

Multiagent Demand-Side Management for Real-World Energy Cooperatives

by

Charilaos Akasiadis

Submitted to the School of Electrical and Computer Engineering
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Electrical and Computer Engineering

at the

TECHNICAL UNIVERSITY OF CRETE

October 2017

© Charilaos Akasiadis, 2017. All rights reserved.

The author hereby grants to TUC permission to reproduce and to distribute publicly paper and electronic copies of this dissertation in whole or in part in any medium now known or hereafter created.

Multiagent Demand-Side Management for Real-World Energy Cooperatives

by

Charilaos Akasiadis

Submitted to the School of Electrical and Computer Engineering
on October 2, 2017, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Electrical and Computer Engineering

Abstract

Balancing energy demand and production in modern Smart Grids with increased penetration of intermittent renewable energy resources is a challenging problem. *Demand-Side Management (DSM)*, i.e., the design and application of sophisticated mechanisms for managing and coordinating energy demand, has been hailed as a means to deal with this problem. In this dissertation, we propose mechanisms for the formation of *agent cooperatives* offering large-scale DSM services, and put forward a *complete framework* for their operation. Individuals, being either mere consumers, or even prosumers of electricity, are represented by rational agents and form *coalitions* to offer demand shifting from peak to non-peak intervals.

For *cooperatives of consumers*, we present an effective consumption shifting scheme, equipped with desirable guarantees, such as individual rationality, truthfulness, and (weak) budget balance. Our scheme employs several algorithms to promote the formation of the most effective shifting coalitions. It takes into account the shifting costs of the individuals, and rewards them according to their shifting efficiency. In addition, it employs internal pricing methods that guarantee individual rationality, and allow agents with initially forbidding costs to also contribute to the shifting effort. The truthfulness of agent statements regarding their shifting behaviour is ascertained via the incorporation of a strictly proper scoring rule. We provide a thorough evaluation of our approach on a simulations setting constructed over a real-world dataset. Our simulation results clearly demonstrate the benefits arising from the use of agent cooperatives in this domain.

Moreover, to also allow the decentralized coordination of *cooperatives of pro-*

sumers, we combine, for the first time in the literature, a strictly proper scoring rule with a specialized cryptocurrency framework. Using our approach, prosumers collaborate with the use of a blockchain-oriented framework to manage their demand, in order to make more profits from the selling of their energy. When tested on a simulation setting that uses dynamic electricity pricing to promote the usage of locally generated renewable energy, our approach drives the prosumers to become more engaged in DSM and achieve increased profits; the balancing of demand and local renewable supply is more effective; and dynamic electricity prices are more stable.

Furthermore, we propose a *vehicle-to-grid/grid-to-vehicle (V2G/G2V)* algorithm that balances demand and local renewable supply in environments populated with electric vehicles. The approach promotes new business models that make effective use of the capability of electric vehicles to store energy in their batteries. Additionally, to assess participating agents' uncertainty, and correctly predict their future behaviour regarding power consumption shifting actions, promoting in this way accuracy and effectiveness, we adopt various machine learning techniques, adapt them to fit the problem domain, and use these to effectively monitor the trustworthiness of agent statements regarding their final shifting actions. Simulation results confirm that the adoption of machine learning techniques provides tangible benefits regarding enhanced cooperative performance, and increased financial gains for the participants.

Finally, we provide the methodology for delivering large-scale DSM services in the real world. To this purpose, we devise an IoT service-oriented architecture for DSM applications, through which we test different GUIs and incentive types for managing energy consumption. In this context, we present a "serious game" solution that was tested by real human subjects. Our approach comes complete with the adoption of a statistical analysis methodology to validate reductions in consumption and the promotion of renewable energy usage in real world settings. Our results show that using the proposed methods in real-world large-scale settings can significantly benefit the end-users, the Grid, and the environment. The success of our approach indicates that the combination of methods from multiple fields of Computer Science can deliver high quality human-centered solutions to complex real-world problems.

Abstract (Greek)

Η εξισορρόπηση της ζήτησης και της παραγωγής ενέργειας στα σύγχρονα έξυπνα δίκτυα με αυξημένη διείσδυση των ανανεώσιμων πηγών ενέργειας είναι ένα μεγάλο πρόβλημα, για τη λύση του οποίου απαιτούνται περίπλοκοι μηχανισμοί για την αποτελεσματική διαχείριση της ζήτησης ενέργειας (Demand-Side Management - DSM). Σε αυτή τη διατριβή, προτείνουμε μηχανισμούς για το σχηματισμό συνεταιρισμών ευφών πρακτόρων που προσφέρουν υπηρεσίες DSM μεγάλης κλίμακας, και παρουσιάζουμε ένα ολοκληρωμένο πλαίσιο για τη λειτουργία τους. Τα άτομα, είτε απλώς καταναλωτές, είτε ακόμη και παραγωγοί, ταυτόχρονα, ηλεκτρικής ενέργειας, αντιπροσωπεύονται από ορθολογικούς πράκτορες και σχηματίζουν συνεταιρισμούς για να προσφέρουν τη μετατόπιση της ζήτησης από τα χρονικά διαστήματα με έλλειψη ανανεώσιμης παραγωγής, σε άλλα με περίσσεια.

Για *συνεταιρισμούς καταναλωτών*, σχεδιάσαμε ένα αποτελεσματικό σχήμα μετατόπισης της κατανάλωσης, το οποίο προσφέρει συγκεκριμένες εγγυήσεις, όπως ορθολογισμό κατ' ανεξαρτησία (individually rational), ειλικρίνεια (truthfulness) και ασθενή ισολογισμό του κεφαλαίου (weak budget balance). Το σχήμα μας χρησιμοποιεί διάφορους αλγορίθμους για την προώθηση του σχηματισμού των πιο αποτελεσματικών ομάδων συνεισφοράς. Λαμβάνει υπόψη τα ατομικά κόστη μετατόπισης κατανάλωσης και ανταμείβει τους συνεισφέροντες πράκτορες ανάλογα με την αποδοτικότητα τους. Επιπλέον, χρησιμοποιεί εσωτερικές μεθόδους ανατιμολόγησης που εγγυώνται τον ορθολογισμό κατ'ανεξαρτησία, και επιτρέπουν στους πράκτορες με αρχικώς απαγορευτικά κόστη να συμβάλλουν επίσης στην προσπάθεια μετατόπισης της κατανάλωσης. Η ειλικρίνεια στη συμμετοχή των πρακτόρων εξασφαλίζεται μέσω της ενσωμάτωσης ενός αυστηρά αρμόζοντα κανόνα βαθμολόγησης. Παρέχουμε μια διεξοδική αξιολόγηση της προσέγγισής μας σε μια προσομοίωση διαμορφωμένη βάσει συνόλου δεδομένων από τον πραγματικό κόσμο. Τα αποτελέσματα καταδεικνύουν σαφώς τα οφέλη που προκύπτουν από τη χρήση συνεταιρισμών πρακτόρων σε αυτόν τον τομέα.

Επιπροσθέτως, για να επιτρέψουμε επίσης τον αποκεντρωμένο συντονισμό των συνεταιρισμών καταναλωτών-παραγωγών ενέργειας, συνδυάζουμε για πρώτη φορά έναν αυστηρά αρμόζοντα κανόνα βαθμολόγησης με ένα εξειδικευμένο κρυπτονόμισμα (cryptocurrency). Χρησιμοποιώντας την προσέγγισή μας, οι καταναλωτές-παραγωγοί συντονίζονται με τη χρήση ενός πλαισίου βασισμένου σε "αλυσίδα από μπλοκ" (blockchain), για να διαχειριστούν τη ζήτησή τους προκειμένου να κερδίσουν περισσότερα από την πώληση της παραγωγής τους. Όταν τη δοκιμάσαμε σε ένα περιβάλλον προσομοίωσης που χρησιμοποιεί δυναμικά μεταβαλλόμενη τιμολόγηση ηλεκτρικής ενέργειας για να προωθήσει τη χρήση τοπικά παραγόμενης ανανεώσιμης ενέργειας, η προσέγγισή μας ωθεί τους καταναλωτές-παραγωγούς να εμπλακούν περισσότερο στη διαχείριση της ζήτησης και να επιτύχουν αυξημένα κέρδη. Ακόμη, η εξισορρόπηση της ζήτησης και της παραγωγής ανανεώσιμων πηγών ενέργειας είναι αποτελεσματικότερη, και οι δυναμικά μεταβαλλόμενες τιμές ηλεκτρι-

κής ενέργειας είναι πιο σταθερές.

Επιπλέον, προτείνουμε έναν αλγόριθμο για την φόρτιση και αποφόρτιση ηλεκτρικών οχημάτων (υπηρεσία V2G / G2V) που έχει επίσης ως στόχο την εξισορρόπηση της ζήτησης και της τοπικής ανανεώσιμης προσφοράς, αλλά προωθεί επίσης και νέα επιχειρηματικά μοντέλα που κάνουν αποτελεσματική χρήση της δυνατότητας των ηλεκτρικών οχημάτων να αποθηκεύουν ενέργεια στη μπαταρία τους. Συμπληρωματικά, για να αξιολογήσουμε την αβεβαιότητα των συμμετεχόντων, και να προβλέψουμε σωστά τη μελλοντική τους συμπεριφορά όσον αφορά τη διαχείριση ενεργειακής κατανάλωσης, προωθώντας έτσι την ακρίβεια και την αποτελεσματικότητα, υιοθετούμε διάφορες τεχνικές μηχανικής μάθησης, τις προσαρμόζουμε στα συγκεκριμένα προβλήματα, και τις χρησιμοποιούμε για την αποτελεσματική παρακολούθηση της αξιοπιστίας των δηλώσεων του κάθε πράκτορα σχετικά με τις πραγματικές του ενέργειες. Τα πειραματικά αποτελέσματα επιβεβαιώνουν ότι η υιοθέτηση τεχνικών μηχανικής μάθησης παρέχει απτά οφέλη όσον αφορά την ενισχυμένη απόδοση των συνεταιρισμών και αυξημένα οικονομικά οφέλη για τους συμμετέχοντες.

Τέλος, παρουσιάζουμε μια μεθοδολογία για την παροχή υπηρεσιών διαχείρισης ζήτησης (DSM) μεγάλης κλίμακας στον πραγματικό κόσμο. Προωθούμε μια αρχιτεκτονική *διαδικτύου των πραγμάτων (IoT)* για εφαρμογές DSM, μέσω της οποίας δοκιμάσαμε διαφορετικές γραφικές διεπαφές (GUI) και τύπους κινήτρων για τη διαχείριση της κατανάλωσης ενέργειας. Οι λύσεις αυτές οργανώθηκαν στην μορφή ενός "σοβαρού παιγνίου" (serious game), το οποίο παρουσιάστηκε σε, και χρησιμοποιήθηκε από πραγματικούς ανθρώπους. Η προσέγγιση μας σε αυτή τη διατριβή συνοδεύεται από μια μεθοδολογία στατιστικής ανάλυσης για την επικύρωση των μεγεθών μείωσης της κατανάλωσης και της προώθησης της χρήσης ανανεώσιμης ενέργειας σε πραγματικές συνθήκες. Τα αποτελέσματα μας δείχνουν ότι η χρήση των προτεινόμενων μεθόδων σε συνθήκες μεγάλης κλίμακας πραγματικού κόσμου μπορεί να ωφελήσει σημαντικά τους τελικούς χρήστες, το δίκτυο, και το περιβάλλον. Η επιτυχία της προσέγγισής μας δείχνει ότι ο συνδυασμός μεθόδων από πολλούς τομείς της Επιστήμης των Υπολογιστών μπορεί να προσφέρει υψηλής ποιότητας ανθρωποκεντρικές λύσεις σε σύνθετα προβλήματα του πραγματικού κόσμου.

Jury

3-member Committee

Georgios Chalkiadakis
Michail G. Lagoudakis
Ioannis A. Vetsikas

Examiners

Eftichios Koutroulis
Sarvapali D. Ramchurn
Constantine D. Spyropoulos
Georgios Stavrakakis

*To my family
for their love,
endless support,
and patience*

Acknowledgments

Over the past four years I have received support and encouragement from a number of individuals for whom I am really thankful. First of all, Professor Georgios Chalkiadakis has been a mentor, colleague, and friend, to whom I express my sincerest gratitude and appreciation. His guidance and motivation made this a very thoughtful and rewarding journey, which has opened many doors for further research and collaborations.

I would like to express my gratitude to the members of my advisory committee, Professor Michail Lagoudakis, and Dr. Ioannis Vetsikas for their help, support, and lessons they offered to me. Also, to the Professors of my examination committee, Eftichios Koutroulis, George Stavrakakis, and Sarvapali Ramchurn for helping me completing my studies. Finally, to Dr. Constantine Spyropoulos, for also giving me the opportunity to work and conduct research at the Institute of Informatics and Telecommunications of NCSR ‘Demokritos’ since 2014.

Contents

1	Introduction	27
1.1	Multiagent Demand-Side Management	30
1.2	Scaling-up: Smart Cities and The Internet-of-Things	33
1.3	Providing the Incentives for Participation	34
1.4	Contributions	37
1.5	Thesis Outline	42
2	Background	45
2.1	Basic Assumptions	46
2.2	Real-World Coops, VPPs, and Aggregators	47
2.3	Demand-Side Management and Mechanism Design	50
2.4	Incentive Compatibility and Related Properties	56
2.5	Scoring rules	60
2.5.1	The Continuously Ranked Probability Score	62
2.6	Uncertainty in Demand-Side Management	63
2.7	Constrained Optimization for DSM	64
2.8	Cryptocurrencies and Blockchain Technology	66
2.9	Designing the front end of DSM applications	70

2.10	The Grid and the Electricity Markets	73
3	Consumer Cooperatives	77
3.1	Related Work	82
3.2	A Generic Electricity Consumption Shifting Model	88
3.2.1	Scheme overview	91
3.2.2	Constraints	92
3.2.3	Agent incentives	94
3.2.4	Group price	96
3.2.5	Continuously ranked probability score	97
3.3	Forming Effective Demand Shifting Coalitions	99
3.3.1	Choosing the acting coalitions	102
3.3.2	Cooperative bidding and billing	109
3.4	Internal Pricing	111
3.4.1	Heuristic internal price balancing	112
3.4.2	Internal pricing as a constrained optimization problem	115
3.5	Mechanism Properties	118
3.6	Experimental Evaluation	122
3.6.1	The simulations dataset	122
3.6.2	Evaluation of the proposed coalition formation methods	124
3.6.3	Assessing the effect of different group pricing slopes	130
3.6.4	Coalition size vs. group price range	131
3.6.5	Assessing the CRPS effect	133
3.6.6	Experimenting with different internal price balancing techniques	136

<i>CONTENTS</i>	13
3.6.7 Participant availability	140
3.7 Conclusions	141
4 Prosumer Cooperatives	143
4.1 A Generic Prosumer Consumption Shifting Model	147
4.1.1 Promoting Demand-Side Management	148
4.1.2 Shifting to profitable time intervals	150
4.1.3 Shifting without coordination	153
4.2 Distributed Shifting and Reward Sharing	154
4.2.1 Cooperative balance increase	160
4.2.2 COOPcoin for prosumer cooperatives	161
4.2.3 Selection of contributors	165
4.3 Experimental Evaluation	167
4.3.1 Simulations setting	167
4.3.2 Individual vs. cooperative action	170
4.3.3 Evaluating contributor selection methods	171
4.3.4 Shifting capacity sensitivity test	172
4.3.5 Reward sharing methods evaluation	176
4.4 Conclusions	177
5 Incorporating EVs and Storage Systems	179
5.1 Related Work	185
5.2 The V2G/G2V Extension	192
5.3 Experimental Evaluation	197
5.3.1 Simulation over a one-year time horizon	199

5.3.2	EV reservation prices sensitivity test	202
5.3.3	Performance comparison	204
5.4	Conclusions	206
6	ML for Enhanced Performance	207
6.1	Related Work	211
6.2	Method Description	214
6.2.1	Histogram filter	216
6.2.2	Gaussian process filter	217
6.2.3	k -nearest neighbours regression	219
6.2.4	Kernel regression	220
6.3	Experimental Evaluation	222
6.3.1	Evaluation of monitoring techniques	224
6.3.2	Experiments with agents of the same behavior	229
6.3.3	Trial with mixed agent behavior classes	231
6.4	Conclusions and Future Work	234
7	Supporting REScoops and Aiding End Users	235
7.1	Related Work	239
7.2	System Architecture for DSM in IoT Ecosystems	244
7.3	A Serious Game for Demand Shifting	246
7.3.1	Game controls	248
7.3.2	Game feedback	249
7.3.3	The game with economic incentives	251
7.3.4	The game with social incentives	252

<i>CONTENTS</i>	15
7.3.5 Results from trials with real users	252
7.4 Methodologies for REScoop Data Analysis	256
7.5 Conclusions	258
8 Conclusions	261
8.1 Summary	262
8.2 Future Work and Research Directions	267
A Datasets Used in the Simulations	275
A.1 Agent shifting capacities	276
A.2 Agent statements and final shifting actions	277
A.3 Shifting costs	279
A.4 Renewable energy sources production levels	283
A.5 Electric vehicles demand levels	284

THIS PAGE INTENTIONALLY LEFT BLANK

List of Figures

1-1	Overview of scientific fields, and our contributions.	40
2-1	Aggregate curves of Greece.	47
2-2	Simplified blockchain of Bitcoin.	68
3-1	Overview of scientific fields, and our contributions in this chapter. . .	81
3-2	Scheme objective representation. Portions of peak-load at t_h intervals, are shifted to lower demand intervals t_l	94
3-3	Forms of p_g for different values of κ	131
3-4	Average coalition size vs. p_g increase for the coalition formation meth- ods <i>CF1</i> and <i>CF3</i>	132
3-5	Average losses in gain (increase in bill) due to <i>CRPS</i> , as induced by increasing participant inaccuracies, across all intervals.	135
3-6	Prices (€/kWh shifted) assigned to each individual for a sample peak interval as a result of employing our pricing methods.	138
3-7	Expected gains for each participant for a sample peak interval, when applying different reward sharing approaches.	139

3-8	Cooperative effectiveness versus contributor availability.	140
4-1	Overview of scientific fields, and our contributions in this chapter. . .	146
4-2	Prosumer network and interactions.	155
4-3	Time-series of B_t^{buy} and B_t^{sell} used in the simulation as the Grid's pricing mechanism.	169
4-4	Average individual gains during 2012, and average shifting coalition sizes vs. shifting capacity changes for all three selection methods. . .	173
4-5	Cooperative balance in 2012 vs. shifting capacity changes for all three selection methods.	175
4-6	COOPcoin total wealth difference between accurate and inaccurate actors after 344 days.	176
5-1	Participant types and interactions of the assumed microgrid with V2G/G2V services.	181
5-2	Overview of scientific fields, and our contributions in this chapter. . .	184
5-3	Flowchart of algorithm for V2G/G2V services.	194
5-4	Hourly aggregate values for a sample day that is produced from the simulator.	197
5-5	Hourly aggregate values for a sample day that is produced from the simulator after V2G/G2V is applied.	198
5-6	Daily mean and standard deviation values from the 365 days simulation.	200
5-7	Daily mean and standard deviation from EV demand values of the 365 days simulation.	200
5-8	Typical daily imbalance curve during the year.	202

