,s}$. For all numerical results we assume $\lambda_{el} = 1$. As we can see, uniform bidding outperforms Λ_s-bidding for all cases. As the number of auctions increases, uniform bidding further improves. This positive effect of adding auctions appears to be stronger for loads that have higher demand, as indicated by the steeper slope for $s = 5$ compared to $s = 1$.

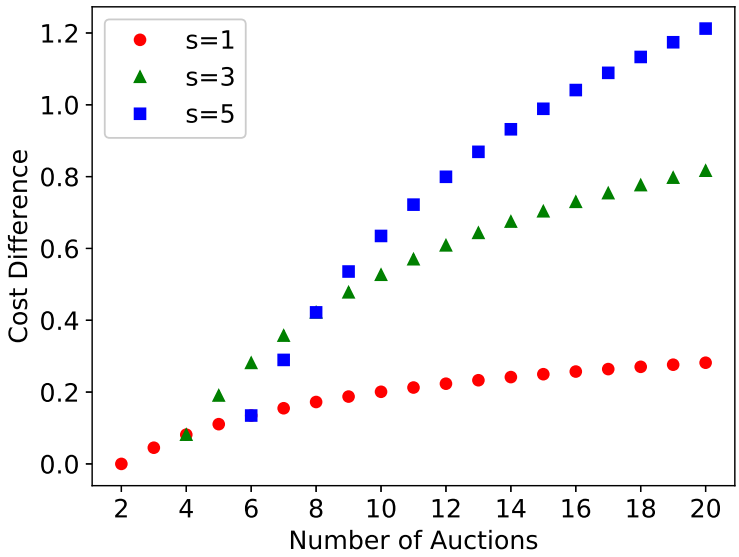


Figure 3.1: Cost Difference of uniform bidding to Λ_s-Bidding

It is noteworthy here that Theorem 3.2 and Figure 3.1 seem to directly contradict some of the results by Mohsenian-Rad (2014). In particular Theorem 3 in (Mohsenian-Rad, 2014) states that the optimal bidding strategy is to acquire all energy from a single auction with a bid equal the expected cost on the intra-day market. In section 3.4.1 we associated this bidding strategy with Λ_s-bidding in our setting. To resolve this seeming contradiction, we need to look towards the assumptions made by Mohsenian-Rad. In particular the constraint of equation (6) in (Mohsenian-Rad, 2014) restricts the bidding strategy to those that never run the risk of obtaining more energy than needed, while we assume free disposal of additional energy, see Assumption 3.2. Our uniform bidding strategy runs the risk of – in the worst case – winning all auctions. Furthermore, Figure 3.2 shows the expected number of units the uniform bidding strategy wins. Figure 3.2 indicates that we not only run the risk of winning more units than necessary but that for most cases we expect to obtain more units than demanded. This over-consumption can be explained by the backup cost being weakly larger than market clearing prices. Therefore obtaining too few units causes a relatively large penalty of λ_{el}, while an agent winning too many auctions pays an often much lower market clearing price.

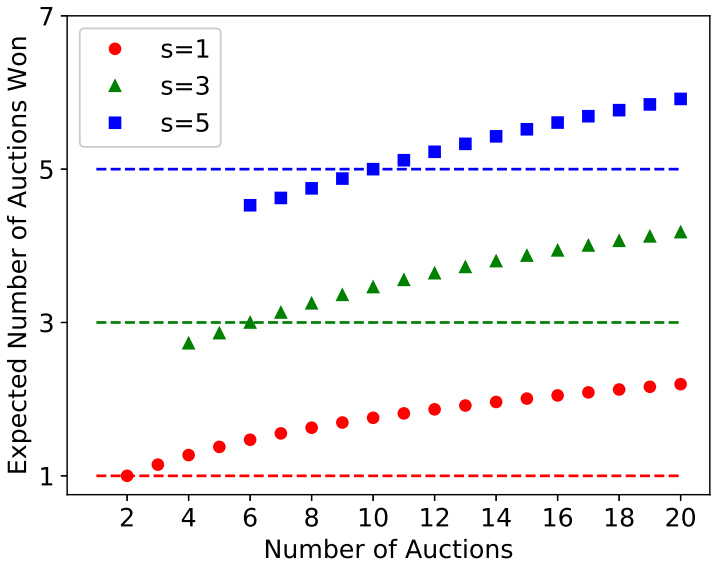


Figure 3.2: Expected number of units won

3.4.3 Non-Identical Auctions

We consider the problem of acquiring s units from a set of non-identical auctions T, which follow truncated normal distributions. We assume that the clearing prices for auctions $t \in T$ are distributed according to $\mathcal{N}_{0,\lambda_{el}}(\mu_t, \sigma^2)$, i.e. the un-truncated distributions differ in mean but not standard deviation. Let the set of mean values be $\boldsymbol{\mu} = \{\mu_t | t \in T\}$. Without loss of generality, we assume that $\mu_i \leq \mu_j$ for $i < j$.

Given this setting we can rewrite the condition in (3.4) as a root finding problem of a system of $|T|$ non-linear equations:

$$b_k - Prob(x < s | T \setminus \{k\}, B_{-k})\lambda_{el} = 0 \quad k \in T$$

For finding a solution, we use the Minpack's (Cowell, 1984) 'hybrj' method, which is an adaptation of Powell hybrid method. This method requires us to provide a Jacobian. The diagonal elements of the Jacobian are $J_{k,k} = 1$ and the off-diagonal elements are $J_{k,j} = f_j(b_j)Prob(x = s-1 | T \setminus \{j,k\})\lambda_{el}$. The entries for the Jacobian only differ by a factor of $f_k(b_k)$ from the second derivative and the reader is therefore referred to Appendix A.1.3 for a derivation.

This approach exhibits two challenges. First, evaluating the function and Jacobian requires repeated calculation of a term of the form $Prob(x = i|T \setminus \{k\}, B_{-k})$, which if done explicitly is composed of $\binom{|T|-1}{i}$ terms. This makes finding a solution for larger systems computationally infeasible. The second challenge is convergence to non-trivial solutions. Since any bid $B_s \in \Lambda_s$, which can be constructed without the need for any solver, satisfies the system of equations, a badly initialized solver often converges to this solution.

The first challenge is solved by recursion. Let D be a set of auctions and let B_D be a corresponding bid vector. Let $d = |D|$ be the corresponding auction for bid $b_{|D|}$ – the last bid in bid vector B_D. We calculate $Prob(x = i|D, B_D)$ recursively as follows

$$\begin{aligned} Prob(x = i|D,B_D) = F_d(b_d)Prob(x = i - 1|D\backslash\{d\}, B_{D,-d}) \\ + [1 - F_d(b_d)]\, Prob(x = i|D\backslash\{d\}, B_{D,-d}), \end{aligned}$$

where $B_{D,-d}$ is bid vector B_D with the last entry removed. Note that F_d and $[1 - F_d(b_d)]$ are the probabilities of winning and losing auction d, respectively. Going through the recursion, we can calculate $Prob(x = i|D, B_D)$ by calculating the entries in a table of size $|D| \times i$.

The second problem, creating an appropriate initial guess, is addressed by iteratively solving approximate versions of the problem. To do so we partition $T = \uplus G$. For each group $g \in G$, a group auction price distribution $\mathcal{N}(\bar{\mu}_g, \sigma^2)$, where $\bar{\mu}_g$ is the average mean of the auctions in g, is defined.

To create the approximate problem, we replace the clearing price distribution $f_t(\tau)$ for each auction $t \in T$ with the corresponding group auction price distribution $\mathcal{N}(\bar{\mu}_g, \sigma^2)$. The partitioning starts as a singleton $G = \{T\}$, i.e. all individual price distributions are replaced with the same group price distribution. Therefore, the first approximate problem is an identical auction setting for which we find the uniform bid as in Section 3.4.2 by interval halving. The grouping is then refined and a new approximate problem is constructed. The solution to the prior solution is sorted (to ensure that higher bids go to cheaper auctions) and used as the initial guess for the new approximate problem. We refine the grouping of auctions until every group consists of a single auction at which point we reached the original problem.

The refinement of the partitioning can be done in one of two ways. Either, we split every group in two approximately equal sized subgroups. This means that we require $log(|T|)$ iterations of refinement to reach the original problem and therefore add a factor of $log(|T|)$ to the computational complexity. Alternatively, we increase the number of groups by one every time we refine the grouping, adding a factor of $|T|$ to the computational complexity. When increasing the number of groups, we use k-means clustering on the set of mean values $\boldsymbol{\mu}$ to find the grouping.

Results and Discussion

We consider the problem of acquiring a certain number of units from a set of non-identical auctions, where the clearing prices for auctions $t \in T$ are distributed according to $\mathcal{N}_{0,\lambda_{el}}(\mu_t, \sigma^2)$. Let μ_t be drawn from a uniform distribution on the interval $[0, \lambda_{el}]$.

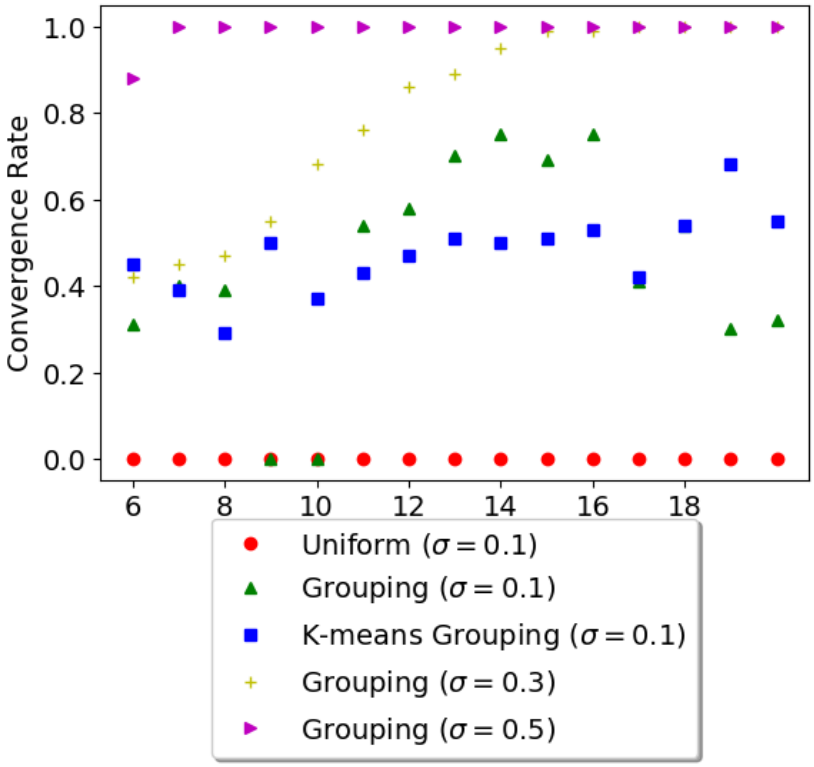


Figure 3.3: Convergence-Rate to non-trivial solutions over the number of auctions

We first look at the convergence rate. We consider the solver to have successfully converged, when it finds a non-trivial solution, i.e. a bid $B \notin \Lambda_s$ that satisfies (3.4). Figure 3.3 compares the convergence success rate of different initializations. The demand of the agent is set to $s = 5$. When initialized with a uniform bid, the solver virtually never converges. Note that we do not mean that the solver did not converge but rather that the solver converged to a Λ_s-bidding strategy, which can be found easily without the

need of a solver. When increasing the number of groups by one (k-means Grouping) the convergence rate is similar to when groups are split in half at every iteration (Grouping). This more refined increase of groups appears to add little to the chance of converging to a non-trivial solution, while adding a factor of $|T|$ instead of $log(|T|)$ to the computational complexity. Increasing the standard deviation σ from 0.1 to 0.3 and 0.5 improve the chances of the solver converging to a non-trivial solution. Similarly, as the number of auctions increases the success rate of the solver increases. For large number of auctions and for high uncertainty settings, i.e. $\sigma = 0.3$ or $\sigma = 0.5$, we reach a near 100% success rate.

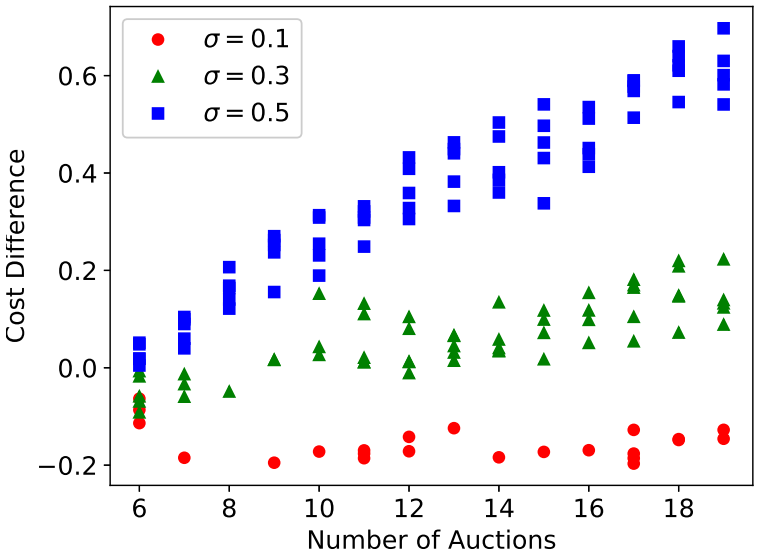


Figure 3.4: Cost Difference of Λ_S to the proposed approach

Next, we compare the performance of our approach to the Λ_s-bidding, which we associated in Section 3.4.1 to the work of Mohsenian-Rad. Figure 3.4 shows the cost difference between Λ_s-bidding and the solution provided by our solver approach. For visual clarity, results where the solver did not converge to $B \notin \Lambda_s$ were omitted. While, under low uncertainty, $\sigma = 0.1$, Λ_S is the better choice, as the uncertainty in the prices grows ($\sigma = 0.3$, $\sigma = 0.5$) our approach appears to improve. The solution provided by us further improves as the number of auctions increases.

3.5 Conclusion and Future Work

This chapter considered the problem of an interruptible time-shiftable electric load in acquiring multiple units of electricity from a set of parallel auctions. We derived a condition for optimality and established a direct comparison to Mohsenian-Rad (2014). First, for the identical auction setting, we show that uniform bidding is an approach that satisfies our necessary optimality condition. For identical uniform price distributions, we show that uniform bidding, i.e. participating in all rather than a single auction, is optimal. In comparing our results to Mohsenian-Rad (2014), we show that assuming free disposal yields significantly reduced cost.

Second, for non-identical auctions, we use approximate problem formulations to guide a non-linear-solver and provide a dynamic programming approach to make solving the set of non-linear equations computationally feasible. Numerical results show that under high price uncertainty our approach again outperforms the literature.

Regarding future work, we would like to extend this line of research to include budget constraints of the agent. A second line of extension is to not only assume free disposal, but to consider the value that is added by providing excess energy to the district heating network.

4 Fair Pricing for Time-Flexible Smart Energy Markets

Preface The adoption of new market mechanisms, vital to the better integration of flexible assets, depends on the fairness and non-discrimination of the pricing rules. We consider a market setting with time-flexible unit buyers and sellers, that additionally submit their availability in time. The temporal flexibility of the agents allows the market operator to choose between differing schedules that are, however, equivalent with regard to social welfare. This arbitrary choice of schedule introduces the potential for price discrimination.

We first demonstrate that, in contrast to other settings, non-discriminatory prices are not trivially defined in market settings with time-flexible agents. Second, we propose an extension of the concept of non-discriminatory pricing to time-flexible systems, as prices that are consistent over equivalent outcomes, i.e., prices that are independent of the particular choice of schedule. Third, for our setting, we provide a mathematical representation that allows for such pricing to be calculated quickly. Forth, we show that such prices can exist without compromising individual rationality, and provide a suitable mechanism. Lastly, we put our work in context by relating it to broader concepts of fairness from economic psychology and comparing it to envy-freeness as well as relaxations thereof. [1]

4.1 Introduction

The climate crises is one of the defining challenges of the 21st century. To curb the exhaust of greenhouse gases, massive investments in renewable energies have been made (Waseem et al., 2019; Zeng et al., 2017) and the electrification of the transportation sector is well on its way (Shaffer, 2019; Waseem et al., 2019; Zheng et al., 2020). However, both these trends are putting enormous stress on the electricity grid (Koltsaklis and Dagoumas, 2018). To handle this burden without unrealistic investment in grid reinforcement, research has

[1] This chapter is based on publication (Saur et al., 2023)

pushed towards a tighter integration of flexible assets into electricity systems (Algarvio et al., 2025; Aurangzeb et al., 2021; Eltohamy et al., 2026; Esmat et al., 2018; Fridgen et al., 2014). To this end electricity systems have seen the integration of many homogeneous flexible in time devices. These devices come in the form of fleets of electric vehicles (Dutta et al., 2026; Sepetanc and Pandzic, 2021; Zheng et al., 2023), groups of thermostatically controlled loads (Radaideh and Ajjarapu, 2016) or household batteries (Adika and Wang, 2014) and are often aggregated or coordinated in some way.

One of the most promising approaches of integration is the introduction of new market mechanisms (Meir et al., 2017; Palensky and Dietrich, 2011; Shen et al., 2015; Zhongjing and Long, 2013). Early approaches of market based integration of flexibility were based on offering time dependent dynamic prices(Limmer, 2019). However, such price signals are not sufficient to efficiently coordinate individual agents (Bitar and Low, 2012). It was, therefore, argued for mechanisms in which agents instead of deciding for themselves when to operate simply provide their flexibility information to a market operator (Bitar and Low, 2012), which then schedules them.

Depending on the exact market definition flexibility information is communicated in different ways. In duration-differentiated energy services agents report a time window in the form of arrival and deadline times as well as the duration of their energy demand. Within the deadline the agents are indifferent to the exact delivery time (Chen et al., 2015; Mo et al., 2016). Other works such as (Esmat et al., 2018; Hayakawa et al., 2015) allow the load of the agents to be freely rescheduled within a defined time window. In (Zaidi and Hong, 2018) flexibility information is expressed in the form of XOR bids. As the goal of our work is to discuss a fundamental property, we will consider a stylized market in which agents report a time window for which they demand/supply a single time step of energy. This can be viewed as a kind of XOR bid and differs from the duration-differentiated energy services only in that we additionally restrict the duration to a single time step.

Whatever the particular market mechanism, the effectiveness of such new markets hinges on a critical mass of participants: acceptance of smart energy systems is vital to their adoption and success (Ellabban and Abu-Rub, 2016). *Fairness* is one of the key aspects that might hinder the social acceptance and therefore implementation of smart energy systems (Milchram et al., 2018; Wildt et al., 2019). Further, psychology surveys have shown that a perception of fairness can improve the acceptance of autonomous smart energy systems (Guerreiro et al., 2015).

A perception of fairness is often tied to how agents in the market are priced (Maxwell, 2002). Literature on fair pricing in this field has mostly revolved around fairness in real time (Mamounakis et al., 2019) or time-of-

use pricing systems (Aurangzeb et al., 2021; Ma et al., 2019) with a focus on congestion management. Fair price sharing for agents that express flexibility in their demand profile has been investigated in (Perrault and Boutilier, 2015). However, few have specifically focused on flexibility in time with regard to fair pricing schemes. One exception is (Limmer and Dietrich, 2018), who consider the pricing problem of flexible-in-time electric vehicles. The authors define fairness as similar prices for electric vehicles that charge at similar time slots. However, their work relies on a quantitative measure, i.e., similarity of price, but does not consider the qualitative property of fair pricing, i.e., the same price for the same item. We explore this qualitative concept of fair pricing – *non-discriminatory pricing* – in the context of *time-flexible markets*, by studying a stylized market in which unit buyers and sellers can define their availability in the form of a time window.

This concept of non-discriminatory or anonymous pricing has been studied in the mechanism design literature. In (Sandholm and Suri, 2001) non-discriminatory and discriminatory pricing have been compared with regard to market clear-ability. The question under which conditions assigning prices to items or bundles leads to prices that support an optimal allocation has been explored in (Conen and Sandholm, 2004). Approximation guarantees of posted prices have been established in (Feldman et al., 2015), while in (Alaei et al., 2015) guarantees for revenue maximization of anonymous pricing have been derived.

To summarize the engineering issue: we want a fair market setting in which time-shiftable consumers and producers, such as electric vehicles, industrial processes or batteries, can report their flexibility directly via their bids. This allows flexible assets to provide their flexiblity without the need for sophisticated bidding strategies. The market mechanism takes care of utilizing the reported flexiblity to maximize social welfare.

The related conceptual challenge this work addresses, arises when applying common concepts of fairness to such a flexibility market. We will show that the concept of price discrimination is ill defined in flexibility markets that allow for different but equivalent outcomes. We then provide consistent pricing over equivalent outcomes as a well defined fairness concept for such flexibility markets and show that this concept can be achieved without sacrificing other market design goals.

The remainder of this chapter is structured as follows. We first define the market setting in Section 4.2 and outline the desired pricing properties in Section 4.3. Section 4.4 contains three parts. First, we take a closer look at price discrimination in time and establish our concept of consistent pricing, see Subsection 4.4.1. In Subsections 4.4.2 and 4.4.3, we show how to obtain such prices and that they do not conflict with Individual Rationality. Section 4.5

provides a suitable mechanism. Finally in Section 4.6, we discuss related fairness concepts.

We contribute to the literature in the following ways:

1. We show that non-discriminatory pricing in a time-flexible setting is not trivial to define and provide the concept of consistent pricing over equivalent outcomes as a suitable property.
2. We show that in our setting consistent and Individually Rational prices always exist for efficient allocations (Theorem 4.1).
3. We provide for our setting a tractable, polynomial-time mechanism that is Budget Balanced, Individually Rational and Consistent (Theorem 4.2).
4. We put our work in context by relating it to fairness concepts in economic psychology and comparing it to envy-freeness (Section 4.6).

4.2 Market Setting

We consider a set of flexible unit buyer and seller agents $N = B \cup S$. Agents $n \in N$ have a time window τ_n in which they can be scheduled. Every agent operates for a single time step. Every buyer $b \in B$ assigns valuation ν_b to being allocated, while every seller $s \in S$ incurs cost c_s when allocated. The market collects reports Θ_n from agents $n \in N$. A buyer b reports $\Theta_b = (\tau_b, \nu_b)$, which contains their time window τ_b and their valuation ν_b. Similarly, a seller s reports $\Theta_s = (\tau_s, c_s)$. We define a buyer–seller pair $\{b, s\}$ and associate it with an operating window $\tau_{b,s} = \tau_b \cap \tau_s$. We say a buyer–seller pair $\{b, s\}$ is feasible if the buyer and seller time windows overlap, i.e. $\tau_b \cap \tau_s \neq \varnothing$. Further, we say that two buyer–seller pairs, $\{b_1, s_1\}$ and $\{b_2, s_2\}$ are in conflict if one agent appears in both pairs, i.e., $b_1 = b_2 \vee s_1 = s_2$.

A matching M is a set of buyer–seller pairs. A matching M is feasible if all buyer–seller pairs in M are feasible and no pairs are conflicting. We refer to an allocation as the set $N^M = B^M \cup S^M$ of buyers and sellers that occur in a buyer–seller pair of the matching M. Note that different matchings can result in the same allocation.

Social welfare $SW(M)$ of a matching is defined as the sum of gains from trade. An allocation or one of its underlying matchings is *efficient* if it maximizes Social Welfare.

$$\max_{M} \left\{ \overbrace{\sum_{\{b,s\}\in M} (\nu_b - c_s)}^{SW} \right\} \qquad (4.1)$$
$$\text{s.t.} \quad \tau_b \cap \tau_s \neq \varnothing \quad \forall \{b,s\} \in M$$

An allocated buyer b incurs a price $p(b)$, while an allocated seller is paid a price $p(s)$. Non allocated agents are not priced. We assume that agents possess quasi-linear utility functions, i.e. when allocated their utility is equal to the absolute difference between their valuation and their assigned price and zero otherwise.

The allocation problem, see equation (4.1), is a constrained assignment problem where the value of seller s servicing buyer b is $(\nu_b - c_s)$ if they form a feasible buyer–seller pair, i.e., $\tau_b \cap \tau_s \neq \varnothing$. This problem can be solved in polynomial time using the Hungarian method (Kuhn, 1955).

4.3 Desired Properties of Pricing

Besides a matching M, we need a pricing function p that assigns prices to agents. As the designer of the market, we would like the mechanism to be financially sustainable, i.e., the mechanism is budget balanced.

Definition 4.1 (Budget Balance)**.** A mechanism is Budget Balance (BB) if the payments received are equal to the payments given out.

$$\sum_{b\in B} p(b) - \sum_{s\in S} p(s) = 0$$

The point of a market is to coordinate the actions of self-interested agents for their mutual benefit. This crucially depends on agents actually participating. To encourage participation, the market needs to be beneficial and fair to the agents. To ensure that agents see the market as beneficial, we want no agent to end up with negative utility, i.e. the mechanism is *Individually Rational* (IR).

Definition 4.2 (Individual Rationality)**.** A mechanism is Individual Rationality (IR) if an honestly participating agent always receives at least zero utility.

To ensure, on the other hand, that agents do not see the market as arbitrary or favouring one agent over another, we want our market to assign non-discriminatory prices.

Definition 4.3 (Non-discriminatory Prices)**.**
The "same" price assigned to the "same" item.

This is a rather vague description of non-discriminatory pricing. The first challenge in our time-flexible setting is to properly define what this informal definition of non-discriminatory pricing means in exact mathematical terms.

4.4 Price Discrimination in Time

Here, we will elaborate on our informal definition of non-discriminatory prices. To do so, we need to determine when price differences are justified and therefore answer the following question: What does it mean for two items to be the same? When both buyers and sellers are flexible in time is energy exchanged at 3 pm the same as energy exchanged at 5 pm? Can we establish a fair, non-arbitrary way of distinguishing when these two time slots should be priced equally and when they should be priced differently.

In the following, we will first go into detail on the concept of non-discriminatory pricing in systems with flexible in time agents in Subsection 4.4.1 , then provide a mathematical representation that will allow us to obtain such prices more quickly in Subsection 4.4.2 and finally show that our concept does not conflict with individual rationality in Subsection 4.4.3.

4.4.1 Consistent Pricing over Equivalent Outcomes

The key insight of this section is that due to the time flexibility of the agents it makes little sense to differentiate items based on time itself. If an agent provides a time window and is arbitrarily scheduled at a point with worse pricing, it will perceive this as unfair. A more reasonable approach is to state non-discriminatory prices as the "same price for the same *outcome*". Here, we will establish what this idea specifically means by going through several examples and pricing attempts.

Given a matching M, the simplest approach is to price agents by the pair. For any allocated agent the price is set half way between its valuation/cost and the valuation/cost of the agent they are paired with in M. We call this pricing approach *simple pricing*.

Definition 4.4 (Simple Pricing)**.** Simple pricing assigns prices by the pair: For every buyer–seller pair in the matching M, the price is set half way between the buyer and seller, i.e.,

$$p_M(n) = \frac{\nu_b + c_s}{2} \quad n \in \{b, s\} \in M$$

We begin with the simplest of cases when there are two buyers b_1, b_2 and two sellers s_1, s_2, all of whom possess the same availability time window, see Figure 4.1. Let the matching be

$$M_1 = \{\{b_1, s_1\}, \{b_2, s_2\}\}\,.$$

When using simple pricing, we obtain prices $p(1)$ and $p(2)$ for b_1, s_1 and b_2, s_2, respectively. Given that all pairs exchange the same units over the same time intervals, prices $p(1)$ and $p(2)$ appear to be discriminatory.

But why specifically do these prices feel discriminatory? Price discrimination is about assigning different prices to different agents for the same item. It is not yet clear that the item traded between $\{b_1, s_1\}$ is the *same* as the one between $\{b_2, s_2\}$. To clearly expose such price differences, in our example, let us consider an alternative matching that results in the same allocation. Let this alternative matching be

$$M_2 = \{\{b_1, s_2\}, \{b_2, s_1\}\}\,.$$

Now consider the pair $\{b_1, s_2\}$. Buyer b_1 buys one unit of energy from seller s_2. However, under the original pricing, buyer b_1 pays a lower price, $p(1)$, compared to the price $p(2)$ that seller s_2 receives. We end up with an exchange of energy, where the buyer and the seller are assigned different prices for exactly the same item, i.e., the prices are discriminatory. Note, that matchings M_2 and M_1 are essentially the same outcome. Both allocate the same agents and therefore result in the same overall SW. The only difference emerges once simple pricing is applied and an arbitrary choice between two equivalent matchings results in some agents being advantaged and some agents being disadvantaged.

Therefore, a better approach to thinking about non-discrimination in pricing is to consider prices that are the same over equivalent outcomes. To explore this idea we introduce the concept of equivalent matchings.

Definition 4.5 (Equivalent Matchings)**.** We call two matchings M_1 and M_2 equivalent if they result in the same allocation, $N^{M_1} = N^{M_2}$. Further, let $\mathbb{M}_M$ be the set of matchings that are equivalent to M; this includes $M \in \mathbb{M}_M$.

Note that any two equivalent matchings result in the same Social Welfare.