5-9	Daily microgrid gain vs. EV reservation prices offset.	203
5-10	Average time for the calculation of daily consumption rescheduling as the number of incorporated EVs rises.	203
6-1	Overview of scientific fields, and our contribution in this chapter. . .	211
6-2	Probability density functions of agent behaviors.	223
6-3	Histogram with 8 bins and corresponding roulette wheel.	225
6-4	<i>GP</i> fits. Illustration of <i>GP</i> training for two non-linear cases.	226
6-5	<i>k-NN</i> and <i>KR</i> fits. Illustration of the regression functions for two non-linear cases, with $k = \sqrt{n}$ for the <i>k-NN</i> and $h = 1$ for the <i>KR</i> . . .	227
6-6	<i>HF</i> , <i>GP</i> , <i>k-NN</i> , and <i>KR</i> prediction examples for the various agent behavior scenarios.	228
7-1	Overview of scientific fields, and our contributions in this chapter. . .	239
7-2	IoT based architecture for large-scale DSM mechanisms.	245
7-3	The two different interfaces.	248
7-4	Typical day of the game rounds.	254
A-1	Probability density functions of agent behaviors.	279
A-2	Probability density function of $\mathcal{B}(0.20865, 3.76572)$	280
A-3	Model of agent costs.	283

THIS PAGE INTENTIONALLY LEFT BLANK

List of Tables

3.1	Notation.	89
3.2	Kissamos 2012: Size and corresponding individual average consumption and bills for each consumption contract type.	123
3.3	Average results over a 100 days simulation period.	125
3.4	Shifting performance of the most active participants per consumer class and coalition formation method, over 100 days.	127
3.5	Monthly financial gains for the four most active participants.	128
3.6	Average gains for participants with 15 or more participations per month.	128
3.7	Average results over a 100 days simulation period for the dataset with only industrial consumers.	129
3.8	Results from an 100 days simulation for different values of κ in the group pricing function.	130
3.9	Average differences between expected and final gain per contribution, with and without internal <i>CRPS</i> penalization.	137
4.1	Performance of individual and cooperative action.	171

4.2	Total cooperative gains and balance in 2012 for each selection method (NRG coins).	172
5.1	V2G participations—energy from batteries only.	201
5.2	Imports and exports of energy.	201
5.3	Shifting performance of the most active participants per consumer class and coalition formation method, over 100 days.	205
6.1	Behavioural classes of the scheme participants used in the simulation.	224
6.2	Average daily gain (€) and accuracy in aggregate amount of <i>kWh</i> shifted (%), from 100 days simulations, one for each agent behavior class.	230
6.3	Average results from an 100 days simulation scenario including all agent accuracy classes.	233
7.1	Acceptance ratio over shifting different appliances per type of game. .	254
A.1	Consumption contract types and their shifting capacities (as percentages of their total demand).	276
A.2	Behavioural classes of the scheme participants.	278

Own publications related to this dissertation

1. Charilaos Akasiadis and Georgios Chalkiadakis. *Agent Cooperatives for Effective Power Consumption Shifting*. In Proceedings of the 27th AAAI Conference on Artificial Intelligence, (AAAI-2013), pp. 1263—1269, Bellevue, WA, USA, July 2013.
2. Charilaos Akasiadis and Georgios Chalkiadakis. *Stochastic Filtering Methods for Predicting Agent Performance in the Smart Grid*. In Proceedings of the 21st European Conference on Artificial Intelligence- Including Prestigious Applications of Intelligent Systems (ECAI/PAIS 2014), pp. 1205—1206, Prague, Czech Republic, August 2014.
3. Charilaos Akasiadis, Kasia Panagidi, Nikolaos Panagiotou, Paolo Sernani, April Morton, Ioannis A. Vetsikas, Lora Mavrouli, Konstantinos Goutsias. *Incentives for Rescheduling Residential Electricity Consumption to Promote Renewable Energy Usage*. In Proceedings of SAI Intelligent Systems Conference (IntelliSys), pp. 328—337, IEEE, London, UK, November 2015.

4. Charilaos Akasiadis and Georgios Chalkiadakis. *Decentralized Large-Scale Electricity Consumption Shifting by Prosumer Cooperatives*. In the 22nd European Conference on Artificial Intelligence (ECAI - 2016), pp. 175—183, The Hague, Netherlands, August-September 2016.
5. Charilaos Akasiadis and Georgios Chalkiadakis. *Mechanism Design for Demand-Side Management*. IEEE Intelligent Systems 32:1, pp. 24—31, IEEE, 2017.
6. Charilaos Akasiadis and Georgios Chalkiadakis. *Cooperative Electricity Consumption Shifting*. Sustainable Energy, Grids and Networks, 9C, pp. 38—58, Elsevier, 2017.
7. Charilaos Akasiadis and Alexandros Georgogiannis. *Predicting Agent Performance in Large-Scale Electricity Demand Shifting*. Advances in Building Energy Research, pp. 1—22, Taylor and Francis, 2017 (in press).
8. Charilaos Akasiadis, Georgios Chalkiadakis, Michail Mamakos, Nikolaos Savvakis, Theocharis Tsoutsos, Thomas Hoppe, Frans Coenen. *Analyzing Statistically the Energy Consumption and Production Patterns of European REScoop Members: Results from the H2020 project REScoop Plus*. In the 9th International Exergy, Energy and Environment Symposium (IEEES - 9), pp. 372—378, Split, Croatia, May 2017.
9. Frans Coenen, Thomas Hoppe, Georgios Chalkiadakis, Theocharis Tsoutsos, Charilaos Akasiadis. *Exploring energy saving policy measures by renewable energy supplying cooperatives (REScoops)*. In ECEEE 2017 Summer Study on Energy Efficiency, pp. 1—12, Presqu'île de Giens, France, May-June 2017.

List of Abbreviations

AI	Artificial Intelligence
CDF	Cumulative Density Function
CRPS	Continuously Ranked Probability Score
DR	Demand Response
DSM	Demand-Side Management
DSO	Distribution System Operator
DT	Decision Theory
EU	European Union
EV	Electric Vehicle
G2V	Grid to Vehicle
GT	Game Theory
IoT	Internet of Things
ISO	Independent System Operator
MAS	Multi-agent System
MD	Mechanism Design
PDF	Probability Density Function
RES	Renewable Energy Sources
RTP	Real-time Pricing
TOU	Time of Use
TSO	Transmission System Operator
V2G	Vehicle to Grid
VPP	Virtual Power Plant

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 1

Introduction

Electricity is undoubtedly one of the most important commodities in our world, affecting almost every aspect of daily life, from industrial production processes and commercialism, to people’s heating, well-being and recreation. As a consequence of this fact, governmental policies have naturally driven energy generation companies to expand their production infrastructure, so as to meet the end-users’ increasing energy demand. However, legacy systems for electricity generation mainly produce electricity by the burning of fossil fuels; apart from the fact that their sources are depleting, their use is harmful to the environment as their extraction might harm surrounding areas, and their burning produces gases which help exacerbate the so-called “greenhouse effect”.

As a remedy for these concerns, recent trends propose “greener” approaches that will help future electricity production become less polluting, introducing the hope for a more sustainable development [60, 71, 147]. The emerging *renewable energy* generation sources can be organized in a non-industrial and decentralized manner,

allowing the average household to contribute and benefit from its participation to the electricity production process [5, 140]. Despite the positive effects from using renewable energy sources, new challenges arise for electricity production and grids management. This is because weather-dependent electricity sources are by definition intermittent, and potentially unreliable regarding their output size.

To overcome the difficulties of effective renewable energy sources integration, “Smart Grid”-related research has received much attention in the last few years. Its general objective is to create a more secure, reliable and efficient electricity networks infrastructure, with affordable energy produced mostly by green sources, production costs minimized, and energy savings maximized [53]. Smart Grids offer communication infrastructure for all levels of the power grid, from generators, to transmission systems, and, finally, to the end-users. The information that is exchanged at real-time can be used to build prediction models, and balance demand and supply proactively, resulting to the so-called *Demand-Side Management (DSM)*. In DSM, the end-users alter their demand profiles, instead of the producer side controlling the production levels, since this might be hard or very expensive, e.g., because renewable generation output is often uncontrollable, or due to the need for turning on emergency generators. DSM in Smart Grid environments generally aims to induce changes to the consumers’ demand curves, so as for the total demand to match the production [68, 192, 177].

Now, to achieve successful DSM, we must make sure that two conditions hold: Firstly, energy consumption management must be performed in a coordinated fashion so as to make the small individual demand portions count in the face of the

large aggregate consumption levels, and also guarantee that certain Grid¹ stability constraints will hold. Secondly, DSM participants must come in very large numbers, overcoming this way a participants “critical mass”, one that is able to contribute in DSM at any time interval of the day, offering adequate quantities of load for shifting, despite the uncertainty that exists.

Naturally, the complex data communication and calculations required, can all be performed by *autonomous agents* [147, 202]: as such, this field provides a grand challenge for the fields of Artificial Intelligence (AI) and Multi-Agent Systems (MAS). Note that *whether* an agent decides to contribute and *how much load* will actually be shifted to other intervals (i.e., DSM participation and action), are both linked to every day human activities, which, usually, are hard to be postponed or rescheduled. This fact induces uncertainty regarding the actual shifting of load, a procedure that will take place in the future. Thus, sophisticated mechanisms that provide incentives and rewards must be incorporated, so that participants cooperate and become engaged in the DSM efforts [181]. More specifically, a mechanism must ensure that each participant will be granted back proportions of the profit achieved by DSM in a personalized manner, and, moreover that, no participant must be able to “game” the mechanism, e.g. via collusions, misleading statements, and unethical practices aiming for manipulation [172].

In the work presented in this thesis, we engineer an effective DSM scheme that can be easily applied to existing Smart Grids. The main mechanism combines methods and techniques inspired by multiple fields of computer science that is, multiagent

¹By the term *Grid*, we refer to the entities that actually manage the electricity grid, e.g. by setting the prices, and managing the distribution of electricity in general. Such entities can be individual utility companies, or nationwide independent system operators (ISO).

systems, artificial intelligence, optimization, and game theory and mechanism design. Here, individual preferences are taken into account, and personalized rewards are given back to the participants, according to their distinct interactions with the scheme. We model Grid constraints and formulate truthful mechanisms that guarantee security and reliability from the Grid side, as well as participation incentives for the end-user side. Apart from the core of the mechanism that is designed to deliver *large-scale* cooperative DSM operations, additional modules, such as machine learning techniques for regression, contemporary Internet of Things [70] architectures for easy application deployment and user interaction, as well as blockchain technology [131] for decentralized coordination, are also incorporated forming a complete, functional, and effective DSM framework. By using such an approach, all stakeholders can benefit—and so does the environment—due to the increased utilization of renewable energy resources, resulting this way to a “win-win” scenario, both for the Grid side and for the end-users.

In the rest of this chapter, we provide a thorough explanation of the motivation behind our work, and highlight the difficulties that exist in the highly complex domain of energy efficiency in the Smart Grid.

1.1 Multiagent Demand-Side Management

As mentioned above, research towards sustainable approaches and solutions, especially in the energy domain, has flourished in recent years (partly due to the observed climate change [2, 4]). In order to reach a sustainable energy plan, new Grid architecture designs are needed, which will incorporate recent technological advancements,

and deliver a green and secure infrastructure [102]. Of course this is a complex venture, due to the nature of electricity, and requires specific handling. This is because, first, the storage of electricity is still rather expensive, and, second, the energy output of green sources is intermittent. Hence, supply might easily become unstable and limit the penetration of “green” sources into the electricity grid [31].

The contemporary *Smart Grid* agenda of research aims to create a more secure, reliable and efficient electricity networks infrastructure, with energy produced mostly by “green” sources, production costs minimized, and affordable electricity made easily and reliably available to the public [51, 53, 60].

Due to the scale and complexity of electrical networks management, AI and MAS solutions are in high demand in the emerging markets involving business entities providing Smart Grid services [23, 147]. Many such entities have already adopted a business model that pulls together the resources and abilities of multiple economically-minded individuals. Specifically, the emergence of *virtual power plants (VPPs)* or *cooperatives*² of small-to-medium size electricity producers, consumers, or even prosumers³—that operate as a virtually single entity—has been hailed as a means to create large, efficient, trustworthy providers of renewable energy production or electricity consumption reduction (peak-trimming) services [23, 40, 96, 100, 147, 155]. VPPs can also deliver a range of *DSM services*. As argued in [129], due to the uncertain nature of the most widespread generators (solar and wind), DSM is expected to play an important role for the reliability of the future smart grids. Note that, frequently, in related literature, DSM is also termed as Demand Response (DR).

²The term cooperative refers to conglomerations of end-users that organize and operate largely on a democratic manner [1]; which is not necessarily the case for Virtual Power Plants.

³A prosumer is an entity that both produces and consumes energy.

However, DR usually indicates that participants might be called in for last minute action, while in DSM participation is requested long in advance, thus the two notions must not be confused [23].

In DSM, consumers contribute to the stability of the electricity grid in exchange for certain, usually monetary, rewards. Furthermore, load-management schemes are an alternative to *electricity storage*—a problem difficult by its nature, and the tackling of which requires the use of expensive equipment [125]. In light of these facts, a reliable and sustainable energy infrastructure should conduct *proactive balancing*, between the amount of electricity produced and consumed, at all times. As technology and metering devices' capabilities advance, new possibilities appear. For example, even though the smart grid is at a relatively early stage, it could be enhanced by using multi-agent technologies; each agent represents a consumer, a producer, or even a prosumer, has its own goals and preferences, is interconnected to the other agents, and must comply with rules, such as paying for electricity, in order to reach these goals. Moreover, these autonomous and intelligent agents can attempt to learn the users' preferences and interests, and control their energy consumption appropriately. For the cooperative shifting to be successful at *large-scale* shifting effort, it is obvious that *coordinated joint shifting efforts* have to take place, carried out by *demand shifting coalitions*.

1.2 Scaling-up: Smart Cities and The Internet-of-Things

Despite the fact that massive participation is required for DSM to work, this often leads to herding effects. As such, the estimated consumption curve could significantly change, both endangering the Grid's stability, but also leading to substantially different economic outcomes [195]. In order for agent cooperatives to be functional, efficient, and profitable, they need to make business decisions regarding which members to include in consumption shifting coalitions, from a very large set of available end-users. These decisions naturally depend on the abilities (e.g., electricity production or consumption reduction capacities) of individual agents. These abilities need to be either monitored by some central cooperative-managing agent, or need to be truthfully and accurately communicated to it. However, it is clear that in the large and dynamically changing scene of the Smart Grid, trust between selfish agents is not implied, and must be guaranteed.

Continuing with the technological advancements of recent years, everyday physical objects have been modified with the embodiment of short-range and energy efficient mobile transceivers and have been enhanced with unique identifiers. The networking of such objects, originally not meant to be "computerized", has led to the so-called "Internet-of-Things" (IoT) [24], which is considered as the next industrial revolution [70]. IoT is expected to find applications in many areas, such as industry, logistics, building and home automation, smart cities, smart manufacturing, health-care, automotive, etc. The aforementioned are only a small fragment of the areas that benefit from IoT technologies [73, 26]. In our case, one can take advantage

of the easily connected devices and cloud services offered by IoT [98] to: construct complex DSM applications; manage the large sets of sensors; data streams and user preferences; and, finally, deliver the infrastructure for DSM that enables large-scale user participation.

Moreover, in large settings where trust is an issue and where centralized control is unlikely, there is much need for tools for trusted decentralized cooperation. *Cryptocurrencies* and *blockchain-oriented algorithms* is a recent technology developed to cover this need [185]. Blockchain algorithms run distributedly, and are transparent by using encryption methods, which guarantee that the transactions are secure, and that no external third-parties need to take part in the exchanges [62]. In our work, to overcome the potential herding effects and the need for cooperative manager determination, we champion the use of a purpose-designed *cryptocurrency protocol* for *distributed prosumer cooperative coordination*.

1.3 Providing the Incentives for Participation

In a Smart Grid setting, *trust* is very important between the participants, thus appropriate incentives must be given to truthfully and accurately report their intended future actions, along with their corresponding uncertainty regarding those actions [7, 40, 96, 155]. Generally, the smooth operation of the Smart Grid and related DSM schemes is aided by the existence of *rules* and *incentives* that lead potentially selfish individuals to adopt a cooperative behaviour and coordinate their actions. This is exactly the problem studied by *mechanism design (MD)* [133], a sub-field of *game theory (GT)* that explores how to design a setting (viewed as a *game*)

so that *rational actors* (or *players*) adopt a behaviour that helps meet the designer's objectives. In other words, MD schemes seek to offer incentives or counter-incentives for achieving desired social outcomes, to individuals that aim to maximize their own utility. Typically, such schemes strive to be *incentive compatible*, meaning that participants are incentivized to be *truthful* regarding their private preferences, and that "gaming" the scheme leads to worse outcomes for "misbehaving" actors. As such, MD can be used to create DSM schemes that promote more *efficient network operation*, by granting *economic and/or social gains* (or, reversely, threaten to inflict similar losses) to the participating individuals—e.g. monetary prizes, and promotions in social networks. Of course, scheme participation is determined by each actor individually; and the incentives must be sophisticated, so as to drive changes to human behavior and habits (with which electricity consumption is strongly correlated), while maintaining the profitability of energy sector businesses. Thus, apart from the technological point of view, a designer should also consider the social aspect, one that incentivizes people to participate in demand side management (DSM) efforts—for instance by shifting their consumption to match the availability of green energy.

In order to provide incentives for consumption rescheduling to the actors, variable pricing techniques are often used. This means that instead of applying a flat pricing scheme, time-of-use (TOU), or real-time pricing (RTP) are employed [27, 146]. By setting higher electricity price values for buying energy during intervals of high demand, and lower values during intervals of low demand, it is possible for an electricity consumer to reduce her expenses by rescheduling her energy usage to the most profitable intervals [7]. This is a task that becomes even more important (and

challenging) when it comes to *electricity prosumers*. As prosumers both produce and consume energy [21, 184], they can take advantage of fluctuations in prices, and generate even more profit [111].

Now, it is conceivable that the Grid would be willing to promise significantly lower electricity rates for considerable shifting efforts only, which cannot normally be undertaken by any consumer alone (due to small reduction capacity or high shifting costs). As a result, agents are motivated to join forces in a cooperative, to coordinate their actions so as to reach the expected reduction levels and make their participation in the scheme worthwhile. This is similar to *group buying* in e-marketplaces, where some members can obtain items that cost more than they are able to pay for alone, but due to group internal price fluctuations set by corresponding mechanisms, the purchase finally becomes advantageous to all [105, 203]. Inspired by work in that domain, and also work on scoring rules [67, 155, 153], in this thesis we devise an *individually rational, incentive compatible, and budget-balanced* reward sharing mechanism which determines variable reduced electricity prices for coalescing agents via internal money transfers, and incentivizes them to participate in the consumption shifting scheme.