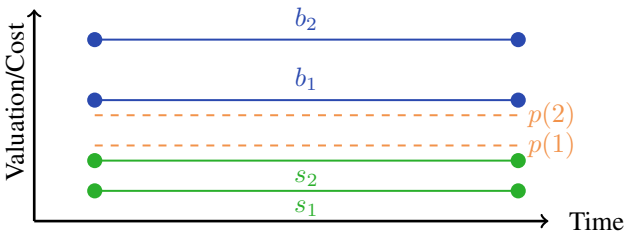
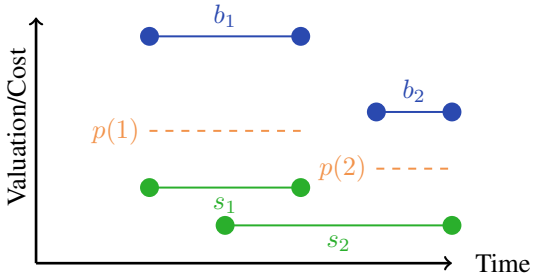
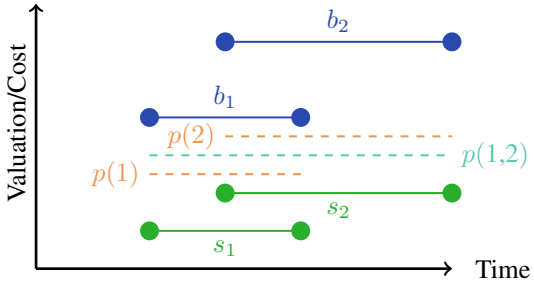


Figure 4.1: Time Intervals and valuation/cost of the Agents



(a) Justified Price Difference



(b) Price-Discrimination

Figure 4.2: Price Difference in Time under Simple Pricing

To further explore this idea, we contrast two examples depicted in Figures 4.2a and 4.2b. Both examples use matching $M_1 = \{\{b_1, s_1\}, \{b_2, s_2\}\}$. Prices $p(1)$ and $p(2)$ for the first and second pair are based on simple pricing and also indicate the pairs' operating window. The visually striking difference between the two cases is that the operating windows of the two pairs $\{b_1, s_1\}$, $\{b_2, s_2\}$ overlap in case 4.2b but not in case 4.2a. The overlapping operating windows in case 4.2b allows us to construct an equivalent matching $M_2 = \{\{b_1, s_2\}, \{b_2, s_1\}\}$ in which the pairs are assigned incongruous prices: b_1 pays a lower price than s_2 receives, i.e. $p(1) < p(2)$. While this pair, $\{b_1, s_2\}$, is also possible in case 4.2a, it comes at the cost of social welfare because $\{b_2, s_1\}$ is not a feasible pair. Case 4.2a, unlike case 4.2b, has no equivalent matching that exposes incongruous prices within a pair.

These examples naturally suggest a different but as we will show still discriminatory pricing approach: pricing based on overlapping operating windows. For this, we note that for any two buyer–seller pairs $\{b_1, s_1\}$, $\{b_2, s_2\}$ that overlap in their operating window, $\tau_{b_1,s_1} \cap \tau_{b_2,s_2} \neq \varnothing$, we can create an

equivalent matching with pairs $\{b_1, s_2\}$, $\{b_2, s_1\}$. Any buyer–seller pair that overlaps constitutes an alternative trading partner. Therefore, instead of only considering the pair itself we will additionally consider all buyer–seller pairs that overlap with its operating window.

Definition 4.6 (Basic Pricing)**.** Basic pricing considers the buyer–seller pair $\{b, s\} \in M$ and also all buyer–seller pairs that overlap in their operating window. For this let

$$S_{b,s} = \left\{ s^\dagger \middle| \tau_{b,s} \cap \tau_{b^\dagger, s^\dagger} \neq \varnothing \ \wedge \{b^\dagger, s^\dagger\} \in M \right\}$$

be the set of sellers that can operate at a common time slot as $\{b, s\}$ given a particular matching M. Similarly, $B_{b,s}$ are such buyers. Then,

$$p_M(n) = \frac{\min_{b^\dagger \in B_{b,s}} \nu_{b^\dagger} + \max_{s^\dagger \in S_{b,s}} c_{s^\dagger}}{2} \quad n \in \{b, s\}.$$

Reconsidering the cases depicted in Figure 4.2 using basic pricing, we see no difference when the price differences were justified. However, where simple pricing created discriminatory prices $p(1)$ and $p(2)$ for the agents, basic pricing generates one non-discriminatory price $p(1, 2)$ for all agents.

While basic pricing might appear to be non-discriminatory, we will now show that it can still lead to price discrimination. In Figure 4.3, we see an example with three buyer–seller pairs and prices based on basic pricing using matching

$$M_1 = \{\{b_1, s_1\}, \{b_2, s_2\}, \{b_3, s_3\}\}.$$

The first two buyer–seller pairs $\{b_1, s_1\}, \{b_2, s_2\}$ overlap in their operating window and are therefore considered together under basic pricing. However, neither $\{b_1, s_1\}$ nor $\{b_2, s_2\}$ overlap with pair $\{b_3, s_3\}$ in their operating window and therefore $\{b_3, s_3\}$ is priced separately. But there also exists the equivalent matching

$$M_2 = \{\{b_1, s_3\}, \{b_2, s_1\}, \{b_3, s_2\}\}.$$

In this matching pair $\{b_3, s_2\}$ is priced incongruously. Seller s_2 is paid price $p(1, 2)$ for providing energy to buyer b_3, who pays a lower price $p(3)$, see Figure 4.3. Therefore, basic pricing yields discriminatory prices in this example.

Truly non-discriminatory pricing would consider *all* pairs in the example in Figure 4.3, in order to price any one of the pairs. As we can see, we need to consider more than simply overlapping operating windows. We need to consider all possible operating windows over all equivalent matchings.

To achieve non-discriminatory prices, we will introduce prices that are consistent over all equivalent matchings. For this we define the *feasible operating*

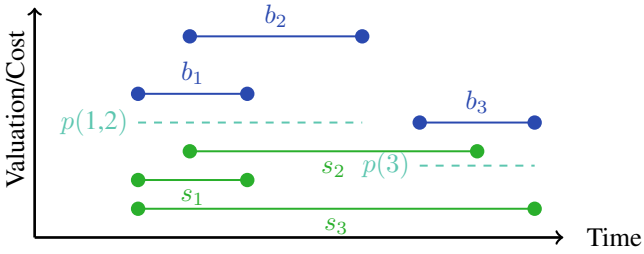


Figure 4.3: Price Difference in Time

range of an agent as all time steps at which an agent can operate considering all equivalent matchings. Note that by definition if and only if a buyer and seller overlap in their feasible operating ranges, there exists an alternative matching that pairs them. In the example depicted in Figure 4.3, there are equivalent matchings that pair seller s_3 with b_1, b_2 or b_3. The feasible operating range of agent s_3 is therefore the union of the resulting three operating windows:

$$\tau_{b_1,s_3} \cup \tau_{b_2,s_3} \cup \tau_{b_3,s_3}$$

Note, we distinguish three different types of time intervals: τ_i denotes the time window of agent i itself, while $\tau_{i,j}$ is the operating window of agent i when paired with agent j. Finally we define the feasible operating range τ_i^M as follows:

Definition 4.7 (Feasible Operating Range)**.** The feasible operating range τ_i^M is the union of operating windows over all M-equivalent matchings:

$$\tau_i^M = \bigcup_{\{i,x\} \in \bigcup \mathbb{M}_M} \tau_{i,x}.$$

Note that, thus τ_i^M does not depend on a particular matching M but only on the resulting allocation.

Definition 4.8 (Consistent Prices)**.** A pricing function p is consistent given a matching M, if it assigns the same price to any agents that overlap in their feasible operating ranges, i.e.,

$$\tau_i^M \cap \tau_j^M \neq \varnothing \implies p(i) = p(j) \quad \forall i, j \in N$$

As consistent prices are defined using all equivalent matchings, they make it impossible to re-schedule the agents in such a way that two trading partners are assigned different prices. The arbitrary choice the market operator makes

about equivalent matchings no longer affects the pricing and therefore utility of the agents. These prices are, therefore, non-discriminatory.

4.4.2 Obtaining Consistent Pricing

The usefulness of the link between non-discrimination and consistent prices is limited by the fact that calculating all equivalent matchings is intractable. Instead of calculating all equivalent matchings, we introduce a labelled directed graph $H_M = (V, E)$ that can express all M-equivalent matchings (Lemma 4.1& 4.2) and show that a consistent pricing function assigns the same price to all agents represented by a strongly connected component in H_M (Lemma 4.4), which can be obtained in polynomial time.

Based on the matching M, $H_M = (V, E)$ is a labelled directed graph that represents every pair in M as a node and every potential pair as an edge.

Definition 4.9 ($H_M = (V, E)$)**.** Let M be a matching, then $H_M = (V, E)$ is a labelled directed graph with as many nodes as there are pairs in the matching M, i.e., $|M| = |V|$. The nodes $v_i \in V$ are labelled with the buyer–seller pairs in the matching M. W.l.o.g., we assume that the node v_i is labelled with the buyer–seller pair $\{b_i, s_i\}$.

Let there be a directed edge $(v_i, v_j) \in E$ whenever the buyer at outgoing node v_i and seller at incoming node v_j overlap in their operating windows, i.e., $\tau_{b_i} \cap \tau_{s_j} \neq \varnothing$. The edge $(v_i, v_j) \in E$ is labelled by the alternative feasible pair $\{b_i, s_j\}$.

We can use the labels assigned to the nodes and edges in H_M to describe new matchings M^*. To do so we consider a subset of nodes $V^* \subseteq V$ (blue circles in Figure 4.4) and a subset of directed edges $E^* \subseteq E$ (green edges in Figure 4.4).

Definition 4.10 (Translation $(V^*, E^*) \mapsto M^*$)**.** Let M be a matching, with associated graph $H_M = (V, E)$ by Definition 4.9, and let $V^* \subseteq V$, $E^* \subseteq E$. The labels assigned to subsets (V^*, E^*) describe a set of buyer–seller pairs. We will refer to this set of buyer-seller pairs as the matching M^*. For short we will say that (V^*, E^*) forms M^*. (Note: (V^*, E^*) is a set of pairs and should not be seen as a sub-graph of H_M and while M^* is a matching it is not necessarily feasible.)

In the example shown in Figure 4.4,

$$V^* = \{v_1, v_5\} \quad \text{and} \quad E^* = \{(v_2, v_4), (v_4, v_6)\}.$$

Pairs $\{\{b_1, s_1\}, \{b_5, s_5\}\}$ are represented by nodes in H_M and are therefore both in the original matching M as well as the newly formed matching M^*. Edges $E^* = \{(v_2, v_4), (v_4, v_6)\}$ are labelled with the pairs $\{b_2, s_4\}$ and

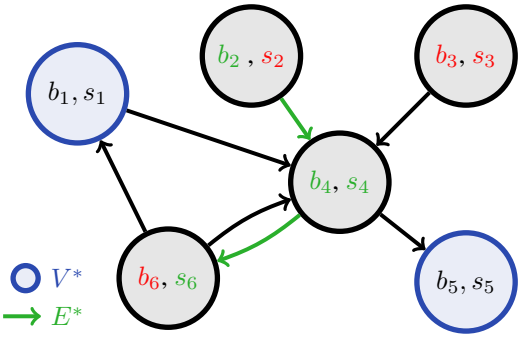


Figure 4.4: Create new matching from V^* and E^*

$\{b_4, s_6\}$. By Definition 4.10 these pairs are in M^* but not in this form in M. Agents $\{s_2, b_3, s_3, b_6\}$ are allocated in M but not in M^*.

Note 4.1. (V^*, E^*) with $V^* \subseteq V, E^* \subseteq E$ form a feasible matching M^* if and only if no edge $(v_i, v_j) \in E^*$ has an endpoint in V^*, i.e., $v_i, v_j \notin V^*$, and no node $v_k \in V \setminus V^*$ is source/destination of two or more distinct edges in E^*.

Proof. To see why Note 4.1 is true, we need to recall that all pairs described by the labels on V and E are feasible. Therefore matching M^*, formed by (V^*, E^*), is only infeasible if there is a buyer or seller assigned to more than one label in V^* or E^*. Note that by Definition 4.9, a buyer b_i only appears in the graph H_M in node $v_i = \{b_i, s_i\}$ and its outgoing edges, while a seller s_i only appears in node $v_i = \{b_i, s_i\}$ and its incoming edges. Therefore, buyer b_i appears in two labels of M^*, only if two or more outgoing edges of v_i are in E^* or if there is one or more outgoing edges of v_i in E^* and v_i is in V^*. Similar holds for sellers and incoming edges. □

To see how this graph representation of the allocation is useful, we need a more abstract representation of the construction process. We will move from considering E^* as a set of edges to thinking of E^* as a set of paths and ultimately a set of cycles.

Definition 4.11 (Path)**.** A path p is a series of nodes, where there is an edge between any two consecutive nodes: $p = (v_k, v_j, ..., v_l)$. We will say that these nodes are the nodes that the path *visits*.

We say a path is maximal if no set of edges can be added such that the result is still a simple path. A path is simple if every node is visited at most once.

Definition 4.12 (Cycle)**.** A cycle z is a non-empty path that ends at the same node at which it starts, i.e., the vertex representation is as follows: $(v_k, v_j, ..., v_k)$

We say a cycle is simple, if any node except the endpoint is visited at most once and the end point is visited exactly twice. We say two or more cycles/paths are *non-overlapping* if no node is visited by more than one cycle/path.

Using cycles in the graph H_M, we can not only describe any alternative matching but more specifically all equivalent matchings. For this, let Z be a set of non-overlapping cycles in H_M and $E_Z \subseteq E$ be the set of directed edges used by the cycles in Z. By slight abuse of notation, we say $\bigcup E_Z$ is the set of nodes in H_M visited by any cycle in E_z. All not visited nodes in H_M are then described by $V_{-Z} = V \setminus \bigcup E_Z$.

Noteworthy here are two things. First, every node in H_M is either in V_{-Z} or is an endpoint of an edge in E_Z. Second, the translation defined in Definition 4.10 for a simple cycle retains the agents in the simple cycle because it pairs every buyer in a node with the seller at the next node in the cycle. Therefore, we can establish a one-to-one correspondence between any equivalent matching $M^\dagger \in \mathbb{M}_M$ and a set of cycles Z in H_M.

Lemma 4.1 (Cycles form equivalent matchings)**.** A matching $M^\dagger$ is equivalent to M if and only if there exists a set of non-overlapping simple cycles Z in H_M such that (V_{-Z}, E_Z) forms $M^\dagger$ (Definition 4.10), where E_Z is the set of edges in the cycles Z and $V_{-Z} = V \setminus \bigcup E_Z$.

Proof. First, we assume that $M^\dagger$ is equivalent to M and will show that there exists a set of non-overlapping simple cycles Z in H_M such that (V_{-Z}, E_Z) forms $M^\dagger$.

By how the graph is defined (Definition 4.9), we know that for every agent pair in $M^\dagger \cap M$ there exists a node in H_M and for every agent pair in $M^\dagger \setminus M$ there exists a directed edge, in H_M, labelled with the respective pair. Let V_{-Z} be the nodes in H_M labelled with the pairs in $M^\dagger \cap M$ and E_Z be the edges, in H_M, labelled with the pairs $M^\dagger \setminus M$. By definition, (V_{-Z}, E_Z) forms $M^\dagger$ and we only need to show that the edges in E_Z describe simple non-overlapping cycles. Contrarily, if E_Z contained a maximal (with regard to E_Z) non-cyclical path, there would be a last node v_k in this path that had an incoming but no outgoing edge in E_Z. The lack of an outgoing edge from v_k results in buyer b_k not being in a matching pair of $M^\dagger$, contradicting our assumption that $M^\dagger$ is equivalent to M. Also, the cycles have to be simple and non-overlapping, as otherwise there would be a node with more than one incoming/outgoing edge making the allocation infeasible, see Note 4.1.

Next, we assume a set of simple non-overlapping cycles Z in H_M such that (V_{-Z}, E_Z) forms $M^\dagger$ and will show that $M^\dagger$ is equivalent to M. Labels

assigned to nodes and edges in H_M only contain buyers and sellers from M. Therefore, $M^\dagger$ cannot include any agent that is not also in M, i.e., $N^{M^\dagger} \subseteq N^M$. To, on the other hand, show that M cannot contain an agent that is not also in $M^\dagger$, $N^M \subseteq N^{M^\dagger}$, we assume the contrary. W.l.o.g., assume there is a buyer $b_k \in \{b_k, s_k\} \in M$ such that $b_k \in N^M$ but $b_k \notin N^{M^\dagger}$. Because H_M is built from M, there exists a node v_k labelled $\{b_k, s_k\}$. This node, v_k is either visited by none of the cycles in Z, $v_k \in V_{-Z}$, or by one of the cycles, $v_k \notin V_{-Z}$. If $v_k \in V_{-Z}$, then $\{b_k, s_k\} \in M^\dagger$ contradicting our assumption that $b_k \notin N^{M^\dagger}$.

If $v_k \notin V_{-Z}$, then by how (V_{-Z}, E_Z) is defined we know there is an edge in E_Z incoming to v_k. However, v_k has no outgoing edge in E_Z since $b_k \notin N^{M\dagger}$ Hence, v_k has an incoming but no outgoing edge in E_Z contradicting our assumption that E_Z is based on cycles Z. □

Having established this one-to-one correspondence, it is easy to see that for a matching M any alternative pair not in M can only be part of an equivalent matching of M if there is a simple cycle in H_M that contains an edge labelled with that alternative pair.

Lemma 4.2 (Possible buyer–seller pairs)**.** Given two buyer–seller pairs $\{b_x, s_x\} \in M$ and $\{b_y, s_y\} \in M$ $(x \neq y)$, there exists an equivalent matching $M^\dagger$ such that $\{b_x, s_y\} \in M^\dagger$, if and only if there exists a simple directed cycle in H_M that includes the directed edge from $v_x = \{b_x, s_x\}$ to $v_y = \{b_y, s_y\}$.

Proof. First, we assume there is a simple cycle $z = (v_x, v_y, ..., v_x)$ in H_M and will show that there is an equivalent matching that includes the pair $\{b_x, s_y\}$. Let $Z = \{z\}$, then by Lemma 4.1, (V_{-Z}, E_Z) forms an M-equivalent matching. The edge (v_x, v_y) with label $\{b_x, s_y\}$ is in the cycle z and by extension E_Z, making the pair $\{b_x, s_y\}$ part of the matching formed by (V_{-Z}, E_Z).

Next, we assume that there exists an M-equivalent matching $M^\dagger$ such that $\{b_x, s_y\} \in M^\dagger$ and will show that there is a directed cycle in H_M containing edge (v_x, v_y). By Lemma 4.1, we know that $M^\dagger$ is only M-equivalent if there exists a set of simple directed non-overlapping cycles Z such that (V_{-Z}, E_Z) form $M^\dagger$. The only edge labelled $\{b_x, s_y\}$ is the edge (v_x, v_y), resulting in $(v_x, v_y) \in E_Z$; so (v_x, v_y) is part of a cycle in H_M. □

Given Lemma 4.2, we can make a more general statement about the feasible operating ranges of buyers and sellers.

Lemma 4.3. A buyer b_k and seller s_l overlap in their feasible operating ranges if and only if there exists an M-equivalent matching that contains the pair $\{b_k, s_l\}$.

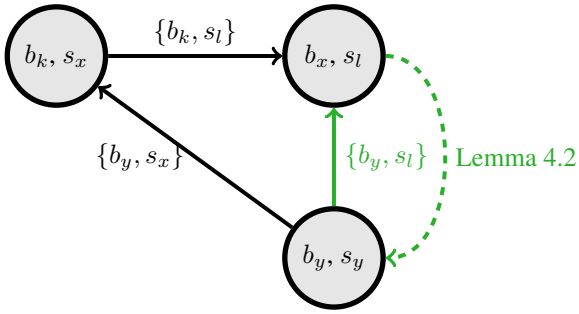


Figure 4.5: Excerpt of $H_{M'}$ with $\{b_k, s_x\}$ not on the cycle z (indicated in green)

Proof. If there exists M-equivalent matching $M^* \in \mathbb{M}_M$ such that $\{b_k, s_l\} \in M^*$, then b_k and s_l have overlapping feasible operating ranges by Definition 4.7.

If b_k and s_l have overlapping operating ranges, we can show that $\exists M^* \in \mathbb{M}_M$ such that $\{b_k, s_l\} \in M^*$ by contradiction. We assume that such an M^* does not exist. Given that b_k and s_l overlap in their operating ranges, there is a common time step $t \in \tau_{b_k} \cap \tau_{s_l}$. Given that M^* does not exist and Definition 4.7, there has to exist two M-equivalent matchings M' and M'' that enable b_k and s_l to operate at time t, i.e.

$$\begin{aligned} \exists M' \in \mathbb{M}_M \quad s.t. \quad \{b_k, s_x\} \in M' \wedge t \in \tau_{b_k, s_x} \\ \exists M'' \in \mathbb{M}_M \quad s.t. \quad \{b_y, s_l\} \in M'' \wedge t \in \tau_{b_y, s_l} \end{aligned}$$

If $\{b_y, s_l\} \in M'$, then $M' \setminus \{\{b_k, s_x\}, \{b_y, s_l\}\} \cup \{\{b_k, s_l\}, \{b_y, s_x\}\}$ is a feasible M-equivalent matching containing $\{b_k, s_l\}$ contradicting our assumption.

Otherwise, the graph $H_{M'}$ contains an edge labelled $\{b_y, s_l\}$, from node $\{b_y, s_y\}$ to node $\{b_x, s_l\}$, that is part of a simple directed cycle z in $H_{M'}$, since $\{b_y, s_l\} \in M'' \in \mathbb{M}_M$, see Lemma 4.2.

We will alter this cycle slightly to include the edge $\{b_k, s_l\}$. Because $\{b_k, s_x\} \in M'$ there exists a node $\{b_k, s_x\}$ in $H_{M'}$. This node is either part of the directed cycle z or not.

If $\{b_k, s_x\}$ is part of z, then there exists a simple directed path from $\{b_x, s_l\}$ to $\{b_k, s_x\}$, which can be extended to form a simple cycle z' by adding the edge labelled $\{b_k, s_l\}$ in $H_{M'}$. This edge exists, because $\{b_k, s_l\}$ is a feasible pair.

If $\{b_k, s_x\}$ is not part of z, see Figure 4.5, the following holds. Since b_k and s_l overlap in their feasible operating ranges, $\{b_k, s_l\}$ is feasible, which means there exists an edge labelled $\{b_k, s_l\}$ from node $\{b_k, s_x\}$ to node $\{b_x, s_l\}$ in $H_{M'}$. Since $t \in \tau_{b_y,s_l}$ and $t \in \tau_{b_k,s_x}$, see definition of M' and M'', we know that $\{b_y, s_x\}$ is feasible and therefore there is an edge from node $\{b_y, s_y\}$ to $\{b_k, s_x\}$ in $H_{M'}$.

We can construct a simple directed cycle z'' by replacing the edge labelled $\{b_y, s_l\}$ in cycle z with the edges labelled $\{b_y, s_x\}$ and $\{b_k, s_l\}$, see Figure 4.5.

By Lemma 4.1, either cycle z' or z'' represent an M-equivalent matching that contains $\{b_k, s_l\}$, contradicting our assumption and concluding the proof. □

Next, we show that the sets of agents that are assigned the same price can be expressed by strongly connected components in graph H_M. This follows because any two agents in a pair are assigned the same price and Lemma 4.2 propagates this price throughout cycles.

Definition 4.13 (Strongly Connected Component)**.** We say a subset $C \subseteq V$ is a strongly connected component of $H_M = (V, E)$, if for any two nodes in C there exists a directed cycle involving these two nodes.

Lemma 4.4 (Price sharing agents)**.** Given a matching M, a consistent pricing function p assigns the same price to two buyer–seller pairs $\{b_x, s_x\}$ and $\{b_y, s_y\}$ if and only if v_x and v_y are part of the same strongly connected component in the graph H_M.

Proof. For convenience, we say that two nodes share a price if the underlying buyer–seller pairs are assigned the same price by a consistent pricing function. Recall, that we defined consistent pricing based on overlapping feasible operating ranges, which are constructed from all equivalent matchings $\mathbb{M}_M$. From Lemma 4.3, we know that a buyer b_k and seller s_l overlap in their feasible operating ranges if and only if there exists an equivalent matching that contains the pair $\{b_k, s_l\}$. An equivalent matching containing pair $\{b_k, s_l\}$ exists if and only if the edge (v_k, v_l) is part of a simple directed cycle in H_M, see Lemma 4.2. Therefore, two nodes v_k, v_l with edge (v_k, v_l) in H_M share a price if and only if(v_k, v_l) is part of a directed simple cycle in H_M. Because the sharing of prices between nodes is transitive, any two nodes v_x, v_y share a price if and only if there is a cycle connecting the two nodes, which is equivalent to v_x, v_y being part of the same strongly connected component (Definition 4.13). □

4.4.3 Consistent and Individually Rational Prices

While consistent prices convey a sense of fairness, they would still not be acceptable to agents if they left some agents with negative utility, i.e. if the pricing broke Individual Rationality. In this section, we show that consistent and IR prices exist (Theorem 4.1) for efficient allocations. For this we only need to realise that consistent and IR prices are only impossible if there exists a seller that, because of consistency, has to be assigned the same price as a buyer that has a lower valuation. However, removing such two agents would lead to an improved allocation, indicating that the original allocation was not efficient.

Theorem 4.1 (Existence of consistent, IR Prices)**.** *Given a SW-maximizing matching M there exist consistent prices that are Individually Rational.*

Proof. Contrary to the statement, we assume that there exists a social welfare maximizing matching M for which no consistent prices are IR. We will come to a contradiction by constructing a new feasible matching with higher social welfare.

By Lemma 4.4, the different subsets of agents that are assigned the same price can be represented by strongly connected components in H_M. Our assumption that there cannot exist consistent and IR prices for matching M, therefore means there is a strongly connected component $C \subseteq V$ for which no single price is IR for all agents in it. This means there exists at least one buyer b_x with node v_x in the connected component C, $v_x \in C$, and one seller s_y with node v_y in the same connected component C, $v_y \in C$, such that $\nu_{b_x} < c_{s_y}$, because otherwise we can pick any price below the lowest valuation of any buyer in C and above the highest cost of any seller in C and obtain a single individually rational price for that strongly connected component.

Because C is a strongly connected component there is a directed cycle involving v_x and v_y. More specifically, there is a directed path starting at v_y and ending in v_x. Let $E_{y,x}$ be the set of edges in that path and $V_{-(y,x)}$ all the nodes that are not visited by the path described by $E_{y,x}$. $(V_{-(y,x)}, E_{y,x})$ forms a new matching $M_{y,x}$. Any buyer b or seller s in N^M, with the exception of buyer b_x and seller s_y is part of a buyer–seller pair that is labelled to either a node in $V_{-(y,x)}$ or an edge in $E_{y,x}$. s_y and b_x are the exceptions because $v_y \notin V_{-(y,x)}$ and $v_x \notin V_{-(y,x)}$ are the start and endpoints of the path described by $E_{y,x}$ and therefore have no incoming/outgoing edge in $E_{y,x}$, respectively.

By only deallocating s_y and b_x, we found a new matching that increased social welfare by $-(\nu_{b_x} - c_{s_y}) > 0$. Therefore, M cannot be a social welfare maximizing matching, contradicting our initial assumption. □

4.5 Pricing Mechanism

Having the concept of non-discriminatory pricing in time-flexible systems, we will describe one particular mechanism that provides such prices. Generally speaking a mechanism takes reports Θ, finds a matching M and assigns prices $p(n)$ to agents $n \in N$.

The mechanism uses an efficient matching M. To find prices, our mechanism builds graph H_M and identifies the set of maximal strongly connected components $\mathbb{C}_M$. For a strongly connected component $C_X \in \mathbb{C}_M$, let $N^{C_x} = S^{C_x} \cup B^{C_x}$ be the set of sellers and buyers in that component. We assign a price half way between the costliest seller and cheapest buyer in the component to every agent in the component:

$$p(n) = \frac{\min\left\{\nu_b | b \in B^{C_x}\right\} + \max\left\{c_s | s \in S^{C_x}\right\}}{2} \forall n \in N^{C_x} \tag{4.2}$$

The reason why this produces Individually Rational and consistent prices is similar to the proof of Theorem 4.1, while Budget Balance is ensured by assigning prices to buyer–seller pairs. The algorithm runs in polynomial time as identifying strongly components can be done in linear time (Tarjan, 1972).

Theorem 4.2 (IR, BB and Consistent)**.** *The above mechanism is Individually Rational, Budget Balanced, prices agents consistently, and runs in polynomial time.*

Proof. Because of Lemma 4.4 we know that the pricing scheme is consistent. In Theorem 4.1 we saw that no buyer in a connected component can have a valuation lower than any seller in the same connected component, i.e. $\min\left\{\nu_b | b \in B^{C_x}\right\} \geq \max\left\{c_s | s \in S^{C_x}\right\}$ and therefore the assigned prices are IR. The mechanism is BB because it assigns prices to connected components which consist of buyer–seller pairs. The price paid to a seller is exactly compensated by the payment received from the buyer at that same node (buyer–seller pair). Regarding the run-time, the matching problem can be solved in polynomial time (Kuhn, 1955). For pricing, the graph H_M can be built in quadratic time, since we have a potential edge between any two nodes, while identifying the strongly connected components can be done in linear time in the number of nodes and edges of the graph (Tarjan, 1972). Therefore both allocation and pricing can be done in polynomial time. □

This consistent pricing mechanism has two advantages. It assures the agents that the market operator did not make a deliberate or accidental choice to disadvantage them. Second, it allows the market operator to clear the market and communicate fixed prices before deciding on a particular schedule/matching.