Unfortunately, even if the participating agents are perfectly truthful regarding their abilities and corresponding uncertainty, their reports and estimates can still be highly inaccurate. This can be due to, for example, communication problems, malfunctioning equipment, or prejudiced beliefs and private assumptions—e.g., a truthful reporting agent might be overly pessimistic or optimistic. To tackle this issue, we also propose the application of generic prediction methods, which are nevertheless able to adapt to a specific agent’s behaviour regarding the promised and

final consumption shifting actions.

1.4 Contributions

In this thesis, we examine how methods and techniques from multiple computer science fields can be combined and augmented, to deliver effective DSM schemes that can allow for broad user participation, i.e. thousands, hundreds of thousands, and even more end-users, and nevertheless examine the preferences of each individual. As the reader may have suspected, in order to establish DSM mechanisms that would actually work in realistic large-scale environments, the designer should examine many aspects, and most probably be called to utilize solutions from various fields of engineering and computer science.

To the best of our knowledge, this is the first time that a complete framework for effective large-scale cooperative demand shifting is provided in the literature. The contributions of our work can be summarized as follows:

1. We coin the term *Multiagent Demand-Side Management*, where large sets of end-users of various types (consumers, prosumers, electricity storage facilities) join forces in *cooperatives* to contribute to the balancing between demand and supply of electricity, by *altering their demand profile* in exchange of rewards, either monetary or social. Though this has been envisaged by prior work, e.g., [96, 146], in this thesis we explicitly propose mechanisms and algorithms to offer a solution that can be applied effectively to large, real-world environments.
2. Inspired by work in the *cooperative games* and related MAS literature [39, 173],

we propose several methods for the formation of shifting coalitions. These *coalition formation* methods group together agents based on criteria such as their perceived shifting contribution potential, and their expected economic gains from participation (Chapter 3).

3. Our mechanisms come with individual rationality, truthfulness, and (weak) budget balancedness guarantees. As such, the design of our mechanism promotes *broad participation opportunities*, guaranteeing that consumers of any category or type have strong economic incentives for participation in the scheme, as verified theoretically and also confirmed by our simulations (Chapters 3, 4, and 5).
4. We also devise *internal pricing* mechanisms that determine *variable*, and *individual specific* reduced electricity prices for our agents, via implementing *expected gain transfers* among the coalescing consumers. The resulting *internal price balancing* incentivizes even agents with initially forbidding shifting costs to participate in the cooperative effort. We put forward several such internal price balancing techniques: a heuristic mechanism; and five alternative ones. All of our proposed internal pricing methods satisfy *budget-balancedness* (Chapter 3).
5. To promote truth-telling and efficiency in load shifting, we employ a *strictly proper scoring rule*, the *continuously ranked probability score (CRPS)*, proposed in the mechanism design literature [67]. The use of *CRPS* incentivizes agents to truthfully and precisely report their predicted shifting capabilities, helping in this way to model and tackle related uncertainty (Chapters 3 and 4).

6. We apply a *novel cryptocurrency model* for the coordination and management of the cooperative shifting actions, and combine it for the first time with a scoring rule, resulting on a large-scale DSM scheme for prosumers that promotes accuracy and engagement (Chapter 4).
7. We propose a V2G/G2V mechanism that can take advantage of fleets of electric vehicles, a load type that has different constraints for the shifting of consumption. Additionally, power that is stored in the electric vehicles batteries can be used as a source of electricity to reduce external electricity imports and better utilize locally generated renewable energy (Chapter 5).
8. The methods that we adopt in this work for *agent behavior prediction*, originate in Machine Learning, are very generic, and have wide areas of application. Their employment in the power consumption shifting domain ensures that member agents can be ranked by the cooperative according to their perceived consumption shifting capabilities; and thus untruthful or inaccurate agent statements regarding their capacity and corresponding uncertainty will not be able to jeopardize the stability and effectiveness of the overall mechanism governing the cooperative business decisions (e.g., which agents to select for consumption shifting at a given point in time). This is key for the economic viability of any such cooperative (Chapter 6).
9. We have developed a *serious game* with *intuitive Graphical User Interfaces (GUIs)* and researched system architectures that are used for engaging the end-users in DSM participation; prototypes have been tested in trials with real people (Chapter 7).

MAS APPLICATION: LARGE-SCALE DEMAND-SIDE MANAGEMENT

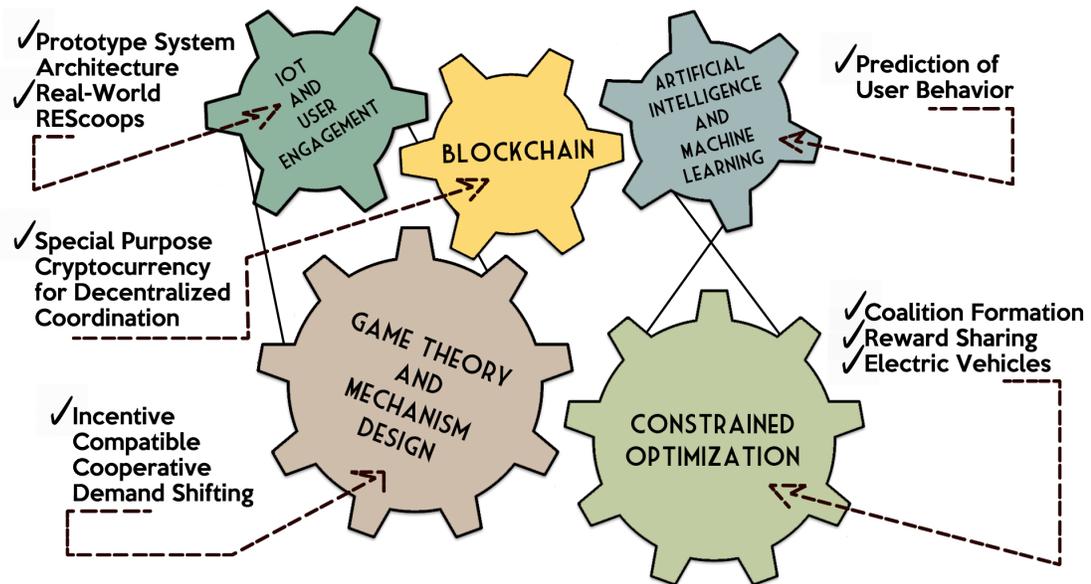


Figure 1-1: Overview of scientific fields, and our contributions.

- Finally, our proposed scheme is easy-to-use and directly applicable, as it requires no legislature changes whatsoever—it only requires the willingness of national authorities (or perhaps even utility companies that adopt DSM as part of their business) to provide better prices for joint, large-scale demand shifting.

Figure 1-1 provides an overview of the scientific fields that inspired this research, and shows the main links of our thesis components and contributions with these fields. The core of our demand shifting mechanism is composed of a constrained optimization process, which applies constraints calculated based on the Grid’s parameters to select the most appropriate participants for the DSM procedure, i.e. to

form coalitions of actors from the large pool of the available agents. Classic coalition formation theory studies [39] the participation incentives and behavior of groups of agents that join forces and act collectively. Here, to form the acting coalitions, further constraints are explicitly examined, that concern the individual preferences of the agents. Our mechanism takes into account agent availability and shifting capability, and offers back tangible benefits for engaged and truthful contributors, as dictated by Game Theory and Mechanism Design practices. These fields offer guidelines for solutions regarding resource allocation, electronic marketplaces, preferences aggregation, and so on, aspects which are found in the domain of DSM in Smart Grids. In our mechanism, after the DSM procedure takes place and utility is generated, this is distributed back to agents in a fair and personalized manner, rewarding accuracy and truthfulness the most with the use of a strictly proper scoring rule [67], thus making the mechanism incentive compatible.

Also, by incorporating a special purpose, blockchain-based cryptocurrency protocol [117], we have extended the proposed DSM mechanism to include electricity prosumers that coordinate without the need of a, potentially third-party, central entity. This fact removes the single point of failure of the cooperative manager, and paves the way to energy democracy and self-production, always guaranteeing the Grid's reliability. To enhance the DSM mechanism even more, we have studied a number of prediction methods from the fields of Artificial Intelligence and Machine Learning, that are able to monitor and learn agent behavior in the DSM procedures, helping this way to better evaluate each agent's capabilities, and manage to select appropriate agent sets for future participation, and guaranteeing effectiveness in the collective balancing procedures. Additionally, we coin and evaluate a constrained

optimization-based algorithm that, once again, forms *coalitions* of EVs that offer V2G/G2V services, and that are rewarded back based on the energy they feed to the Grid and on the driver-set prices for offering this service. Moreover, and in light of the recent technological advancements in the Internet of Things, we have studied system and agent architectures for delivering DSM schemes in the real world, and, in particular, how intuitive graphical user interfaces for aiding the end-user can be developed. Finally, in the scope of a Horizon 2020⁴ project, in this thesis, we have worked closely together with several European real-world energy cooperatives, and aim to expand our collaboration to examine how the proposed DSM scheme can be incorporated into their business plans.

1.5 Thesis Outline

The structure of the rest of this dissertation is as follows. Chapter 2 presents the general placement of our work against the current state-of-the-art. Note that at the beginning of each chapter we also provide links to past related work on corresponding fields.

Chapter 3 introduces a cooperative DSM scheme suitable for small consumers of electricity. The scheme comes along with desirable properties that guarantee effectiveness and strong incentives for user participation.

Chapter 4 builds on the same principles, but extends to electricity *prosumers* settings, where a cryptocurrency protocol is also included for agent coordination and

⁴The Horizon 2020 programme was established by the European Union (EU) in order to serve as the financial instrument that would help secure Europe's global competitiveness in research and innovation (see <https://ec.europa.eu/programmes/horizon2020/>).

reward sharing, without the need of centralized entities.

Chapter 5 provides an algorithm for the incorporation of electricity storage systems, and more specifically Electric Vehicles, where the user and system constraints require specific handling.

Chapter 6 presents agent behavior monitoring and prediction techniques, and shows the increased accuracy and effectiveness of DSM, when we add these in our scheme.

Chapter 7 discusses the new architecture trends of the Internet of Things, and showcases a serious game scenario, where human users interact with intelligent agents inside their residences, to achieve profitable and meaningful demand shifting.

Chapter 8 concludes our work, and draws on future work directions, discussing, apart from the Smart Grid, different application fields of our approach, as well as numerous extensions that can be performed.

Finally, Appendix A includes sections that describe the realistic simulation datasets that were used in our experiments.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 2

Background

The scope of the work presented in this thesis, combines and involves many different areas at various levels of the DSM process. Starting from a higher level, in this chapter we discuss existing notions regarding *cooperatives*, *virtual power plants*, and *aggregators*. Then, we set the theoretical framework, originating from mechanism design, that is used to analyze the properties and guarantees of our approach. We also highlight the *uncertainty* of agent behavior and introduce the difficulties that the designer of DSM schemes would confront, as well as the means to effectively drive agent behavior, that is, *truthful* and *incentive compatible* solutions. This chapter also introduces some essential background on optimization and blockchain-oriented cryptographic applications, since they can be used to enable decentralized DSM. Finally, we present some essential principles for the design of agent-based systems and front-ends for the DSM process, which motivate user engagement, and define some key notions that are present in the Smart Grid.

2.1 Basic Assumptions

To set the basis of our model and introduce some essential notation, consider a set of actors (agents), \mathcal{A} :

$$\mathcal{A} = \{a_1, a_2, \dots, a_N\} \quad (2.1)$$

each a_i of which is characterized by her electricity demand $q_{i,t}^-$ during a time interval t , as well as her electricity production during t , $q_{i,t}^+$, both measured in kWh. For mere electricity consumers, $q_{i,t}^+$ is considered to be zero, while for “prosumers” (i.e. consumers and producers of electricity) it can take (real) positive values. Each measurement $q_{i,t}^-$ is composed of the power consumption of the electrical and electronic appliances that are placed in the premises of agent i ’s building, or infrastructure. The same holds for $q_{i,t}^+$, where part of the amount may originate from photovoltaic (PV) panels, or wind generators. It is possible to disaggregate loads to identify the kind of appliance that is enabled simply by observing the aggregate curve [139]. In our work however, we do not focus on such matters, and assume (without loss of generality) that each agent has full control and perfect knowledge regarding their own electrical and electronic appliances consumption, time intervals of initiation and turning off, etc. Importantly, we acknowledge that electricity consumption is strongly related to each inhabitant’s daily routine and personal preferences, and that specific patterns exist in the aggregate hourly demand curve. This is a natural assumption, that exists also in the real-world, as can be seen also in an example of real hourly demand and production curves of Greece, shown in Figure 2-1. As concerns the RES production curve, this cannot be easily controlled, since it strongly depends on time and weather

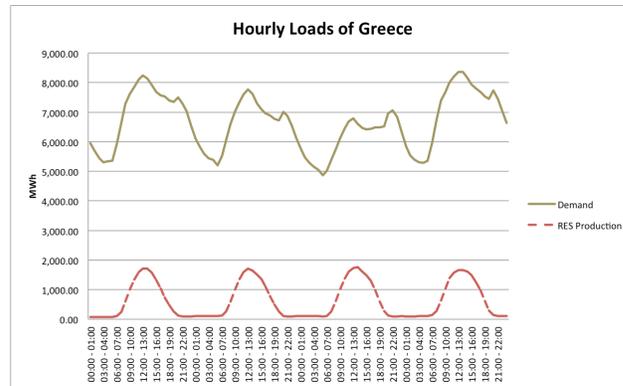


Figure 2-1: Aggregate demand and RES supply curves of Greece (1/7/2016-4/7/2016). [Source: <https://www.entsoe.eu>, <http://www.admie.gr>]

conditions.¹

2.2 Cooperatives, Virtual Power Plants, and Aggregators in the Real-world

The emergence of effective, large-scale DSM is not a thing for the future. User conglomerates, such as cooperatives, VPPs and aggregators, have already been formed in various countries of the world. Such notions are often used interchangeably to describe large sets of consumers, prosumers, or generators, which, although small in size as individuals, coordinate and interact with the electricity grid as a single entity in order to have much larger impact on the market [58, 66, 161].

According to the 2015-2016 report of the American Public Power Association, there are 877 cooperative providers (26.5% of all providers) in total in the U.S., offer-

¹In the case of Greece, most RES generators are photovoltaic panels, thus they produce energy only during daytime.

ing energy to the 12.8% of the total customers. In Europe, the requirement for large-scale, coordinated consumer action, lies at the foundations of the fledging industry of renewable energy sources cooperatives (REScoops, see, e.g., <https://rescoop.eu>). These are cooperatives of electricity consumers, renewable energy producers, or “prosumers” which form in the emerging Smart Grid, and, so far, have exceeded 2500 business entities in numbers. Their activities include producing and trading renewable energy, and providing DSM services. The operations of such real-world cooperatives and related Smart Grid businesses can naturally benefit from incorporating schemes and methods such as ours, in order to select the most appropriate participants for DSM actions; take individual preferences into account; and redistribute profits in a “fair” manner, rewarding truthful and accurate participants more than unhelpful and misbehaving ones [158]. Further elaboration and analysis of the REScoops profiles and activities can be found in [12], in [50], and in [142].

To be more specific, we provide the following definitions that can be found in the literature:

- **Cooperatives** are autonomous associations of people who join voluntarily to meet their common economic, social, and cultural needs and aspirations through jointly owned and democratically controlled businesses. Cooperative businesses carry with them underlying social values and ethical principles. Those principles are: *voluntary* and *open membership*; *democratic member control*, *economic participation* by members; *autonomy and independence*; *education, training* and *information dissemination*; *cooperation among cooperatives*; and *concern for community* [197]. In the Smart Grid, cooperative members are electricity end-users, or even small scale producers. An example is *coop-*

erative virtual power plants [40], a broad term that intuitively represents the aggregated capabilities of a set of distributed energy resources (DER). For example, it can be thought of as a portfolio of DERs, as an independent entity or agent that coordinates DERs pooling their resources together, or as an external aggregator that “hires” DERs to profit from their exploitation.

- **Virtual Power Plants** are aggregations of different types of distributed resources, which may be dispersed in different points of medium voltage distribution networks. A VPP is composed of a number of various technologies with various operating patterns and availability, which can connect to different points of the distribution network. It actually aggregates the capacity of many diverse DERs, it creates a single operating profile from a composite of the parameters characterizing each of the DERs and can incorporate the impact of the network on aggregate DERs output. A VPP is a flexible representation of a portfolio of DERs that can be used to make contracts in the wholesale market and to offer services to the system operator [161]. Virtual power plants rely upon software systems to remotely and automatically dispatch and optimize generation or demand-side or storage resources in a single, secure web-connected system [23].
- **Aggregators** act as intermediates to several small assets, such as home users, in order to allow them to interact with the market and have increased negotiation power. The aggregator has two main goals, first to provide actual services (e.g. DSM) to the system operator, and secondly to guarantee profits (e.g. via reduced electricity bills) to the end users [66].

Note that, although these notions may be built upon different backgrounds, the final result from the scope of the Grid is the same in every case: they can provide substantial control on the traded energy, and help increase the stability and safety of the network. In our work we promote the coordination of the individuals via such schemes. We choose to use the term *cooperative* because it implies commitment and engagement from the participants side, features that are imperative for the success of decentralized DSM operations.

Under the “umbrella” of the cooperative, every $q_{i,t}^-$ and $q_{i,t}^+$ of the individuals of \mathcal{A} , are summed up to the aggregate cooperative demand and supply for each time interval t , measured in kWh:

$$Q_t^- = \sum_{i=1}^N q_{i,t}^-, \forall i \in \mathcal{A} \quad (2.2)$$

$$Q_t^+ = \sum_{i=1}^N q_{i,t}^+, \forall i \in \mathcal{A} \quad (2.3)$$

2.3 Demand-Side Management and Mechanism Design

DSM services can be loosely divided into three categories: consumption reduction programs, load management programs, and energy conservation programs [3, 96, 196]. Now, although reduction and conservation can obviously contribute to the reliability and sustainability of the power system, they require large investments for equipment and infrastructure upgrade [68]. Moreover, these efforts can reach a

maximum effectiveness level, beyond which further reductions cannot be tolerated, because such actions begin to interfere with consumer comfort and well being [168]. In any case, for any DSM scheme to be easily adopted, it should be made sure that consumer needs are accommodated and consumer tasks are eventually completed—even at earlier or later times than originally scheduled.

In mathematical terms, agent i has the capability to reduce its demand $q_{i,t}^-$ during a time interval, by an amount

$$r_{i,t} \leq q_{i,t}^-$$

and, as concerns the cooperative,

$$R_t = \sum_{i=1}^N r_{i,t} \leq Q_t^-, \forall i \in \mathcal{A}$$

Assuming that the agent will have to consume $r_{i,t}$ at an earlier or later time interval, $r_{i,t}$ is referred to as the *shifting capacity* of agent i during t . However, in order for a rational, self-interested agent to reduce and shift its demand, it expects to be rewarded, at least as much as the anticipated *shifting cost*, $c_i^{t \rightarrow t'}$, measured in euros.