4.6 Discussion: consistent pricing, envy-freeness and fairness

Fairness is a vague intuitive human term. However, in order to design systems, we require precise definitions and have therefore come up with specific mathematical concepts to represent our human intuition.

One of the strongest fairness concepts is envy-freeness. An envy-free outcome is one in which no agent would want to switch places – be assigned the allocation and prices of another – with any other agent. However, due to the difficulty of achieving full envy-freeness, several relaxations have been proposed. Our concept of consistent pricing can be viewed as a new kind of relaxation of envy-freeness.

In the following we will first look at how our work compares to envy-freeness by considering how the field of economic psychology views fairness (Section 4.6.1) and then compare our concept of consistent pricing to other relaxations of envy-freeness (Section 4.6.2).

4.6.1 Fairness: envy-free or consistent pricing

Here, we compare our concept of consistent pricing directly to envy-freeness. In particular, we are interested in how the mathematical concepts relate to people's perception of fairness.

First, we show that consistent pricing does not necessarily lead to envy-freeness. Let us consider the example depicted in Figure 4.6, with non-discriminatory prices derived as described in Section 4.5. The prices are non-discriminatory but we can also see that seller s_3 envies seller s_2, because $p(1,2)$ is a higher price than $p(3)$. Unlike pure envy-freeness, we allow this envy and justify the price difference because pair $\{b_2, s_3\}$ is not part of an equivalent matching and therefore agents b_1, b_2, s_1, s_2 and agents b_3, s_3 are priced separately.

In economic psychology, there exist two notions of fairness in pricing. Distributed justice states that agents perceive prices as fair, if they are in line with their reference price (Chapuis, 2012; Maxwell, 2002), e.g. historical prices or prices that other agents are paying. Procedural justice, on the other hand, says that prices are seen as fair by agents, if the agents know that the process of obtaining the price was fair itself (Chapuis, 2012; Maxwell, 2002).

In the context of economic psychology, envy-freeness can, therefore, be viewed as distributed justice, no agent sees another agent being paid better or worse, while our concept of consistent prices reflect the procedural justice notion of fairness, the main idea behind the process of obtaining the prices was non-discrimination, i.e. a form of fairness. It should be noted here that these concepts are not necessarily in conflict. Distributed and procedural justice are

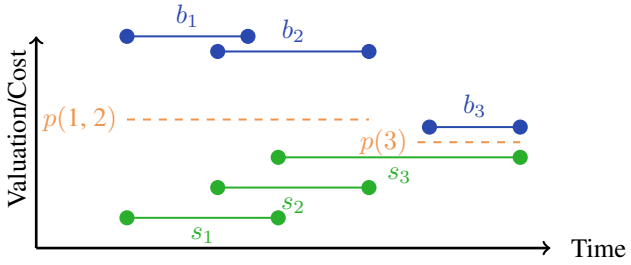


Figure 4.6: Not envy-free but non-discriminatory

simply two aspects of how humans perceive fairness in pricing, while envy-freeness and consistent pricing are mathematical concepts that reflect these aspects.

4.6.2 Comparison to relaxations of envy-freeness

Relaxations on the concept of envy-freeness allow either for a particular kind or magnitude of envy. The concept put forth in this chapter allows for a particular kind rather than magnitude of envy within social welfare maximizing matchings. The only way an agent b_x of pair $\{b_x, s_x\}$ envies an agent b_y of pair $\{b_y, s_y\}$ is if the pair $\{b_x, s_y\}$ is feasible and the price for b_y differs from that for b_x. However, under consistent pricing, the prices can only differ if the pair $\{b_x, s_y\}$ is not part of any equivalent, i.e. social welfare maximizing matching. Thus, b_x can only switch places with b_y if a non social welfare maximizing matching is chosen. Therefore, consistent pricing, as put forth here, relaxes envy-freeness only to maintain maximum social welfare.

There are several distinct differences between our concept of consistent pricing and relaxations of envy-freeness found in literature.

From the field of fair division, relaxed envy-freeness that allows for a certain magnitude of envy has emerged. "Envy-freeness up to any good" is the idea that given a certain bundle allocation, an agent envying another agent would cease to do so, if we removed just one item from the other agents bundle (Chaudhury et al., 2020; Livanos et al., 2022; Plaut and Roughgarden, 2018). This kind of relaxation can be viewed as trying to approximate the distributed justice notion of fairness. Our concept of consistent prices, on the other hand, invokes fairness through procedural justice, i.e. explaining why the process of price formation was fair.

One way to allow a particular kind of envy is to restrict the way in which envy is directed. One such restriction is expressed in the concept of group envy-freeness (Aleksandrov and Walsh, 2018; Todo et al., 2019), in which only

the envy between groups of agents is considered. Group envy-free outcomes ensure that no group of agents would want to switch with another group of agents while individual agents can still envy each other. However, this approach is limited to settings where agents can be seen as groups with regard to fairness.

Another way is to use a social graph, which represents awareness of others between the agents. (Alon et al., 2013) introduces a non-directed graph and limits the price differences between agents on adjacent nodes, effectively limiting envy between neighbors. More recently, in (Flammini et al., 2020, 2021) this idea is extended to directed social graphs representing the agents' awareness of other prices in the market. However, these approaches, while solving the problem of envy-freeness, do not translate well back to the initial intention of fairness. An agent, who has arbitrarily been given a worse deal is not treated fairly simply because they are not aware of the better deal they could have been given. This is also corroborate by studies in economic psychology. Psychological studies indicate that especially in the absence of price difference a sense of fairness is established via procedural justice (Chapuis, 2012). The fairness perception of an agent who is not aware of other prices in the system is more strongly determined by the procedural way prices are derived.

In comparison, an agent under consistent pricing can rest assured that even if they are not aware of a better deal, the reasoning behind the prices was fair and that if they received a worse deal, that it was not arbitrary but was rooted in a notion of consistency.

4.7 Conclusion and Future Work

Integration of flexible assets via new markets is a cheap alternative to costly grid reinforcement. This integration, however, depends on the adoption by users and therefore on a perception of fairness. This chapter took a deeper look at the concept of non-discriminatory pricing itself. While often simply assumed in literature, we showed that non-discriminatory pricing has no trivial definition in time flexible systems. For such settings, we provided a stronger foundation for non-discriminatory pricing as prices that are consistent over different but equivalent outcomes. In a stylized market setting, we showed that such consistent prices can exist without breaking individual rationality or budget balance and provided a mechanism, which runs in polynomial time. Furthermore, we put our notion of fairness into context by relating it to the concept of procedural justice from economic psychology, comparing it to envy-freeness and showing that it can be viewed as a particular relaxation of envy-freeness.

Building upon this work, we would like to extend our stylized market setting to incorporate a more complex bidding language and realistic constraints. Beyond markets with time-flexible agents, the extended concept of non-discriminatory prices as presented here can be applied to any setting in which outcomes can be different but equivalent. It might be particularly interesting to multi-energy systems in which the same outcome can be obtained using a different mix of energy carriers. In such a multi-energy system an agent expressing flexibility over several energy carriers might be reassured by knowing that the particular mix of energy that they were allocated was not more expensive than any equivalent alternative.

5 Incentives for Accurate Energy Predictions: How to Reduce Epistemic Uncertainty

Preface Accurate predictions of power fluctuations are pivotal to the operation of flexibility markets. While the design of flexibility markets is an active and ongoing field of research, the question of how to elicit high quality predictions in a non-cooperative setting is often overlooked. We contribute the concept of best prediction incentivizing contracts. Under such contracts the best response of an agent is to report the true distribution of its power fluctuation. This concept differs from Incentive Compatibility by explicitly taking epistemic uncertainty into account: while Incentive Compatible mechanisms often assume agents possess perfect knowledge of their own valuation, our concept incentivizes agents to reduce their epistemic uncertainty about the world. In practical terms, we present generic closed form solutions for polynomial distributions and show they can be used to approximate realistic Gaussian distributions. Lastly, placing our work in a larger context, we show that third party agents can profit from providing improved predictions via arbitrage. [1]

5.1 Introduction

Renewable energy sources (RES) are generally seen as the essential solution to global warming. Excessive use of fossil fuels and the accompanying output of carbon dioxide has caused global temperatures to rise (Lee et al., 2023). RES have become a carbon-neutral alternative to fossil fuels through advancements in technology as well as investments in mass production (IRENA, 2023; Timilsina, 2021).

[1] This chapter is based on publication (Saur et al., 2024). Supporting code can be found under doi.org/10.4121/1c4efcd7-1aca-498d-9337-091c8ae49db2

However, large scale integration of RES has brought with it some challenges for electricity systems. In contrast to conventional power plants, most types of RES cannot control their power output. As the share of fluctuating RES increases so does the magnitude and frequency of power imbalances, which need to be addressed by grid operators (Kuiken and Más, 2019; Lund et al., 2019).

To effectively deal with these power imbalances, flexibility has to be integrated into electricity grids (Auer and Haas, 2016; Kuiken and Más, 2019; Shahzad and Jasińska, 2024). While there is no agreed upon definition of flexibility, it is generally seen as the ability of power producing or consuming devices to quickly adjust their energy output/input (Lannoye et al., 2012).

Flexibility can be provided via several different technologies. Conventional gas turbine generators are well established and can provide flexibility by adjusting their operating point (Oree and Hassen, 2016). Besides adjustments on the production side, demand response can provide flexibility by adjusting electricity demand (Korkas et al., 2016; Kuiken and Más, 2019; Palensky and Dietrich, 2011; Zhang et al., 2025). However, flexibility can not only be provided by adjusting production or consumption, but also by shifting energy in different dimensions. Such flexibility can be provided by for instance electrical storage devices (Engels, 2020; Lüth et al., 2018), which effectively shift energy from times of excess to when energy production is falling short. Given the high cost of electrical storage (Denholm et al., 2010), the idea of shifting energy by utilizing several energy carriers in multi-energy systems has emerged. Flexibility providing multi-energy systems range from combined heat and power generators (Johnson and Mayfield, 2022; Su et al., 2026; Wang et al., 2019) and the latent heat in residential buildings (Golmohamadi, 2022; Korkas et al., 2016; Lizana et al., 2018) to tightly integrated industrial processes (Chen et al., 2019; Wang et al., 2017; Zhang et al., 2025) and even compute workloads in data centers(Liu et al., 2024).

While there exists a variety of potential flexibility providers as well as a clear need for it, efficiently integrating this flexibility is still a topic of active research.

Flexibility market mechanisms have emerged as a viable approach for large scale flexibility integration. Direct integration, i.e. equipping fluctuating RES with flexible assets (Papavasiliou and Oren, 2013) has allowed RES to better participate in conventional electricity markets (Aghamohamadi and Mahmoudi, 2019; Qiu et al., 2022; Saur et al., 2021; Yang et al., 2023). However, such integration ties the benefit to one particular fluctuating energy source limiting the overall usefulness of the flexible asset. This limitation can somewhat be addressed by integration of flexibility in larger multi-energy systems (Kleinschmidt et al., 2020). While this adds to the usefulness of

flexible assets, it still requires rather complex control systems (Kleinschmidt et al., 2020), which cannot be scaled to larger systems, e.g. national electricity markets. Market mechanisms on the other hand are able to coordinate between disparate energy assets leading to a more efficient use of available flexibility.

The operation of flexibility markets rely on accurate predictions. Flexibility trading has been approached both via centralized (Valarezo et al., 2021) as well as peer-to-peer markets (Valarezo et al., 2021; Zhao et al., 2015). It is noteworthy that all of these approaches implicitly assume accurate probability models. This assumption is complicated two-fold: by the difficulty of the task itself as well as the non-cooperative nature of market settings.

The task of forecasting power fluctuations is non-trivial involving different kinds of uncertainty. The need for accurate predictions in RES integration has been noted early on (Wan and Parsons, 1993) and steady progress has been made over the years on providing improved probabilistic models (Hanifi et al., 2020). However, only in recent years, the power forecasting literature has started addressing the issue of epistemic uncertainty (Ding et al., 2022; Liu et al., 2023; Wen et al., 2022) i.e. the uncertainty that exists about the probabilistic models.

Uncertainty can be grouped into aleatoric and epistemic uncertainty. Aleatoric uncertainty is the known uncertainty that can be modeled by probabilistic models. Epistemic uncertainty, on the other hand, represents the unknown uncertainty or lack of knowledge (Der Kiureghian and Ditlevsen, 2009; Murphy, 2022). The epistemic component of the uncertainty can be reduced by learning from additional data (Hüllermeier and Waegeman, 2021). However, since, epistemic uncertainty represents a lack of knowledge it is very difficult to capture in mathematical models, as has been discussed in (Hüllermeier and Waegeman, 2021). We will circumvent this issue by creating a mechanism that incentivizes the reduction of epistemic uncertainty rather than attempting to model it.

The Forecasting issue is further complicated in flexibility markets by the non-cooperative nature of market settings. Within a market setting the question is not just how to generate accurate forecasts but why a self-interested agent should reveal them to the system. This has been explicitly noted by (Zhao et al., 2015) as an open challenge. Recent years have seen the development of data markets (Goncalves et al., 2020; Han et al., 2022; Sun et al., 2023) as well as the use of prediction markets (Shamsi and Cuffe, 2021a,b,c) in the energy sector to address this challenge. However, such information markets are disjoint from the physical operation of the energy system and can therefore be more susceptible to speculation.

To summarize the engineering issue: consider a setting in which a fluctuating renewable energy source, such as a wind farm, seeks an insurance contract

from a flexibility provider. Such an insurance contract can be mutually beneficial as the fluctuating energy source is shielded from high imbalance costs, while the flexibility provider is paid for a service they provide. The issue when designing such an insurance contract starts with incentivizing the fluctuating energy source to provide accurate forecasts of their energy output.

The conceptual challenge, this work addresses, stems from the fact that the fluctuating energy source cannot have a perfect model of their forecast. The concept of Incentive Compatibility incentivizes agents to be honest, but implicitly assumes that the agent knows the *true* value to report. However, when reporting a forecast, the agent cannot be assumed to know the *true* distribution. The gap between what the agent believes to be the true forecast and the actually true forecast is called epistemic uncertainty and is what this work aims to address in a peer-to-peer flexibility trading setting.

The state of the art is extended in the following way.
We

- properly define our peer-to-peer flexibility trading setting and provide the concept of best prediction incentivizing contracts as a conceptual means to incentivizing the reduction of epistemic uncertainty (Section 5.2).
- provide best prediction incentivizing contract functions for uniform (Theorem 5.1) and single term polynomial (Theorem 5.2) probability distributions.
- show that our single-term polynomial solution is suitable for modeling real world applications (Section 5.3.3).
- demonstrate how the arbitrage potential in our solution incentivizes third party agents to contribute improved predictions (Section 5.4).
- put our work into context by comparing our approach to data markets (Section 5.5.1) as well as prediction markets (Section 5.5.2).

5.2 Problem Formulation

We consider a wind farm (WF) that is contractually obliged to produce a certain amount of power but in reality deviates from this obligation by an amount Δq. We say the actual power deviation Δq is described by a true probability density function (pdf) f_T.

To deal with the fluctuation, the WF can obtain an insurance contract $\Theta = (p_0, p_\Delta)$, which will protect it from costly balancing markets. The insurance contract Θ is composed of an upfront cost p_0 and a payment $p_\Delta(\Delta q)$, which depends on the actual power deviation Δq. The balancing cost to the WF

of power deviation Δq given a contract Θ can be expressed as:

$$c_B = p_0 + p_\Delta(\Delta q)$$

However, the contract Θ has to be provided by an insurance provider. The contracts this insurance provider is willing to provide depend on the event that the insurance is for, i.e. Θ will depend on f_T. To make this connection explicit, we say the insurance provider uses a contract creating function Ω that accepts a probability distribution f and returns a contract Θ. Let $p_0(.|\Omega(f))$ be the up-front cost p_0 of the contract returned by $\Omega(f)$, while $p_\Delta(\Delta q|\Omega(f))$ is the cost of power deviation Δq given the contract returned by $\Omega(f)$.

Given that the insurance contract Ω depends on f_T, the main research question is how to elicit best approximation of f_T from the agent.

5.2.1 Best Prediction Incentivizing

In a first step we assume that the WF knows the true probability distribution f_T, while the insurance provider has no information on f_T. The WF reports probability distribution f to contract creating function Ω to generate a contract Θ.

To answer our research question we are looking for a contract creating function Ω to make reporting f_T the best approach for the WF.

The expected cost to the WF of the deviation given that the insurance provider uses contract creating function Ω

$$C_B(f|\Omega, f_T) = p_0(.|\Omega(f)) + \mathbb{E}_{f_T}\left[p_\Delta(\Delta q|\Omega(f))\right] \tag{5.1}$$

should be minimized when the WF provides the true distribution f_T, i.e.

$$C_B(f_T|\Omega, f_T) \leq C_B(f|\Omega, f_T) \quad \forall f$$

Definition 5.1 (Incentivizing Best Predictions)**.** We call a contract creating function Ω or the related pricing functions $(p_0(.|\Omega(f)), p_\Delta(\Delta q|\Omega(f)))$ best prediction incentivizing if

$$C_B(f_T|\Omega, f_T) \leq C_B(f|\Omega, f_T) \quad \forall f$$

5.2.2 Best Prediction Incentizing under epistemic uncertainty

So far we assumed that the WF knows the true distribution f_T. However, in reality the WF only has an approximation f of the true distribution f_T. This is because there are two kinds of uncertainty at play here. The true uncertainty described by f_T that is inherent in the event itself. The second kind of uncertainty is called epistemic (Murphy, 2022) and is caused not by the

event but by a lack of knowledge that is reflected in the discrepancy between f and f_T.

Given that reporting f_T is optimal, there is a certain cost associated with reporting an approximate distribution f. We will refer to this cost as the cost of epistemic uncertainty.

$$CEU(f|\Omega, f_T) = C_B(f|\Omega, f_T) - C_B(f_T|\Omega, f_T)$$

This epistemic uncertainty about the true distribution f_T and its associated cost complicate our earlier definition of best prediction incentivizing contract functions. The earlier definition specifically incentivizes reporting the true distribution f_T. However, it is impossible to incentivize a WF to reliably report something it does not know.

We therefore extend our definition of best prediction incentivizing contract functions to settings with epistemic uncertainty. To do so we will not only require that reporting f_T is optimal, but that reporting a strictly better approximation always results in a lower cost of epistemic uncertainty. This way improving the approximation will always result in lower costs, in essence incentivizing the WF to report better approximations of f_T.

First we need to define what it means for one approximation to be better than the other. To do so we introduce an ordering $\succ_{f_T}$ with respect to the true distribution f_T, where f_T is the dominant element, i.e. $f_T \succ_{f_T} f_1 \quad \forall f_1 \neq f_T$.

Definition 5.2. We say probability density function f_1 is a better approximation of f_T than probability density function f_2 if $f_1 \succ_{f_T} f_2$. Several different orderings can reasonable be considered for $\succ_{f_T}$. We will provide two here:

a) Given an interval $[-d_T, d_T]$ and a class $\mathcal{F}$ of probability density functions on said interval, we say that $f_1 \succ_{f_T} f_2$ if

$$f_1(\Delta q) \in [min(f_T(\Delta q), f_2(\Delta q)), max(f_T(\Delta q), f_2(\Delta q))] \, \forall \Delta q$$

as well as

$$\exists \Delta q \text{ s.t. } f_1(\Delta q) \neq f_2(\Delta q)$$

b) Given a parameter d, which defines a unique function in a class $\mathcal{F}$ on the interval $[-d, d]$, we say that $f(d_1) \succ_{f_T} f(d_2)$ if

$$d_1 \in (d_2, d_T] \quad \text{or} \quad d_1 \in [d_T, d_2)$$

With this definition, we can extend our definition of Incentivizing Best Predictions to settings with epistemic uncertainty.

Definition 5.3 (Incentivizing Best Predictions under Epistemic Uncertainty)**.** We call a contract creating function Ω or the related pricing functions $(p_0(.|\Omega(f)), p_\Delta(\Delta q|\Omega(f)))$ best prediction incentivizing under epistemic uncertainty if reporting a better approximation of f_T always results in a lower epistemic cost, i.e.

$$CEU(f_1|\Omega, f_T) < CEU(f_2|\Omega, f_T) \qquad \forall f_1 \succ_{f_T} f_2$$

Corollary 5.1. Given that Ω is Incentivizing Best Predictions under Epistemic Uncertainty, a contract Θ_1 created by providing f_1 will have a lower expected cost than a contract Θ_2 created by providing f_2 if f_1 is a better approximation of the true distribution f_T, i.e.

$$C_B(f_1|\Omega, f_T) < C_B(f_2|\Omega, f_T) \qquad \forall f_1 \succ_{f_T} f_2$$

Note 5.1. Any pricing function that is Incentivizing Best Predictions under Epistemic Uncertainty is also Incentivizing Best Predictions under perfect knowledge.

Going forward, we will be referring to best prediction incentivizing contracts/pricing functions (Definition 5.3) as *BPI*.

5.2.3 Comparison to Incentive Compatibility

In this section we explore the difference between our concept and the related concept of Incentive Compatibility (IC). In an IC mechanism an agent's best response is to reveal their true valuation to the market (Nisan et al., 2007). Regarding the information revelation part, IC is similar to BPI contracts. However, in our setting the true valuation of the agent is obscured by epistemic uncertainty. The concept presented here differs from IC by explicitly accounting for the epistemic uncertainty that the agent has about the world.

The epistemic uncertainty stems from the kind of information we are interested in revealing: internal versus external information.

In settings with perfectly known valuations, IC mechanisms incentivizes the revelation of private/internal information to value a good or service. For instance, the value of a vase being sold at an art auction is entirely determined by the private appreciation of the participants. This holds similarly for any kind of commodity or service contract. The assumption, here, being that the agent knows their own true valuation.

Pricing in BPI contracts explicitly aims at revealing external information about the world itself, i.e. the true future power fluctuation. Instead of valuing an object or service, BPI aims to arrive at the true probability distribution of a wind farm. While the argument can be made that the agent knows their true internal valuation, it is very unlikely that they possess perfect knowledge about

the world. This inability to perfectly model the future is what we refer to as epistemic uncertainty.

This difference in revealing external instead of internal information has implications for the purpose of pricing. In both concepts, pricing is used as a means to an end. In IC mechanisms without epistemic uncertainty, the end is to value a good or service. In the concept of BPI, pricing is directly linked to finding the correct predictions of an uncertain future event. This connection between pricing and prediction can also be seen in Corollary 5.1, which states that better predictions result in lower costs. The prices in BPI contracts is therefore directly connected to the amount of epistemic uncertainty.

However, taking epistemic uncertainty into account also comes with challenges for the agent. Honest reports are easy to compute. Therefore, when the true valuation is known to the agent, truth-telling is an easy strategy to implement. However, correctly predicting an external event is a non-trivial task. We will return to this challenge in Section 5.4, when discussing the arbitrage potential of BPI contracts.

5.3 Applications of BPI contracts

So far we have provided BPI as a novel concept and discussed how it relevantly differs from IC. However, a concept is only useful if it can feasibly be applied. To this end, we provide BPI contracts for uniform (Section 5.3.1) and single term polynomial distributions (Section 5.3.2) and show that the latter can be used in more real-world settings (Section 5.3.3). For the particular cases in this section, we will be using a parameter d instead of the entire function f as the input to the contract generating function Ω. Given this input, we will consider the particular ordering of Definition 5.2b.

5.3.1 Solutions for uniform distribution

To begin with, we consider a very simple setting.

Setting 1 (Uniform). The true power deviation of a fluctuating energy source is described by a uniform probability density distribution on the interval $[-d_T, d_T]$. The task of the energy source operator is to provide the correct parameter d_T.

We want to design prices, $p_0(.|d)$ and $p_\Delta(\Delta q|d)$ for Setting 1 such that the operator minimizes its cost by submitting $d = d_T$, i.e. they are best predictions incentivizing.

Theorem 5.1. *Contract generating function $\Omega_u(d)$ with pricing functions ($p_0(.|d) = 0.5ln(d)$, $p_\Delta(\Delta q|d) = |\Delta q|\frac{1}{d}$) are best BPI functions under Epistemic Uncertainty for Setting 1.*

Proof. To show that the epistemic cost always decreases by moving from a worse approximation d_2 to a better one d_1, we only need to show that epistemic costs are falling for d less than d_T and are rising for d larger than d_T, i.e.:

$$\frac{\partial CEU(d|\Omega_u, d_T)}{\partial d} < 0 \quad \text{if } d \in (0, d_T)$$
$$\frac{\partial CEU(d|\Omega_u, d_T)}{\partial d} > 0 \quad \text{if } d \in (d_T, \infty).$$

This can easily be shown:

$$\begin{aligned}\frac{\partial CEU(d|\Omega_u, d_T)}{\partial d} &= \frac{\partial}{\partial d} C_B(d|d_T, \Omega_u) - \underbrace{\frac{\partial}{\partial d} C_B(d_T|d_T, \Omega_u)}_{=0} \\ &= \frac{\partial}{\partial d}\left[0.5ln(d) + \frac{1}{2}d^{-1}d_T\right] \\ &= 0.5d^{-1}\left[1 - \frac{d_T}{d}\right]\end{aligned}$$

Note that $\left[1 - \frac{d_T}{d}\right] > 0$ for $d > d_T$ and $\left[1 - \frac{d_T}{d}\right] < 0$ for $d < d_T$, which concludes the proof.

□

5.3.2 Solutions for single term polynomial distributions

In this section, we will be deriving a family of contract generating functions $\Omega(d|k)$ parameterized on a shape parameter k. We consider a single term polynomial as the pdf $f(\Delta q|d)$ on the symmetric interval $[-d, d]$. The general form of f on the positive interval is:

$$f(\Delta q) = -m(\Delta q)^k + t$$

Furthermore, we assume that $f(\Delta q|d)$ is symmetric around $\Delta q = 0$ and $f(d) = 0$. With these two assumptions $f(\Delta q|d)$ is uniquely defined given an agent's report d.

From $f(d) = 0$ it follows that

$$t = md^k$$

Because f is a probability density function and symmetric, we know that:

$$\int_0^d f(\Delta q)\partial\Delta q = \frac{1}{2} = \left[-\frac{m}{k+1}(\Delta q)^{k+1} + t\Delta Q\right]_0^d$$
$$= -\frac{m}{k+1}d^{k+1} + t \cdot d$$

Replacing $t = md^k$ we obtain:

$$\frac{1}{2} = -\frac{m}{k+1}d^{k+1} + md^{k+1}$$

From which follows:

$$m = \frac{1}{2}\frac{k+1}{k}d^{-(k+1)} \wedge t = \frac{1}{2}\frac{k+1}{k}\frac{1}{d}$$

This results in the general from of

$$f_p(\Delta q|d,k) = \frac{1}{2}\frac{k+1}{k}\left[-d^{-(k+1)}(\Delta q)^k + \frac{1}{d}\right] \quad \text{if } \Delta q \in [0,d]$$
$$f_p(\Delta q|d,k) = f(-\Delta q) \quad \text{for } \Delta q \in [-d,0] \tag{5.2}$$

Setting 2 (Single Term Polynomial). The true power deviation of a fluctuating energy source is described by a probability density distribution $f_p(\Delta q|d,k)$ of the shape of equation 5.2 on the interval $[-d_T, d_T]$. The task of the energy source operator is to provide the correct parameter d_T.

For Setting 2, we derive the following BPI pricing functions.

Theorem 5.2. *Given a power deviation that can be described by a single term polynomial $f_p(\Delta q|d,k)$ of specified degree k as described in equation* (5.2)*, contract function Ω_p with pricing functions $p_0(.|d) = \frac{k+1}{2(k+2)}log(d)$ and $p_\Delta(\Delta q) = |\Delta q|\frac{1}{d}$ are BPI pricing functions according to Definition 5.3.*

Proof. To show that p_0 and p_Δ is BPI pricing functions, we need to show that the expected cost C_B contains a global minimum at $d = d_T$, which we will do by showing that the expected cost increases for $d > d_T$ and decreases for $d < d_T$. We denote f_T as the true distribution $f(\Delta q|d_T,k)$.

First, we calculate the expected cost:

$$
\begin{aligned}
C_B(f,|\Omega f_T) &= \frac{k+1}{2(k+2)} log(d) + \int_{-\infty}^{\infty} f_T(\Delta q)|\Delta q| \frac{1}{d} \partial \Delta q \\
&= \frac{k+1}{2(k+2)} log(d) + 2\int_{0}^{d_T} f_T(\Delta q)\Delta q \frac{1}{d} \partial \Delta q \\
&= \frac{k+1}{2(k+2)} log(d) \\
&\quad + \frac{k+1}{k}\frac{1}{d}\int_{0}^{d_T} -\frac{1}{d_T^{k+t}}(\Delta q)^{k+1} + \frac{1}{d_T}(\Delta q)\partial \Delta q \\
&= \frac{k+1}{2(k+2)} log(d) \\
&\quad + \frac{k+1}{k}\frac{1}{d}\left[-\frac{1}{k+2}\frac{1}{d_T^{k+1}}(\Delta q)^{k+2} + \frac{1}{2d_T}(\Delta q)^2\right]_0^{d_T} \\
&= \frac{k+1}{2(k+2)} log(d) + \frac{k+1}{k}\frac{1}{d}\left[-\frac{1}{k+2}d_T + \frac{1}{2}d_T\right] \\
&= \frac{k+1}{2(k+2)} log(d) + \frac{k+1}{k}\frac{d_T}{d}\left[-\frac{2}{2(k+2)} + \frac{k+2}{2(k+2)}\right] \\
&= \frac{k+1}{2(k+2)} log(d) + \frac{k+1}{2(k+2)}\frac{d_T}{d} \\
&= \frac{k+1}{2(k+2)}\left[log(d) + \frac{d_T}{d}\right]
\end{aligned}
$$

Taking the first derivative of the expected cost C_B with respect to the reported interval d results in:

$$
\begin{aligned}
\frac{\partial}{\partial d} C_B &= \frac{k+1}{2(k+2)}\left[\frac{1}{d} - \frac{d_T}{d^2}\right] \\
&= \underbrace{\frac{k+1}{2(k+2)}\frac{1}{d}}_{>0}\left[1 - \frac{d_T}{d}\right]
\end{aligned}
$$

It is easy to see that $\left[1 - \frac{d_T}{d}\right] > 0$ for all $d > d_T$, while also $\left[1 - \frac{d_T}{d}\right] < 0$ for all $d < d_T$ and therefore the cost increases while $d > d_T$ and decreases while $d < d_T$. This results in a global minimum at $d = d_T$. □

5.3.3 Real World Applications

While Theorem 5.2 provides us with a solution for a wide range of polynomial distributions, on first glance it seems to be limited when it comes to real-world

applications. First, Theorem 5.2 assumes that any market operator knows the particular degree k of the polynomial up front. Secondly, not many problems regarding uncertainty in the energy domain are modeled as single term polynomial distributions.

However, we can fine tune our theoretical solution to be applicable to real world problems. Uncertainty in the energy domain is often modeled as Gaussian distributions(Ge et al., 2018; Trombe et al., 2012; Zhang et al., 2019). We will, therefore, consider in Setting 3 a truncated Gaussian distribution as the true model of reality. Fortunately, Theorem 5.2 is parameterized on k and, thus, provides an entire family of solutions. Using parameter k, we can choose the solution that best approximates a truncated Gaussian distribution, i.e. we are looking for k^* such that $\Omega_P(d|k^*)$ is best suited for the following Gaussian setting.

Setting 3 (Gaussian). We will consider reality to be described by a truncated Gaussian distribution $\mathcal{N}(\Delta q|\mu, \sigma)$ (Robert, 1995) with mean $\mu = 0$ and standard deviation σ on a symmetric interval $[-d, d]$.

This truncated Gaussian is, therefore, completely described by the input parameters d and σ. We scale the standard deviation σ by $0.5d$ to reduce the number of input parameters. The Gaussian setting contains not only the interval length d but also the standard deviation σ as an input parameter. This input parameter does not appear in Setting 2, which we will use for approximation. By scaling $\sigma = 0.5d$, the distribution can be described solely by the input parameter d. Therefore, the input for both this setting as well as Setting 2 is the same, i.e. the interval length d.

This results in a truncated normal probability distribution $\mathcal{N}(\Delta q|d)$ that can be described solely by the interval size d in the following way:

$$\frac{1}{0.5d}\frac{1}{\Phi^*(d|0, 0.5d) - \Phi^*(-d|0, 0.5d)}\mathcal{N}^*(\Delta q|0, 0.5d), \tag{5.3}$$

where $\mathcal{N}^*(\Delta q|0, 0.5d)$ is the untruncated normal distribution of mean $\mu = 0$ and standard deviation $\sigma = 0.5d$ with its corresponding cumulative $\Phi^*(\Delta q|0, 0.5d)$ and $\Delta q \in [-d, d]$.

Note, that reducing the number of input parameters does not mean that this setting can only describe a single shape of Gaussian distributions. Figure 5.1 shows Gaussians from this setting for different values of d.

In the following sub-sections, we will show that the solution from Theorem 5.2 is appropriate for Setting 3. While setting 2 and 3 differ in specific ways – one is a polynomial the other a Gaussian, one is zero at the ends of the interval the other is not – we will see that they are similar enough in practical terms. In particular, we will show that our solution can provide a good approximation of a Gaussian distribution (Section 5.3.3) and that the introduced

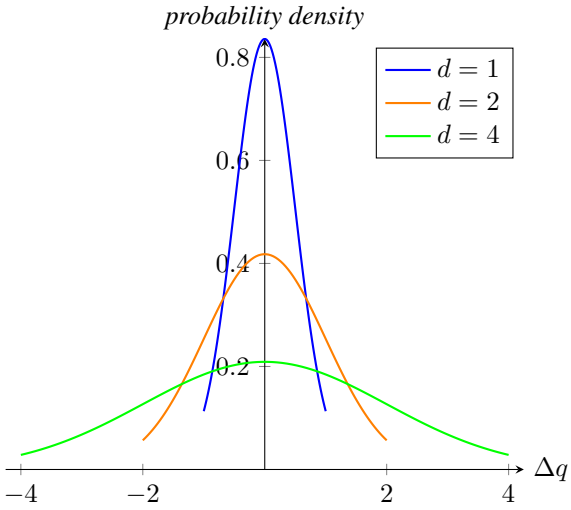


Figure 5.1: Truncated Gaussians for different reports of d

modeling error is especially small when comparing it to the error caused by epistemic uncertainty (Section 5.3.3). Lastly, we consider the possibility of the agent to gain a profit by reporting a value other than d_T and conclude that both the ability and the effect of such misreporting is very limited (Section 5.3.3).

Approximating Gaussian Distributions

While we do not have BPI pricing functions for Setting 3 in closed form, we will leverage our previous result in Theorem 5.2 to derive pricing functions that specifically aim to approximate truncated Gaussian distribution as described in Setting 3. Specifically, we are looking for a value of k such that the probability distribution described in equation (5.2) best approximates the truncated Gaussian. We measure the quality of our approximation by the root mean-square error (Lane et al., 2003) over the interval $[-d_T, d_T]$. In this subsection, our task is to find k^* to minimize the error between $\mathcal{N}(\Delta q|d_T)$ and the single term polynomial $f_p(\Delta q|d_T, k)$ as defined in equation (5.2).

$$Error(\mathcal{N}, f) = \sqrt{\frac{1}{2d_T}\int_{-d_T}^{d_T} \left[\mathcal{N}(\Delta q|d_T) - f(\Delta q|d_T, k)\right]^2 \partial\Delta q} \qquad (5.4)$$

Figure 5.2 shows the average error between our approximate solution and the Gaussian distribution over k. The error appears to be minimal for approx-

imately $k = 1.4$. Note that this value is independent of the actual size of the interval d_T as both the Gaussian $\mathcal{N}(\Delta q|d_T)$ and $f_p(\Delta q|d_T, k)$ are being scaled by d_T.

$$k^* \approx \arg\min_k \{Error(\mathcal{N}, f)\} \tag{5.5}$$

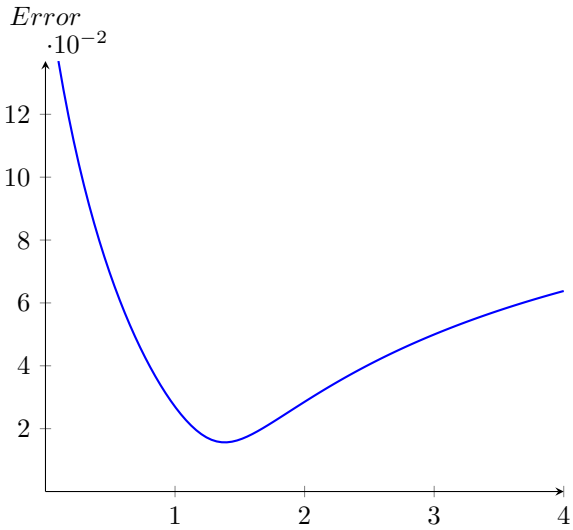


Figure 5.2: Average Error over k

Figure 5.3 shows both Gaussian distributions describing the original problem as well as the approximation f_p provided by our approach. It appears that our approach approximates the truncated Gaussian well for a range of values of d.

In subsequent sections, we will use the solution from Theorem 5.2 with $k^* = 1.4$.

Approximation under Epistemic Uncertainty

In this section, we explore the effect of epistemic uncertainty on our approximation and conclude that the error we introduce by using an approximate model of the problem is often insignificant compared to the error caused by epistemic uncertainty. We again consider Setting 3 and the approximation from Theorem 5.2 with $k^* = 1.4$.

In this section we will deal with two kinds of errors: the modeling error and the epistemic error. The modeling error is caused by using an approximate

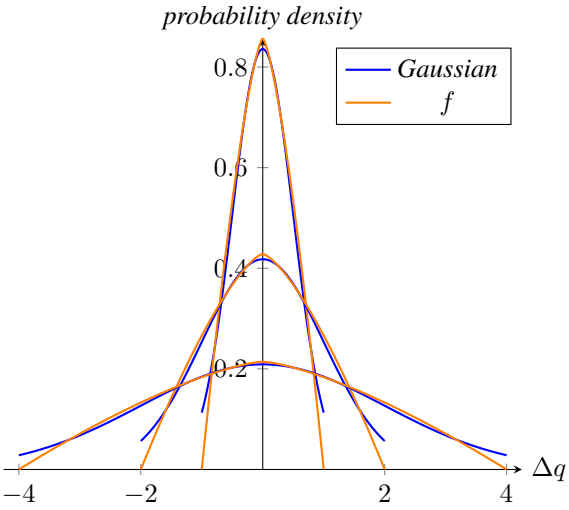


Figure 5.3: Gaussian and its approximate distribution for $k^* = 1.4$

model, i.e. single term polynomial instead of a Gaussian. The epistemic error is due to an inaccurate report by the agent. We assume that the agent only knows an approximate value $d \approx d_T$ instead of the true value d_T.

To more clearly highlight the different (modeling and epistemic) contributions to the overall approximation error, we introduce several notions of errors. The total error $Error_{total}$ is simply the overall error caused by providing an approximate value d to function $f_p(\Delta q|d, k^*)$, while the true distribution is the Gaussian $\mathcal{N}(\Delta q|d_T)$.

$$Error_{total} = Error(\mathcal{N}(\Delta q|d_T), f_p(\Delta q|d, k^*)) \tag{5.6}$$

We define the purely epistemic error as the result of providing an approximate value d to the true model $\mathcal{N}(\Delta q|d))$.

$$Error_{epistemic} = Error(\mathcal{N}(\Delta q|d_T), \mathcal{N}(\Delta q|d)) \tag{5.7}$$

Given the total error and the purely epistemic error, we introduce the *relative modeling error* as an indication of how much of the total error is due to incorrect modeling. While the true distribution is Gaussian, we are using the polynomial f_p, see equation (5.2), to approximate it. The error caused by this discrepancy we refer to as the modeling error. We, then, define the relative modeling error $Error_{model}$ as the relative difference between the total and the

epistemic error.

$$Error_{model} = \frac{Error_{total} - Error_{epistemic}}{Error_{total}} \tag{5.8}$$

Figure 5.4 shows the relative modeling error over reports d, while Figure 5.5 displays the total error. As is to be expected, the relative modeling error is 1 (i.e. the error is entirely caused by the incorrect model) when $d = d_T$ because then the epistemic error is zero and the entire error is due to the polynomial approximation. However, the relative modeling error quickly drops off as d deviates from the true value of d_T. It is worth noting that this narrow band around d_T, where the modeling error makes a significant contribution, also mostly contains the values of d for which the overall error is small, see Figure 5.5. Therefore, whenever there is a significant difference between the true distribution and the approximate distribution provided by our solution (Figure 5.5), most of the error is caused by epistemic uncertainty rather than incorrect model assumptions (Figure 5.4).

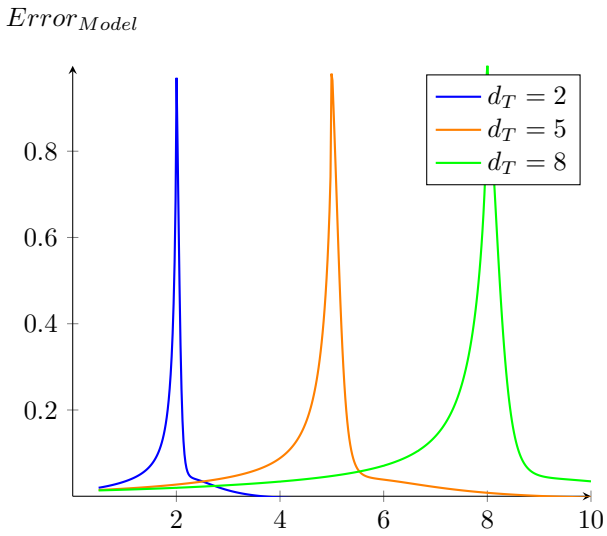


Figure 5.4: Relative Modeling Error over report d for $k^* = 1.4$

In this section, we showed that our solution not only provides a good approximation of Gaussian distributions, but also that the introduced modeling error is relatively small once epistemic uncertainty is present. However, so far we assumed that the agent is still incentivized to provide accurate predictions, which we will address in the next section.

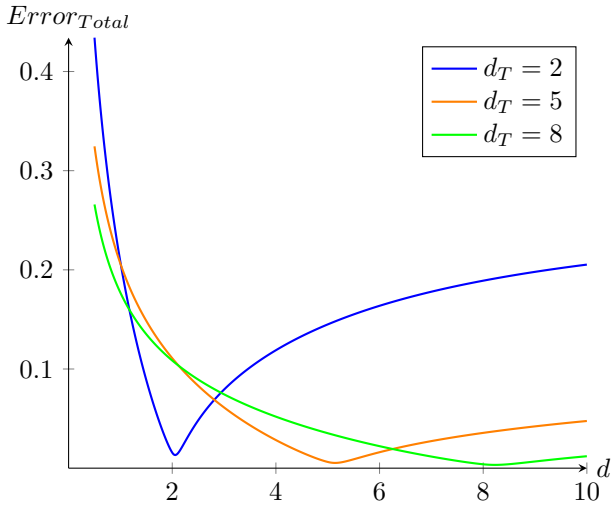


Figure 5.5: Total Error over Report d

Incentives under Epistemic Uncertainty

While Section 5.3.3 showed that the modeling error is small compared to the epistemic error, we, implicitly, assumed that the incentive to provide the best prediction of d_T still holds when the underlying true distribution is Gaussian. In this section, we consider Setting 3 and the approximation, i.e. Ω_P and f_p, from Theorem 5.2 with $k^* = 1.4$.

Specifically, we will explore to which extend the agent is still incentivized to provide an accurate estimation of d_T. Theorem 5.2 only specifies pricing functions that are BPI when the true distribution is a single term polynomial of specified degree k. This raises the question of incentives when these pricing functions are used in Setting 3. Reporting the true interval length d_T is no longer proven to be the optimal strategy. Therefore, there is a cost minimizing d^* that may not be equal to d_T.

$$d^* = \arg\min_d C_B\left(d|\Omega_p(d|k^*), \mathcal{N}(\Delta q|d_T)\right) \tag{5.9}$$

We explore the ramifications of this deviation from truthfully reporting d_T by looking at three metrics related to the strategy, agent cost and the resulting distribution. Regarding the strategy, we consider the absolute value of the relative distance between d^* and d_T, $\left|\frac{d^* - d_T}{d_T}\right|$. For the cost of the agent,

we look at the relative cost reduction gained by misreporting, i.e. $\left|\frac{\Delta Cost}{Cost_T}\right|$. Finally, we look at the quality of the resulting distribution by calculating the total error caused by optimal misreporting d^* when using the approximation $f_p(\Delta q|d, k^*)$ as defined in equation (5.2). This error we will refer to as the error of misreporting $Error_{mis}$.

$$Error_{mis} = Error[\mathcal{N}(\Delta q|d_T), f_p(\Delta q|d^*, k^*)] \tag{5.10}$$

However, any exploration of these effects is contingent on the ability to obtain the optimal report d^*. Unfortunately, an analytical solution to d^* is difficult to obtain as there is no closed form of the integral of a Gaussian. Luckily, the expected balancing cost $C_B\left(d|\Omega_p(d|k), \mathcal{N}(\Delta q|d_T)\right)$ of reporting d to the polynomial contract function $\Omega_p(d|k)$ in Gaussian Setting 3 has a single local minimum with regard to the reported parameter d, see Theorem 5.3. This means that approaches such as gradient descent and bisection method converge quickly to the global minimum d^*.

Theorem 5.3. *When applying contract function $\Omega_p(d|k)$ (Theorem 5.2), to Setting 3, the expected balancing cost $C_B\left(d|\Omega_p(d|k), \mathcal{N}(\Delta q|d_T)\right)$ has a single local minimum with regard to the report d.*

Proof. Note that while Setting 3 defines the real probability distribution to be on the interval $[-d, d]$, the agent can report any $d > 0$ and cost function $\Omega_p(d|k)$ is defined on $d > 0$. Therefore, to show that $C_B\left(d|\Omega_p(d|k), \mathcal{N}(\Delta q|d_T)\right)$ has a single local minimum, it suffices to show that its derivative $\frac{\partial}{\partial d} C_B\left(d|\Omega_p(d|k), \mathcal{N}(\Delta q|d_T)\right)$ changes sign from negative to positive at a single point for $d > 0$. Recall the general cost function, equation (5.1) as well as the explicit pricing functions in Theorem 5.2.

$$\begin{aligned}
&\frac{\partial}{\partial d} C_B(d|\Omega_P, \mathcal{N}_T) \\
&\quad = \frac{\partial}{\partial d}\left[\frac{k+1}{2(k+2)} log(d) + \int_{-\infty}^{\infty} \mathcal{N}_T(\Delta q)|\Delta q| \frac{1}{d} \partial \Delta q\right] \\
&\quad = \frac{k+1}{2(k+2)} \frac{1}{d} + \frac{\partial}{\partial d} \int_{-\infty}^{\infty} \mathcal{N}_T(\Delta q)|\Delta q| \frac{1}{d} \partial \Delta q \\
&\quad = \frac{k+1}{2(k+2)} \frac{1}{d} + \frac{\partial}{\partial d} \frac{1}{d} \underbrace{\int_{-\infty}^{\infty} \mathcal{N}_T(\Delta q)|\Delta q| \partial \Delta q}_{const > 0} \\
&\quad = \frac{k+1}{2(k+2)} \frac{1}{d} - \frac{1}{d^2} \int_{-\infty}^{\infty} \mathcal{N}_T(\Delta q)|\Delta q| \partial \Delta q
\end{aligned}$$

$$= \frac{1}{d}\left[\frac{k+1}{2(k+2)} - \frac{1}{d}\underbrace{\int_{-\infty}^{\infty} \mathcal{N}_T(\Delta q)|\Delta q|\partial\Delta q}_{>0}\right]$$

It is easy to see that $\left[\frac{k+1}{2(k+2)} - \frac{1}{d}\underbrace{\int_{-\infty}^{\infty} \mathcal{N}_T(\Delta q)|\Delta q|\partial\Delta q}_{>0}\right]$ changes sign at a single point, while $\frac{1}{d} > 0$ for all $d > 0$, which concludes the proof. □

We numerically explore the effect of reporting optimal d^* as in equation (5.10) instead of d_T in the context of using $\Omega_p(d|k^*)$ (Theorem 5.2) in Setting 3. Similar to Section 5.3.3, we assume a value of $k^* = 1.4$. We use the scipy optimize package (Virtanen et al., 2020) to calculate d^*.

First we consider the relative distance between the optimal report d^* and the true value d_T. Figure 5.6 displays the relative distance between optimal d^* and true d_T over different values of d_T. We see that the deviation from the true value d_T is generally low and less than three percentage points for a range of values. Note, that in order to reliably report d such that the cost is less than when reporting d_T, the agent needs to know d^* within this three percentage point window. Given that we can assume that the reporting agent does not have perfect knowledge of the future, it stands to reason that they will not be able to take advantage of this difference.

Next, we consider the effect that reporting an untrue but optimal value $d^* \neq d_T$ has on both the agent's cost and the quality of the resulting prediction. In Figure 5.7a, we see that the relative difference in cost for reporting d^* instead of d_T is very small. This indicates that any effort invested in optimizing for d^* over d_T might ultimately be larger than its gain. Figure 5.7b depicts the misreporting error as defined in equation 5.10 when assuming a single term polynomial and optimal report d^* compared to a Gaussian distribution on the true interval $[-d_T, d_T]$. As we can see, the effect on the prediction is less than $3 \cdot 10^{-2}$, indicating that even if the agent reports d^*, the flexibility provider buying the contract is still left with a virtually unchanged prediction.

Therefore, we believe our approach to still incentivize honest behavior: First, effectively deviating from the true report is difficult in practical settings, see Figure 5.6; Second, the gains of deviating are small, see Figure 5.7a; And third, the error-effects on the prediction are similarly small, see Figure 5.7b.

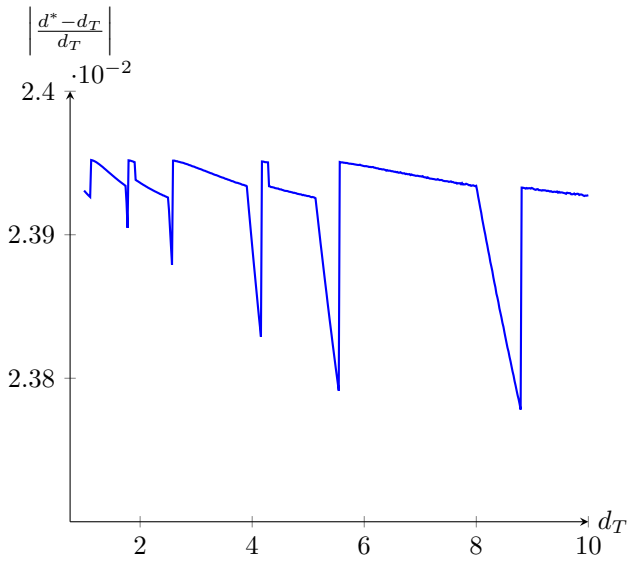
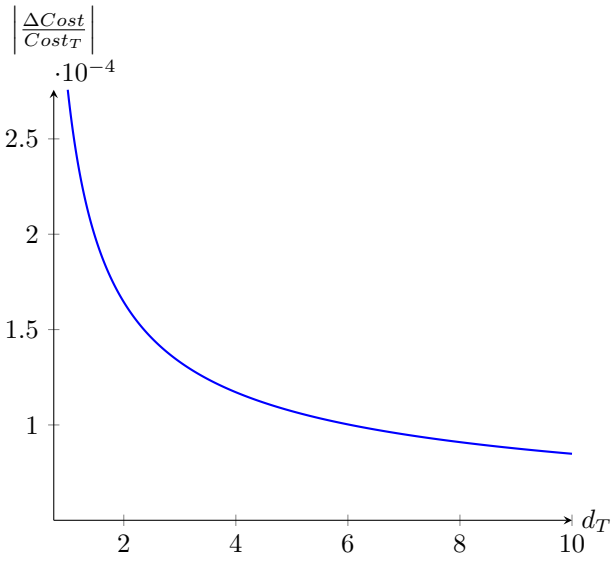
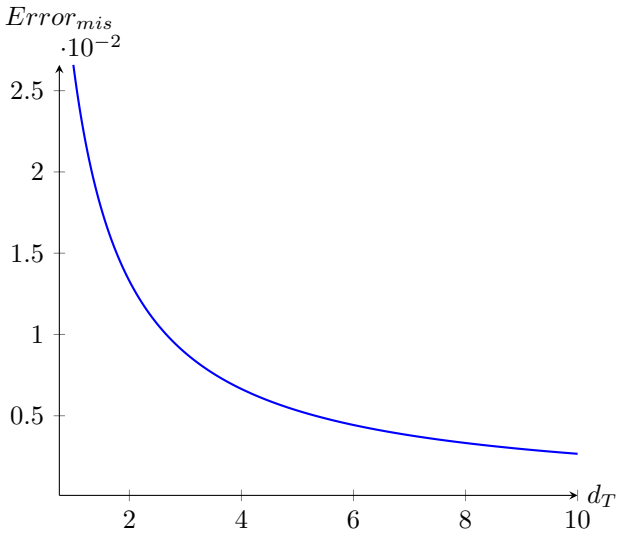


Figure 5.6: Relative distance between optimal and true report



(a) Effect on cost



(b) Effect on Error of misreporting

Figure 5.7: Effect of reporting optimal instead of true intervals

Real World Applications: Concluding Remarks

In conclusion, we demonstrated in this section the applicability of our theoretical results to real world problems. We showed the versatility of the results from Theorem 5.2. From the family of BPI contracts provided by Theorem 5.2, we produced a solution that is suitable for a setting with Gaussian distributions.

It is important to note that the solution from Theorem 5.2 does not only fit a Gaussian well, as can be seen in Section 5.3.3. Beyond a good fit, the modeling error is often less significant than the error caused by epistemic uncertainty, i.e. the agent only knowing an approximate value of d, see Section 5.3.3. And even under perfect knowledge, the effect of misreporting is small both for the cost of the agent as well as the resulting prediction, see Section 5.3.3. All together this produces a solution that incentivizes agents to provide improved predictions most of the time.

5.4 Arbitrage: Reselling improved predictions

In this section, we show how reselling contracts can both benefit the reseller and lead to improved predictions.

In particular, we revisit the issue that the agents do not know their true valuation, see Section 5.2.3. But, we will frame the issue in the context of arbitrage. So far we have shown how our solution can incentivize a fluctuating energy source to produce accurate predictions. However, this incentive might not actually translate into better predictions, because, it is not just a matter of willingness but also ability to report the true forecast. Unlike providing private information, the task of providing accurate predictions is non-trivial and might not be within the expertise of the operator of a fluctuating energy source. However, providing such predictions might very well be the expertise of other agents.

The challenge of enabling third party agents to participate in the prediction process is not necessarily a technical one. In most cases, the sources of uncertainty in renewable energies are events with publicly available data, such as weather forecasts. This means from a data availability standpoint, the task of providing accurate predictions could be done by any third party. The challenge lies more with how to incentivize a third party to participate in the first place.

This is where arbitrage becomes relevant. Arbitrage is the act of generating a profit by reselling a good or service. The arbitrage of BPI contracts addresses the above mentioned challenge, because the contract creating function connects prices directly to the quality of the prediction, as we have argued in Section 5.2.3. This direct connection between prices and quality of prediction

means that arbitrage of such contracts rewards agents that provide improved predictions with favorable prices.

We will illustrate this potential in the following example.

5.4.1 Arbitrage Example: Middleman providing an improved Prediction

To illustrate how a third party agent can benefit from arbitraging contracts created from BPI pricing functions, we will consider an example setting consisting of a Wind Farm (WF), flexibility provider (FP) and middleman (M) as depicted in Figure 5.8. For this example we assume the true day-ahead wind power fluctuating is described by a truncated Gaussian on the interval $[-d_T, d_T]$. For the contract generating function $\Omega_p(d|k^*)$, we use the cost function as outlined in Theorem 5.2 with parameter $k = 1.4$, as this has been shown to be a good approximation for a truncated Gaussian, see Section 5.3.3. Figure 5.9 depicts the expected cost the Wind farm incurs when providing an estimation d.

Initially, there only exists the rather inaccurate estimate d_1 that the WF produced – indicated in Figure 5.9. The WF uses this estimate to generate a contract Θ_1 from contract generating function $\Omega_p(d|k^*)$. It offers the contract Θ_1 to the market. Contract Θ_1, the input parameter d_1 as well as the wind farm that the contract refers to are public knowledge.

A third party agent M, observes contract Θ_1 as well as input d_1 and believes it has a better estimate d_2 – indicated in Figure 5.9. Therefore M decides to buy contract Θ_1 and generates a new contract using its improved estimate d_2. It then offers this improved contract Θ_2 to the market, where it gets sold to the flexibility provider FP.

It is important to note that middleman M has no energy obligation and only engages in arbitrage. On a physical level both contract Θ_1 and Θ_2 refer to the same uncertain future energy production. This means that both contracts insure the same event and therefore the energy Δq that is exchanged is the same in both contracts, leaving M with no obligation.

With regard to prices, however, Θ_1 and Θ_2 differ in favor of M. Recall, Corollary 5.1 from Section 5.2.2, which states that better predictions lead to a lower expected cost when using BPI contract functions. Since d_2 is closer to the true interval size d_T than d_1, the cost paid for Θ_1 is higher than the cost of Θ_2. This way the third party agent M has an expected profit of $C_B(\Theta_1) - C_B(\Theta_2) > 0$ by providing an improved prediction to the flexibility provider. This is visually indicated in Figure 5.9 as the difference between $C_1 = C_B(\Theta_1)$ and $C_2 = C_B(\Theta_2)$.

This process can also lead to better predictions over time. When it becomes public knowledge that M is generating a profit from providing improved predictions, it will both attract others to compete with M and put pressure

on WF to improve their predictions to lower cost. It creates a competitive environment that rewards the agent with the best predictions.