From a strict game theoretic scope that follows the descriptions of [138] and [173], each agent has a *strategy*, i.e. *a complete contingent plan, or decision rule, that defines the action an agent will select in every distinguishable state of the world*, $s_i(\theta_i) \in \Sigma_i$, which, in our case, is the choice of an agent to participate in DSM or not (Σ_i denotes the set of all strategies available to agent i). The $\theta_i \in \Theta$ denotes the agents preference type. Here, it is a tuple containing all the necessary values required to describe the

behavior of each agent i , e.g.

$$\theta_i = \langle q_{i,t}^-, q_{i,t}^+, r_{i,t}, c_i^{t \rightarrow t'} \rangle$$

Now, a *self-interested, rational* agent seeks to maximize its *utility*, $u_i(o, \theta_i)$, which numerically corresponds to reward received, as a result of a selected action that leads to an outcome $o \in \mathcal{O}$. Due to the fact that the utility is awarded after each (unknown) agent's action takes place, the speculation regarding the outcome for each agent is termed as *expected utility*. This means that the expected utility can be used to determine the preference of an agent towards selecting a specific strategy, given its type. We say that, if

$$u_i(o_1, \theta_i) > u_i(o_2, \theta_i)$$

then agent i prefers outcome o_1 over o_2 .

Game-theoretic *solution concepts*, correspond to formal “rules” for predicting how a game will be played. To continue, we must define the *strategy profile*, the strategies used by all agents, $s = \{s_1, s_2, \dots, s_N\}$. A subscript of $-i$ denotes the set that includes every agent apart from i . This way, the set of all agent strategies apart from i 's would be $s_{-i} = \{s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_N\}$. Also, s'_i denotes any other strategy profile that agent i could choose to use. We now briefly describe some fundamental GT solution concepts [138].

Definition 2.3.1. [Nash equilibrium] The **Nash equilibrium** describes the strategy profiles s , for which every agent maximizes its expected utility:

$$u_i((s_i(\theta_i), s_{-i}(\theta_{-i})), \theta_i) > u_i((s'_i(\theta_i), s_{-i}(\theta_{-i})), \theta_i), \forall s'_i \neq s_i$$

This means, that if a strategy profile is in Nash Equilibrium, no participant has an incentive to select another strategy, due to the fact the the expected utility decreases when choosing alternatives. Although Nash Equilibrium is a very strong concept, its calculation requires that each agent must have perfect information regarding the types and strategies of all other agents. In practice, this is rarely realizable (plus the calculations cost can be prohibitive) [52], thus more robust solution concepts are used in settings with uncertainty, such as the *Bayes-Nash Equilibrium*.

Definition 2.3.2. [Bayes-Nash Equilibrium] A strategy profile $s = \{s_1(\cdot), s_2(\cdot), \dots, s_N(\cdot)\}$ is in **Bayes-Nash equilibrium** if for every agent i and all preferences $\theta_i \in \Theta_i$:

$$u_i((s_i(\theta_i), s_{-i}(\cdot)), \theta_i) > u_i((s'_i(\theta_i), s_{-i}(\cdot)), \theta_i), \forall s'_i(\cdot) \neq s_i(\cdot)$$

Now, in some settings, one might be lucky enough to discover *Dominant Strategies*:

Definition 2.3.3. [Dominant strategy] **Dominant strategy** is a strategy s_i that weakly maximizes an agent's expected utility for all possible strategies of the other agents:

$$u_i((s_i, s_{-i}), \theta_i) > u_i((s'_i, s_{-i}), \theta_i), \forall s'_i \neq s_i, s_{-i} \in \Sigma_{-i}$$

The Dominant Strategy equilibrium signifies that *regardless of the strategies of every other agent*, selecting s_i will grant the maximum utility. In strategic dominance, the personal choices and preferences of the rest of the agents do not have to be known, and are irrelevant.

In the *cooperative games* domain—i.e., games in which participants cooperate and try to maximize the *collectively achieved utility*, also termed as *social welfare*—a

quite popular solution concept is the *Shapley Value* [39]. In cooperative games agents play by forming *coalitions* and each possible coalition is mapped to a real number with the use of a *characteristic function*, $v : 2^N \rightarrow \mathbb{R}$. Such games belong to the class of *characteristic function games*. The result of the characteristic function can be thought of as the collectively achieved utility of a potential coalition of players. Here, there is an additional requirement that is, to divide the social welfare and reward back the individual participants. This is a quite complex task, since, frequently, each agent contributes in different degrees. The Shapley value aims to reward each agent proportionally to their average *marginal contribution*, that is the collective utility increase actually caused by the agent's participation. Let $c \in \mathcal{C}$ denote a coalition from the set of all possible coalitions between players in \mathcal{A} , and c_{-i} , the same coalition, but without agent i .

Definition 2.3.4. Shapley value The **Shapley Value** in a characteristic function game $G = (\mathcal{A}, v)$ for an agent i is given by:

$$\phi_i(G) = \frac{1}{N!} \sum_{c \in \mathcal{C}} (v(c) - v(c_{-i}))$$

As we can see, the Shapley value requires a vast number of calculations that grows super exponentially as N , the number of participating agents, increases. This is why in this thesis we do not employ the Shapley value to distribute utility to the agents.

To persuade candidates for DSM participation to actually contribute, the outcomes should be such that, despite the shifting costs, the agent finally profits after contributing. At the same time, there should be counter-incentives for agents that

do not “play fair”, that is, for those that select actions such that they jeopardize the total outcome for all agents, to increase their own. The subfield of game theory that sets the rules of a game, as well as the game itself, is called *Mechanism Design*(MD), or “inverse game theory” [133]. MD is defined by [138] as “the sub-field of microeconomics and game theory that considers how to implement good system-wide solutions to problems that involve multiple self-interested agents, each with private information about their preferences.” These solutions are modelled by a single function, the *social choice function*.

Definition 2.3.5. [Social choice function] The social choice function $f : \Theta_1 \times \dots \times \Theta_N \rightarrow \mathcal{O}$ chooses an outcome $f(\theta) \in \mathcal{O}$, given agent types $\theta = (\theta_1, \dots, \theta_N)$.

This function selects the optimal choice for agents that have given types. The task of the mechanism designer is to set the rules of a game, e.g. which are the available actions, how the outcomes are determined, so that the solution of the social function is solved, despite the individual self interests of each agent.

We can now give a formal definition of what a *mechanism* is:

Definition 2.3.6. [Mechanism] A mechanism $\mathcal{M} = (\Sigma_1, \dots, \Sigma_N, g(\cdot))$ defines a set of strategies Σ_i available to each agent, and an outcome rule $g : \Sigma_1 \times \dots \times \Sigma_N \rightarrow \mathcal{O}$ such that $g(s)$ is the outcome implemented by the mechanism for strategy profile $s = (s_1, \dots, s_N)$.

What a mechanism really does, is to define the strategies that are available to players, and, furthermore, the method by which the final outcomes are determined, as a result of agents selecting their strategies. Now, suppose that agents select their strategies such that they are in equilibrium (e.g. Nash, Dominant strategy, etc.), and

compute the outcome. If this outcome is the solution of the social choice function, for all possible agent preferences, then we say that the mechanism \mathcal{M} with outcome function $g(\cdot)$ implements the social choice function $f(\theta)$:

Definition 2.3.7. [Mechanism implementation] A mechanism $\mathcal{M} = (\Sigma_1, \dots, \Sigma_N, g(\cdot))$ implements social choice function $f(\theta)$ if $g(s_1^*(\theta_1), \dots, s_N^*(\theta_N)) = f(\theta)$, for all $(\theta_1, \dots, \theta_N) \in \Theta_1 \times \dots \times \Theta_N$ where strategy profile $s_1^*(\theta_1), \dots, s_N^*(\theta_N)$ is an equilibrium solution to the game induced by \mathcal{M} .

2.4 Incentive Compatibility and Related Properties

Usually, a key property sought after by MD approaches is *incentive compatibility* [133]. This means that the design must be such, that actors finally choose willingly to follow the social choice function. Additionally, this guarantees that they are better-off being *truthful* regarding the private preferences they reveal to the scheme, while untruthful ones suffer penalties, exclusions, and are in general subject to negative incentives. These properties are very important in mechanism design in general, but in the DSM domain as well: since agents need to coordinate in the present, and commit for collective consumption shifting in the future, coordination must be based on accurate and truthful statements of the participants. Otherwise, the success of collective DSM actions is jeopardized because there will be no guarantees for the satisfaction of the necessary constraints for effective and profitable DSM.

Incentive compatible mechanisms are mainly built upon the *direct revelation* prin-

principle. This principle dictates that the only action available to agents by a mechanism, is to state its *preferences* [138].

Definition 2.4.1. [Direct-revelation mechanism] A direct revelation mechanism $\mathcal{M} = (\Theta_1, \dots, \Theta_N, g(\cdot))$ restricts the strategy set $\Sigma_i = \Theta_i$ for all i , and has outcome rule $g : \Theta_1 \times \dots \times \Theta_N \rightarrow \mathcal{O}$, which selects an outcome $g(\hat{\theta})$ based on reported preferences $\hat{\theta} = (\hat{\theta}_1, \dots, \hat{\theta}_N)$.

The next step is to guarantee that the preferences reported by the agents to the mechanism are their true ones. This is performed by first asking for their *true preferences*, and, secondly, guaranteeing that the agents will finally *choose* to report them [138].

Definition 2.4.2. [Truth-revelation] A strategy $s_i \in \Sigma_i$ is truth-revealing if $s_i(\theta_i) = \theta_i$ for all $\theta_i \in \Theta_i$.

Definition 2.4.3. [Strategy-proof] A direct revelation mechanism is strategy-proof if truth revelation is a dominant strategy equilibrium.

Finally, we need an *incentive compatible* mechanism consisting of a social choice function that is solved for truth-revealing agent strategies [138].

Definition 2.4.4. [Incentive compatible implementation] An incentive compatible direct revelation mechanism \mathcal{M} implements a social choice function $f(\theta) = g(\theta)$, where $g(\theta)$ is the outcome rule of the mechanism.

A complementary definition to the above is the *quasi-linear agent preference* [138].

Definition 2.4.5. [Quasi-linear preference] A quasi-linear utility function for agent i with type θ_i is of the form:

$$u_i(o, \theta_i) = v_i(x, \theta_i) \Leftrightarrow p_i$$

where outcome o defines a choice x from a discrete choice set and a payment p_i by the agent.

The quasi-linear preferences imply that the utility of each agent can be translated to a value function i.e., a price that is paid (or claimed), making this way the utility transferable among agents.

Another desired property of a mechanism, or, more accurately, of the social choice function implemented by a mechanism, is *budget balance*. Budget-balanced mechanisms guarantee that the utility that is to be redistributed among participants is generated by their own participation, and no external funding is required [138].

Definition 2.4.6. [Budget-balance] A social choice function $f(\theta) = (x(\theta), p(\theta))$ is budget balanced if for all preferences $\theta = (\theta_1, \dots, \theta_N)$

$$\sum_i^N p_i(\theta) > 0$$

A more relaxed case is that of *weak budget balance*, where, although again the mechanism will not require external funding, some profit might be generated that is not claimed by the participants:

Definition 2.4.7. [Weak budget-balance] A social choice function $f(\theta) = (x(\theta), p(\theta))$

is weakly budget balanced if for all preferences $\theta = (\theta_1, \dots, \theta_N)$

$$\sum_i^N p_i(\theta) \geq 0$$

A third property that is a pre-requisite to *any* solution mechanism is *individual rationality*: since individuals are assumed to participate in the mechanism only if it is rational for them to do so, the mechanism ensures that participation incentives do in fact exist [138].

Definition 2.4.8. [Individual rationality] A mechanism \mathcal{M} is individually rational if for all preferences θ_i it implements a social choice function $f(\theta)$ with

$$u_i(f(\theta, \theta_{-i})) \geq \bar{u}_i(\theta_i)$$

where $u_i(f(\theta, \theta_{-i}))$ is the expected utility for agent i at the outcome, given distributional information about the preferences θ_{-i} of other agents, and $\bar{u}_i(\theta_i)$ is the expected utility for non-participation.

This means that when a mechanism is individually rational, then participating agents achieve *at least* as much expected utility as from not participating, given their prior beliefs regarding other agents' preferences.

It should be clear to the reader by now, that in the multi-agent DSM mechanism that we propose, agents should truthfully report their shifting preference types, and maximize their utilities/payoff functions by doing so. This can be achieved by combining *a priori* statements—these can be thought of as agents' *commitments* regarding their participation—with the *actual actions* that take place after DSM ac-

tualization. This is achieved with the incorporation of *scoring rules*, which we will be discussing next.

2.5 Scoring rules

Scoring rules [67] are functions that can be used to mathematically evaluate the quality of probabilistic forecasts. Such functions assign a numerical value—termed as score—given a *predictive distribution*, and the event that finally materializes. The scoring rule function is referred to as

$$S(P, x) : f_X, \mathbb{R} \rightarrow \mathbb{R}$$

where P is the probability density function (PDF) of the forecast that belongs to the set of the probability density functions. The sample x can be generated from another, or the same PDF, noted as Q . In general, scoring rule functions have a range in \mathbb{R} , but the extended $\bar{\mathbb{R}} = [-\infty, \infty]$ can also be used.

Now, we say that a scoring rule is *proper* if the following hold [67]:

$$S(P, P) \geq S(P, Q), \forall P, Q$$

This means that if the materialized value originates from the stated predictive distribution, then the score value is at least as high as if x originated by any other PDF Q . Although this concept is appropriate for some fields, in our case we would like the score to generate even *higher* values when x is generated by the predictive PDF P , than from any other Q . As such, the participant will be guaranteed that

the maximum profit is earned by being truthful. This is the case for *strictly proper* scoring rules:

$$S(P, P) > S(P, Q), \forall P, Q, P \neq Q$$

Here, the equality must hold if and only if $P = Q$, which means that the forecasters are better off stating *their true beliefs* accurately, since this is the only value that gives the best score. As such, a mechanism can exploit a rule's strict propriety property, to ensure that scores that are lower than the best, result to lower returns (e.g., via the use of penalties) to the participants—and, by so doing, ensure the *truthfulness* [133] of the agent reports. Scoring rules are already being applied in various scientific fields, such as probabilistic weather forecasts that use ensembles, economic and financial forecasting methods, wind turbine generation forecasts, etc.

For example, in the work of [35] Boutilier discusses proper scoring rules that incentivize self-interested forecasting experts to report predictions truthfully to decision makers. The main idea is that the agent might have increased inherent utility by misreporting the forecast to the aggregator. Thus, the author proposes the use of a *compensation rule*, that actually rewards the agent the loss in her utility caused by truthful reports. In this way, the agent has no incentive to report untruthfully. In our research, we are concerned about truthful prediction statements, and thus incorporate such scoring rules, in particular the Continuously Ranked Probability score. However, it is also desirable to study prediction methods' accuracy and efficiency (Chapter 6).

2.5.1 The Continuously Ranked Probability Score

In the mechanism proposed in this thesis, the incentivization for truthful reporting is made possible with the incorporation of the strictly proper scoring rule, the *Continuously Ranked Probability Score (CRPS)* [67], that actually performs an opposite reward assignment; as the participant's actions deviate from the committed ones, fines apply that punish such deviations. Thus, it is of the participants' best interests to be truthful.

CRPS has also been recently used in [155] to incentivize renewable energy-dependent electricity *producers* to accurately state their estimated output when participating in a cooperative. The evaluation of producers', but also consumers' forecasts, is performed using CRPS in other works as well, e.g. [186, 19, 65, 33].

The use of *CRPS* allows us to directly evaluate probabilistic forecasts, and the score is given by:

$$CRPS(\mathcal{N}(\mu, \sigma^2), x) = \sigma \left[\frac{1}{\sqrt{\pi}} - 2\phi\left(\frac{x - \mu}{\sigma}\right) - \frac{x - \mu}{\sigma} \left(2\Phi\left(\frac{x - \mu}{\sigma}\right) - 1 \right) \right] \quad (2.4)$$

In our setting, $\mathcal{N}(\mu, \sigma^2)$ is the uncertainty stated over the expected *absolute relative errors* regarding the reduction capacity, as reported by an agent; while x is the actually observed error, ϕ the PDF, and Φ the CDF of a standard Gaussian variable. The mean μ and variance σ^2 of the predictive distribution can be estimated by each agent through private knowledge of its consumption requirements and business needs.

2.6 Uncertainty in Demand-Side Management

The incorporation of CRPS and proper scoring rules in general, implies that the mechanism also takes into account *uncertainty factors* that might be present in a setting. As we have already discussed, large-scale consumption shifting is in general accompanied with high uncertainty, since human factors affect the operations. In our work, we not only model uncertainty in agent behavior, but we also take it into account into the mechanism; and provide functional tools to tackle it.

The work of [74] describes the highly uncertain nature of such environments. Some key points regarding the residential behaviour characteristics of DR participants are the following:

- Houses of similar sizes, demographically similar occupants, similar set of appliances, and under the same geographical condition, vary in consumption by as much as 200%.
- Houses where occupants have moved, new occupants consumption cannot be estimated.
- Individuals do not make consistent rational decisions, according to economists.
- People tend to overestimate the amount of energy they use, and are optimistic with commitments in energy efficiency applications.
- The effects of *Price Unresponsiveness*.
- Users might not be able to respond to price signals, e.g., due to misunderstandings or the lack of immediate feedback.

- Equity issues that is, differences between the real value of assets and the costs of liability.

A potential solution to the equity issues is providing *personalized* suggestions, based on the individual capabilities and profile of each participant, where they receive different suggestions for consumption shifting and pricing according to their own needs and interests. This is a main feature of the mechanism, as described in Chapter 3.

2.7 Constrained Optimization for DSM

We term DSM approaches that aim to *optimize consumption schedules* given specific user preferences *optimization* or *decision-theoretic (DT)* approaches, since they largely use techniques originating in optimization and decision theory. Typically, these methods aggregate individual *preferences*, turn them into *constraints*, and solve the resulting joint optimization problem *centrally*, to produce altered demand profiles for the individuals. Their usual objective is to *minimize a cost* (or to *maximize a payoff*) *function*. Costs (payoffs) can be monetary or associated with participant discomfort minimization, maximizing the use of renewable or locally produced energy, and similar concerns.

DT methods can be deterministic or stochastic, and focus on either real-time or longer-term planning [29]. Regardless, an assumption they commonly make is that stated individual preferences are exactly as provided, and not subject to change. Moreover, unlike their GT counterparts, these methods do not tackle conflicts arising among users due to optimization (e.g., due to unfair sharing of DSM profits), nor do

they explicitly incentivize users to participate or cooperate in the DSM process [29].