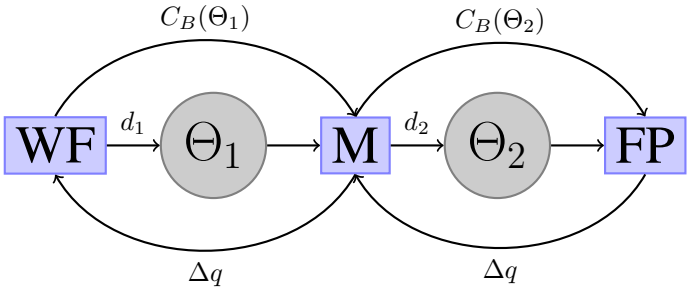


Figure 5.8: Arbitrage Setup

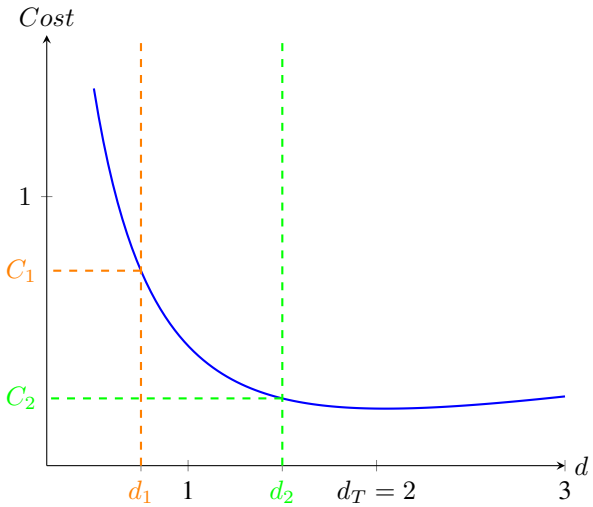


Figure 5.9: Arbitrage by Cost Difference

5.5 Discussion

Having shown the efficacy of our solution in incentivizing good predictions (Section 5.3) as well as the ability of our approach to accommodate third party

prediction providers (Section 5.4) we compare our solution to two other information theoretical approaches: data markets (Section 5.5.1) and prediction markets (Section 5.5.2)

5.5.1 Comparison to Data Markets

Data markets are mechanisms for the sharing of data between non-cooperative agents. For the use in renewable energy forecasting, data markets enable different RES owners to purchase each other's data/forecasts in order to improve their own forecasts. A market system for a single buyer and several data sellers has been proposed in (Han et al., 2022; Sun et al., 2023), while a two sided mechanism with several data buyers and sellers has been designed in (Goncalves et al., 2020).

Both data markets and our approach have the aim of improving forecasts. However, they differ conceptually in two distinct ways.

First data markets are designed as secondary markets that operate exclusively on the informational layer. The concept of BPI contracts as presented here, on the other hand, are power markets that come with balancing responsibilities.

Secondly, data markets unlike BPI contracts lack a clear incentive structure to reduce epistemic uncertainty. Purchasing additional data may lead to an improved prediction, effectively reducing epistemic uncertainty. However, this is not guaranteed and more importantly, there is no clear incentive structure to provide the most useful data. In contrast, BPI contracts directly incentivize agents to provide improved forecasts by reducing epistemic uncertainty.

5.5.2 Comparison to Prediction Markets

Predictions markets aim to to predict outcomes of future events (Wolfers and Zitzewitz, 2004) and have been applied in a wide range of applications (Horn et al., 2014). Contracts, traded on prediction markets, yield a pay-off depending on the outcome of a pre-defined event. From the discrepancy between trading price and pay-off, the probabilities of specific outcomes can be derived. Through the continuous trading process the distributed information of the participants is being aggregated.

While prediction markets generally have been around for several decades, their use in the energy literature has been a more recent development. Initial research has focused on utilizing prediction markets for hedging purposes (Shamsi and Cuffe, 2021a,c). Recently, a prediction market for complete probabilistic forecasts has been proposed by Shamsi and Cuffe (2021b).

Prediction markets are a suitable tool for providing energy forecasts, because their aim is to reveal information about a future event. Such markets, however, differ from BPI in the way they seek to achieve this goal. While prediction markets try to leverage the wisdom of the crowds to provide an aggre-

gate prediction, BPI contracts directly incentivize accurate predictions within the contract structure.

Similar to data markets, prediction markets differ from the concepts presented here by operating only on the informational layer. In the case of prediction markets, this disconnect from the operation of the grid, makes them more susceptible to speculation. BPI contracts on the other hand are power contracts first and come with balancing obligations.

5.6 Conclusion

The work presented here contributes to the energy flexibility literature by providing the concept of Best Prediction Incentivizing contracts. The concept addresses the often overlooked issue of how to incentivize agents to accurately predict their fluctuating future energy production. The concept differs from Incentive Compatibility by focusing on reducing epistemic uncertainty about the real world rather than just incentivizing the revelation of private information.

Furthermore, we showed the applicability of this concept to the real world. Starting by providing BPI pricing functions for simple cases – uniform and single term polynomial distributions – we then showed that these simple solutions can be tuned to approximate more realistic Gaussian distributions.

Finally, we put our results into a larger context. We explored how third party agents are incentivized by the arbitrage potential to contribute improved predictions and compared our solution to other information theoretical approaches.

Going forward, there still remain some open questions around the integration of BPI contracts into larger energy systems. For instance, it is unclear how wide-spread existence of such contracts would affect balancing markets. Further research also has to be done on how inter-dependencies of different events affect their corresponding contracts.

6 Conclusion

The last four chapters have given us a deeper look at the questions posed at the beginning of this work. Following this, we will connect our findings to the specific research questions (Section 6.1), provide some broad research directions (Section 6.2) and end with more general remarks on the energy transition and global climate change (Section 6.3).

6.1 Addressing Research Questions

In the following, we will address the individual research questions based on the insights we have gained from the previous chapters and thereby shed some light on the main question of this thesis.

6.1.1 Research Question 1

The first research question is about the issue of how to balance the need for incentives with other desired properties.

Research Question 1
How can we incentivize honesty in multi-energy markets?

This question was explored in the context of a combined heat and power unit demand/supply double sided auction, in which also power to heat devices can participate.

At the heart of the issue lies the basic limitation that a market cannot simultaneously be optimal, IC, IR and weakly budget balanced. Given that IC and IR aligns the optimization objective of the market with actual social welfare and weak budget balance ensures that the market is financially sustainable, we are only left with optimality to compromise on. As we have shown in Chapter 2, trade reduction mechanisms can be used to maintain IC, IR and weak budget balance by removing the least valuable trades from the allocation and use them for pricing the still allocated agents. By singling out the least valuable trades, the mechanism only has a small negative effect on social welfare.

Future research on the topic could explore the effect of allowing a wider range of bundles on the social welfare loss as well as the application of trade reduction to other domains.

6.1.2 Research Question 2

The second research question addresses the issue of how multi-energy flexibility can be integrated in the current energy system in the absence of dedicated multi-energy flexibility markets.

Research Question 2
How can access to a district heating network be used for energy procurement in current electricity markets?

This question was explored in the context of a flexible in time electricity consumer with a District Heating Network (DHN) participating in simultaneous day-ahead auctions.

The essential motivation behind this question is that flexible in time energy-systems will need to be integrated before dedicated flexibility markets have been widely established. It, therefore, becomes important to understand how flexibility can be made useful within traditional energy markets. While flexibility in time enables an agent to shift their consumption to when electricity is cheapest, access to a DHN allows for excess energy to be balanced without incurring imbalance costs. It is not clear how this combination of time flexibility and energy carrier flexibility affect the optimal bidding strategy of agents.

In Chapter 3, we derive an optimal bidding criteria for a setting of multi-unit demand and provide closed form solutions for identical auctions as well as a numerical approach for non-identical auctions. The results show that the ability to balance excess energy generally affects the optimal bidding strategy. In the absence of such a balancing ability, the optimal bidding strategy is to participate in n auctions and bid zero in all other auctions. On the other hand, when there is such a balancing ability available, the agent is better off participating in all auctions to varying degrees. In most cases this leads to a more aggressive bidding strategy in which the agent is expected to win more units of energy than they need.

It is important to note that we assume the energy used in the DHN has no value. We, thereby, show that a connection to the DHN adds value simply by its ability to absorb excess energy. The addition of the connection to the DHN simply prevents high imbalance costs in case the agent obtained too many units of energy. This ability to avoid an imbalance charge in itself is enough to cause a change in bidding behaviour.

Further research should focus on the real world applicability of our approach with regard to both technical and economic feasibility. In regard to economic feasibility, the most important assumption to re-evaluate is that the

energy used in the DHN has no value itself. In this sense this work provided a lower bound on the usefulness of power-to-heat connections by assuming heat to have no intrinsic value at all. However, usually heating is not free and therefore has some economic value. In regard to technical feasibility, it is important to explore the assumption that the district heating network is able to absorb any excess energy. In practical settings a DHN does have limitations in the amount of energy and maximum power it can absorb. Whether these technical limitations limit its operational ability to absorb excess energy from a flexible energy consumer, will vary from case to case and depends on the relative size of DHN to energy consumer. To see to which degree the foundational results of this work are applicable to a particular scenario, this question of feasibility has to be addressed.

6.1.3 Research Question 3

The third research question is the first of two questions related to broader conceptual challenges with flexibility markets: the question of what is considered fair in flexible energy systems.

Research Question 3
What is fairness in flexible energy systems and how can it be achieved?

This question was explored in the context of a double sided auction with unit supply/demand agents that are flexible within a continuous time window.

What distinguishes flexible from regular energy systems and thereby makes this question interesting is their ability to yield outcomes that differ in details yet result in the same overall welfare. This sameness while also being different can be expressed in time-flexible systems as the question:"When is electricity at 3pm the same as at 5pm?". This ambiguity causes complications in the traditional fairness concept of non-discrimination: the same price for the same item is no longer an adequate definition if there is uncertainty about which units of energy are the same.

We explored this ambiguity in Chapter 4, by making the connection between the psychological foundations and mathematical formulations of fairness. We then used this connection to extend the definition of non-discrimination to a new concept called *consistent pricing over equivalent outcomes*. The application of this concept requires finding a computationally efficient representation of the set of all equivalent outcomes. For our setting we found such a representation in the form of a directed graph in which simple cycles indicate an alternative yet equivalent matching. Connected components within this directed graph represent agents that should be priced equally under consistent pricing over equivalent outcomes. To better understand our approach, we put the proposed concept in the context of other fairness metrics as well as marketing psychology. In particular, we show that our fairness

approach can be viewed as a relaxation of envy-freeness. In terms of marketing psychology, our fairness approach follows the idea of procedural justice, which emphasizes the process being fair. This contrasts with distributed justice, which considers prices fair if they are in line with reference prices – in the case of envy-freeness the reference is other agents.

The broader implication of our findings is that usually well defined concepts such as non-discrimination can become ill-defined in flexible energy systems. As flexibility across both time and energy carriers adds several dimensions of how to deal with a problem, it unfortunately also introduces ambiguity about which approach is optimal. In order to put the concept back on proper footing, we have to reconsider the initial intent of the concept rather than its mathematical formulation.

While the concept as presented here can generally be applied to any system in which outcomes can be different yet equivalent, it does require a computationally efficient representation of equivalent outcomes. Any future application of the concept will have to face this challenge. For some settings it might simply not be possible to provide such a representation. In such cases, it might be useful to consider bounded fairness constraints. Instead of aiming for complete fairness, a system can ensure that there is a limit on the favour or disadvantage given to one agent over another.

6.1.4 Research Question 4

The fourth Research Question deals with the challenge of sharing accurate prediction models between non-cooperative agents. Specifically it asks:

Research Question 4
How to incentivize accurate predictions in flexibility markets?

This question was explored in the context of a peer-to-peer flexibility market with insurance contracts.

What makes the question interesting is the non-cooperative nature of markets further complicated by the existence of epistemic uncertainty. In a market setting, agents strive to optimize for their own utility. An agent seeking to purchase flexibility to compensate for a fluctuating renewable energy source will attempt to obtain that flexibility at the lowest price possible. Most likely the agent has some model describing their renewable energy fluctuation. However, the only reason why an agent would share their model is if it is in their own best interest. In flexibility markets, the challenge of creating incentives for the agents is further complicated by epistemic uncertainty. Even when incentivized to report truthfully, the agent does not have perfect knowledge and is, therefore, simply not able to provide the true model.

We addressed this specific issue in Chapter 5 by showing that Incentive Compatibility is no longer unambiguous or at the very least achievable when

there is epistemic uncertainty. Truth revelation is not enforceable, when the agent does not know the truth. We then provided a broader concept: Best Prediction Incentivizing under epistemic uncertainty. It addresses the complication introduced by epistemic uncertainty by rewarding the revelation of improved models rather than just the true model. This way the agent does not need to know the true model but can operate more profitably by improving their model of their energy fluctuation. Since concepts are only useful if they can reasonably be applied to the real world, we provided solutions that can be fine-tuned to fit specific settings.

The broader insight we gained is that for incentives in flexibility markets there needs to be a more elaborate coupling between the physical and market layer. Unlike regular energy markets in which the price is solely determined by competition, i.e. on the market level, the resulting payments in flexibility markets need to be informed by the outcome on the physical level, since that is the only place that can provide feedback about the accuracy of the prediction.

The broader implication of our findings is that trading goods that relate to uncertainties requires a different approach to incentives. While we specifically looked at the incentive to improve the reported prediction, we did not consider the incentive to participate in the first place. Since payments depend on an uncertain outcome and are only incentivizing in expectation, there can still be indidivual outcomes that put an agent at a disadvantage for having participated. Given the uncertainty of the outcome, it becomes important to ensure that the worst case scenarios do not discourage agents from participating.

With regards to applying our findings to actual systems, it is important to keep in mind that neither our conceptual approach nor the specific contracts exist in a vacuum. Our solution approach of peer-to-peer contracts has to compete with alternative approaches such as prediction markets, which might be easier to implement in some situations but have a less tight coupling between the information and physical level. Furthermore, when applying the approach, we presented in Chapter 5, it is currently unclear how the interdependencies between energy fluctuations affects the flexibility contracts.

6.1.5 Main Research Question

So far we have examined specific questions related to multi-dimensional energy systems. These questions have ranged from the practical issues of day-ahead bidding to the more conceptual inquiries into the meaning of fairness. Equipped with these insights, we can look towards the main question of this work: the interaction between multi-dimensional energy and energy markets.

Main Research Question
What are ways to integrate multi-dimensional flexibility into current and future energy markets.

Broadly speaking, the overall insights presented in this work can be grouped into short term operational and long term conceptual insights.

Regarding short term insights, we have seen that the interconnected nature of modern energy systems affects the behaviour of individual participants with yet unknown long term effects on the operation of traditional energy markets. The ability to shift energy allows market participants to over-commit to their energy purchases as excess energy can be shuffled around later with the overall benefit of lower operational cost. This can push market agents to participate more broadly and aggressively across energy markets. In the end, the interconnected nature of modern energy systems will affect the behaviour of conventional energy markets in the future.

Regarding the long term implications, we have seen the importance of rethinking basic concepts such as honesty, uncertainty and fairness. Even though, we have shown that via trade reduction multi-energy markets can retain the most important mechanism design properties, we have also explored the unforeseen challenges that come with introducing flexibility to the market itself.

Future energy markets addressing the issue of renewable energy fluctuations by taking into account the interconnectedness of time and different energy carriers will have to re-evaluate basic concepts such as fairness and honesty. As we have shown in the last two research questions, the introduction of flexibility and uncertainty have rendered previously well defined concepts ill-defined. Ultimately, the implication of multi-dimensional energy systems for future energy market design is that well established concepts have to be re-examined to ensure that their definitions are still appropriate for whichever new market setting they are being applied to.

6.2 Further Research Directions

From the insights we have gained so far, we now look towards further research directions that might benefit the broader issue of market design for energy integration. The suggestions cover both mechanism design itself as well as an interdisciplinary perspective.

Regarding mechanism design it might be useful to consider weaker versions of Incentive Compatibility for applications in the energy domain. Incentive Compatibility is difficult to achieve, especially as market settings are becoming more complicated. However, as we have already argued in Chapter 5, when agents do not have perfect knowledge of themselves, incentives might hold even in the absence of Incentive Compatibility. Under the assumption that an agent has some epistemic uncertainty about their own valuation, it might be

possible to design incentives that work in practical terms for settings in which Incentive Compatibility is difficult or impossible to achieve.

There are two aspects of interdisciplinary research where the field of energy market integration could stand to benefit from widening its scope: market psychology to guide design goals and financial mathematics to find better solutions.

The connection from market design towards market psychology might yield relevant insights into which properties of a market are important to the adoption of a market or technology more general. For instance, while mechanism design research generally assumes individual rationality to be an important property, it might actually not be in practice. Companies might see an entrance fee simply as the cost of doing business and, therefore, be satisfied by a mechanism that is not individually rational. It is important to realize that the question of how companies within the energy domain perceive such trade-offs is important to the the field of mechanism design, but it is not a question that should be answered by it. Market psychology, on the other hand, can explore these questions and provide design goals or indicate which aspects of a new market design to highlight.

Financial mathematics should be viewed as a potential source of solution approaches. The new challenges that will emerge within the energy domain might look nothing like the issues that engineering fields are used to, while resembling problems from the financial sector. Financial tools such as insurances and options are well established mechanisms for dealing with fluctuations and risk. Their application to energy systems with uncertain energy production might yield solutions without the need to reinvent the wheel. Exploring this connection is the purview of the field of energy economics.

6.3 Concluding Remarks

Coming to the end of this thesis, we leave behind the technical discussions on energy integration to make some general remarks on the topic of global climate change.

According to the most recent IPCC report (Lee et al., 2023), tackling global climate change still requires significant changes to most economic activities. A sharp reduction in greenhouse gas emissions is needed, which will require rethinking entire industrial sectors, in order to achieve the $2\ ^{\circ}C$ goal. However, there are also positive signs. It should be noted that most of the issues discussed in this work are the result of a world that is rapidly adopting green energies. The need for flexibility is a reaction to a massive integration of fluctuating renewable energy sources. With the cost of electricity generation from both wind and solar being lower than that of conventional fossil fuels, the IEA

expects renewables to not only grow in absolute terms but their share of the overall electricity production to significantly increase (IEA, 2023a). We are looking towards 2030s in which half of global electricity production is carbon neutral.

However, while there are some positive signs towards a carbon free future, the energy transition also has a socio-economic dimension to it. As it stands it seems that the burden of this transition is not equally distributed, neither on the global nor the local level.

Historically, North America and Western Europe becoming rich and prosperous has been a very energy and thereby carbon intensive process. While de-carbonizing these countries is important, the much bigger challenge will be for currently poor countries to become wealthy nations without a massive carbon footprint. The fact that carbon free technologies often come with high up front investments, further complicates this. For the poorer nations of this world, green technologies not only have to prevent carbon emissions but also provide a clear path towards prosperity.

Even within industrialized nations, the economic burden of change is distributed somewhat unevenly. In order to phase out fossil fuels, carbon certificates and thereby conventional energy will become more expensive. Such an increase in basic energy cost hits more vulnerable citizens disproportionately hard, while at the same time the technologies to alleviate this cost – electric vehicles, heat pumps, improved housing insulation – require high upfront investment.

These socio-economic challenges will require direct financial aid as well as some fine tuning of market mechanisms. But, ultimately, with the right incentives as well as some level of economic aid, we are stepping towards a sustainable, more prosperous future for all.

A Appendix

A.1 Regarding Chapter 3

A.1.1 Proofs

Proof of Lemma 3.1.

Proof. Any bid of an optimal bidding strategy is either on the edge of the domain, i.e. some bid $b_k \in \{0, \lambda_{el}\}$, or the first derivative of the cost function is zero. First, we set the first derivative of the cost function (3.1) equal to zero.

$$\frac{\partial Cost(B|T,s)}{\partial b_k} = \frac{\partial Cost_M(B|T,s)}{\partial b_k} + \frac{\partial Cost_B(B|T,s)}{\partial b_k} = 0 \tag{A.1}$$

The derivative of the market cost is simply

$$\frac{\partial Cost_M(B|T,s)}{\partial b_k} = f_k(b_k)b_k \tag{A.2}$$

For the backup cost, we have to consider the inner sum over subsets w in (3.3). We split the sum into the terms for which $k \in w$, see second line in (A.3) and into those for which $k \notin w$, see third line in A.3.

$$\begin{aligned}\frac{\partial Cost_B(B|T,s)}{\partial b_k} = f_k(b_k) \Bigg[& \\ \sum_{j\in[1,s-1]} \sum_{\substack{w\subset T\setminus\{k\} \\ s.t.|w|=j-1}} \prod_{t\in w} F_t(b_t) \prod_{\tau\in T\setminus(w\cup\{k\})} [1-F_\tau(b_\tau)](s-j)\lambda_{el} & \\ - \underbrace{\sum_{j\in[0,s-1]} \sum_{\substack{w\subset T\setminus\{k\} \\ s.t.|w|=j}} \prod_{t\in w} F_t(b_t) \prod_{\tau\in T\setminus(w\cup\{k\})} [1-F_\tau(b_\tau)](s-j)\lambda_{el}}_{Prob(x<s|T\setminus\{k\},B_{-k})} & \Bigg]\end{aligned} \tag{A.3}$$

Indeed, since $k \in w$, second line in (A.3), we know that $|w| \geq 1$ and therefore the outer sum is over $j \in [1, s-1]$. The derivative of $F_k(b_k)$ is $f_k(b_k)$. Since $k \in w$, the inner sum then sums over the subsets of w of size $(j-1)$. When $k \notin w$, third line in (A.3), we take the derivative of $[1 - F_k(b_k)]$ resulting in $-f_k(b_k)$, while the outer and inner sum stay unchanged.

Next we combine the terms of the second and third line in (A.3) into one summation. To do so, we first extend and shift the summation in the second line of (A.3):

$$\sum_{j \in [1, s-1]} \sum_{\substack{w \subset T \setminus \{k\} \\ s.t. |w| = j-1}} \prod_{t \in w} F_t(b_t) \prod_{\tau \in T \setminus (w \cup \{k\})} [1 - F_\tau(b_\tau)](s-j)\lambda_{el}.$$

We extended the summation to iterate over the range $[1, s]$, which can be done as $(s-s)\lambda_{el} = 0$. We then shift the index j by one so that both sums iterate over $j \in [0, s-1]$. This only changes the tail of the second line from $(s-j)\lambda_{el}$ to $(s-j-1)\lambda_{el}$. These operations alter the second line in (A.3) to

$$\underbrace{\sum_{j \in [0, s-1]} \sum_{\substack{w \subset T \setminus \{k\} \\ s.t. |w| = j}} \prod_{t \in w} F_t(b_t) \prod_{\tau \in T \setminus (w \cup \{k\})} [1 - F_\tau(b_\tau)]}_{Prob(x < s | T \setminus \{k\}, B_{-k})} (s-j-1)\lambda_{el}$$

Next, we combine the two terms in (A.3) to obtain

$$\frac{\partial Cost_B(B|T,s)}{\partial b_k} = -f_k(b_k) Prob(x < s | T \setminus \{k\}, B_{-k})\lambda_{el}. \tag{A.4}$$

Given (A.2) and (A.4) we can rewrite (A.1) in the following way:

$$\frac{\partial Cost(b)}{\partial b_k} = f_k(b_k) \left[b_k - Prob(x < s | T \setminus \{k\}, B_{-k})\lambda_{el}\right] = 0 \tag{A.5}$$

This equality holds if the following equation holds:

$$b_k = Prob(x < s | T \setminus \{k\}, B_{-k})\lambda_{el} \tag{A.6}$$

Next, we will show that any possibly optimal bidding strategy with bids on the boundary of the domain also satisfies Equation A.6. To see that this holds, first note that any Λ_s strategy as defined in Definition 3.1 satisfies Equation A.6. Indeed, given Λ_s bidding, for any $b_l = 0$ we know $Prob(x < s | T \setminus \{l\}, B_{-l}) = 0$, while for any $b_k = \lambda_{el}$, we know that $Prob(x < s | T \setminus \{k\}, B_{-k}) = 1$ and therefore Equation A.6 holds.

Regarding other boundary points of the domain, we will show that these strategies can easily be improved upon. If the bidding strategy contains at least s bids of λ_{el}, then it is either in Λ_s, i.e. all other bids are zero, or the strategy

can be improved by adjusting all bids, other than the s bids of λ_{el}, to zero. This holds because the penalty of the adjusted strategy is still zero while the payment for the items is strictly less. When the bidding strategy contains less than s bids of λ_{el}, then for any $b_k = \lambda_{el}$, $Prob(x < s|T \setminus \{k\}, B_{-k}) < 1$ and therefore $\frac{\partial Cost(b)}{\partial b_k}\Big|_{b_k=\lambda_{el}} > 0$, see Equation A.5. This means that reducing b_k leads to lower cost. As we can see, if there is a bid of λ_{el}, the bidding strategy can only be optimal if it is in Λ_s. The only other boundary points to consider is when there is no bid of λ_{el} and some bid of $b_k = 0$. However, since there are no λ_{el}-bids, it is not ensured that the bidding strategy wins s auctions, i.e. $Prob(x < s|T \setminus \{k\}, B_{-k}) > 0$ and therefore $\frac{\partial Cost(b)}{\partial b_k} < 0\Big|_{b_k=0}$, see Equation A.5. This means the cost can be reduced by increasing b_k, making this strategy not optimal.

This shows that an optimal bidding strategy on the boundary still satisfies Equation A.6.

□

Proof of Theorem 3.1.

Proof. Let $B = (b_1, b_2, ..., b_{|T|})$ be a bid vector which satisfies condition (3.4). Let $b_i = 0$ for some $i \in T$. Since B satisfies (3.4), we know that $Prob(x < s|T \setminus \{i\}, B_{-i}) = 0$. This implies that it is certain that the agent wins s auctions from $T \setminus \{i\}$, which can only be true if we submit λ_{el} to at least s auctions. However, submitting λ_{el} in more than s auctions cannot satisfy (3.4) as it would follow that for every auction $k \in T$ the agent wins s auctions from the remaining $T \setminus \{k\}$ auctions with certainty. Therefore $Prob(x < s|T \setminus \{k\}, B_{-k}) = 0$ holds, causing a bid of 0, for every auction $k \in T$ contradicting the previous statement of bidding λ_{el} in more than s auctions. Thus the agent bids λ_{el} in exactly s auctions. By (3.4) the agent bids 0 in the remaining $(|T| - s)$. The proof works similarly for $b_i = \lambda_{el}$. □

Proof of Lemma 3.2.

Proof. Given a uniform bid $B_u = (b_u, b_u, ..., b_u)$, the set of equations in (3.4) becomes a set of the same equation and $Prob(x < s|T \setminus \{k\}, B_{-k})\lambda_{el}$ becomes a function of b_u. Therefore, we only need to find the value of b_u for which this one equation holds:

$$b_u = Prob(x < s|T \setminus \{k\}, B_{-k})\lambda_{el}.$$

As $Prob(x < s|T \setminus \{k\}, B_{-k})\lambda_{el}$ is λ_{el} for $b_u = 0$, 0 for $b_u = \lambda_{el}$ and strictly decreasing and continuous in between, there can only exist one crossover point where $b_u = Prob(x < s|T \setminus \{k\}, B_{-k})\lambda_{el}$. □

Proof of Lemma 3.3

Proof. Let B be a bid vector that satisfies (3.4) and let b_l and b_m be the bids on auctions l and m, respectively. The proof will show that either $B \in \Lambda_s$ or $b_l = b_m$. Let $B_{-l,m}$ be the bid vector B without the bids b_l and b_m. We assume $B_{-l,m}$ to be fixed and solve (3.4)for b_l and b_m, respectively (derivation can be found in Appendix A.1.2):

$$\begin{aligned} b_l &= Prob(x < s | T \setminus \{l,m\}, B_{-l,m})\lambda_{el} \\ &\quad - Prob(x = s-1 | T \setminus \{l,m\}, B_{-l,m})\lambda_{el} F(b_m) \\ b_m &= \underbrace{Prob(x < s | T \setminus \{l,m\}, B_{-l,m})\lambda_{el}}_{a} \\ &\quad - \underbrace{Prob(x = s-1 | T \setminus \{l,m\}, B_{-l,m})}_{c} \lambda_{el} F(b_l) \end{aligned} \tag{A.7}$$

Given the constants a and c and the fact that $F(\tau) = \frac{\tau}{\lambda_{el}}$, (A.7) can be rewritten as follows:

$$\begin{aligned} b_l &= a - c b_m \\ b_m &= a - c b_l \end{aligned}$$

Subtracting the second from the first equation leaves us with $b_l - b_m = c(b_l - b_m)$, which can only be true if $b_l = b_m$ or $c = 1$. However, $c = 1$, i.e. $Prob(x = s-1 | T \setminus \{l,m\}, B_{-l,m}) = 1$, means winning exactly $s-1$ units from auctions $T \setminus \{l,m\}$ with certainty. This can only be true when $B \in \Lambda_s$ and $(b_l = 0, b_m = \lambda_{el})$ or $(b_l = \lambda_{el}, b_m = 0)$. Therefore either $B \in \Lambda_s$ or B is a uniform bid over all auctions. □

Proof of Lemma 3.4

Proof. If the uniform bid $b_u = \lambda_{el}$, then $Prob(x < s | T \setminus \{k\}, B_{-k}) = 1 \quad \forall k \in T$. This can only be the case if the requested number of units s is equal the number of available auctions $n = |T|$. In this case $B \in \Lambda_s$. Note that $b_u = 0$ is impossible as it would require $Prob(x < s | T \setminus \{k\}, B_{-k})\lambda_{el} = 0 \quad \forall k \in T$, i.e. winning s auctions with certainty while submitting a bid of 0. Going forward, we will only consider non-boundary bids, i.e. $0 < b_u < \lambda_{el}$.

The general condition for a local minimum is that the Hessian matrix M is positive definite(Lax and Terrell, 2017), i.e. $x^T M x > 0 \quad \forall x \in \mathbb{R}^{|T|}$. Given are the uniform bid B_u and identical uniform price distributions. The Diagonal elements of the Hessian M, are identical d, with $d = \frac{\partial^2 Cost(b)}{(\partial b_k)^2}\Big|_{B_u} = f_k(b_u)$,

while all off-diagonal elements are

$$p = \left.\frac{\partial^2 Cost(b)}{\partial b_l \partial b_k}\right|_{B_u} = f_k(b_u) f_l(b_u) Prob(x = s-1 | T \setminus \{k,l\}) \lambda_{el}.$$

A derivation for these second order derivatives can be found in Appendix A.1.3. Note that for uniform price distributions $f_l(b_u)\lambda_{el} = 1$ and that for uniform bidding $Prob(x = s-1|T \setminus \{k,l\}) < 1$. Therefore $d > p$.

$$\begin{aligned} x^T M x &= \left[\sum_{t_1 \in T} \left(x_{t_1} \sum_{t_2 \in T \setminus \{t_1\}} x_{t_2} \right)\right] p + \sum_{t_1 \in T} x_{t_1}^2 d \\ &> \left[\sum_{t_1 \in T} \left(x_{t_1} \sum_{t_2 \in T \setminus \{t_1\}} x_{t_2} \right)\right] p + \sum_{t_1 \in T} x_{t_1}^2 p = (x^T A x) p. \end{aligned}$$

Matrix A is 1 at every entry, making all rows linearly dependent and therefore A has rank 1 and only one non-zero eigenvalue(Aleskerov et al., 2011) exists. Any eigenvalue, eigenvector (λ, x) pair has to satisfy $Ax = \lambda x$, which holds for $x = (x_1, x_1, ..., x_1)$ and $\lambda = |T| > 0$. Since all eigenvalues of A are non-negative, A is positive semi-definite(Aleskerov et al., 2011), i.e. $x^T A x \geq 0 \quad \forall x \in \mathbb{R}^{|T|}$, and therefore $x^T M x > 0 \quad \forall x \in \mathbb{R}^{|T|}$. □

Proof of Theorem 3.2

Proof. Given Lemma 3.3 we know that uniform bidding and Λ_s-bidding are the only two strategies satisfying Equation (3.4). Given these two options, we only need to show that the cost of uniform bidding is always at most the cost of Λ_s-bidding. We define $B^u_{n,s}$ as the uniform bid vector with bid value $b^u_{n,s}$ that satisfies Equation (3.4) for requiring s units out of n auctions. Similarly, let $B^\lambda_{n,s}$ be the Λ_s-bidding vector for needing s units out of n auctions.

We will, first, derive the general cost of bidding $B^\lambda_{n,s}$. When bidding $B^\Lambda_{n,s}$, a bid of λ_{el} is submitted to s auctions and a bid of 0 is submitted otherwise. This means that s auctions are certainly won and all others certainly lost. The expected cost of each auction won is $0.5\lambda_{el}$, while the expected backup cost is 0, since s auctions are won. Therefore,

$$Cost_u(B^\lambda_{n,s}|n,s) = s \cdot 0.5\lambda_{el}. \tag{A.8}$$

We start by showing that the theorem holds for a more restricted case of demanding $n-1$ units out of n auctions. [1] We do so via proof by induction.
Base Case:

[1] For convenience of notation, we say that $n = |T|$ is the number of auction.

Let $n = 2$ and $s = 1$. Note that for this edge case uniform bidding while not a unique optimal solution still minimises the overall cost, see Appendix A.1.4.
Inductive Step: Given that the theorem holds for $n = 2$ and $s = 1$, we will now show that it holds for general n and $s = n - 1$. Assuming

$$Cost_u(B^u_{n,n-1}|n, n-1) \leq Cost_u(B^\lambda_{n,n-1}|n, n-1),$$

we show that

$$Cost_u(B^u_{n+1,n}|n+1, n) \leq Cost_u(B^\lambda_{n+1,n}|n+1, n)$$

by constructing $\tilde{B}_{n+1,n} = (B^u_{n,n-1}, \lambda_{el})$, i.e. bidding in n auctions according to the uniform bid $B^u_{n,n-1}$ and submitting λ_{el} in the last auction. Since $\tilde{B}_{n+1,n}$ wins the last auction with certainty, $F(\lambda_{el}) = 1$, $Cost(\tilde{B}_{n+1,n}|n+1, n)$ can be expressed as

$$Cost\left(\tilde{B}_{n+1,n}\middle|n+1, n\right) = Cost_u\left(B^u_{n,n-1}\middle|n, n-1\right) + 0.5\lambda_{el}.$$

Given Equation (A.8), we can express the cost of submitting $B^\lambda_{n+1,n}$ as

$$Cost_u\left(B^\lambda_{n+1,n}\middle|n+1, n\right) = Cost_u\left(B^\lambda_{n,n-1}\middle|n, n-1\right) + 0.5\lambda_{el}.$$

Using the inductive assumption, we get

$$Cost_u\left(\tilde{B}_{n+1,n}\middle|n+1, n\right) \leq Cost_u\left(B^\lambda_{n+1,n}\middle|n+1, n\right). \tag{A.9}$$

This shows that the cost of the constructed bid $\tilde{B}_{n+1,n}$ is at most that of $B^\lambda_{n+1,n} \in \Lambda_s$.

Next, we show that the cost of uniform bid $B^u_{n+1,n}$ is at most that of $\tilde{B}_{n+1,n}$. Note, that since $s < n$, $B^u_{n,n-1} \notin \Lambda_s$, see the beginning of the proof of Lemma 3.4. This means that the constructed bid $\tilde{B}_{n+1,n}$ is neither uniform nor in Λ_{s+1}. Lemma 3.3 then tells us that $\tilde{B}_{n+1,n}$ is not a local minimum. Therefore its cost has to be strictly larger than the cost of at least one local minimum, which by (A.9) cannot be $B^\lambda_{n+1,n}$. The only alternative according to Lemma 3.3 is the uniform bid $B^u_{n+1,n}$, which by Lemma 3.4 is always a local minimum. Therefore,

$$\begin{aligned} Cost_u\left(B^u_{n+1,n}\middle|n+1, n\right) &< Cost_u\left(\tilde{B}_{n+1,n}\middle|n+1, n\right) \\ &\leq Cost_u\left(B^\lambda_{n+1,n}\middle|n+1, n\right). \end{aligned}$$

Having shown that the theorem holds for $s = n - 1$ out of n auctions, we next generalise this result to the setting of s units out of n auctions. To do so, we introduce $\tilde{B}_{n,s} = (B^u_{s+1,s}, 0)$, i.e. bidding uniformly as if there only

existed $s+1$ auctions and submitting 0 to all other auctions. Note that adding auctions in which the agent bids 0 has no effect on the cost, because a bid of 0 incurs no market cost and also does not change the probability of needing the backup generator. Therefore,

$$Cost_u\left(\tilde{B}_{n,s}\middle|n,s\right) = Cost_u\left(B^u_{s+1,s}\middle|s+1,s\right)$$

as well as

$$Cost_u\left(B^\lambda_{n,s}\middle|n,s\right) = Cost_u\left(B^\lambda_{s+1,s}\middle|s+1,s\right).$$

Since we already established the validity of the theorem for $s = n-1$ out of n auctions, we know

$$Cost_u\left(B^u_{s+1,s}\middle|s+1,s\right) \leq Cost_u\left(B^\lambda_{s+1,s}\middle|s+1,s\right)$$

and using this with the previous two equations we get

$$Cost_u\left(\tilde{B}_{n,s}\middle|n,s\right) \leq Cost_u\left(B^\lambda_{n,s}\middle|n,s\right) \tag{A.10}$$

Note that $\tilde{B}_{n,s}$ is neither a uniform bid nor in Λ_s and, therefore, by Lemma 3.3 is not a local minimum. This means that $\tilde{B}_{n,s}$ has to result in a cost that is larger than a local minimum. By Lemma 3.3, the only local minima are uniform bidding or in Λ_s. However, given the inequality in (A.10), Λ_s-bidding cannot be cheaper than $\tilde{B}_{n,s}$.

This leads to the following inequality, concluding the proof:

$$Cost_u\left(B^u_{n,s}\middle|n,s\right) < Cost_u\left(\tilde{B}_{n,s}\middle|n,s\right) \leq Cost_u\left(B^\lambda_{n,s}\middle|n,s\right).$$

□

A.1.2 Solving $B_{-l,m}$ for a Pair of Bids in Lemma 3.3

We assume a fixed bid vector $B_{-l,m}$ and try to find the bids b_l and b_m such that (3.4) is satisfied. The procedure is the same for b_l and b_m, we will therefore only consider b_l. Recall (3.4).

$$b_l = Prob(x < s | T \setminus \{l\}, B_{-l})\lambda_{el}$$

. Winning fewer than s units from $T \setminus \{l\}$ is the probability of winning fewer than s units when losing auction m plus the probability of winning fewer than $s-1$ units when winning auction m.

$$\begin{aligned} b_l = {} & F(b_m)Prob(x < s | T \setminus \{l,m\}, B_{-l,m})\lambda_{el} \\ & + [1 - F(b_m)]Prob(x < s-1 | T \setminus \{l,m\}, B_{-l,m})\lambda_{el} \end{aligned}$$

This can be rewritten as

$$b_l = Prob(x < s|T \setminus \{l,m\}, B_{-l,m})\lambda_{el} \\ - Prob(x = s-1|T \setminus \{l,m\}, B_{-l,m})\lambda_{el}F(b_m)$$

A.1.3 Second Derivatives in Lemma 3.4

We derive the second derivative for when submitting a uniform bid that satisfies (3.4). Recall that the first order derivative is

$$\frac{\partial Cost(b)}{\partial b_k} = f_k(b_k)\left[b_k - Prob(x < s|T \setminus \{k\}, B_{-k})\lambda_{el}\right]$$

Therefore,

$$\frac{\partial^2 Cost}{(\partial b_k)^2} = \frac{\partial f_k(b_k)}{\partial b_k}\overbrace{\left[b_k - Prob(x < s|T \setminus \{k\}, B_{-k})\lambda_{el}\right]}^{=0} \\ + f_k(b_k)\frac{\partial}{\partial b_k}\left[b_k - Prob(x < s|T \setminus \{k\}, B_{-k})\lambda_{el}\right]$$

'Note, that the first term in the above equation is zero because B_u satisfies (3.4) and that $Prob(x < s|T \setminus \{k\}, B_{-k})$ is not a function of b_k. Therefore, $\frac{\partial^2 Cost}{(\partial b_k)^2}\Big|_{B_u} = f_k(b_k)$.

Similarly for distinct k and l ($k \neq l$), the following holds

$$\frac{\partial^2 Cost}{\partial b_k \partial b_l} = -f_k(b_k)\frac{\partial}{\partial b_l}Prob(x < s|T \setminus \{k\}, B_{-k})\lambda_{el} \tag{A.11}$$

in which $Prob(x < s|T \setminus \{k\}, B_{-k})$ can be expressed more verbosely as

$$\sum_{j\in[0,s-1]} \sum_{\substack{w\subset T\setminus\{k\} \\ s.t.|w|=j}} \prod_{t\in w} F_t(b_t) \prod_{t\in T\setminus(w\cup\{k\})} [1 - F_t(b_t)]$$

The following derivation works similar to what has been done in the proof for Lemma 3.1. For the derivative of the above mentioned term we need to separately consider when $l \in w$ and when $l \notin w$. When $l \in w$, we know that $|w| \geq 1$, causing the outer sum to run over the interval $[1, s-1]$, we take the derivative of $F_l(b_l)$ and obtain the following

$$f_l(b_l) \sum_{j\in[1,s-1]} \sum_{\substack{w\subset T\setminus\{k,l\} \\ s.t.|w|=j-1}} \prod_{t\in w} F_t(b_t) \prod_{t\in T\setminus(w\cup\{k,l\})} [1 - F_t(b_t)].$$

We shift the outer sum by one resulting in

$$f_l(b_l) \underbrace{\sum_{j\in[0,s-2]} \sum_{\substack{w\subset T\setminus\{k,l\} \\ s.t.|w|=j}} \prod_{t\in w} F_t(b_t) \prod_{t\in T\setminus(w\cup\{k,l\})} [1-F_t(b_t)]}_{Prob(x<s-1|T\setminus\{k,l\},B_{-k,l})}. \quad \text{(A.12)}$$

When $l \notin w$, the outer sum is not affected and we take the derivative of $[1 - F_l(b_l)]$ and therefore obtain

$$-f_l(b_l) \underbrace{\sum_{j\in[0,s-1]} \sum_{\substack{w\subset T\setminus\{k,l\} \\ s.t.|w|=j}} \prod_{t\in w} F_t(b_t) \prod_{t\in T\setminus(w\cup\{k,l\})} [1-F_t(b_t)]}_{Prob(x<s|T\setminus\{k,l\},B_{-k,l})}. \quad \text{(A.13)}$$

Combining (A.12) and (A.13) and substituting the term in (A.11) we obtain:

$$\frac{\partial^2 Cost}{\partial b_k \partial b_l} = f_k(b_k) f_l(b_l) Prob(x = s-1|T\setminus\{k,l\}, B_{-k,l})\lambda_{el}.$$

A.1.4 Considering case s=1 and |T|=2

for Theorem 3.2 The case of requiring $s = 1$ units out of $n = |T| = 2$ auctions is an edge case that has to be considered separately. We will show that, while uniform bidding under the assumption of identical uniform distributions is not a unique minima, it still minimises the cost. We do so by showing that uniform bidding in this scenario is one of many cost minimising bidding strategies.

For this scenario, we denote the bid as $B = (b_k, b_l)$. Recall that the first derivative from equation (A.5) is as follows:

$$\frac{\partial Cost(b)}{\partial b_k} = f_k(b_k)\left[b_k - Prob(x < s|T\setminus\{k\}, B_{-k})\lambda_{el}\right] = 0$$

Note that for this case, $Prob(x < s|T\setminus\{k\})$ is reduced to the probability of losing auction l, i.e. $[1 - F_l(b_l)]$.

Therefore, the derivatives become

$$\frac{\partial Cost(b)}{\partial b_k} = f_k(b_k)\left[b_k - [1-F_l(b_l)]\,\lambda_{el}\right]$$
$$\frac{\partial Cost(b)}{\partial b_l} = f_l(b_l)\left[b_l - [1-F_k(b_k)]\,\lambda_{el}\right]$$

Given that we assume uniform identical distributions, i.e. $f_l(b) = f_b(b) = \frac{1}{\lambda_{el}}$ and $F_l(b) = F_l(b) = \frac{b}{\lambda_{el}}$, the derivatives can be simplified to:

$$\frac{\partial Cost(b)}{\partial b_k} = \frac{b_k}{\lambda_{el}} - \left[1 - \frac{b_l}{\lambda_{el}}\right]$$
$$\frac{\partial Cost(b)}{\partial b_l} = \frac{b_l}{\lambda_{el}} - \left[1 - \frac{b_k}{\lambda_{el}}\right]$$

The optimality condition of the first derivative being 0 becomes:

$$b_k - \lambda_{el} + b_l = 0$$
$$b_l - \lambda_{el} + b_k = 0$$

Note that these two conditions are the same and we are left with an underdetermined system of equations leading to an infinite number of potential solutions. These solutions are described by:

$$B = (b_k, \lambda_{el} - b_k) \quad b_k \in [0, \lambda_{el}] \tag{A.14}$$

This set of solutions also contains the uniform bid $B = (\frac{\lambda_{el}}{2}, \frac{\lambda_{el}}{2})$.

The second degree derivatives are easily calculated:

$$\frac{\partial^2 Cost(b)}{(\partial b_k)^2} = \frac{\partial^2 Cost(b)}{(\partial b_l)^2} = \frac{\partial^2 Cost(b)}{\partial b_k \partial b_l} = \frac{1}{\lambda_{el}}$$

This leads to a Hessian matrix with a set of eigenvalues $\{0, \frac{2}{\lambda_{el}}\}$. The Hessian is therefore positive semi-definite, which means the cost function is convex (Boyd and Vandenberghe, 2004). The set of bids in Equation (A.14) are therefore extreme points in a convex function.

Next we will show that all these points result in the same cost. The total cost is composed of market and backup costs. For the case of $n = 2$ and $s = 1$, the market cost, see equation (3.2), becomes:

$$\begin{aligned} Cost_M &= \int_0^{b_k} \tau f_k(\tau) \partial\tau + \int_0^{b_l} \tau f_l(\tau) \partial\tau \\ &= \frac{1}{2}\frac{b_k^2}{\lambda_{el}} + \frac{1}{2}\frac{b_l^2}{\lambda_{el}} \\ &= \frac{1}{2\lambda_{el}}\left(b_k^2 + b_l^2\right) \end{aligned}$$

The backup cost, see equation (3.3), is simply the probability of losing both auctions multiplied by λ_{el}:

$$\begin{aligned} Cost_B &= [1 - F_k(b_k)]\,[1 - F_l(b_l)]\,\lambda_{el} \\ &= \left[1 - \frac{b_k}{\lambda_{el}}\right]\left[1 - \frac{b_l}{\lambda_{el}}\right]\lambda_{el} \\ &= \left[1 - \frac{b_k}{\lambda_{el}} - \frac{b_l}{\lambda_{el}} + \frac{b_k b_l}{\lambda_{el}^2}\right]\lambda_{el} \\ &= \lambda_{el} - b_k - b_l + \frac{b_k b_l}{\lambda_{el}} \end{aligned}$$

Combining the market and backup cost, the total cost becomes:

$$Cost = \frac{1}{2\lambda_{el}}\left(b_k^2 + b_l^2\right) + \lambda_{el} - b_k - b_l + \frac{b_k b_l}{\lambda_{el}}$$

Using the previously derived condition for an optimal bid ($b_l = \lambda_{el} - b_k$) see equation (A.14), the total cost for the case of $n = 2$ and $s = 1$ can be simplified to:

$$\begin{aligned} Cost &= \frac{1}{2\lambda_{el}}\left(b_k^2 + [\lambda_{el} - b_k]^2\right) + \lambda_{el} - b_k - [\lambda_{el} - b_k] + \frac{b_k\,[\lambda_{el} - b_k]}{\lambda_{el}} \\ &= \frac{1}{2\lambda_{el}}\left(b_k^2 + [\lambda_{el} - b_k]^2\right) + \frac{b_k\,[\lambda_{el} - b_k]}{\lambda_{el}} \\ &= \frac{1}{2\lambda_{el}}\left(b_k^2 + \lambda_{el}^2 - 2\lambda_{el} b_k + b_k^2\right) + b_k - \frac{b_k^2}{\lambda_{el}} \\ &= \frac{1}{2\lambda_{el}}\left(2b_k^2 + \lambda_{el}^2 - 2\lambda_{el} b_k\right) + b_k - \frac{b_k^2}{\lambda_{el}} \\ &= \frac{b_k^2}{\lambda_{el}} + \frac{1}{2}\lambda_{el} - b_k + b_k - \frac{b_k^2}{\lambda_{el}} \\ &= \frac{1}{2}\lambda_{el} \end{aligned}$$

This means that independent of the exact choice of b_k or b_l, the resulting cost is $0.5\lambda_{el}$ as long as bid B satisfies condition (A.14). We therefore have a set of cost minimising bids, one of which is the uniform bid $B = (0.5\lambda_{el}, 0.5\lambda_{el})$.

Bibliography

Mohammad Hosein Abbasi, Badr Abdullah, Muhammad Waseem Ahmad, Ali Rostami, and Jeff Cullen. Heat transition in the european building sector: Overview of the heat decarbonisation practices through heat pump technology. *Sustainable Energy Technologies and Assessments*, 48:101630, 2021.

Christopher O. Adika and Lingfeng Wang. Non-cooperative decentralized charging of homogeneous households' batteries in a smart grid. *IEEE Transactions on Smart Grid*, 5(4):1855–1863, jul 2014. doi: 10.1109/tsg.2014.2302449.

Mehrdad Aghamohamadi and Amin Mahmoudi. From bidding strategy in smart grid toward integrated bidding strategy in smart multi-energy systems, an adaptive robust solution approach. *Energy*, 183:75–91, 2019.

Alessandro Agnetis, Gianluca De Pascale, Paolo Detti, and Antonio Vicino. Load scheduling for household energy consumption optimization. *IEEE Transactions on Smart Grid*, 4(4):2364–2373, 2013.

S. Alaei, J. Hartline, Rad Niazadeh, Emmanouil Pountourakis, and Y. Yuan. Optimal auctions vs. anonymous pricing. *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, pages 1446–1463, 2015.

Martin Aleksandrov and Toby Walsh. Group envy freeness and group pareto efficiency in fair division with indivisible items. In *Joint German/Austrian Conference on Artificial Intelligence (Künstliche Intelligenz)*, pages 57–72. Springer, 2018.

Fuad Aleskerov, Hasan Ersel, and Dmitri Piontkovski. *Linear Algebra for Economists*. Springer Science & Business Media, 2011.

Hugo Algarvio, António Couto, Fernando Lopes, Ana Estanqueiro, Ricardo Faia, Gabriel Santos, Rui Carvalho, Pedro Faria, and Zita Vale. The role of demand and vres flexibility in carbon-neutral power systems: Insights from portugal and spain in prospective 2050 scenarios. In *2025 21st International Conference on the European Energy Market (EEM)*, pages 1–6. IEEE, 2025.

Noga Alon, Yishay Mansour, and Moshe Tenneholtz. Differential pricing with inequity aversion in social networks. In *Proceedings of the fourteenth ACM conference on Electronic commerce*, pages 9–24, 2013.

Hans Auer and Reinhard Haas. On integrating large shares of variable renewables into the electricity system. *Energy*, 115:1592–1601, 2016.

Khursheed Aurangzeb, Sheraz Aslam, Syed Muhammad Mohsin, and Musaed Alhussein. A fair pricing mechanism in smart grids for low energy consumption users. *IEEE Access*, 9:22035–22044, 2021.

Lukas Barth, Veit Hagenmeyer, Nicole Ludwig, and Dorothea Wagner. How much demand side flexibility do we need? Analyzing where to exploit flexibility in industrial processes. In *Proceedings of the Ninth International Conference on Future Energy Systems*, pages 43–62, 2018.

Ricardo J Bessa, Manuel A Matos, Filipe Joel Soares, and João A Peças Lopes. Optimized bidding of a ev aggregation agent in the electricity market. *IEEE Transactions on Smart Grid*, 3(1):443–452, 2011.

Eilyan Bitar and Steven Low. Deadline differentiated pricing of deferrable electric power service, 2012.

Diana Böttger, Mario Götz, Nelly Lehr, Hendrik Kondziella, and Thomas Bruckner. Potential of the power-to-heat technology in district heating grids in Germany. *Energy Procedia*, 46:246–253, 2014.

Stephen Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.

Song Hwa Chae, Sang Hun Kim, Sung-Geun Yoon, and Sunwon Park. Optimization of a waste heat utilization network in an eco-industrial park. *Applied Energy*, 87(6): 1978–1988, 2010.

Jean Michel Chapuis. Price fairness versus pricing fairness. 2012.

Bhaskar Ray Chaudhury, T. Kavitha, Kurt Mehlhorn, and Alkmini Sgouritsa. A little charity guarantees almost envy-freeness. *ArXiv*, abs/1907.04596, 2020.

Tianmei Chen, Yi Jin, Hanyu Lv, Antao Yang, Meiyi Liu, Bing Chen, Ying Xie, and Qiang Chen. Applications of lithium-ion batteries in grid-scale energy storage systems. *Transactions of Tianjin University*, 26(3):208–217, 2020.

Wei Chen, Li Qiu, and Pravin Varaiya. Duration-deadline jointly differentiated energy services. In *2015 54th IEEE Conference on Decision and Control (CDC)*. IEEE, dec 2015. doi: 10.1109/cdc.2015.7403358.

Xinyu Chen, Chongqing Kang, Mark O'Malley, Qing Xia, Jianhua Bai, Chun Liu, Rongfu Sun, Weizhou Wang, and Hui Li. Increasing the flexibility of combined heat and power for wind power integration in china: Modeling and implications. *IEEE Transactions on Power Systems*, 30(4):1848–1857, 2014.

Zhengqi Chen, Yingyun Sun, Xin Ai, Sarmad Majeed Malik, and Liping Yang. Integrated demand response characteristics of industrial park: a review. *Journal of Modern Power Systems and Clean Energy*, 8(1):15–26, 2019.

Leon Yang Chu. Truthful bundle/multiunit double auctions. *Management Science*, 55 (7):1184–1198, 2009.

W. Conen and T. Sandholm. Anonymous pricing of efficient allocations in combinatorial economies. *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems, 2004. AAMAS 2004.*, pages 254–260, 2004.

Carlos Adrian Correa-Florez, Andrea Michiorri, and George Kariniotakis. Optimal participation of residential aggregators in energy and local flexibility markets. *IEEE Transactions on Smart Grid*, 11(2):1644–1656, 2019.

Wayne R. Cowell, editor. *Sources and Development of Mathematical Software*. Prentice-Hall Series in Computational Mathematics. 1984. ISBN 0-13-823501-5.

Rajdeep K. Dash, Nicholas R. Jennings, and David C. Parkes. Computational-mechanism design: A call to arms. *IEEE Intelligent Systems*, 18:40–47, 2003.

Paul L Denholm, Erik Ela, Brendan Kirby, and Michael R. Milligan. Role of energy storage with renewable electricity generation. 2010. URL https://www.osti.gov/biblio/972169.

Armen Der Kiureghian and Ove Ditlevsen. Aleatory or epistemic? does it matter? *Structural Safety*, 31(2):105–112, 2009.

Ilias Dimoulkas and Mikael Amelin. Constructing bidding curves for a CHP producer in day-ahead electricity markets. In *Energy Conf. (ENERGYCON)*, pages 487–494, 2014.

Ilias Dimoulkas, Mikael Amelin, and Fabian Levihn. District heating system operation in power systems with high share of wind power. *Journal of Modern Power Systems and Clean Energy*, 5(6):850–862, 2017.

Jinfeng Ding, Kaigui Xie, Bo Hu, Changzheng Shao, Tao Niu, Chunyan Li, and Congcong Pan. Mixed aleatory-epistemic uncertainty modeling of wind power forecast errors in operation reliability evaluation of power systems. *Journal of Modern Power Systems and Clean Energy*, 10(5):1174–1183, 2022.

Ayan Kumar Dutta, Sivaneasan Balakrishnan, and Pallavi Joshi. Optimal scheduling of electric-vehicle fleets for vehicle-to-grid services under soft-and hard-demand-limit constraints. *Applied Energy*, 423:128416, 2026.

Jacob Ehrlich, Maximilian Moll, and Stefan Pickl. A generalized trade reduction mechanism: J. ehrlich et al. *Central European Journal of Operations Research*, 33 (2):429–448, 2025.

O. Ellabban and H. Abu-Rub. Smart grid customers' acceptance and engagement: An overview. *Renewable & Sustainable Energy Reviews*, 65:1285–1298, 2016.

M Saber Eltohamy, Muhayya S Aldawsari, Amir Raouf, Hossam Youssef, Ijaz Ahmed, Theyab R Alsenani, Yun-Su Kim, and Mohamed Mostafa Ramadan Ahmed. Enhancing electric grid flexibility for the integration of variable renewable energy: Challenges, innovations, and future directions. *IEEE Access*, 2026.

Jonas Engels. *Integration of flexibility from battery storage in the electricity market*. PhD thesis, KU Leuven, 2020.

Ayman Esmat, J. Usaola, and M. A. Moreno. Distribution-level flexibility market for congestion management. *Energies*, 11:1056, 2018.

M. Feldman, N. Gravin, and B. Lucier. Combinatorial auctions via posted prices. In *SODA*, 2015.

Michele Flammini, Manuel Mauro, Matteo Tonelli, and Cosimo Vinci. Inequity aversion pricing in multi-unit markets. In *24th European Conference on Artificial Intelligence (ECAI)*, volume 325, pages 91–98. IOS Press, 2020.

Michele Flammini, Manuel Mauro, and Matteo Tonelli. On fair price discrimination in multi-unit markets. *Artificial Intelligence*, 290:103388, 2021.

Kyriaki Foteinaki, Rongling Li, Alfred Heller, and Carsten Rode. Heating system energy flexibility of low-energy residential buildings. *Energy and Buildings*, 180: 95–108, 2018.

Gilbert Fridgen, Philipp Mette, and Markus Thimmel. The value of information exchange in electric vehicle charging. In *ICIS*, 2014.

Yang Fu, Jie Shan, Zhenkun Li, Bolin Xie, and Jeng-Shyang Pan. The optimal bidding strategy for multi-energy prosumers in the double auction electricity-heat market: A bidding space model. *Energy*, 314:134167, 2025.

Shuang Gao, Hailong Li, Yichen Hou, and Jinyue Yan. Benefits of integrating power-to-heat assets in chps. *Applied Energy*, 335:120763, 2023.

Fuchao Ge, Yuntao Ju, Zhinan Qi, and Yi Lin. Parameter estimation of a gaussian mixture model for wind power forecast error by riemann l-bfgs optimization. *IEEE Access*, 6:38892–38899, 2018.