Now, as is well-known, optimization problems are formulated as a function f that is to be minimized (or maximized), subject to a set of constraints, that can include both equality, and inequality constraints. In mathematical terms this is expressed as:

$$\begin{aligned} & \text{minimize } f_0(x) \\ & \text{subject to } f_i(x) \leq 0, i = 1, \dots, m \\ & \quad \quad \quad h_i(x) = 0, i = 1, \dots, l \end{aligned} \tag{2.5}$$

The above statement describes the problem of finding an x that minimizes the function $f_0(x)$, by choosing from the x 's that also satisfy the conditions $f_i(x) \leq 0$ and $h_i(x) = 0$. The x is called *optimization variable*, and the function $f_0 : \mathbb{R}^n \rightarrow \mathbb{R}$ the *objective function*. The inequalities $f_i(x) \leq 0$ are termed as *inequality constraint functions*, while $h_i(x) = 0$ are the *equality constraint functions* [36]. For the solution of *convex problems*, i.e. for an x to be the optimal x^* given convex $f_0(x)$ and affine inequality constraints, a set of conditions must hold, the Karush-Kuhn-Tucker (KKT) conditions. Lastly, unconstrained optimization problems can be formulated as well, by setting m , and l to zero.

In DSM, $f_0(x)$ represents the costs of participating parties, and the constraint functions model the preferences and limitations that exist. The main problem with the usual approaches however, is that the problem is solved in a centralized manner, by a party that needs to have perfect knowledge regarding agent personal preferences—and the Grid's as well, so that to be able to formulate the problem.

Even worse, after the global solution is found, then every agent is obligated to follow it and alter their consumption profiles, thus leading to inflexible solutions that are rarely successful in practice for large-scale applications. In the core of the mechanism that we propose we also solve constrained optimization, but we combine it with incentivization schemes (by using GT techniques), and also incorporate a cryptocurrency protocol to overcome centralizations.

2.8 Cryptocurrencies and Blockchain Technology

A field that is of interest to us due to its potential for decentralized problem solving, is the *blockchain-oriented technology* with the most famous application being that of digital currencies, or *cryptocurrencies*. To help explain the way blockchain and cryptocurrencies work, we will use the example of *Bitcoin*. Introduced by the work of [131] in 2008, Bitcoin is the first, and most widespread cryptocurrency. Bitcoin consists of a number of interconnected peers that execute a special purpose, open-source, and decentralized software. Peers exchange messages that contain information regarding the bitcoin transactions, i.e., who sent electronic coins to whom. They also aggregate such messages, and simultaneously try to solve a “computationally hard” cryptographic problem, termed as *proof-of-work*. The first peer that manages to solve this problem combines information regarding (a) a proof of the solution, (b) the so far received transactions, and (c) additional data required by the ciphering algorithms. This, along with its hash, is propagated to the rest of the network as a message, which constitutes the next block of the blockchain. The rest of the network that receives this block, can validate its legitimacy by performing a set of crypto-

graphic checks. If the block is indeed valid, peers add it into the blockchain, and, as a whole, listen to new transaction messages, and try to solve another instance of the cryptographic problem, this time to generate the next block of the blockchain. The process of creating blocks is called *mining*. Mining also guarantees that the peer who managed to create the new block will be rewarded, either by newly created coins, or by the transaction fees that are set by the “transactees”, or by a combination of both; it depends on the state of the network, i.e. if the maximum amount of bitcoins has already been mined, etc.

Now, if a malicious peer tries to modify past transactions, e.g. in order to double-spend an amount that has already been spent in a past transaction, it will have to change part of the data included into previous blocks. However, such an attempt would fail, since each block includes the hash of the previous block in its header, as shown in Figure 2-2. Locally changing a past block, and propagating it to the peer network would cause inconsistencies in the hashes, thus the other peers would reject such changes through the validation process. Another type of attack, would be to try and alter data in the current block, thus no changes in history should be required. In such cases, the blockchain is “forked” and is split into two different paths, one with the most recent block holding the true state of the blockchain, and another one that has the most recent block altered. Here, the path that finally prevails and is accepted by the peer network is the longest one. This way, highjacking the blockchain and the peer network might be *possible*, however, to achieve this, one should be able to control a majority of peers, which should also have much larger computational power than the rest of the network. With the current technology, and due to the presence of large “fair-playing” mining pools controlled by different stakeholders and

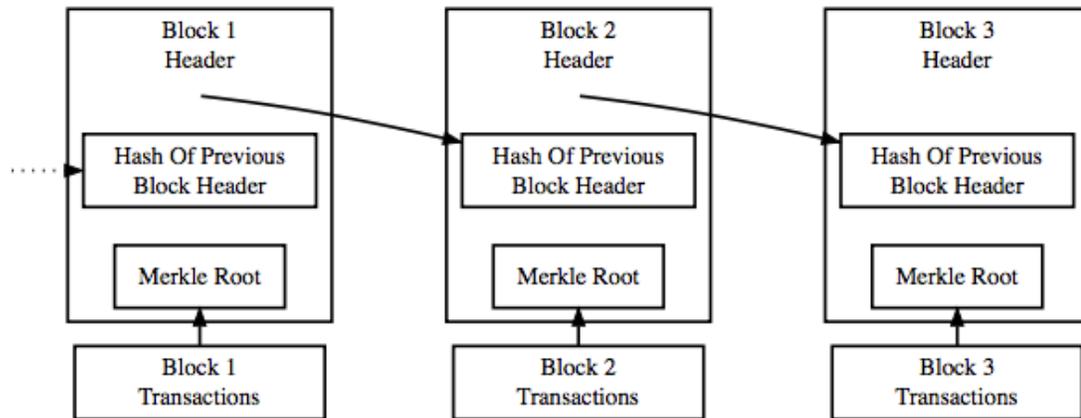


Figure 2-2: Simplified blockchain of Bitcoin. [Source: <https://bitcoin.org/en/developer-guide>]

placed in various physical locations on the globe, this is so far considered as rather impossible [6]. As a minor note, transactions and related data are compressed into Merkle trees [118], for communication and computational efficiency.

So far, as the reader might have already noticed, there is no reference to a *trusted third party* that is required to make the transactions realize. This is why Bitcoin has become so popular by now; as a means of currency, bitcoins are not regulated by governments or central banks, and thus, less fees are required to be payed in general. More importantly, no party needs to perfectly trust anyone in order for transactions to take place. *Trust*, in particular, is a fundamental issue to consider when designing large-scale distributed systems that are open to participants [145], such as a DSM scheme. Especially in cases where cooperative schemes are used, trusted participants guarantee the correct and smooth operation of the system, thus their existence is very important.

Apart from Bitcoin, more than 800 alternative cryptocurrencies are currently fully functional and available to individuals via exchanges in the real world. A great number of them are pure clones of the Bitcoin technology, and are introduced to serve other communities, however there exist examples that differentiate from the original, e.g. by adopting alternative proofs, including anonymity and privacy features, etc.

Furthermore, cryptocurrency protocols are also enabling *smart contracts* [49]. Smart contracts are automated procedures, which can be also considered as autonomous agents. Their goal is to execute a contract, when the required actions have been performed. In our work, as we explain later in Chapter 4, we utilize a cryptocurrency protocol that is able to execute smart contracts for DSM purposes. Here, the inputs are smart meter measurements that indicate effective participation in electricity demand shifting efforts, offered by prosumer cooperatives. The scheme is mainly aiming to evaluate how effective and “helpful” each individual prosumer is in flattening the difference between renewable supply and demand. Additionally, the exchangeability of the proposed cryptocurrency for fiat currencies i.e., real money, is also possible constituting it this way as even more engaging for participants. Additionally, by incorporating such a protocol we offer a decentralized solution that has no need for central cooperative managers, and removes single points of failure from the mechanism.

More interestingly, only recently, innovative business cases that utilize blockchain-oriented algorithms have begun to appear.² Specific blockchain-based networks are proposed, which give to small prosumers—i.e. residences, or small companies, the

²<https://hbr.org/2017/03/how-utilities-are-using-blockchain-to-modernize-the-grid>,
<https://www.theguardian.com/sustainable-business/2017/jul/13/could-a-blockchain-based-electricity-network-change-the-energy-market>,
<https://www.euractiv.com/section/electricity/news/is-blockchain-about-to-disrupt-the-electricity-sector/>

ability to trade energy on their own, change energy suppliers, etc. More specifically in the work of [121], the concept of an “NRGcoin” was incorporated to balance demand and renewable production in a Smart Grid setting that included prosumers. There, given the levels of demand during a time interval, the producer was awarded more NRGcoins for importing energy to the Grid at intervals when renewable supply was weak, and fewer at intervals when renewable supply was in excess. This way, prosumers were incentivized to produce when renewable energy was really needed. However, given the intermittent and uncontrollable nature of the most common renewable generators (i.e. wind and solar), such approaches must be enhanced with carefully designed coordination mechanisms in order to effectively promote DSM. The scheme that we describe in Chapter 4 is the first approach that combines cryptocurrency technology with mechanism design for large-scale decentralized DSM, allowing also for prosumers with uncontrollable renewable generators to benefit from such actions.

2.9 Designing the front end of DSM applications

Since electricity usage and participation in DSM highly relies on human activities and interventions, one critical step for delivering successful and effective DSM schemes is to design appropriate *graphical user interfaces (GUI)*, which could be considered as the front-end of the Smart Grid’s mechanisms and applications [147]. Through the visual presentation of data, also termed as *data visualization*, human end-users are informed regarding the opportunities for DSM participation, and can easily weigh the potential impacts. This helps to plan their actions, recognise underlying threats,

and maximize the gain of each user according to her personalized needs [20].

In the DSM field, there are various data sources that influence both the user decisions, as well as the system efficiency. Such data sources are electricity prices, either retail or wholesale, periodical consumption measurements and forecasts, renewable production levels, monetary gains (cumulative and per-device/participation). Additionally, stored historical data from the aforementioned data sources can play a very important role in a user's final decisions, and help provide an overview of her overall DSM efficiency [193].

In order for data visualization approaches to become effective, three principles must be taken into account, namely *accuracy*, *utility*, and *efficiency* [208]:

1. *Accuracy* states that the actual data items that are interpreted to visual items must have matching attributes. This means that both the placement and the color of a visual item must adhere to the relationship that the original data items had—e.g. use same color for the same group, specify respective positions according to the relationship the original data items have, etc.
2. *Utility* implies that by using visualizations, users should be directed on how to implement certain tasks and achieve their goals. This includes the clear presentation of the current tasks, as well as the subsequent ones that should follow, in order to achieve a certain goal.
3. *Efficiency* defines the ease of comprehension and usage. Data visualization should make clear statements regarding the current state of situation, thus requiring minimum effort by the user to understand and take action. Here, an important factor is also a task's time of completion, as well as the learning

curve for the user to be able to use the data visualization system.

Moreover, the work of [95] discusses an additional categorization of visualizations, the *pragmatic* and the *artistic*.

Pragmatic visualization includes the accuracy principle and is defined as the technical application of visualization techniques to analyze data. From mere numerical data presentation to more sophisticated graphs and scatter plots, pragmatic visualizations are used to directly communicate scientific results in an objective way.

Artistic visualization, on the other hand, is usually used to communicate intuitions and drive inspiration, using the data as the basis for the presentation. Here, efficiency is not a concern, since raw data is not the centre of attention, however it is used to prove that the presented concern is actually real.

In systems that are highly user-centric, it is important to combine multiple views and deliver charts and graphs that update at real-time when a user changes parameters [165]. This technique is also known as *coordinated views*. Coordinated views are crucial for *what-if* scenarios in GUIs for DSM. By utilizing what-if scenarios, the GUI can be used to “evaluate fine-grained effects of behavior changes, such as rescheduling and appliance’s running time” [20].

It should be clear by now that, to achieve the design of useful and effective GUIs in a DSM setting, combining all the above techniques is crucial. These visualization techniques help the user to immediately understand how her preferences or changes in decisions affect multiple aspects of the system. In Chapter 7, we present a prototype GUI design for DSM, which incorporates multiple visualization techniques, and can be utilized by real-world applications. The proposed design was evaluated by human users in a *serious game* setting that simulated day-ahead planning of residential users

in DSM, both for electrical appliances consumption shifting and heating/cooling equipment operation optimization.

Serious games [120] can be defined as any form of interactive computer-based game software for one or multiple players that has been developed with the intention to offer more than entertainment [152]. For example, serious games have been developed for educational and e-learning purposes, as well as for health and safety training [113]. Such games promote the engagement of the players via various means, such as by configuring the difficulty according to each player's capabilities, via virtual rewards, the tracking of the player's skill improvement, etc. [32].

2.10 The Grid and the Electricity Markets

In this section we describe some key notions and authorities that are present in contemporary electricity markets, e.g. in Europe and in North America, relating to the establishment of the Smart Grid.

To begin, typically there exist three different markets, the *wholesale market*, the *retail* (customer market), and the *balancing market*. The wholesale market, also termed as the “day-ahead market” allows generating and utility companies to buy and sell energy for future delivery [90]. Such markets are regional or country-wide and aim to provide liquidity to the energy market.

The customer market is placed one level below, and includes the utility companies that sell electricity to the end users. In customer markets, prices for buying (from local producers) or selling electricity (to the consumers) are set according to contracts and the portfolio of each market actor. Such prices are often set by the various power

authorities, or, in many cases, by a nationwide independent system operator (ISO), managing the electricity grid [93]. ISO authorities are also the *Transmission System Operator (TSO)* and the *Distribution System Operator (DSO)*. TSOs are responsible for the operation and maintenance of the electricity transmission network in large-scale markets, while DSOs for the distribution of the energy from generators to the end-users in local markets. In our work here, as is also common in the literature, we term such authorities as “the Grid” for convenience.

Next, the balancing market, is placed at the top level of the Grid hierarchy, and is provided by the ISO and TSO authorities. The balancing market is responsible for the balancing between supply and demand in real-time [90]. In general, in this kind of market, prices are much worse for the participants, giving them this way incentives to match production and demand beforehand, to the degree possible.

The methods proposed in this dissertation, although requiring participants in large numbers being this way closer to the customer market, can be incorporated into the two other markets as well; this can be decided and mutually agreed by the cooperative members.

Finally, we describe the pricing schemes that are commonly used in electricity markets around the world. We note that in recent years economists have been advocating the use of *dynamic pricing* schemes in the electricity market as a means to avoid market inefficiencies and the aforementioned shortcomings of existing demand reduction schemes [38].

- **Flat pricing:** This is the most simple and common pricing scheme, where a fixed price is assigned to every time interval for consuming or producing electricity.

- **Time-of-Use (TOU):** To apply TOU rates, one should first divide each day in consecutive time slots and then apply different (although fixed) rates for producing or consuming during each time interval.
- **Real-time Pricing (RTP):** RTP can be defined as energy prices that are set for a specific time period on an advance or forward basis and which may change according to price changes in the market. Prices paid for energy consumed during these periods are typically established and known to consumers a day ahead or an hour ahead in advance of such consumption, allowing them to vary their demand and usage in response to such prices and manage their energy costs by shifting usage to a lower cost period, or reducing consumption overall [125]. However, RTP has been strongly criticized for promoting the complete liberalization of household energy pricing. In addition, due to increased levels of consumer uncertainty regarding imminent price fluctuations, it may also require user manual response or the continuous monitoring of smart meters, leading to difficulties in application. Moreover, recent work shows that RTP mechanisms do not necessarily lead to peak-to-average ratio reduction, because large portions of load may be shifted from a typical peak hour to a typical non-peak hour [126]. By contrast, a scheme proposed in this dissertation is a TOU one, and explicitly takes into account the Grid's perspective on which time intervals are preferable for shifting consumption to, and imposes the necessary constraints to avoid—to the extent possible—the event of new peaks arising.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 3

Collective Electricity Consumption Shifting by Consumer Cooperatives

In this chapter, we take the provision of electricity load management solutions one step forward, by proposing the formation of *agent cooperatives for demand shifting services*. In particular, this chapter demonstrates how to perform *large-scale, collective electricity consumption shifting*. To the best of our knowledge, ours is the first to date complete mathematical framework for collective electricity consumption shifting, which comes with several desirable properties and guarantees, and which is also evaluated extensively via simulations on a real-world consumption dataset. A mechanism is designed to promote truthful DSM participation, which also takes uncertainty into account, and forms effective demand shifting coalitions. The shifting of these coalitions is such, that reliability of the Grid is ascertained, and agents do not loose from participation. Simulation results on realistic data imply that the mechanism is “win-win”, and results to flatter aggregate daily demand curves. Our

scheme motivates self-interested business units, represented by autonomous agents to join forces in a cooperative and shift power consumption from peak intervals to others with lower demand, in order to receive lower electricity price rates for their contribution. This is similar to economy of scale approaches, where groups of buyers join together to finally buy goods in larger quantities at a better price each [143, 203]. Consumption re-scheduling can be performed a day-ahead, thus avoiding the dangers and risks of last-minute action. In this way, our scheme *proactively balances future supply and demand*, without any losses affected to the utility of the contributors.

Now, for the cooperative to be successful at *large-scale* shifting efforts, it is obvious that *coordinated joint shifting efforts* have to take place, carried out by *demand shifting coalitions*. Inspired by work in the *cooperative games* and related MAS literature [39], we propose several methods for the formation of shifting coalitions. These *coalition formation* methods group together agents based on criteria such as their perceived shifting contribution potential, and their expected economic gains from participation. We also devise *internal pricing* mechanisms that determine *variable*, and *individual-specific* reduced electricity prices for our agents, via implementing *expected gain transfers* among the coalescing consumers. The resulting *internal price balancing* incentivizes even agents with initially forbidding shifting costs to participate in the cooperative effort.¹ We put forward several such internal price balancing techniques: a heuristic mechanism and five alternative ones. All of our proposed internal pricing methods satisfy budget-balance.

¹This is similar to group purchasing in e-marketplaces, where agents collectively get better prices for their purchases; and where, due to group-internal price fluctuations set by corresponding mechanisms, the purchase finally becomes advantageous to all—even though some members would not be able to obtain the items even at the better rate promised [105, 203].

Thus, our DSM mechanism employs coalition formation and internal pricing techniques in order to facilitate the shifting of sizeable electricity consumption amounts from peak to non-peak time intervals. The effectiveness of the joint coalitional shifting actions naturally depends on the accuracy of the members' statements regarding their shifting capabilities, and their confidence about meeting their forecasted goals. These statements, however, might not be accurate forecasts, as consumers might not be able to accurately predict their shifting capabilities, or are not truthful (due to low trust towards their partners, or similar concerns). Therefore, to promote truth-telling and efficiency in load shifting, we employ a *strictly proper scoring rule*, the *continuously ranked probability score (CRPS)*, proposed in the mechanism design literature [67]. The use of *CRPS incentivizes* agents to truthfully and precisely report their predicted shifting capabilities.

An extensive experimental evaluation of the proposed mechanisms and methods was performed on a large dataset containing real consumption patterns from the Kissamos district at western Crete, Greece. Our experiments confirm that granting a low enough price to the consumers incentivizes cooperative consumption shifting—as long as this price allows the agents to (collectively) overcome their shifting costs. Moreover, our results indicate that the coalition formation method that yields the best results with respect to load shifting effectiveness and actual monetary gains for the participants, is one that creates agent coalitions that maximize in expectation both the eligible amount for shifting, and the expected gains of their members. We also assess the behaviour of the various internal pricing methods we propose, and conclude that our heuristic price balancing approach, in particular, is the most appropriate for use in this domain. Further, we demonstrate that employing *CRPS* is

effective: inaccurate agents suffer penalties that are higher than those of their more accurate counterparts, and individual agents and coalitions alike are incentivized to be truthful and accurate regarding their stated shifting capacities.

Summing up, our work here provides several contributions, also depicted schematically in Figure 3-1:

- This is the first time that a complete framework for effective large-scale demand shifting is provided in the literature.
- We propose a variety of novel coalition formation methods that can be used by consumers wishing to join forces for cooperative demand shifting, along with several novel internal pricing techniques, the application of which incentivizes even agents with low shifting capacities to offer their services in the scheme.
- Our mechanisms come with individual rationality, truthfulness, and (weak) budget balancedness guarantees. As such, the design of our mechanism promotes *broad participation opportunities*, guaranteeing that consumers of any category or type, have strong economic incentives for participation in the scheme, as verified theoretically and also proven by our simulations.
- DSM participation is planned in a proactive manner. The proposed scheme aims to shift power consumption to non-peak intervals and flatten the consumption curve on a day-ahead basis. This allows more careful planing and removes the difficulties of last-minute call for action, which is a serious issue in real world DSM applications.
- Finally, our proposed scheme is easy-to-use and directly applicable, as it re-

SCOPE OF CHAPTER 3

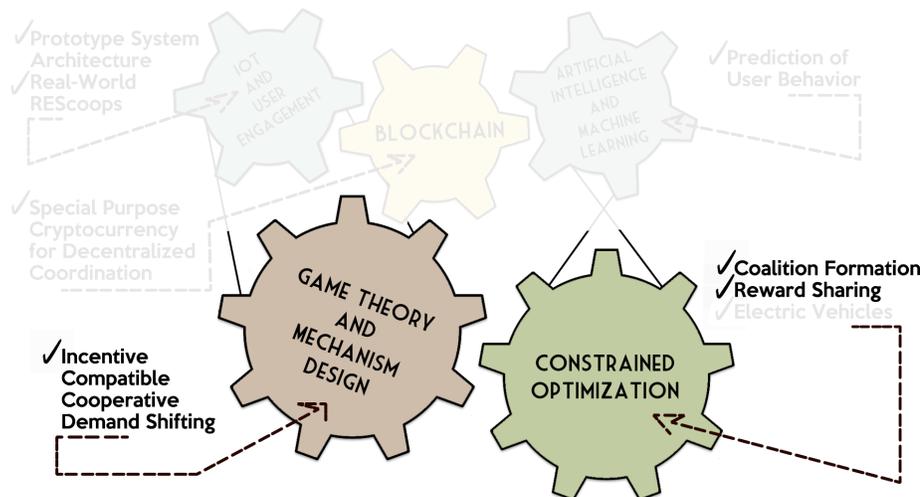


Figure 3-1: Overview of the scientific fields, and our contributions in Chapter 3.

quires no legislature changes whatsoever—it only requires the willingness of national authorities (or perhaps even utility companies that adopt DSM as part of their business) to provide better prices for joint, large-scale demand shifting.

The remainder of this chapter is structured as follows. Section 3.1 reviews related work. In Section 3.2 we give a detailed description of our problem and setting. Section 3.3 then discusses the details of methods for forming effective shifting coalitions. Section 3.4 describes our internal price balancing methods, and Section 3.5 discusses the scheme’s properties and complexity. Section 3.6 describes our experiments and corresponding results; and, finally, Section 3.7 concludes this chapter.

Parts of the research described in this chapter appeared originally in [7] and in [10].

3.1 Related Work

Here we review related work and outline its connections to our approach in this chapter of the dissertation. To begin, the beneficial nature of cooperative *producer* VPPs is demonstrated in [40], where the benefits arising from distributed energy resources coalescing to profitably sell energy to the Smart Grid. Our approach also advocates the creation of cooperative VPPs, but on the *consumer* side—and details a complete framework for the cooperative shifting of consumption loads, to achieve the proactive and effective large-scale electricity demand curve balancing.

Regarding existing approaches that have explicitly attempted to provide solutions to the problem of aggregating distributed energy resources, Chalkiadakis and co-authors have in several research papers in recent years proposed mechanisms for the effective coordination of decentralized energy resources (DER), equipping these with the tools and incentives to participate in DER cooperatives (which could also be viewed as VPPs or aggregators). Such cooperatives could either be cooperatives of energy producers (as in [40, 155]), energy consumers (as in [7, 96, 10]), or energy prosumers (as in [9]). These approaches are, in broad terms, viewing the coordination/aggregation problem as a cooperative game played among the individual agents, and mechanism design tools are often used to provide incentives to the players to report their true capabilities. These are also aided by machine learning and statistical tools to enable (a) decision-making under uncertainty [22, 21], and (b) the detection of errors and the correction of individual consumption/production estimates to the degree possible [10, 136].

Then, in [135] the authors suggest a MAS-based aggregator architecture to man-

age different kind of assets playing a role in a hypothetical smart grid environment. The aggregator, or curtailment service power, undertakes the responsibilities to manage DR events, identify curtailable load at customer levels, enrol new customers, and calculate payments/penalties. Players behavioural strategies are stochastically assigned. Six typical consumers profiles are simulated (two residential, two industrial, and two commercial).

The work of [156], presents a demand-response exchange platform where aggregators can buy and sell flexibility. The architecture focuses exclusively on residential consumers (PHEV, appliances, air conditioners, water heater). The objective is to reduce the overall aggregator load under a threshold for a given duration. One day before an event the ISO call for a bidding session at the DR exchange then the aggregators inform their consumers. The architecture is simulated in Jade. There are 5,555 residential consumers simulated with a 15 minutes time stamp, and no utility price scheme is assumed.

Large-scale and agent-based DSM has been presented in the work of [200]. Authors provide an informationally decentralized framework for managing micro-storage facilities. They perform a game-theoretic analysis and discover optimal charging strategies, as well as optimal levels of micro-storage infrastructure adoption for participants. On the other hand, [97] focuses on predicting the demand-response potential of prosumers. A centralized optimization algorithm is used as a planning mechanism. An incentives-based contribution system is theorized, for which at a certain incentive level a certain shedding potential is expected. There is a one-day simulation provided in the paper, using a 15 min time frame, and considering 100 arbitrarily chosen residential load profiles. Different combination of price level and

load reduction are supposed.

Our work provides both centralized and decentralized solutions for agent-based large-scale DSM. In particular, we focus on a specific load type for DSM, that is the *flexible*, or *shiftable* load type. This type of load may refer to the operation of appliances directly connected to human daily activities, for example, heating of water, cooking, clothes washing, charging of electronics devices, etc., involving this way user shifting preferences that vary according to each case, and, importantly, high levels of uncertainty regarding the future planning and DSM performance. Another characteristic that is present is that, although such loads can be curtailed during a time interval, it is imperative to be consumed at later, or earlier time intervals, i.e., it cannot be reduced from the aggregate daily load. Additionally, the aggregation of the expected consumption profiles of cooperative or community members, allows for a more precise planning regarding the ordered DSM from the Grid side, since we apply specific constraints regarding the aggregate hourly load. This fact helps with the elimination of potential herding effects, which can emerge as a result of dynamic pricing of electricity [146].

Kota *et al.* [96] were actually the first to propose a demand side management scheme involving electricity consumer cooperatives. Their scheme comes complete with certain incentive compatibility guarantees, but differs to ours in many important ways. In their scheme, consumers represented by agents form cooperatives with the purpose to participate in the (wholesale) electricity markets *as if they were producers*, essentially selling energy *nega-watts* in the form of reduction services. However, their approach would, in most countries, require legislature changes in order to be applied in real life. Moreover, agents in their scheme have to essentially sign strict contracts

with the Grid to participate in the market, and cooperative members risk the danger of being significantly “punished” for not meeting their obligations through what might appear to a small-scale, household consumer as a complicated protocol. Thus, real consumers might prove reluctant to join cooperatives and participate in their scheme. By comparison, in our scheme, agents simply run the danger of being granted less profit for their actions than originally promised. Importantly, they are also capable of minimizing this risk, due to the fact that they are guided a priori (and have agreed) to the time slots where they can actually *shift* consumption to. Indeed, no guidelines whatsoever as to where to shift consumption to are provided in [96], and deals agreed there involve *reduction promises* only. Our approach allows for a more relaxed agent-Grid interaction, for explicit power consumption *shifting* to targeted time intervals.

Other approaches, like those in [80] and [127], aim to optimize consumption schedules via searching for Nash equilibria in specific game settings. We chose not to follow this line of research, because it requires that every “player” retains a specific and fixed strategy. This cannot be realistically assumed to hold—let alone be guaranteed—in any large, open multiagent environment. In the work of [163], authors propose a Vickrey-Clarke-Groves (VCG) mechanism for DSM. The use of their mechanism is shown via simulations to benefit both the provider side, as well as the demand side. The main difference with our approach is that this is centralized since rescheduling of loads is calculated by the energy provider, and, furthermore, it calls for last minute action, which can induce decreased participation.

Several simple reduction schemes, that promise reduced flat electricity rates for lower consumption levels over prolonged periods of time, are already in place in the

real world [23]. Unfortunately, most of those schemes can be easily manipulated in “unethical” ways by individuals. For instance, they have no means to exclude consumers that simply happened to be able to not demand electricity over some period; that is, an individual could go away on holiday for a month, and collect a cash reward for doing so. Our scheme does not suffer such problems, as it (a) rewards consumption reduction—and, importantly, promotes *consumption shifting*—on essentially an hour-to-hour basis (planned a day ahead), and (b) rewards these “short-term” services based on how successfully they were delivered.

Examining the work of [119], it addresses the issue of realistic mechanism design for sharing the cost of electricity between strategic, rational, and selfish household agents with private information, in day-ahead planning. Actually, an aggregator undertakes to buy electricity from a wholesale market. Next, agents provide the aggregator with their consumption needs (appliance types, desired hours of initiation, and expected consumption). Then, the aggregator solves the aggregate optimization problem *wrt.* monetary cost of energy for each interval, and communicates the solutions (i.e. consumption profiles for each home and appliance) back to the agents. On the second stage of the mechanism, agents try to meet the promised behavior, and are finally charged based on their actual actions, after the application of the spherical scoring rule, which is strictly proper. Authors provide mathematical proofs of asymptotic incentive compatibility, and ex-ante weak budget balancedness. This model is applied only in residential cases. In [82] coalition formation is utilized to combine intermittent renewable generators with energy stores, so that the coalitions always provide what promised, regardless of the uncertainty due to, e.g., weather conditions.

In [199], the authors present a *tariff model* where the price for electricity depends on *agent predictions of demand*. This approach forms a concave game where participants are called to report their future baseline consumption. The final price that agents pay for consumption is either better for the accurate ones, or for those that joined the game earlier in time. The *Shapley value* solution concept, used widely in cooperative game theory for computing *fair* payments among coalition members [39], is utilized to distinguish rewards among good and poor predictors; and it is shown that no agent has an incentive to consume electricity in excess of its actual requirements, regardless of the baseline it reported.

The work of Gottwalt *et al.* [68], assumes variable electricity prices and presents an environment that simulates the behaviour of house tenants participating in demand side management efforts. However, it makes no references to the potential discrepancies between expected and final actions. By contrast, our model captures such uncertainties, and uses specific techniques to promote efficiency, as we will explain in Section 3.3.

Finally, various *scoring rules* have been widely used to incentivize agents to act by the rules and not deviate intentionally. The work of [35] proposes the use of scoring rules that compensate *forecasting experts* for the utility loss they suffer by stating truthful forecasts. Other recent works in the Smart Grid domain, use the *spherical* scoring rule to incentivize participants to accurately predict aggregate electricity demand or demand response actions [119, 159]. We chose to use the *CRPS* instead of the spherical scoring rule, because the *CRPS* is indifferent between the sign of the error, while the spherical is not [114]: our goal here was to incentivize precise forecasts, and deviations to either side—i.e., either lower, or higher than stated—

should be penalized equally.

3.2 A Generic Electricity Consumption Shifting Model

In this section we describe our electricity consumption shifting model. We begin by describing the problem setting and the required notation. After defining the complete model, we proceed to discuss the necessary constraints that must hold in order to guarantee the feasibility of collective demand shifting efforts. Next, we analyze the incentives of participation for each agent and identify the relationships between shifting costs and prices offered for consumption. Based on these, we propose a pricing scheme that grants lower prices for consuming the (possibly collectively) shifted demand during non-peak intervals; and equip our mechanism with a strictly proper scoring rule to promote accuracy and effectiveness in the coordinated efforts.

To begin, power supply must continuously meet demand that varies between time intervals. To meet this need and in order to provide incentives for consumption at times where production and power supply are cheap, the electricity pricing scheme used in many countries consists of two different pricing rates, one for day-time and one for night-time consumption. In our model, we also assume that there exist exactly two different price levels $p_h > p_l$. These, however characterize *each specific time interval* t , based on a demand threshold τ^t under which electricity generation costs are lower, e.g. when renewable energy levels are higher:

Table 3.1: Notation.

Symbol	Meaning
i	Agent ID
t	Time interval
t_h	Peak time intervals
t_l	Non-peak time intervals
p^t	Price of electricity during t
p_h	Peak time price
p_l	Non-peak time price
p_g	Better price granted for consumption shifting efforts
τ^t	Demand threshold that characterizes peak and non-peak intervals
sl	Safety limit $\leq \tau$
D^t	Electricity demand during t
$q_\tau^{t_h}$	Amount of load whose removal can allow for a better price
$Q_{max}^{t_h}$	Maximum amount of load eligible for a better price when shifted from t_h
$q_{min}^{t_h}$	Minimum amount of load eligible for a better price when shifted from t_h
$r_i^{t_h}$	Load reduced by i at t_h
$q_i^{t_l}$	Load shifted by i to t_l
$q_{sl}^{t_l}$	Load quantity available under sl during a t_l
$\hat{r}_i^{t_h}$	Stated agent shifting capacity (amount of load pledged to be removed from t_h)
$\tilde{r}_i^{t_h}$	Cooperative estimate over the agent's reduction capacity
$c_i^{t_h \rightarrow t_l}$	Agent cost for shifting one kWh from t_h to t_l
$\hat{\sigma}_i$	Stated agent uncertainty regarding stated $\hat{r}_i^{t_h}$
\hat{p}_i	Agent i 's reservation price for shifting to specific t_l
ξ_i	Agent i 's contribution potential
p_i^{eff}	Agent i 's "effective price"—i.e., price eventually paid by i at t_l
$\hat{r}_C^{t_h}$	Stated cooperative shifting capacity
$\hat{\sigma}_C$	Stated cooperative uncertainty regarding stated $\hat{r}_C^{t_h}$
G_C	Cooperative expected gain
p_C	Price awarded for cooperative action
B_i	Electricity bill for i
b_i	Cooperative contributor's bill

Note: the key parameters of our scheme are only the prices $p_l, p_h, p_g(\cdot)$, and $Q_{max}^{t_h}$, and $q_{min}^{t_h}$, which are essentially determined by the Grid.

$$p^t = \begin{cases} p_h, & \text{if } D^t \geq \tau^t \\ p_l, & \text{if } D^t < \tau^t \end{cases} \quad (3.1)$$

where D^t is the energy demand during t . The intervals during which $p^t = p_h$ are considered to be peak-intervals, at which consumption needs to be reduced. We note peak intervals as $t_h \in T_H$ and non-peak ones as $t_l \in T_L$.

Now, given the daily consumption pattern known to the Grid, it would ideally like consumption to drop under a *safety limit* that is placed below τ . Dropping below the safety limit would ensure that some low cost generated load is available in case of high uncertainty or an emergency, thus minimizing the risk that high-cost generators would have to be turned on. That is, the Grid would ideally want to reduce consumption by $Q_{max}^{t_h} \geq q_\tau^{t_h}$, where:

1. $Q_{max}^{t_h}$ is the load normally consumed over the safety limit at t_h (that is the maximum load eligible for shifting), and
2. $q_\tau^{t_h}$ is the minimum amount of load whose potential removal can, under the Grid's estimations, allow for a better electricity price² to be offered to contributing reducers.

Intuitively, $q_\tau^{t_h}$ is a sizable load quantity that makes it cost-effective for the Grid to grant a very low electricity rate, in anticipation of reaching a demand level that is close to the safety limit. We denote the load reduced by some agent i at a t_h as $r_i^{t_h}$, and that shifted to each $t_l \in T_L$ as $q_i^{t_l}$.

²The specific nature of the authority maintaining or setting the prices is not relevant to our mechanism.

3.2.1 Scheme overview

We now provide an overview of the proposed shifting scheme. First, the Grid gives information for the time intervals that consumption needs to be reduced at, and those that it is best to shift consumption to. The consumption of the shifted load during these preferred non-peak intervals is granted a better price. Then, the consumer side weighs its costs and potential profits, and chooses to participate in a shifting operation or not. The exact procedure is:

1. The Grid announces peak and non-peak time intervals³ with high and low consumption prices and asks agents to announce their willingness to shift some of their production from peak to non-peak intervals, promising them a better consumption price for doing so.
2. The agents put forward bids to shift specific amounts from peak to non-peak intervals, along with their costs for doing so, and their *uncertainty* (in the form of a probability distribution) regarding their ability to honor their bids. If the agents represent a cooperative, deliberations internal to the cooperative occur, in order to determine its bids, as we detail later.⁴
3. A *clearing process* takes place, determining the accepted agent bids.

³We must note that time intervals can be of any size. In this work, we consider them to be 24 hour intervals per day, that is a pretty common division in the energy domain. However, any frequency could be accepted, as the performance of the algorithm is affected by the number of peak time intervals regardless of the resolution.

⁴Note that, instead of communicating individual agent shifting costs, it would be possible for the cooperative to announce a marginal shifting cost that would allow profitable participation; then the agents would decide whether they contribute or not (see, e.g., Chapter 4). However, this removes the ability to perform internal price balancing, i.e. most profited agents granting a small part of their gain to necessary, but potentially non-profited agents, so that to make shifting profitable for everyone, as we discuss later in Section 3.4.

4. During the next day,⁵ the agreed consumption shifting activities take place.

The clearing process, apart from determining the expected gains from participation, also guarantees the feasibility of the efforts, by applying specific constraints, which we describe in the following subsection. In reality, of course, such efforts are expected to be best conducted by consumers joining forces in *cooperatives*, as this is the only way to actually deliver substantial—and thus effective—power consumption shifting. We explain the formation of the cooperatives in Sec. 3.3 below.

3.2.2 Constraints

We came up with specific constraints that must hold in order to safely shift demand. These constraints are checked during the demand shifting operations and actually limit the eligible load for reduction by our mechanism. Such limits are important, in order to tackle herding effects. Note that the constraints do not attempt to directly control consumers' actions, just characterize which amounts of load are considered eligible for shifting. If the consumer deviates and constraints no longer hold, then the standard fares are applied.

To begin,

$$\sum_i r_i^{t_h} \geq q_\tau^{t_h} \quad (3.2)$$

that is, the amount of load reduced must be larger than the minimum needed at t_h . In order for the Grid to achieve profits, the reduction must be above certain levels, guaranteeing that the peak will be trimmed.