Enrico H Gerding, Rajdeep K Dash, Andrew Byde, and Nicholas R Jennings. Optimal strategies for simultaneous vickrey auctions with perfect substitutes. *Journal of Artificial Intelligence Research*, 32:939–982, 2008.

Mahdi Ghamkhari and Hamed Mohsenian-Rad. Energy and performance management of green data centers: A profit maximization approach. *IEEE Transactions on Smart Grid*, 4(2):1017–1025, 2013.

Gunther Glenk and Stefan Reichelstein. Reversible power-to-gas systems for energy conversion and storage. *Nature communications*, 13(1):2010, 2022.

Hessam Golmohamadi. Demand-side flexibility in power systems: A survey of residential, industrial, commercial, and agricultural sectors. *Sustainability*, 14(13):7916, 2022.

Carla Goncalves, Pierre Pinson, and Ricardo J Bessa. Towards data markets in renewable energy forecasting. *IEEE Transactions on Sustainable Energy*, 12(1):533–542, 2020.

Mira Gonen, Rica Gonen, and Elan Pavlov. Generalized trade reduction mechanisms. In *Proceedings of the 8th ACM Conference on Electronic Commerce*, pages 20–29, 2007.

Noemi González, Pilar Calatayud, Lucía Arcos, Adriana Trujillo, Marta García Pellicer, and Alfredo Quijano-López. Optimal participation of aggregated residential customers in flexibility markets. In *2020 8th International Conference on Smart Grid (icSmartGrid)*, pages 43–47. IEEE, 2020.

Giorgio Graditi, Maria Luisa Di Silvestre, Roberto Gallea, and Eleonora Riva Sanseverino. Heuristic-based shiftable loads optimal management in smart micro-grids. *IEEE Transactions on Industrial Informatics*, 11:271–280, 2015.

Amy Greenwald and Justin Boyan. Bidding under uncertainty: Theory and experiments. *Twentieth Conference on Uncertainty in Artificial Intelligence*, pages 209–216, 2004.

S. Guerreiro, Susana Batel, M. L. Lima, and S. Moreira. Making energy visible: sociopsychological aspects associated with the use of smart meters. *Energy Efficiency*, 8:1149–1167, 2015.

Liyang Han, Pierre Pinson, and Jalal Kazempour. Trading data for wind power forecasting: A regression market with lasso regularization. *Electric Power Systems Research*, 212:108442, 2022.

Shahram Hanifi, Xiaolei Liu, Zi Lin, and Saeid Lotfian. A critical review of wind power forecasting methods—past, present and future. *Energies*, 13(15):3764, 2020.

Keiichiro Hayakawa, Enrico H Gerding, Sebastian Stein, and Takahiro Shiga. Online mechanisms for charging electric vehicles in settings with varying marginal electricity costs. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*, 2015.

Rocio Herranz, Antonio Munoz San Roque, Jose Villar, and Fco Alberto Campos. Optimal demand-side bidding strategies in electricity spot markets. *IEEE Transactions on Power Systems*, 27(3):1204–1213, 2012.

Christian Franz Horn, Bjoern Sven Ivens, Michael Ohneberg, and Alexander Brem. Prediction markets: A literature review 2014. *The Journal of Prediction Markets*, 8 (2):89–126, 2014.

Yimeng Huang and Ju Li. Key challenges for grid-scale lithium-ion battery energy storage. *Advanced Energy Materials*, 12(48):2202197, 2022.

Eyke Hüllermeier and Willem Waegeman. Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods. *Machine Learning*, 110:457–506, 2021.

IEA. Approximately 100 million households rely on rooftop solar PV by 2030. 2022. URL https://www.iea.org/reports/approximately-100-million-households-rely-on-rooftop-solar-pv-by-2030. Licence: CC BY 4.0.

IEA. Electricity market report 2023. 2023a. URL https://www.iea.org/reports/electricity-market-report-2023. Licence: CC BY 4.0.

IEA. Low-emissions sources of electricity. 2023b. URL https://www.iea.org/reports/low-emissions-sources-of-electricity. Licence: CC BY 4.0.

IEA. Global EV outlook 2024. 2024a. URL https://www.iea.org/reports/global-ev-outlook-2024. Licence: CC BY 4.0.

IEA. Electricity 2024. 2024b. URL https://www.iea.org/reports/electricity-2024. Licence: CC BY 4.0.

IEA. Global hydrogen review 2025. 2025a. URL https://www.iea.org/reports/global-hydrogen-review-2025. Licence: CC BY 4.0.

IEA. Renewables 2025: Analysis and forecasts to 2030. 2025b. URL https://www.iea.org/reports/renewables-2025. Licence: CC BY 4.0.

IRENA. Renewable power generation costs in 2022. 2023.

Nima Javanshir, Sanna Syri, Seela Tervo, and Argo Rosin. Operation of district heat network in electricity and balancing markets with the power-to-heat sector coupling. *Energy*, 266:126423, 2023.

Taoli Jiang, Dongyang Shen, Zuodong Zhang, Hongxu Liu, Guili Zhao, Yidi Wang, Shunxin Tan, Ruihao Luo, and Wei Chen. Battery technologies for grid-scale energy storage. *Nature Reviews Clean Technology*, 1(7):474–492, 2025.

Ruiying Jin, Peng Xu, Jiefan Gu, Tong Xiao, Chunhui Li, and Hongxin Wang. Review of optimization control methods for hvac systems in demand response (dr): Transition from model-driven to model-free approaches and challenges. *Building and Environment*, 280:113045, 2025.

RC Johnson and M Mayfield. The effectiveness of decentralised electro-thermal load shifting strategies in low voltage network violation management. *International Journal of Electrical Power & Energy Systems*, 142:108317, 2022.

Verena Kleinschmidt, Thomas Hamacher, Vedran S. Perić, and Mohammad Reza Hesamzadeh. Unlocking flexibility in multi-energy systems: A literature review. *2020 17th International Conference on the European Energy Market (EEM)*, pages 1–6, 2020. URL https://doi.org/10.1109/EEM49802.2020.9221927.

Nikolaos E. Koltsaklis and Athanasios S. Dagoumas. State-of-the-art generation expansion planning: A review. *Applied Energy*, 230:563–589, 2018.

Christos D Korkas, Simone Baldi, Iakovos Michailidis, and Elias B Kosmatopoulos. Occupancy-based demand response and thermal comfort optimization in microgrids with renewable energy sources and energy storage. *Applied Energy*, 163:93–104, 2016.

H. Kuhn. The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2:83–97, 1955.

Dirk Kuiken and Heyd F Más. Integrating demand side management into eu electricity distribution system operation: A dutch example. *Energy Policy*, 129:153–160, 2019.

Pulkit Kumar, Harpreet Kaur Channi, Raman Kumar, Asha Rajiv, Bharti Kumari, Gurpartap Singh, Sehijpal Singh, Issa Farhan Dyab, and Jasmina Lozanović. A comprehensive review of vehicle-to-grid integration in electric vehicles: Powering the future. *Energy Conversion and Management: X*, 25:100864, 2025.

David Lane, David Scott, Mikki Hebl, Rudy Guerra, Dan Osherson, and Heidi Zimmer. *Introduction to statistics*. David Lane, 2003.

Eamonn Lannoye, Damian Flynn, and Mark O'Malley. Evaluation of power system flexibility. *IEEE Transactions on Power Systems*, 27(2):922–931, 2012.

Peter D Lax and Maria Shea Terrell. *Multivariable Calculus with Applications*. Springer, 2017.

Hoesung Lee, Katherine Calvin, Dipak Dasgupta, Gerhard Krinner, Aditi Mukherji, Peter Thorne, Christopher Trisos, José Romero, Paulina Aldunce, Ko Barret, et al. Ipcc, 2023: Climate change 2023: Synthesis report, summary for policymakers. contribution of working groups i, ii and iii to the sixth assessment report of the intergovernmental panel on climate change [core writing team, h. lee and j. romero (eds.)]. ipcc, geneva, switzerland. 2023.

Hengyi Li, Boyu Qin, Shihan Wang, Tao Ding, Jialing Liu, and Hongzhen Wang. Aggregate power flexibility of multi-energy systems supported by dynamic networks. *Applied Energy*, 377:124565, 2025.

Zhengmao Li, Lei Wu, Yan Xu, Luhao Wang, and Nan Yang. Distributed tri-layer risk-averse stochastic game approach for energy trading among multi-energy microgrids. *Applied Energy*, 331:120282, 2023.

Zhigang Li, Wenchuan Wu, Jianhui Wang, Boming Zhang, and Taiyi Zheng. Transmission-constrained unit commitment considering combined electricity and district heating networks. *IEEE Trans. on Sust. Energy*, 7(2):480–492, 2016.

Steffen Limmer. Dynamic pricing for electric vehicle charging—a literature review. *Energies*, 12(18):3574, sep 2019. doi: 10.3390/en12183574.

Steffen Limmer and Manuel Dietrich. Optimization of dynamic prices for electric vehicle charging considering fairness. In *2018 IEEE Symposium Series on Computational Intelligence (SSCI)*, pages 2304–2311. IEEE, 2018.

Lei Liu, Jicheng Liu, Yu Ye, Hui Liu, Kun Chen, Dong Li, Xue Dong, and Mingzhai Sun. Ultra-short-term wind power forecasting based on deep bayesian model with uncertainty. *Renewable Energy*, 205:598–607, 2023.

Luyao Liu, Xinwei Shen, Zhigang Chen, Qie Sun, and Ronald Wennersten. Optimal energy management of data center micro-grid considering computing workloads shift. *IEEE Access*, 12:102061–102075, 2024.

Ya'an Liu and Xiaohong Guan. Purchase allocation and demand bidding in electric power markets. *IEEE Transactions on Power Systems*, 18(1):106–112, 2003.

Vasilis Livanos, Ruta Mehta, and Aniket Murhekar. (almost) envy-free, proportional and efficient allocations of an indivisible mixed manna. In *AAMAS*, 2022.

Jesus Lizana, Daniel Friedrich, Renaldi Renaldi, and Ricardo Chacartegui. Energy flexible building through smart demand-side management and latent heat storage. *Applied Energy*, 230:471–485, 2018.

Fei Lu, Zhenyu Yu, Yu Zou, and Xudong Yang. Energy flexibility assessment of a zero-energy office building with building thermal mass in short-term demand-side management. *Journal of Building Engineering*, 50:104214, 2022.

Katarzyna M Luc, Rongling Li, Luyi Xu, Toke R Nielsen, and Jan LM Hensen. Energy flexibility potential of a small district connected to a district heating system. *Energy and Buildings*, 225:110074, 2020.

Peter D Lund, John Byrne, Reinhard Haas, and Damian Flynn. *Advances in energy systems: The large-scale renewable energy integration challenge*. John Wiley & Sons, 2019.

Alexandra Lüth, Jan Martin Zepter, Pedro Crespo Del Granado, and Ruud Egging. Local electricity market designs for peer-to-peer trading: The role of battery flexibility. *Applied Energy*, 229:1233–1243, 2018.

Ziming Ma, Haiwang Zhong, Qing Xia, and Chongqing Kang. A block-of-use electricity retail pricing approach based on the customer load profile. *IEEE Transactions on Smart Grid*, 11(2):1500–1509, 2019.

Ioannis Mamounakis, Nikolaos Efthymiopoulos, Dimitrios J Vergados, Georgios Tsaousoglou, Prodromos Makris, and Emmanouel Manos Varvarigos. A pricing scheme for electric utility's participation in day-ahead and real-time flexibility energy markets. *Journal of Modern Power Systems and Clean Energy*, 7(5):1294–1306, 2019.

S. Maxwell. Rule-based price fairness and its effect on willingness to purchase. *Journal of Economic Psychology*, 23:191–212, 2002.

R. Preston McAfee. A dominant strategy double auction. *Journal of Economic Theory*, 56(2):434–450, 1992.

Reshef Meir, Hongyao Ma, and Valentin Robu. Contract design for energy demand response. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, pages 1202–1208, 2017. doi: 10.24963/ijcai.2017/167. URL https://doi.org/10.24963/ijcai.2017/167.

Christine Milchram, G. Kaa, N. Doorn, and R. Künneke. Moral values as factors for social acceptance of smart grid technologies. *Sustainability*, 10:1–23, 2018.

Yanfang Mo, Wei Chen, and Li Qiu. Duration-differentiated energy services with peer-to-peer charging. In *2016 IEEE 55th Conference on Decision and Control (CDC)*. IEEE, dec 2016. doi: 10.1109/cdc.2016.7799430.

Hamed Mohsenian-Rad. Optimal demand bidding for time-shiftable loads. *IEEE Transactions on Power Systems*, 30(2):939–951, 2014.

Stine Mueller, Richard Tuth, David Fischer, Bernhard Wille-Haussmann, and Christof Wittwer. Balancing fluctuating renewable energy generation using cogeneration and heat pump systems. *Energy Technology*, 2(1):83–89, 2014.

Kevin P Murphy. *Probabilistic machine learning: an introduction*. MIT press, 2022.

Roger B Myerson and Mark A Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29(2):265–281, 1983.

Daniel Navia Simon and Laura Diaz Anadon. Power price stability and the insurance value of renewable technologies. *Nature Energy*, 10(3):329–341, 2025.

Marcella Nicolini and Massimo Tavoni. Are renewable energy subsidies effective? evidence from europe. *Renewable and Sustainable Energy Reviews*, 74:412–423, 2017.

Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V Vazirani. *Algorithmic Game Theory*. Cambridge Univ. Press, 2007.

Konstantinos Oikonomou and Masood Parvania. Optimal coordination of water distribution energy flexibility with power systems operation. *IEEE Transactions on Smart Grid*, 10(1):1101–1110, 2018.

Vishwamitra Oree and Sayed Z Sayed Hassen. A composite metric for assessing flexibility available in conventional generators of power systems. *Applied Energy*, 177:683–691, 2016.

Peter Palensky and Dietmar Dietrich. Demand side management: Demand response, intelligent energy systems, and smart loads. *IEEE Transactions on Industrial Informatics*, 7(3):381–388, 2011.

Christina Papachristou, Pieter-Jan Hoes, Marcel GLC Loomans, TAJ van Goch, and Jan LM Hensen. Investigating the energy flexibility of dutch office buildings on single building level and building cluster level. *Journal of Building Engineering*, 40: 102687, 2021.

Anthony Papavasiliou and Shmuel S Oren. Large-scale integration of deferrable demand and renewable energy sources. *IEEE Transactions on Power Systems*, 29 (1):489–499, 2013.

Marek Pavlík, František Kurimskỳ, and Kamil Ševc. Renewable energy and price stability: An analysis of volatility and market shifts in the european electricity sector (2015–2025). *Applied Sciences*, 15(12):6397, 2025.

Andrew Perrault and Craig Boutilier. Approximately stable pricing for coordinated purchasing of electricity. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*, 2015.

Andy B Philpott and Erling Pettersen. Optimizing demand-side bids in day-ahead electricity markets. *IEEE Transactions on Power Systems*, 21(2):488–498, 2006.

Benjamin Plaut and Tim Roughgarden. Almost envy-freeness with general valuations. *ArXiv*, abs/1707.04769, 2018.

Zihang Qiu, Wang Zhang, Shuai Lu, Chaojie Li, Jingbo Wang, Ke Meng, and Zhaoyang Dong. Charging-rate-based battery energy storage system in wind farm and battery storage cooperation bidding problem. *CSEE Journal of Power and Energy Systems*, 8(3):659–668, 2022.

Ashraf Radaideh and Venkataramana Ajjarapu. Extracting expedient short term services from homogeneous group of thermostatically controlled loads. In *2016 IEEE Power & Energy Society Innovative Smart Grid Technologies Conference (ISGT)*. IEEE, sep 2016. doi: 10.1109/isgt.2016.7781030.

Christian P Robert. Simulation of truncated normal variables. *Statistics and Computing*, 5:121–125, 1995.

Juan Sebastian Roncancio, José Vuelvas, Diego Patino, and Carlos Adrián Correa-Flórez. Flower greenhouse energy management to offer local flexibility markets. *Energies*, 15(13):4572, 2022.

Michael H Rothkopf. Bidding in simultaneous auctions with a constraint on exposure. *Operations Research*, 25(4):620–629, 1977.

Tuomas Sandholm and Subhash Suri. Market clearability. In *International joint conference on artificial intelligence*, volume 17, pages 1145–1151. ACM, 2001.

Roland Saur, Neil Yorke-Smith, and Han La Poutré. Combined heat and power markets by double-sided auction mechanisms. In *Proceedings of 2019 IEEE PES Innovative Smart Grid Technologies Europe, ISGT-Europe 2019*, pages 1–5, September 2019. doi: 10.1109/ISGTEurope.2019.8905714.

Roland Saur, Han La Poutré, and Neil Yorke-Smith. Towards optimal demand-side bidding in parallel auctions for time-shiftable electrical loads. In *Proceedings of the IEEE International Conference on Industrial Informatics*, pages 340–347, June 2021. doi: 10.1109/INDIN45582.2020.9442077.

Roland Saur, Han La Poutré, and Neil Yorke-Smith. Fair pricing for time-flexible smart energy markets. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '23, page 2703–2705, Richland, SC, May 2023. International Foundation for Autonomous Agents and Multiagent Systems. ISBN 9781450394321.

Roland Saur, Han La Poutré, and Neil Yorke-Smith. Incentives for accurate energy predictions: How to reduce epistemic uncertainty. In *Proceedings of the 15th ACM International Conference on Future and Sustainable Energy Systems*, e-Energy '24, page 192–202, New York, NY, USA, June 2024. Association for Computing Machinery. ISBN 9798400704802. doi: 10.1145/3632775.3661956.

Franziska Schöniger and Ulrich B. Morawetz. What comes down must go up: Why fluctuating renewable energy does not necessarily increase electricity spot price variance in europe. *Energy Economics*, 111:106069, 2022. ISSN 0140-9883. doi: https://doi.org/10.1016/j.eneco.2022.106069. URL https://www.sciencedirect.com/science/article/pii/S0140988322002353.

Joachim Seel, Dev Millstein, Andrew Mills, Mark Bolinger, and Ryan Wiser. Plentiful electricity turns wholesale prices negative. *Advances in Applied Energy*, 4:100073, 2021. ISSN 2666-7924. doi: https://doi.org/10.1016/j.adapen.2021.100073. URL https://www.sciencedirect.com/science/article/pii/S2666792421000652.

Karlo Sepetanc and H. Pandzic. A cluster-based model for charging a single-depot fleet of electric vehicles. *IEEE Transactions on Smart Grid*, 12(4):3339–3352, jul 2021. doi: 10.1109/tsg.2021.3064272.

Blake Shaffer. Energy and environmental policy trends: Will electric vehicle rebates spur widespread adoption? *The School of Public Policy Publications*, 12, 2019.

Sulman Shahzad and Elżbieta Jasińska. Renewable revolution: A review of strategic flexibility in future power systems. *Sustainability*, 16(13):5454, 2024.

Mahdieh Shamsi and Paul Cuffe. A prediction market trading strategy to hedge financial risks of wind power producers in electricity markets. *IEEE Transactions on Power*

Systems, 36(5):4513–4523, 2021a.

Mahdieh Shamsi and Paul Cuffe. Prediction markets for probabilistic forecasting of renewable energy sources. *IEEE Transactions on Sustainable Energy*, 13(2):1244–1253, 2021b.

Mahdieh Shamsi and Paul Cuffe. Using binary prediction markets as hedging instruments: Strategies for renewable generators. *IEEE Transactions on Sustainable Energy*, 13(2):1160–1163, 2021c.

Bochao Shen, Balakrishnan Narayanaswamy, and Ravi Sundaram. Smartshift: expanded load shifting incentive mechanism for risk-averse consumers. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015.

Goran Strbac. Demand side management: Benefits and challenges. *Energy policy*, 36 (12):4419–4426, 2008.

Chenyu Su, Jakub Jurasz, Lei Feng, Reza Fachrizal, Zahra Shahroozi, Anders Avelin, and Hailong Li. Impacts of dynamic behaviours of thermal energy storage on dispatch optimization for combined heat and power plants. *Energy Conversion and Management*, 368:121982, 2026.

Zelin Sun, Leandro Von Krannichfeldt, and Yi Wang. Trading and valuation of day-ahead load forecasts in an ensemble model. *IEEE Transactions on Industry Applications*, 2023.

R. Tarjan. Depth-first search and linear graph algorithms. *SIAM J. Comput.*, 1:146–160, 1972.

Govinda R Timilsina. Are renewable energy technologies cost competitive for electricity generation? *Renewable Energy*, 180:658–672, 2021.

Taiki Todo, Atsushi Iwasaki, and M. Yokoo. Competitive auctions and envy-freeness for group of agents. In *COCOON*, 2019.

Pierre-Julien Trombe, Pierre Pinson, and Henrik Madsen. A general probabilistic forecasting framework for offshore wind power fluctuations. *Energies*, 5(3):621–657, 2012.

Stylianos I Vagropoulos and Anastasios G Bakirtzis. Optimal bidding strategy for electric vehicle aggregators in electricity markets. *IEEE Transactions on Power Systems*, 28(4):4031–4041, 2013.

Orlando Valarezo, Tomás Gómez, José Pablo Chaves-Ávila, Leandro Lind, Mauricio Correa, David Ulrich Ziegler, and Rodrigo A. Escobar. Analysis of new flexibility market models in europe. *Energies*, 2021. URL https://doi.org/10.3390/en14123521.

Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020. doi: 10.1038/s41592-019-0686-2.

Yih-huei Wan and Brian K Parsons. Factors relevant to utility integration of intermittent renewable technologies. Technical report, National Renewable Energy Lab.(NREL),

Golden, CO (United States), 1993.

Jianxiao Wang, Haiwang Zhong, Ziming Ma, Qing Xia, and Chongqing Kang. Review and prospect of integrated demand response in the multi-energy system. *Applied Energy*, 202:772–782, 2017.

Jiawei Wang, Shi You, Yi Zong, Chresten Træholt, Zhao Yang Dong, and You Zhou. Flexibility of combined heat and power plants: A review of technologies and operation strategies. *Applied Energy*, 252:113445, 2019.

Mohammad Waseem, Ahmad Faizan Sherwani, and Mohd Suhaib. Integration of solar energy in electrical, hybrid, autonomous vehicles: a technological review. *SN Applied Sciences*, 1(11):1–14, 2019.

Michael P Wellman, Amy Greenwald, Peter Stone, and Peter R Wurman. The 2001 trading agent competition. *Electronic Markets*, 13(1):4–12, 2003.

Honglin Wen, Jinghuan Ma, Jie Gu, Lyuzerui Yuan, and Zhijian Jin. Sparse variational gaussian process based day-ahead probabilistic wind power forecasting. *IEEE Transactions on Sustainable Energy*, 13(2):957–970, 2022.

T. D. Wildt, E. J. Chappin, G. Kaa, P. Herder, and I. Poel. Conflicting values in the smart electricity grid: A comprehensive overview. *Renewable & Sustainable Energy Reviews*, 111:184–196, 2019.

Justin Wolfers and Eric Zitzewitz. Prediction markets. *Journal of Economic Perspectives*, 18(2):107–126, 2004.

Xiandong Xu, Quan Lyu, Meysam Qadrdan, and Jianzhong Wu. Quantification of flexibility of a district heating system for the power grid. *IEEE Transactions on Sustainable Energy*, 11(4):2617–2630, 2020.

Xiyun Yang, Liwei Fan, Xiangjun Li, and Lingzhuochao Meng. Day-ahead and real-time market bidding and scheduling strategy for wind power participation based on shared energy storage. *Electric Power Systems Research*, 214:108903, 2023.

Bizzat Hussain Zaidi and S. Hong. Combinatorial double auctions for multiple microgrid trading. *Electrical Engineering*, 100:1069–1083, 2018.

Shihong Zeng, Yuchen Liu, Chao Liu, and Xin Nan. A review of renewable energy investment in the BRICS countries: History, models, problems and solutions. *Renewable and Sustainable Energy Reviews*, 74:860–872, 2017.

Jinhua Zhang, Jie Yan, David Infield, Yongqian Liu, and Fue-sang Lien. Short-term forecasting and uncertainty analysis of wind turbine power based on long short-term memory network and gaussian mixture model. *Applied Energy*, 241:229–244, 2019.

Jiubo Zhang, Bowen Zhou, Zhile Yang, Yuanjun Guo, Chen Lv, Xiaofeng Xu, and Jichun Liu. A review of industrial load flexibility enhancement for demand-response interaction. *Sustainability*, 17(11):4938, 2025.

Yue Zhao, Junjie Qin, Ram Rajagopal, Andrea J. Goldsmith, and H. Vincent Poor. Wind aggregation via risky power markets. *IEEE Transactions on Power Systems*, 30: 1571–1581, 2015. URL https://doi.org/10.1109/TPWRS.2014.2359671.

Jihu Zheng, X. Sun, Li-Jie Jia, and Y. Zhou. Electric passenger vehicles sales and carbon dioxide emission reduction potential in China's leading markets. *Journal of Cleaner Production*, 243:118607, 2020.

Yanchong Zheng, Yubin Wang, and Qiang Yang. Bidding strategy design for electric vehicle aggregators in the day-ahead electricity market considering price volatility:

A risk-averse approach. *Energy*, 283:129138, 2023.

Weifeng Zhong, Kan Xie, Yi Liu, Chao Yang, and Shengli Xie. Auction mechanisms for energy trading in multi-energy systems. *IEEE Trans. Industr. Inform.*, 14(4): 1511–1521, 2018a.

Weifeng Zhong, Chao Yang, Kan Xie, Shengli Xie, and Yan Zhang. Admm-based distributed auction mechanism for energy hub scheduling in smart buildings. *IEEE Access*, 6:45635–45645, 2018b.

Ma Zhongjing and R. Long. Optimal distributed charging coordinations of plug-in electric vehicles with market uniform clearing price mechanism. *2013 25th Chinese Control and Decision Conference (CCDC)*, pages 3497–3502, 2013.