⁵Our mechanism can be employed for any future date of our choice.

Second,

$$\sum_{t_l} q_i^{t_l} \leq \sum_{t_h} r_i^{t_h}, \forall i \quad (3.3)$$

meaning that every reducer shifts to a subset of non-peak intervals an aggregate load amount of at most the load reduced over the t_h intervals he participates in. Note that the consumer might consume during a t_l more than reduced during t_h , as we do not explicitly restrict power flow. In such situations, the better price of Subsec. 3.2.4 is charged only for the eligible amount, and the excess is charged according to the original prices that would be charged during t_l .

Moreover,

$$\sum_i \sum_{t_l} q_i^{t_l} \leq Q_{max}^{t_h}, \forall t_h \in T_H \quad (3.4)$$

has to hold, meaning that the sum of all reducing agents shifted load to all non-peak intervals must be at most equal to $Q_{max}^{t_h}$, assuming that the Grid has no interest in further reducing consumption, once it has dropped under τ^{t_h} .

Finally,

$$\sum_i q_i^{t_l} \leq q_{sl}^{t_l}, \forall t_l \in T_L \quad (3.5)$$

Namely, the total shifted load at each t_l must not exceed the $q_{sl}^{t_l}$ quantity which is actually available under the safety limit, in order to avoid the creation of a new “peak” at t_l . The objective is to keep demand close to the safety limit in as many intervals as possible.

As an example, consider the situation in Fig. 3-2, where a peak induced by some load needs to be trimmed (red shaded). Consumers shift that load as an aggregate consumption to earlier or later intervals (green shaded), where demand is below the

safety limit. Note that to avoid the creation of a new peak during those intervals, the addition of shifted load must not make demand exceed the safety limit. The threshold placement denotes the production levels for cheaper production, thus a lower price is charged for consumption.

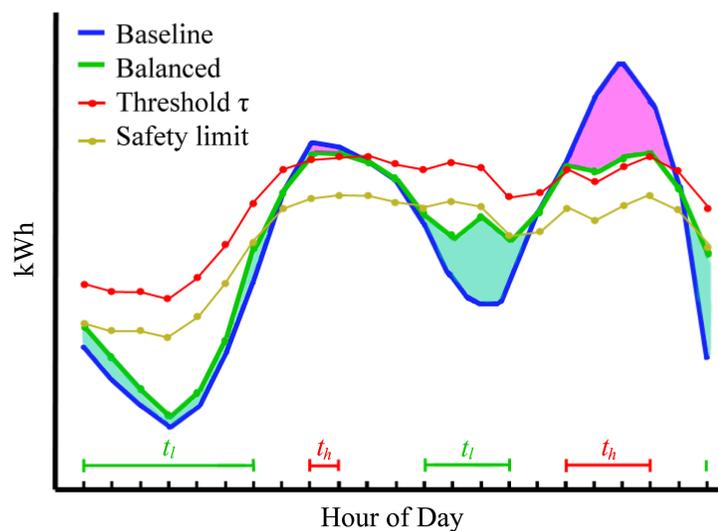


Figure 3-2: Scheme objective representation. Portions of peak-load at t_h intervals, are shifted to lower demand intervals t_l .

3.2.3 Agent incentives

The participation of each agent in the scheme obviously depends on his individual costs and potential gains. Suppose that an agent i ponders the possibility of altering his baseload consumption pattern by shifting some electricity consumption r_i from an interval t_h to t_l . This shifting effort is associated with a *cost* $c_i^{t_h \rightarrow t_l}$ per kWh for

the agent.⁶ The *gain* that an agent has for shifting r_i to t_l given t_l 's lower price p_l , is equal to

$$\text{gain}(i|p_l) = r_i(p_h - p_l - c_i^{t_h \rightarrow t_l}) \quad (3.6)$$

since the agent would be able to consume r_i at t_l for a lower rate. However, under normal circumstances this gain is *negative* for the agent, that is,

$$p_l + c^{t_h \rightarrow t_l} > p_h \quad (3.7)$$

because if not, then the agent would have already been able to make that shift (and its baseload pattern would have been different than its current one).

Now, if the Grid is able to grant an *even lower* rate p_g for consumption of $r_i \geq q_{min}^{t_h}$ at t_l s.t.

$$p_g + c^{t_h \rightarrow t_l} \leq p_h \quad (3.8)$$

then the agent will be incentivized to perform the shift, as his perceived $\text{gain}(i|p_g)$ would now be non-negative.

Lemma 1 (Better Price). The better price must lead to non-negative gains

Proof Assume that p_g rate is better than p_h by at least $c_i^{t_h \rightarrow t_l}$. Then:

$$p_g \leq p_h - c_i^{t_h \rightarrow t_l} \Leftrightarrow \quad (3.9)$$

$$(p_h - p_g - c_i^{t_h \rightarrow t_l}) \geq 0 \quad (3.10)$$

⁶Note that this cost can be elicited by using existing CHI methods, such as [190].

Since r_i denotes the amount of kWh shifted, its value is always positive and multiplication does not change the sign:

$$r_i(p_h - p_g - c_i^{t_h \rightarrow t_l}) \geq 0 \quad (3.11)$$

As defined by Eq. (3.6), this is the gain of an agent for shifting from a peak interval to another non-peak one where a price rate p_g is given, so the following finally holds:

$$gain(i|p_g) \geq 0 \quad (3.12)$$

Thus, the gain for every agent is non-negative, and individual rationality is guaranteed. \square

3.2.4 Group price

Our scheme allows the shifting of sizable load consumption from peak to non-peak intervals. The eligible load for coordinated shifting is granted an *even lower*, “*group*” price $p_g < p_l$, which is a function of the actual load reduced at t_h , $r_i^{t_h}$, in a way that for larger load portions, the price becomes better. We term this price as p_g because such reduction will likely be possible only by groups⁷ of agents. Thus, the group price is given as:

$$p_g(r_i^{t_h}) < p_l \quad (3.13)$$

and it is awarded if *the actual quantity of the load shifted* from t_h exceeds some minimum value $q_{min}^{t_h}$, set by the Grid given its knowledge of $q_\tau^{t_h}$ (e.g., it could be

⁷In our case this reduces to a single and large group, i.e. the cooperative

$q_{min}^{t_h} = q_{\tau}^{t_h}$). The function can actually be linear or non-linear, as long as it is decreasing and complies with the incentive analysis of Subsection 3.2.3.

To elaborate further, the Grid is not willing to offer a lower price to every consumer, as then it would only need to lower the p_l rates. Instead, p_g is awarded only to adequately large participant numbers, those whose coordinated consumption shifting finally reduces the Grid’s generation costs. This is analogous to what happens in *economy of scale* and *group buying* paradigms [143, 203] which takes place for several years in actual real world trades.

3.2.5 Continuously ranked probability score

Now, to promote efficiency in load shifting and (in the face of the global constraints described in our model) avoid Grid interaction with unreliable participants, the agents need to be motivated to precisely report their true reduction capabilities. To achieve this, we employ a *strictly proper scoring rule*, the *continuous ranked probability score (CRPS)* [67], which has also been recently used in [155] to incentivize renewable energy-dependent electricity *producers* to accurately state their estimated output when participating in a cooperative. Recall that a scoring rule $S(\hat{P}, x)$ is a real valued function that assesses the accuracy of probabilistic forecasts, where \hat{P} is the reported prediction in the form of a probability distribution over the occurrence of a future event, and x is the actual occurrence itself. As such, a mechanism can exploit a rule’s strict propriety property, to ensure that scores that are lower than the best result to lower returns (e.g., via the use of penalties) to the participants—and, by so doing, ensure the *truthfulness* and *incentive compatibility* [133] of the agent reports. The use of *CRPS*, in particular, allows us to directly evaluate probabilistic

forecasts, and the score is given by:

$$CRPS(\mathcal{N}(\mu, \sigma^2), x) = \sigma \left[\frac{1}{\sqrt{\pi}} - 2\phi\left(\frac{x - \mu}{\sigma}\right) - \frac{x - \mu}{\sigma} \left(2\Phi\left(\frac{x - \mu}{\sigma}\right) - 1 \right) \right] \quad (3.14)$$

In our setting, $\mathcal{N}(\mu, \sigma^2)$ is the uncertainty stated over the expected *absolute relative errors*⁸ regarding the reduction capacity, as reported by an agent; while x is the actually observed error, ϕ the PDF, and Φ the CDF of a standard Gaussian variable. A *CRPS* value of zero signifies a precise forecast, while a positive value shows the distance between prediction and occurrence. For convenience, we normalize *CRPS* values to $[0, 1]$, with 0 assigned when we have exact forecast, and 1 assigned when the forecast gets far from the occurrence. To improve readability, we also henceforth note $CRPS(\mathcal{N}(\mu, \sigma^2), x)$ as *CRPS* without the arguments and write $CRPS_i$ to denote the *CRPS* rule applied to agent i 's performance, while the stated agent uncertainty over its error is considered to be zero mean, $\mathcal{N}(0, \hat{\sigma}^2)$, so this confidence measure will be sometimes simply referred to as $\hat{\sigma}$. Given this notation, an agent i whose bid to shift some load from t_h to t_l is accepted, is charged an electricity bill B_i ⁹, given its actual contribution $r_i^{t_h}$ ¹⁰, and its additional consumption $q_i^{t_l}$ at the low-cost time interval t_l :

$$B_i^{t_l} = (1 + CRPS_i) q_i^{t_l} p_g(r_i^{t_h}) \quad (3.15)$$

Note also that it can be $q_i^{t_l} < r_i^{t_h}$, since an agent can shift $r_i^{t_h}$ to multiple t_l s.

⁸The mean μ and variance σ^2 of this distribution can be estimated by each agent through private knowledge of its consumption requirements and business needs.

⁹ $B_i^{t_l}$ is the bill the agent receives for the $q_i^{t_l}$ quantity it shifted to t_l . The agent might have been billed a separate amount for any quantity already being consumed at t_l before shifting.

¹⁰Note that $r_i^{t_h}$ is the *actual amount reduced* at t_h , which also determines the electricity price p_g for i at t_l ; while $q_i^{t_l}$ is the quantity shifted from t_h to t_l and consumed there.

Lemma 2. The payment rule $B_i^{t_i}$, as defined in (3.15), is strictly proper.

Proof Equation (3.15) shows that $B_i^{t_i}$ is an affine transformation of $CRPS_i$, which is a strictly proper scoring rule. According to [123] any affine transformation of a strictly proper scoring rule is also strictly proper. Thus, $B_i^{t_i}$ is strictly proper. \square

To summarize, $CRPS$ provides a scoring function for evaluating the accuracy of a forecast, given its actual occurrence. When agent-stated forecasts are off the occurrences, contributors are “fined” proportionally to their $CRPS$ score. However, while this *mechanism design* technique provides the agents with strong incentives to stay truthful (and, indeed, provides theoretical guarantees for statement truthfulness), it does not guarantee agent statements accuracy, as we have already explained in Chapter 1.

3.3 Forming Effective Demand Shifting Coalitions

In the general case, it is very rare for reducers to have shifting capacity greater than $q_{min}^{t_h}$, even for large industrial consumers. Therefore, the agents need to organize into cooperatives in order to coordinate their actions and achieve the better rates promised by the Grid for effective consumption shifting. As we have already mentioned, shifting is meaningful only for substantial quantities, such that the Grid can guarantee better prices, which make participation worthwhile. Consequently, cooperative action is needed. We envisage a cooperative as being composed by hundreds or even thousands of consumers. In what follows, we explain how bids are formed for the cooperative case, and what is actually charged at the end of the shifting efforts

to the participants. In this section we detail different methods for the formation of demand shifting coalitions i.e., that act to carry out specific “shifting contracts”, inside the cooperative, and discuss their final billing, which also incorporates a *CRPS* score.

Naturally, at every given time interval t_h earmarked for potential consumption reduction, only a subset \mathbb{C}^{t_h} of cooperative members might be available for shifting services. We assume that every member agent announces its availability for every such t_h to a cooperative manager agent, along with its reduction (shifting) capacity $\hat{r}_i^{t_h}$; its confidence $\hat{\sigma}_i^2$ on its ability to reduce that amount at t_h (specifically, the agent’s expected relative error with respect to $\hat{r}_i^{t_h}$ reduction is normally distributed according to $\mathcal{N}(0, \hat{\sigma}_i^2)$); and the set of low cost intervals t_l that he pledges to move consumption to.

Even so, more often than not, it is impossible for all agents in \mathbb{C}^{t_h} to participate in the cooperative effort. This is because their shifting costs might be so high that they would not allow their inclusion in any profitable cooperative bid. Therefore, only a subset C of \mathbb{C}^{t_h} will be selected for participation in the bid. Such a shifting bid is composed by the following parts: t_h , the high cost interval to reduce consumption from; \hat{r}_C , the amount C pledges to reduce at t_h ; a pair (T_l, Q_l) that determines the set of low cost intervals t_l to shift consumption to, along with the set of corresponding quantities that will be moved to each t_l ; the per *kWh shifting costs* $c_i^{t_h \rightarrow t_l}$ associated with moving consumption from t_h to t_l ; and a probabilistic estimate of its $\hat{\sigma}_C^2$, in the form of a normal distribution describing the joint relative error on predicted r_C , $\mathcal{N}(0, \hat{\sigma}_C^2)$.

Thus, at a given time interval t_h the cooperative has to select C agents to co-

operate and shift demand to some t_l . Without loss of generality and for ease of presentation, let us assume that t_l is just one specific time interval (though in reality it is just one interval among many in T_l). Assume also that $|C| = n$, and thus the total stated reduction capability of C at t_h is:

$$\hat{r}_C^{t_h} = \sum_{i=1}^n \hat{r}_i^{t_h} \quad (3.16)$$

Given this, the collective expected gain by this shifting operation, assuming the stated quantity $\hat{r}_C^{t_h}$ and corresponding stated confidence $\hat{\sigma}_C$ are *accurate forecasts*, is:¹¹

$$G_C = \hat{r}_C^{t_h} (p_h - p_g(\hat{r}_C^{t_h}) - \bar{c}^{t_h \rightarrow t_l}) \quad (3.17)$$

where $\hat{r}_C^{t_h}$ is the amount of load pledged to be shifted, with $q_{min}^{t_h} \leq \hat{r}_C^{t_h} \leq Q_{max}^{t_h}$ and $p_g(\hat{r}_C^{t_h})$ is the price corresponding to the shifting forecast. Cost $\bar{c}^{t_h \rightarrow t_l}$ denotes the (expected) weighted mean shifting cost:

$$\bar{c}^{t_h \rightarrow t_l} = \frac{\sum_{i=1}^n \hat{r}_i^{t_h} c_i^{t_h \rightarrow t_l}}{\sum_{i=1}^n \hat{r}_i^{t_h}} \quad (3.18)$$

where $\hat{r}_i^{t_l}$ is the load pledged to be shifted by each agent.

Obviously, the cooperative has to select a C so that G_C is non-negative in Eq. 3.17 above. Notice that in order for inequality $\bar{c}^{t_h \rightarrow t_l} < p_h - p_g(\hat{r}_C^{t_h})$ to hold, it is not required that each $c_i^{t_h \rightarrow t_l}$ is less than the price difference too, because the quantity of shifted load also matters. This allows for the possibility that reducers with high

¹¹This is assuming that all $\hat{r}_C^{t_h}$ pledged to be shifted is actually later on consumed at t_l . If only $q_C^{t_l} < \hat{r}_C^{t_h}$ is consumed at t_l , then the gain is $G_C = \hat{r}_C^{t_h} p_h - p_g(\hat{r}_C^{t_h}) q_C^{t_l} - \bar{c}^{t_h \rightarrow t_l} q_C^{t_l}$.

shifting costs but low shifting capacity can also contribute in the cooperative bid.

Moreover, the cooperative needs to be able to estimate its confidence $\hat{\sigma}_C$ over the expected performance of any subset of an acting set of agents C . This confidence is also calculated using the weighted mean formula of Eq. 3.18, where, however, instead of individual costs, individual confidence estimates are now used:

$$\hat{\sigma}_C = \frac{\sum_{i=1}^n \hat{r}_i^{t_h} \hat{\sigma}_i}{\sum_{i=1}^n \hat{r}_i^{t_h}} \quad (3.19)$$

Now, there are several ways to determine the subset C of agents to be selected for shifting. Also, due to the uncertain nature, the estimation of the reduction quantity that C will contribute at t_h is also another issue which can be addressed with different approaches. We proceed to describe the process by which the cooperative determines the subset C of \mathbb{C}^{t_h} to participate in its bid at t_h .

3.3.1 Choosing the acting coalitions

In more detail, suppose a set of n agents that constitute the cooperative. As already explained, in order for the cooperative to place a bid, each contributing agent i must state its reduction capacity, $\hat{r}_i^{t_h}$, at t_h high-consumption (peak) intervals, and corresponding shifting costs $c_i^{t_h \rightarrow t_l}$ for moving consumption to non-peak, t_l , intervals. Agents are also required to state their uncertainty over their *expected relative error* regarding their reduction capacity, in a form of a normal distribution $\mathcal{N}(\mu_i, \hat{\sigma}_i^2)$.¹² Next, the cooperative assigns an estimate of each agent's expected performance $\tilde{r}_i^{t_h}$,

¹² As stated earlier, we set $\mu_i = 0$. This simply assumes that random relative errors, over a long enough time range, will be normally distributed around a mean of 0.

which is a function \mathcal{F} of its stated capacity $\hat{r}_i^{t_h}$ and its error distribution $\mathcal{N}(\mu_i, \hat{\sigma}_i^2)$. Function \mathcal{F} can correspond to various performance prediction methods, as we will be explaining in Chapter 6.

$$\tilde{r}_i^{t_h} = \mathcal{F}(\hat{r}_i^{t_h}, \mathcal{N}(\mu_i, \hat{\sigma}_i^2)) \quad (3.20)$$

One potential estimate is the *conservative* one, i.e.:

$$\tilde{r}_i^{t_h} = \hat{r}_i^{t_h} - \hat{\sigma}_i \hat{r}_i^{t_h} \quad (3.21)$$

With $\tilde{r}_i^{t_h}$ at hand, the absolute relative error between the actual reduction and the expected one is given by:

$$x_i^{t_h} = \frac{|r_i^{t_h} - \tilde{r}_i^{t_h}|}{\tilde{r}_i^{t_h}} \quad (3.22)$$

This value is used for the computation of agent i 's *CRPS* score (cf. Eq. (4.21)).

Then, each agent's *reservation price* \hat{p}_i is calculated as the difference between the price paid for consuming at the “high-cost” t_h intervals, and the agent i 's costs for shifting consumption to t_l intervals with lower electricity charges:

$$\hat{p}_i = p_h - c_i^{t_h \rightarrow t_l} \quad (3.23)$$

Thus, this \hat{p}_i quantity is the highest price i is willing to pay for shifting consumption from t_h to t_l without suffering a monetary loss. Now, the agent's *contribution potential* ξ_i can be calculated as the product of the expected reduction and reservation

price:

$$\xi_i = \tilde{r}_i^{t_h} \hat{p}_i \quad (3.24)$$

The contribution potential is a measure that compactly captures the impact of the individual agents contribution to the coalition. It is higher when more load is offered for shifting and agent reservation prices are high, meaning they are in less need for compensation, and, as such, their employment would cost less to the Grid side; and it is lower for agents with low reduction capacities, and low reservation prices (i.e., greater compensation demands for shifting).