2020 01 Armon Toubman (UL), Calculated Moves: Generating Air Combat Behaviour
02 Marcos de Paula Bueno (UL), Unraveling Temporal Processes using Probabilistic Graphical Models
03 Mostafa Deghani (UvA), Learning with Imperfect Supervision for Language Understanding
04 Maarten van Gompel (RUN), Context as Linguistic Bridges
05 Yulong Pei (TU/e), On local and global structure mining
06 Preethu Rose Anish (UT), Stimulation Architectural Thinking during Requirements Elicitation - An Approach and Tool Support
07 Wim van der Vegt (OU), Towards a software architecture for reusable game components
08 Ali Mirsoleimani (UL),Structured Parallel Programming for Monte Carlo Tree Search
09 Myriam Traub (UU), Measuring Tool Bias and Improving Data Quality for Digital Humanities Research
10 Alifah Syamsiyah (TU/e), In-database Preprocessing for Process Mining
11 Sepideh Mesbah (TUD), Semantic-Enhanced Training Data AugmentationMethods for Long-Tail Entity Recognition Models
12 Ward van Breda (VUA), Predictive Modeling in E-Mental Health: Exploring Applicability in Personalised Depression Treatment
13 Marco Virgolin (CWI), Design and Application of Gene-pool Optimal Mixing Evolutionary Algorithms for Genetic Programming
14 Mark Raasveldt (CWI/UL), Integrating Analytics with Relational Databases
15 Konstantinos Georgiadis (OU), Smart CAT: Machine Learning for Configurable Assessments in Serious Games
16 Ilona Wilmont (RUN), Cognitive Aspects of Conceptual Modelling
17 Daniele Di Mitri (OU), The Multimodal Tutor: Adaptive Feedback from Multimodal Experiences
18 Georgios Methenitis (TUD), Agent Interactions & Mechanisms in Markets with Uncertainties: Electricity Markets in Renewable Energy Systems
19 Guido van Capelleveen (UT), Industrial Symbiosis Recommender Systems
20 Albert Hankel (VUA), Embedding Green ICT Maturity in Organisations
21 Karine da Silva Miras de Araujo (VUA), Where is the robot?: Life as it could be

22 Maryam Masoud Khamis (RUN), Understanding complex systems implementation through a modeling approach: the case of e-government in Zanzibar
23 Rianne Conijn (UT), The Keys to Writing: A writing analytics approach to studying writing processes using keystroke logging
24 Lenin da Nóbrega Medeiros (VUA/RUN), How are you feeling, human? Towards emotionally supportive chatbots
25 Xin Du (TU/e), The Uncertainty in Exceptional Model Mining
26 Krzysztof Leszek Sadowski (UU), GAMBIT: Genetic Algorithm for Model-Based mixed-Integer opTimization
27 Ekaterina Muravyeva (TUD), Personal data and informed consent in an educational context
28 Bibeg Limbu (TUD), Multimodal interaction for deliberate practice: Training complex skills with augmented reality
29 Ioan Gabriel Bucur (RUN), Being Bayesian about Causal Inference
30 Bob Zadok Blok (UL), Creatief, Creatiever, Creatiefst
31 Gongjin Lan (VUA), Learning better – From Baby to Better
32 Jason Rhuggenaath (TU/e), Revenue management in online markets: pricing and online advertising
33 Rick Gilsing (TU/e), Supporting service-dominant business model evaluation in the context of business model innovation
34 Anna Bon (UM), Intervention or Collaboration? Redesigning Information and Communication Technologies for Development
35 Siamak Farshidi (UU), Multi-Criteria Decision-Making in Software Production

---

2021 01 Francisco Xavier Dos Santos Fonseca (TUD),Location-based Games for Social Interaction in Public Space
02 Rijk Mercuur (TUD), Simulating Human Routines: Integrating Social Practice Theory in Agent-Based Models
03 Seyyed Hadi Hashemi (UvA), Modeling Users Interacting with Smart Devices
04 Ioana Jivet (OU), The Dashboard That Loved Me: Designing adaptive learning analytics for self-regulated learning
05 Davide Dell'Anna (UU), Data-Driven Supervision of Autonomous Systems
06 Daniel Davison (UT), "Hey robot, what do you think?" How children learn with a social robot
07 Armel Lefebvre (UU), Research data management for open science
08 Nardie Fanchamps (OU), The Influence of Sense-Reason-Act Programming on Computational Thinking

09 Cristina Zaga (UT), The Design of Robothings. Non-Anthropomorphic and Non-Verbal Robots to Promote Children's Collaboration Through Play
10 Quinten Meertens (UvA), Misclassification Bias in Statistical Learning
11 Anne van Rossum (UL), Nonparametric Bayesian Methods in Robotic Vision
12 Lei Pi (UL), External Knowledge Absorption in Chinese SMEs
13 Bob R. Schadenberg (UT), Robots for Autistic Children: Understanding and Facilitating Predictability for Engagement in Learning
14 Negin Samaeemofrad (UL), Business Incubators: The Impact of Their Support
15 Onat Ege Adali (TU/e), Transformation of Value Propositions into Resource Re-Configurations through the Business Services Paradigm
16 Esam A. H. Ghaleb (UM), Bimodal emotion recognition from audio-visual cues
17 Dario Dotti (UM), Human Behavior Understanding from motion and bodily cues using deep neural networks
18 Remi Wieten (UU), Bridging the Gap Between Informal Sense-Making Tools and Formal Systems - Facilitating the Construction of Bayesian Networks and Argumentation Frameworks
19 Roberto Verdecchia (VUA), Architectural Technical Debt: Identification and Management
20 Masoud Mansoury (TU/e), Understanding and Mitigating Multi-Sided Exposure Bias in Recommender Systems
21 Pedro Thiago Timbó Holanda (CWI), Progressive Indexes
22 Sihang Qiu (TUD), Conversational Crowdsourcing
23 Hugo Manuel Proença (UL), Robust rules for prediction and description
24 Kaijie Zhu (TU/e), On Efficient Temporal Subgraph Query Processing
25 Eoin Martino Grua (VUA), The Future of E-Health is Mobile: Combining AI and Self-Adaptation to Create Adaptive E-Health Mobile Applications
26 Benno Kruit (CWI/VUA), Reading the Grid: Extending Knowledge Bases from Human-readable Tables
27 Jelte van Waterschoot (UT), Personalized and Personal Conversations: Designing Agents Who Want to Connect With You

	28	Christoph Selig (UL), Understanding the Heterogeneity of Corporate Entrepreneurship Programs
2022	01	Judith van Stegeren (UT), Flavor text generation for role-playing video games
	02	Paulo da Costa (TU/e), Data-driven Prognostics and Logistics Optimisation: A Deep Learning Journey
	03	Ali el Hassouni (VUA), A Model A Day Keeps The Doctor Away: Reinforcement Learning For Personalized Healthcare
	04	Ünal Aksu (UU), A Cross-Organizational Process Mining Framework
	05	Shiwei Liu (TU/e), Sparse Neural Network Training with In-Time Over-Parameterization
	06	Reza Refaei Afshar (TU/e), Machine Learning for Ad Publishers in Real Time Bidding
	07	Sambit Praharaj (OU), Measuring the Unmeasurable? Towards Automatic Co-located Collaboration Analytics
	08	Maikel L. van Eck (TU/e), Process Mining for Smart Product Design
	09	Oana Andreea Inel (VUA), Understanding Events: A Diversity-driven Human-Machine Approach
	10	Felipe Moraes Gomes (TUD), Examining the Effectiveness of Collaborative Search Engines
	11	Mirjam de Haas (UT), Staying engaged in child-robot interaction, a quantitative approach to studying preschoolers' engagement with robots and tasks during second-language tutoring
	12	Guanyi Chen (UU), Computational Generation of Chinese Noun Phrases
	13	Xander Wilcke (VUA), Machine Learning on Multimodal Knowledge Graphs: Opportunities, Challenges, and Methods for Learning on Real-World Heterogeneous and Spatially-Oriented Knowledge
	14	Michiel Overeem (UU), Evolution of Low-Code Platforms
	15	Jelmer Jan Koorn (UU), Work in Process: Unearthing Meaning using Process Mining
	16	Pieter Gijsbers (TU/e), Systems for AutoML Research
	17	Laura van der Lubbe (VUA), Empowering vulnerable people with serious games and gamification
	18	Paris Mavromoustakos Blom (TiU), Player Affect Modelling and Video Game Personalisation
	19	Bilge Yigit Ozkan (UU), Cybersecurity Maturity Assessment and Standardisation

20 Fakhra Jabeen (VUA), Dark Side of the Digital Media - Computational Analysis of Negative Human Behaviors on Social Media
21 Seethu Mariyam Christopher (UM), Intelligent Toys for Physical and Cognitive Assessments
22 Alexandra Sierra Rativa (TiU), Virtual Character Design and its potential to foster Empathy, Immersion, and Collaboration Skills in Video Games and Virtual Reality Simulations
23 Ilir Kola (TUD), Enabling Social Situation Awareness in Support Agents
24 Samaneh Heidari (UU), Agents with Social Norms and Values - A framework for agent based social simulations with social norms and personal values
25 Anna L.D. Latour (UL), Optimal decision-making under constraints and uncertainty
26 Anne Dirkson (UL), Knowledge Discovery from Patient Forums: Gaining novel medical insights from patient experiences
27 Christos Athanasiadis (UM), Emotion-aware cross-modal domain adaptation in video sequences
28 Onuralp Ulusoy (UU), Privacy in Collaborative Systems
29 Jan Kolkmeier (UT), From Head Transform to Mind Transplant: Social Interactions in Mixed Reality
30 Dean De Leo (CWI), Analysis of Dynamic Graphs on Sparse Arrays
31 Konstantinos Traganos (TU/e), Tackling Complexity in Smart Manufacturing with Advanced Manufacturing Process Management
32 Cezara Pastrav (UU), Social simulation for socio-ecological systems
33 Brinn Hekkelman (CWI/TUD), Fair Mechanisms for Smart Grid Congestion Management
34 Nimat Ullah (VUA), Mind Your Behaviour: Computational Modelling of Emotion & Desire Regulation for Behaviour Change
35 Mike E.U. Ligthart (VUA), Shaping the Child-Robot Relationship: Interaction Design Patterns for a Sustainable Interaction

---

2023 01 Bojan Simoski (VUA), Untangling the Puzzle of Digital Health Interventions
02 Mariana Rachel Dias da Silva (TiU), Grounded or in flight? What our bodies can tell us about the whereabouts of our thoughts
03 Shabnam Najafian (TUD), User Modeling for Privacy-preserving Explanations in Group Recommendations

04 Gineke Wiggers (UL), The Relevance of Impact: bibliometric-enhanced legal information retrieval
05 Anton Bouter (CWI), Optimal Mixing Evolutionary Algorithms for Large-Scale Real-Valued Optimization, Including Real-World Medical Applications
06 António Pereira Barata (UL), Reliable and Fair Machine Learning for Risk Assessment
07 Tianjin Huang (TU/e), The Roles of Adversarial Examples on Trustworthiness of Deep Learning
08 Lu Yin (TU/e), Knowledge Elicitation using Psychometric Learning
09 Xu Wang (VUA), Scientific Dataset Recommendation with Semantic Techniques
10 Dennis J.N.J. Soemers (UM), Learning State-Action Features for General Game Playing
11 Fawad Taj (VUA), Towards Motivating Machines: Computational Modeling of the Mechanism of Actions for Effective Digital Health Behavior Change Applications
12 Tessel Bogaard (VUA), Using Metadata to Understand Search Behavior in Digital Libraries
13 Injy Sarhan (UU), Open Information Extraction for Knowledge Representation
14 Selma Čaušević (TUD), Energy resilience through self-organization
15 Alvaro Henrique Chaim Correia (TU/e), Insights on Learning Tractable Probabilistic Graphical Models
16 Peter Blomsma (TiU), Building Embodied Conversational Agents: Observations on human nonverbal behaviour as a resource for the development of artificial characters
17 Meike Nauta (UT), Explainable AI and Interpretable Computer Vision – From Oversight to Insight
18 Gustavo Penha (TUD), Designing and Diagnosing Models for Conversational Search and Recommendation
19 George Aalbers (TiU), Digital Traces of the Mind: Using Smartphones to Capture Signals of Well-Being in Individuals
20 Arkadiy Dushatskiy (TUD), Expensive Optimization with Model-Based Evolutionary Algorithms applied to Medical Image Segmentation using Deep Learning
21 Gerrit Jan de Bruin (UL), Network Analysis Methods for Smart Inspection in the Transport Domain
22 Alireza Shojaifar (UU), Volitional Cybersecurity

	23	Theo Theunissen (UU), Documentation in Continuous Software Development
	24	Agathe Balayn (TUD), Practices Towards Hazardous Failure Diagnosis in Machine Learning
	25	Jurian Baas (UU), Entity Resolution on Historical Knowledge Graphs
	26	Loek Tonnaer (TU/e), Linearly Symmetry-Based Disentangled Representations and their Out-of-Distribution Behaviour
	27	Ghada Sokar (TU/e), Learning Continually Under Changing Data Distributions
	28	Floris den Hengst (VUA), Learning to Behave: Reinforcement Learning in Human Contexts
	29	Tim Draws (TUD), Understanding Viewpoint Biases in Web Search Results
2024	01	Daphne Miedema (TU/e), On Learning SQL: Disentangling concepts in data systems education
	02	Emile van Krieken (VUA), Optimisation in Neurosymbolic Learning Systems
	03	Feri Wijayanto (RUN), Automated Model Selection for Rasch and Mediation Analysis
	04	Mike Huisman (UL), Understanding Deep Meta-Learning
	05	Yiyong Gou (UM), Aerial Robotic Operations: Multi-environment Cooperative Inspection & Construction Crack Autonomous Repair
	06	Azqa Nadeem (TUD), Understanding Adversary Behavior via XAI: Leveraging Sequence Clustering to Extract Threat Intelligence
	07	Parisa Shayan (TiU), Modeling User Behavior in Learning Management Systems
	08	Xin Zhou (UvA), From Empowering to Motivating: Enhancing Policy Enforcement through Process Design and Incentive Implementation
	09	Giso Dal (UT), Probabilistic Inference Using Partitioned Bayesian Networks
	10	Cristina-Iulia Bucur (VUA), Linkflows: Towards Genuine Semantic Publishing in Science
	11	withdrawn
	12	Peide Zhu (TUD), Towards Robust Automatic Question Generation For Learning
	13	Enrico Liscio (TUD), Context-Specific Value Inference via Hybrid Intelligence

14 Larissa Capobianco Shimomura (TU/e), On Graph Generating Dependencies and their Applications in Data Profiling
15 Ting Liu (VUA), A Gut Feeling: Biomedical Knowledge Graphs for Interrelating the Gut Microbiome and Mental Health
16 Arthur Barbosa Câmara (TUD), Designing Search-as-Learning Systems
17 Razieh Alidoosti (VUA), Ethics-aware Software Architecture Design
18 Laurens Stoop (UU), Data Driven Understanding of Energy-Meteorological Variability and its Impact on Energy System Operations
19 Azadeh Mozafari Mehr (TU/e), Multi-perspective Conformance Checking: Identifying and Understanding Patterns of Anomalous Behavior
20 Ritsart Anne Plantenga (UL), Omgang met Regels
21 Federica Vinella (UU), Crowdsourcing User-Centered Teams
22 Zeynep Ozturk Yurt (TU/e), Beyond Routine: Extending BPM for Knowledge-Intensive Processes with Controllable Dynamic Contexts
23 Jie Luo (VUA), Lamarck's Revenge: Inheritance of Learned Traits Improves Robot Evolution
24 Nirmal Roy (TUD), Exploring the effects of interactive interfaces on user search behaviour
25 Alisa Rieger (TUD), Striving for Responsible Opinion Formation in Web Search on Debated Topics
26 Tim Gubner (CWI), Adaptively Generating Heterogeneous Execution Strategies using the VOILA Framework
27 Lincen Yang (UL), Information-theoretic Partition-based Models for Interpretable Machine Learning
28 Leon Helwerda (UL), Grip on Software: Understanding development progress of Scrum sprints and backlogs
29 David Wilson Romero Guzman (VUA), The Good, the Efficient and the Inductive Biases: Exploring Efficiency in Deep Learning Through the Use of Inductive Biases
30 Vijanti Ramautar (UU), Model-Driven Sustainability Accounting
31 Ziyu Li (TUD), On the Utility of Metadata to Optimize Machine Learning Workflows
32 Vinicius Stein Dani (UU), The Alpha and Omega of Process Mining
33 Siddharth Mehrotra (TUD), Designing for Appropriate Trust in Human-AI interaction

34 Robert Deckers (VUA), From Smallest Software Particle to System Specification - MuDForM: Multi-Domain Formalization Method
35 Sicui Zhang (TU/e), Methods of Detecting Clinical Deviations with Process Mining: a fuzzy set approach
36 Thomas Mulder (TU/e), Optimization of Recursive Queries on Graphs
37 James Graham Nevin (UvA), The Ramifications of Data Handling for Computational Models
38 Christos Koutras (TUD), Tabular Schema Matching for Modern Settings
39 Paola Lara Machado (TU/e), The Nexus between Business Models and Operating Models: From Conceptual Understanding to Actionable Guidance
40 Montijn van de Ven (TU/e), Guiding the Definition of Key Performance Indicators for Business Models
41 Georgios Siachamis (TUD), Adaptivity for Streaming Dataflow Engines
42 Emmeke Veltmeijer (VUA), Small Groups, Big Insights: Understanding the Crowd through Expressive Subgroup Analysis
43 Cedric Waterschoot (KNAW Meertens Instituut), The Constructive Conundrum: Computational Approaches to Facilitate Constructive Commenting on Online News Platforms
44 Marcel Schmitz (OU), Towards learning analytics-supported learning design
45 Sara Salimzadeh (TUD), Living in the Age of AI: Understanding Contextual Factors that Shape Human-AI Decision-Making
46 Georgios Stathis (Leiden University), Preventing Disputes: Preventive Logic, Law & Technology
47 Daniel Daza (VUA), Exploiting Subgraphs and Attributes for Representation Learning on Knowledge Graphs
48 Ioannis Petros Samiotis (TUD), Crowd-Assisted Annotation of Classical Music Compositions

---

2025 01 Max van Haastrecht (UL), Transdisciplinary Perspectives on Validity: Bridging the Gap Between Design and Implementation for Technology-Enhanced Learning Systems
02 Jurgen van den Hoogen (JADS), Time Series Analysis Using Convolutional Neural Networks
03 Andra-Denis Ionescu (TUD), Feature Discovery for Data-Centric AI
04 Rianne Schouten (TU/e), Exceptional Model Mining for Hierarchical Data

05 Nele Albers (TUD), Psychology-Informed Reinforcement Learning for Situated Virtual Coaching in Smoking Cessation
06 Daniël Vos (TUD), Decision Tree Learning: Algorithms for Robust Prediction and Policy Optimization
07 Ricky Maulana Fajri (TU/e), Towards Safer Active Learning: Dealing with Unwanted Biases, Graph-Structured Data, Adversary, and Data Imbalance
08 Stefan Bloemheuvel (TiU), Spatio-Temporal Analysis Through Graphs: Predictive Modeling and Graph Construction
09 Fadime Kaya (VUA), Decentralized Governance Design - A Model-Based Approach
10 Zhao Yang (UL), Enhancing Autonomy and Efficiency in Goal-Conditioned Reinforcement Learning
11 Shahin Sharifi Noorian (TUD), From Recognition to Understanding: Enriching Visual Models Through Multi-Modal Semantic Integration
12 Lijun Lyu (TUD), Interpretability in Neural Information Retrieval
13 Fuda van Diggelen (VUA), Robots Need Some Education: on the complexity of learning in evolutionary robotics
14 Gennaro Gala (TU/e), Probabilistic Generative Modeling with Latent Variable Hierarchies
15 Michiel van der Meer (UL), Opinion Diversity through Hybrid Intelligence
16 Monika Grewal (TU Delft), Deep Learning for Landmark Detection, Segmentation, and Multi-Objective Deformable Registration in Medical Imaging
17 Matteo De Carlo (VUA), Real Robot Reproduction: Towards Evolving Robotic Ecosystems
18 Anouk Neerincx (UU), Robots That Care: How Social Robots Can Boost Children's Mental Wellbeing
19 Fang Hou (UU), Trust in Software Ecosystems
20 Alexander Melchior (UU), Modelling for Policy is More Than Policy Modelling (The Useful Application of Agent-Based Modelling in Complex Policy Processes)
21 Mandani Ntekouli (UM), Bridging Individual and Group Perspectives in Psychopathology: Computational Modeling Approaches using Ecological Momentary Assessment Data
22 Hilde Weerts (TU/e), Decoding Algorithmic Fairness: Towards Interdisciplinary Understanding of Fairness and Discrimination in Algorithmic Decision-Making

23 Roderick van der Weerdt (VUA), IoT Measurement Knowledge Graphs: Constructing, Working and Learning with IoT Measurement Data as a Knowledge Graph
24 Zhong Li (UL), Trustworthy Anomaly Detection for Smart Manufacturing
25 Kyana van Eijndhoven (TiU), A Breakdown of Breakdowns: Multi-Level Team Coordination Dynamics under Stressful Conditions
26 Tom Pepels (UM), Monte-Carlo Tree Search is Work in Progress
27 Danil Provodin (JADS, TU/e), Sequential Decision Making Under Complex Feedback
28 Jinke He (TU Delft), Exploring Learned Abstract Models for Efficient Planning and Learning
29 Erik van Haeringen (VUA), Mixed Feelings: Simulating Emotion Contagion in Groups
30 Myrthe Reuver (VUA), A Puzzle of Perspectives: Interdisciplinary Language Technology for Responsible News Recommendation
31 Gebrekirstos Gebreselassie Gebremeskel (RUN), Spotlight on Recommender Systems: Contributions to Selected Components in the Recommendation Pipeline
32 Ryan Brate (UU), Words Matter: A Computational Toolkit for Charged Terms
33 Merle Reimann (VUA), Speaking the Same Language: Spoken Capability Communication in Human-Agent and Human-Robot Interaction
34 Eduard C. Groen (UU), Crowd-Based Requirements Engineering
35 Urja Khurana (VUA), From Concept To Impact: Toward More Robust Language Model Deployment
36 Anna Maria Wegmann (UU), Say the Same but Differently: Computational Approaches to Stylistic Variation and Paraphrasing
37 Chris Kamphuis (RUN), Exploring Relations and Graphs for Information Retrieval
38 Valentina Maccatrozzo (VUA), Break the Bubble: Semantic Patterns for Serendipity
39 Dimitrios Alivanistos (VUA), Knowledge Graphs & Transformers for Hypothesis Generation: Accelerating Scientific Discovery in the Era of Artificial Intelligence
40 Stefan Grafberger (UvA), Declarative Machine Learning Pipeline Management via Logical Query Plans
41 Mozhgan Vazifehdoostirani (TU/e), Leveraging Process Flexibility to Improve Process Outcome - From Descriptive Analytics to Actionable Insights

42 Margherita Martorana (VUA), Semantic Interpretation of Dataless Tables: a metadata-driven approach for findable, accessible, interoperable and reusable restricted access data
43 Krist Shingjergji (OU), Sense the Classroom - Using AI to Detect and Respond to Learning-Centered Affective States in Online Education
44 Robbert Reijnen (TU/e), Dynamic Algorithm Configuration for Machine Scheduling Using Deep Reinforcement Learning
45 Anjana Mohandas Sheeladevi (VUA), Occupant-Centric Energy Management: Balancing Privacy, Well-being and Sustainability in Smart Buildings
46 Ya Song (TU/e), Graph Neural Networks for Modeling Temporal and Spatial Dimensions in Industrial Decision-making
47 Tom Kouwenhoven (UL), Collaborative Meaning-Making. The Emergence of Novel Languages in Humans, Machines, and Human-Machine Interactions
48 Evy van Weelden (TiU), Integrating Virtual Reality and Neurophysiology in Flight Training
49 Selene Báez Santamaría (VUA), Knowledge-centered conversational agents with a drive to learn
50 Lea Krause (VUA), Contextualising Conversational AI
51 Jiaxu Zhao (TU/e), Understanding and Mitigating Unwanted Biases in Generative Language Models
52 Qiao Xiao (TU/e), Model, Data and Communication Sparsity for Efficient Training of Neural Networks
53 Gaole He (TUD), Towards Effective Human-AI Collaboration: Promoting Appropriate Reliance on AI Systems
54 Go Sugimoto (VUA), MISSING LINKS Investigating the Quality of Linked Data and its Tools in Cultural Heritage and Digital Humanities
55 Sietze Kai Kuilman (TUD), AI that Glitters is Not Gold: Requirements for Meaningful Control of AI Systems
56 Wijnand van Woerkom (UU), A Fortiori Case-Based Reasoning: Formal Studies with Applications in Artificial Intelligence and Law
57 Syeda Amna Sohail (UT), Privacy-Utility Trade-Off in Healthcare Metadata Sharing and Beyond: A Normative and Empirical Evaluation at Inter and Intra Organizational Levels
58 Junhan Wen (TUD), "From iMage to Market": Machine-Learning-Empowered Fruit Supply
59 Mohsen Abbaspour Onari (TU/e), From Explanation to Trust: Modeling and Measuring Trust in Explainable Decision Support

60 Marcel Jurriaan Robeer (UU), Beyond Trust: A Causal Approach to Explainable AI in Law Enforcement
61 Shuai Wang (VUA), Links in Large Integrated Knowledge Graphs: Analysis, Refinement, and Domain Applications
62 Khaleel Asyraaf Mat Sanusi (OU), Augmenting a learning model within immersive learning environments for psychomotor skills
63 Rashid Zaman (TU/e), Online Conformance Checking on Degraded Data
64 Jens d'Hondt (TU/e), Effective and Efficient Multivariate Similarity Search
65 Aswin Balasubramaniam (UT), Disentangling Runner Drone Interaction Potentialities

2026 01 Pei-Yu Chen (TUD), Human-Agent Alignment Dialogues: Eliciting User Information at Runtime for Personalized Behavior Support
02 Hezha Hassan Mohammedkhan (TiU), Estimating Body Measurements of Children from 2D Images: Towards the Automatic Detection of Malnutrition
03 Kyriakos Psarakis (TUD), Democratizing Scalable Cloud Applications: Transactional Stateful Functions on Streaming Dataflows
04 Boyu Xu (UU), Exploring Indirect Relations Between Topics in Neuroscience Literature Using Augmented Reality to Inform Experimental Design
05 Koen Minartz (TU/e), Stochastic Simulation with Geometric Deep Generative Models
06 Azim Afroozeh (CWI, VUA), FastLanes: A Next-Gen File Format
07 Inès Blin (VUA), Narrative Understanding with Knowledge Graphs
08 Paul van Vulpen (UU), Debating Digital Dominance: Decentralized Technology Governance For Strategic Autonomy
09 Afrizal Doewes (TU/e), Rethinking Automated Essay Scoring: Agreement, Fairness, and Feedback
10 Nikolaos Delapaschos Kondylidis (VUA), Establishing Task-Oriented Understanding between Agents
11 Işıl Baysal Erez (UT), Handling Missing Data with Meta-Learning and Large Language Models
12 Xue Li (UvA), From Fine-tuning to Prompting: A Paradigm Shift in Knowledge Graph Construction
13 Isaac da Silva Torres (VUA), Guidelines To Flux Between Conceptual Models: Understanding Complex Digital Business Ecosystems
14 Philip Lippmann (TUD), Synthetic Data for Robust Language Modelling

15 Rashmi Khazanchi (OU), Artificial Intelligence in Education: Impact of AI-Based Systems on Mathematics Achievement
16 Carolina Ferreira Gomes Centeio Jorge (TUD), Modelling Artificial Trust for Effective Human-AI Teamwork
17 Maria Tsfasman (TUD), Towards Predicting Memory in Multimodal Group Interactions
18 Riccardo Lo Bianco (TU/e), Deep Reinforcement Learning for Automated Decision-Making in Process Management Systems
19 Israel Campero Jurado (TU/e), Innovations in Optimization and Applications in Healthcare
20 Iftitahu Ni'mah (TU/e), Contrastive Learning and Evaluation in Low Resource Scenario of Natural Language Processing
21 Francisco N.F.Q. Simoes (UU), Causality, Information, and Decision-Making
22 Ruben Verhagen (TUD), Transparent and Explainable Agents for Human-Agent Teaming
23 Eduardo Calò (UU), Automatically Expressing the Meaning of Logical Formulae in Natural Language
24 Shuai Han (UU), Improving Sample Efficiency of Reinforcement Learning: Exploiting Structural Knowledge for Decision Making
25 Francis Saa-Dittoh (VUA), From Radio to AI: African Community-Driven Development of Sustainable Information Systems
26 Mohammed Al Owayyed (TUD), Interactive Simulation-Based Learning Tools for Training Children's Helpline Counsellors
27 Bram Grooten (TU/e), Adaptive Reinforcement Learning: Lean and Dynamic Agents for Robust Generalization
28 Anouck Braggaar (TiU), Does This Answer Your question? Evaluating, Replicating, and Improving Task-Oriented Dialogue Systems
29 Afsana Khan (UM), Unlocking the Value of Data with Vertical Federated Learning
30 Yucheng Yang (TU/e), Generalization in Reinforcement Learning: Theories and Algorithms for Downstream Task and Multi-objective Trade-off Adaptation
31 Luis Pedro Silvestrin (VUA), Efficient Machine Learning for Time-Varying Data: Contributions to Time-Series Machine Learning and Transfer Learning
32 Giovanni Varricchione (UU), Declarative Specifications for Efficient and Safe Reinforcement Learning
33 Markus Funke (VUA), Carving Sustainability into Software Architecture

34 Melika Ayoughi (UvA), External and Internal Semantics for Visual Understanding
35 Clinton Cao (TUD), Modeling Behavior Patterns in Microservice Applications: Practical Applications Using Finite State Machines
36 Jonathan Kamp (VUA), Interpreting the Methods that Interpret Language Models
37 Emre Erdogan (UU), Computational Theory of Mind for Effective Hybrid Intelligence
38 Mahmoud Shokrollahi-Far (TiU), Quran Mining: Computational Stylometry of Quranic Texts
39 Aleksandr Chebykin (CWI, TUD), Efficient Evolutionary Hyperparameter Optimization for Deep Learning
40 Kateryna Ihorivna Holubinka (OU), Holistic Integration of Desktop Virtual Reality Technology in Higher Education: A Design-Based Research Approach
41 Hideaki Joko (RUN), Personalization and Evaluation of Conversational Information Access
42 Nedo Alexander Bartels (VUA), Conceptualizing Platform Revenue Models: A Taxonomy-Driven Design Approach
43 Maksim Gladyshev (UU), "Who Is to Blame?" and "What Is to Be Done?" A Formal Study of Counterfactual Approaches to Responsibility and Causation in (Multi-Agent) AI Systems
44 Iffat Fatima (VUA), Software Architecture Evaluation for Sustainability
45 Tristan Tomilin (TU/e), Scalable Benchmarking of Continual and Safe Embodied Reinforcement Learning
46 Savvina Daniil (CWI, VUA), Bias in Book Recommendation
47 Sabine Molenaar (UU), Requirements Engineering Artifacts in Agile Development
48 Joeri Peters (UU), Bridging the Gap: on the Role of Domain Knowledge in Reconciling Machine Learning with Expert Reasoning in High-Risk Domains
49 Dimitar Rangelov (UT), Enhancing Crime Scene Investigations: From Technology to Evidence
50 W.J. van Leeuwen (TU/e), On Estimating the Cardinality of Conjunctive Queries over Property Graph Databases for Query Planning
51 Roland Saur (TUD, CWI), Multi-Dimensional Flexibility Integration Practical and Conceptual Challenges

Curriculum Vitae



Roland Saur graduated in July 2013 with a Bachelor's of Science from the University of Stuttgart, Germany. In August 2015 he graduated with a Master's of Science in Sustainable Energy Technology from Delft University of Technology. Afterwards, he spent a year as a researcher in South Korea, at the Busan Campus of the Friedrich-Alexander University, Erlangen-Nürnberg. In September 2017, he began his PhD research at the Centrum Wiskunde & Informatica (CWI) and Delft University of Technology. The result of this research is presented in this Dissertation.

The widespread adoption of renewable energy sources has greatly contributed to the fight against climate change. This success, however, has put a significant strain on power grids.

Further integration of renewable energy sources, while necessary, will require a smarter use of existing infrastructure. Power-to-gas, electric vehicles and heat pumps already allow us to provide flexibility by shifting energy use across domains. This thesis explores how new market designs can better facilitate such multi-energy flexibility.

This work discusses both foundational questions of multi-dimensional energy trading and conceptual questions of applying fairness and honesty to multi-dimensional energy markets: What are fair unit-prices when energy can be shuffled across time? How can an agent be incentivized to be honest about its forecast when it does not know it perfectly itself?