The agents are then ranked by some ξ_i -related *ranking criterion* of the cooperative's choosing (cf. below), and *shifting coalitions* are formed by the number of top agents that meet any set of requirements determined by the cooperative (e.g., “maximize the shifting capacity”, etc.). Selected coalition agents are awarded low, *variable prices* for shifting to t_l , determined by the *group price* $p_g \leq p_l$ which is guaranteed by the Grid, and by monetary *gain transfers* that make it worthwhile for everyone selected to participate in the shift, as described in Section 3.4 below.

Thus, in order to form the acting coalitions, one must first check, for every i , whether $\hat{p}_i \leq 0$. If that holds for all i , we stop; the problem is infeasible (as *all* agents need to *be paid* with a rate equal at least \hat{p}_i in order to participate). Furthermore, if for all i it holds that $\hat{p}_i \leq p_g^{\min}$ (where p_g^{\min} is the best possible quantity-dependent group price that can be granted), then the problem is again infeasible and we stop (as *all* agents need to be *paid* with a rate at least equal to $p_g^{\min} - \hat{p}_i$ in order to participate). If that is not the case, then there exist some agents in \mathbb{C}^{t_h} for which there is a price they can accept to pay so as to move some of their consumption to t_l

without suffering a loss. We now present four different coalition formation methods that one can use in order to identify such agents.¹³

Rank by contribution potential and maximize expected capacity (Method CF1)

The algorithm ranks all agents in \mathbb{C}^{t_h} by their *contribution potential*, i.e., $\xi_i = \tilde{r}_i^{t_h} \hat{p}_i$, in decreasing order. Then, starting from the agent with the highest ξ_i value, we sum these values up in decreasing order, and add the respective agents in a group C . Intuitively, the algorithm attempts to add in the coalition members with high “potential” to contribute to reduction—that is, members with potentially high \tilde{r}_i to contribute, while being able to accept a relatively high (though reduced) energy price \hat{p}_i . This process continues until the following conditions are met for the *maximum possible* group of agents C :

$$\sum_{i \in C} \xi_i \geq \tilde{r}_C p_C \quad (3.25)$$

$$\tilde{r}_C \geq q_{min}^{t_h} \quad (3.26)$$

$$\tilde{r}_C \leq Q_{max}^{t_h} \quad (3.27)$$

where $q_{min}^{t_h}$ and $Q_{max}^{t_h}$ are the minimum and the maximum quantity admitting a

¹³In what follows, we relax the notation somewhat, by dropping time indices where these are clearly implied.

group price at t_h , respectively;

$$\tilde{r}_C = \sum_{i \in C} \tilde{r}_i \quad (3.28)$$

; and

$$p_C = p_g(\tilde{r}_C) \quad (3.29)$$

is the price rate offered by the Grid for reduction \tilde{r}_C .

To provide further intuition, note that the expected gain of every agent in some group C given p_C is:

$$gain(j|p_C) = \tilde{r}_j(\hat{p}_j - p_C) \quad (3.30)$$

If we were simply given a C for which this gain was positive for every member, then each agent would have been able to just pay p_C and enjoy the corresponding gain. However, the reducing set C has to be dynamically determined by the cooperative—and, in order to guarantee individual rationality, so is the price paid by each one of its members. Note also that by progressively adding agents with lower ξ_i values to C , the agents j preceding them can only become better-off, in terms of expected gain, as the group price p_C expected to be attained drops.

Now, if *all* agents in \mathbb{C}^{t_h} are inserted in C and \tilde{r}_C is still lower than $q_{min}^{t_h}$ (i.e., there is a violation of the constraint of Eq. (3.26)), the problem is infeasible and we stop. Likewise, if all agents are in C and $\sum_{i \in C} \xi_i - \tilde{r}_C p_C < 0$ (violation of the constraint of Eq. (3.25)), the problem is again infeasible and we have to stop.

Assume that this has not happened, and both conditions have been met for *maximal* C .¹⁴ This means that there is *at least* one agent j in C that has a positive

¹⁴That is, assume that C was actually constructed so that, after a subset of agents L has met $\tilde{r}_L \geq q_{min}^{t_h}$, we kept adding agents to L until by adding an agent k we constructed some C' so that

gain, given p_C . That is,

$$\text{gain}(j|p_C) = \tilde{r}_j(p_h - c_i - p_C) = \xi_j - \tilde{r}_j p_C > 0 .$$

If not, then no agent has a positive gain, and thus $\sum_{i \in C} \xi_i - \tilde{r}_i p_C \leq 0$, leading to $\sum_{i \in C} \xi_i \leq \tilde{r}_C p_C$, contradicting condition of Eq. (3.25) above. This also means that agents in C are collectively willing to pay a total amount for moving their \tilde{r}_i consumptions to t_l , which is greater than what their group will be asked to pay for (assuming accurate forecasts), given the Grid's offer p_C for \tilde{r}_C reduction.

Thus we have ended up with the maximal C so that Eq. (3.25) and Eq. (3.26) hold, and which contains some agents with positive and some agents with negative gain given p_C . This means that the cooperative bid is actually formed as \tilde{r}_C . We will later (in Section 3.4) use these agents to implement a gain transfer scheme so that all individual agents in C end up with non-negative gain themselves.

Rank by contribution potential and meet minimum expected capacity requirements (Method CF2)

When using *CF2*, we form coalitions in the same manner we did when employing *CF1*, with the difference that we stop adding agents in the coalition at the point when the reduction capacity \tilde{r}_C becomes equal or more than the minimum amount of load eligible for a better price q_{min}^{th} . This way, coalitions with the minimum shifting capacity are formed. Thus, *CF2* runs the risk of proving overly "optimistic" wrt. its final shifting performance estimates. This fact was actually verified by our

$\sum_{i \in C'} \xi_i - \tilde{r}'_C p'_C$ has turned to negative, in which case we remove k from the list C' so that we end up with a C that has $\sum_{i \in C} \xi_i - \tilde{r}_C p_C \geq 0$.

experiments in Section 3.6.2.

Rank by expected gain and maximize expected capacity (Method CF3)

Here we rank the available agents by their *expected gain* (cf. Eq. (3.30)) with respect to the p_g offered at the very moment they are to be included in the coalition. Then, the agents are added in the coalition under formation given this ranking, as long as the “feasibility constraints” of Eqs. (3.25), (3.26), and (3.27) hold.

Note that the ranking used in this particular method, encompasses an insight regarding the expected gain given the better price offered (Eq. (3.30)). By contrast, *CF1* and *CF2* rank the agents by contribution potential. However, having a high contribution potential does not necessarily mean that the gain of an agent is also high. For instance, assume two agents, the first of which with $\tilde{r}_1 = 100$, $\hat{p}_1 = 0.05$, and thus $\xi_1 = 5$; and the second one with $\tilde{r}_2 = 20$, $\hat{p}_2 = 0.06$, and thus $\xi_2 = 1.2$. Assume $p_C = 0.055$. Then, we get (expected) gains $gain(1|p_C) = -0.5$ and $gain(2|p_C) = 0.1$ respectively. In this case, *CF1* would rank agent 1 before agent 2, while *CF3* would do the opposite.

This results to the potential exclusion from the acting coalition of agents that could contribute much to the joint gain—because by the time their inclusion is examined by *CF1*, the feasibility constraints¹⁵ have already been violated, and formation has stopped. For this reason, one would expect that using *CF3* would give rise to coalitions with an increased gain potential and a higher number of agents, which is naturally linked to increased shifting ability. This was verified by our experiments in Section 3.6.2.

¹⁵Usually Eq. (3.27), due to the addition of agents with high capacity but low gain potential.

However, employing *CF3* is computationally more expensive. This is because whenever an agent enters the coalition, the total reduction capacity increases, and consequently p_g changes. This price change induces further changes in the expected gain of the agents, the criterion by which agents get ranked, and thus the respective values have to be recalculated with every new addition. The added complexity is of order $\mathcal{O}((n - k)^2)$, where n is the number of the possible contributors, and k is the number of agents that form the minimum coalition.

Random selection of contributors and maximize expected capacity (Method Random)

Here we *do not* rank agents by any criterion whatsoever, but those included in the coalitions are simply chosen randomly. The random inclusion continues until the maximum capacity is reached, just as it was the case in *CF1*.

At this point we must note that any simpler approaches than the ones we mention, comes with no incentive compatibility guarantees. The method that ranks agents by their contribution potential (*CF1*) and the one that randomly selects them (*Random*) are our baseline approaches.

3.3.2 Cooperative bidding and billing

Above we defined coalition formation methods to determine which agents to actually include in the acting coalitions, and we showed how to compute specific values that characterize the coalition in terms of expected performance. Given these, the actual bid that the cooperative submits for a peak interval t_h , consists of:

1. The sets of non-peak intervals that the coalescing agents agree to shift to, namely the t_l s.
2. The total expected reduction capacity \tilde{r}_C .
3. The total confidence of the coalition $\hat{\sigma}_C$, as calculated using Eq. (3.19).

Next, the shifting actions take place, and the cooperative is billed according to its shifting performance, as we now explain. Assuming that a coalition C was selected for shifting and eventually acted, the total cost (or, the electricity bill) charged by the Grid to the cooperative for consuming $q_C^{t_l}$ after curtailing $r_C^{t_h}$ is given as:

$$B_C = (1 + CRPS_C)q_C^{t_l}p_g(r_C^{t_h}) \quad (3.31)$$

with $CRPS_C$ being calculated using the cooperative confidence (Eq (3.19)), and the collective absolute relative error realized after shifting. Note that *strict propriety* is maintained in this rule, since the only factor depending on agent forecasts is $(1 + CRPS_C)$.

However, even if G_C (Eq. 3.17) is positive, it is not certain that all individual agents in C have a positive gain (and thus an incentive to participate) as well. Nevertheless, with $G_C \geq 0$, the possibility of allowing *all* agents in C to make a non-negative gain arises, via “price balancing” and the use of internal *gain transfers*. These transfers also have to be performed in such a way so that the budget-balance of any cooperative bid is ensured—that is, the sum of the n members’ bills will have to be equal to B_C (or, if not, definitely not be less than it). We now proceed to discuss potential ways that achieve these goals.

3.4 Internal Pricing

As a next step, the cooperative must pre-assign different *effective price* rates p_i^{eff} to each contributor, producing bills that must sum up *at least* to B_C , i.e., the bill charged to the shifting team C . This is done with the understanding that a member’s final effective price will eventually be weighted according to its individual contribution, given also that an acting coalition C of agents will receive an actual price rate that is dependent on its *CRPS* score. Moreover, in order to allow even negative gain agents to be included in the coalition, their prices must be reassigned in order to ascertain that no contributor suffers a gain loss, and maintain individual rationality.

In short, an *internal price balancing* or “*gain balancing*” process is employed, assigning p_i^{eff} effective prices that guarantee that each agent will be granted non-negative gains from participation. More specifically, we propose certain (weakly) budget-balanced internal pricing mechanisms, which meet the goals above, and also ensure that the gains derived from cooperative shifting efforts are shared among its members in intuitively fair ways (e.g., larger and more accurate contributors can still expect after internal price balancing to rank higher than smaller and less accurate ones in terms of gain). The first pricing method we propose is a heuristic gain balancing algorithm. That is then followed by five additional pricing methods based on different mathematical programming formulations. Note that gain balancing takes place only if the collective (expected) gain is positive, and under no circumstances leads to negative (expected) gains for any of the individual participants.

3.4.1 Heuristic internal price balancing

The heuristic algorithm actually balances the gains of the participants, such that nobody ends up with negative expected gain, and that the most favored ones grant some of their gain to achieve this. However, *the new ranking maintains the same order*, i.e. the one who used to have the most gain, ends up also with the most after the balancing takes place. The procedure is described in detail by Algorithm 1. To begin, the cooperative initially computes individual reservation prices and corresponding gains and proceeds to rank agents in C in decreasing order with respect to their *expected gain* (Lines 1-3), that is:

$$gain(i|p_i^{eff} = p_C) = \tilde{r}_i(\hat{p}_i - p_C) \quad (3.32)$$

If all agents already have non-negative *gain*, then everyone pays p_C and expects to achieve $gain(i|p_C)$ without need of balancing (Line 1). If negativities exist (Line 4), then we must rearrange p_i^{eff} such that agents with the highest gain provide some of their surplus to those with negative, to make their participation individually rational. The first step is to count the total negative gain existing and assign negative gain agents a reduced p_i^{eff} so that their gain becomes exactly zero (Lines 5-6). These agents are added in the set of negative gain participants, denoted as G^- . In Lines 7-32, an iterative process takes place, where the gain of the top agents is reduced until it reaches the gain of those below. This is achieved by the following procedure: having ranked the contributors based on their expected gains, we increase p_i^{eff} of the top agent until its gain drops to the point that it is equal to the g_j gain of the

$j = i + 1$ agent below (as long as $g_j \geq 0$). The value of p_i^{eff} is calculated by:

$$p_i^{eff} = \frac{\xi_i - g_j}{\tilde{r}_i} \quad (3.33)$$

with g_j being the target gain i.e., the gain of the agent below (Lines 20-22).

Then we do the same for the second top agent, until its gain reaches that of the third. We continue in this manner until all requested gain is transferred (Lines 16-17, 24-25), or one's gain reaches zero (Lines 12-14). If the latter happens, we move to the top again and repeat (Line 8).

The p_i^{eff} prices thus determined represent internally pre-specified prices, agreed upon by all the agents, and set ahead of the actual shifting operations. The actual bill b_i that an agent $i \in C$ will be called to pay, however, is determined *after* the actual shifting operations have taken place, and depends on its actual performance wrt. the performance of other agents also, as follows:

$$b_i = \frac{(1 + CRPS_i)p_i^{eff}q_i}{(\sum_{j \in C \setminus \{i\}} (1 + CRPS_j)p_j^{eff}q_j) + p_i^{eff}q_i} B_C \quad (3.34)$$

Strict propriety is ensured by this rule, as it is an affine transformation of a member's $CRPS_i$ score; if the sum in the denominator was over all agents including i , i.e. if $\sum_{j \in C} (1 + CRPS_j)p_j^{eff}q_j$ then Eq. (3.34) would not have been an affine transformation of i 's $CRPS$ score. By contrast, Eq. (3.34) is an affine transformation of $CRPS_i$. As such, its strict propriety is maintained [123]. Moreover, the sum of the b_i bills is always at least as much as the overall bill B_C charged to C , making the mechanism *weakly budget balanced*, and generating some small *cooperative surplus*

Algorithm 1 The Heuristic Internal Price Balancing Method

Input: $p_C, \tilde{r}_i, c_i \forall agent_i \in C$ **Output:** $p_i^{eff} \forall agent_i \in C$

```

1: Assign  $p_i^{eff} = p_C, \forall i \in C$ 
2: Compute reservation prices  $\hat{p}_i$  and gain  $g_i = gain_i, \forall i \in C$ 
3: Sort agents by gain in decreasing order
4: if Negative expected gain agents exist (i.e., the set  $G^-$  is non empty) then
5:   Count total negative gain  $L = \sum g_j, \forall j \in G^-$ 
6:   Assign agents  $\forall j \in G^-, p_j^{eff} = \hat{p}_j$ 
7:    $donation := 0$ 
8:   while  $donation < L$  do
9:     for all Positive gain agents do
10:      if  $donation < L$  then
11:        if  $agent_i$  is the last positive gain agent in sorted gain list then
12:          if  $donation + gain_i \leq L$  then
13:             $donation = donation + gain_i$ 
14:             $p_i^{eff} = \hat{p}_i$ 
15:          else
16:             $donation = L$ 
17:            Assign  $p_i^{eff}$  s.t. only the remaining gain needed is transferred
18:          end if
19:        else
20:          if  $donation + (gain_i - gain_{i+1}) \leq L$  then
21:             $donation = donation + (gain_i - gain_{i+1})$ 
22:            Assign  $p_i^{eff}$  s.t. the amount of  $i$ 's gain is equal to that of  $i + 1$ 's
23:          else
24:             $donation = L$ 
25:            Assign  $p_i^{eff}$  s.t. only the remaining gain needed is transferred
26:          end if
27:        end if
28:      else
29:        Assign  $p_i^{eff} = p_C$ 
30:      end if
31:    end for
32:  end while
33: end if
34: return  $p_i^{eff}, \forall i \in C$ 

```

(which could be used for cooperative administration expenses, maintenance, or other similar purposes).

In this approach, the amount of gain that needs to be transferred in order for required but high-shifting cost agents to contribute with no losses, is provided from the most profitable ones. However, the heuristic internal price balancing method is such, that after the application, the ranking with respect to gains does not change. This means that (a) no contributor suffers losses and also (b) that those who gained more than other agents, still gain more than those specific agents. Summarizing, the internal price balancing informally provides incentive compatibility guarantees for the stated shifting cost values.

3.4.2 Internal pricing as a constrained optimization problem

It is often desired to formulate problems in a form of constrained minimization or maximization, in order to be able to apply generic solving methods. In our case, the problem of dividing the aggregate profit of the cooperative to the individuals could be optimized according to various criteria. The initial approach that was discussed in the previous subsection, had as a criterion to change as few prices as possible, without violating any constraints. Also, when the algorithm had to change the prices due to the existence of negative gains of participants, it chose to alter those of the *highest gain* agents first. Here, by adding specific constraints we can guarantee that the outcomes, despite the optimization criterion, will maintain the individual rationality and budget balancedness properties. Furthermore, each cooperative might set different optimization criteria, according to its own interests. Then, off-the-shelf, practically efficient methods can be applied to solve the optimization problems. Later, in Section 3.6.6, we study the effects that each of these optimization criteria has on the agents payoffs. In what follows, we present the general form of our optimization

problem, and five different criteria for gain balancing.

Equation (3.35) describes the general constrained optimization problem that we will be solving, but for each of the following cases, different objective functions $f(\mathbf{p}^{eff})$ ¹⁶ are going to be used:

$$\begin{aligned}
 \text{opt}_{\mathbf{p}^{eff}} \quad & f(\mathbf{p}^{eff}) \\
 \text{s. t.} \quad & \mathbf{p}^{eff} \leq \hat{\mathbf{p}} \\
 & \text{gain}(\hat{\mathbf{p}}|\mathbf{p}^{eff}) \geq \mathbf{0} \\
 & \mathbf{r}^\top \mathbf{p}^{eff} = r_C p_C
 \end{aligned} \tag{3.35}$$

The opt operator can mean either minimize, or maximize. The first constraint guarantees that each agent will pay at most its reservation price, and the second constraint, that its gain will be non-negative (*individual rationality*). The third constraint dictates that the sum of individual monetary charges will be equal to the total amount that the cooperative is charged for the shifted consumption of electricity. After using an off-the-shelf method to solve the problem described by Eq. 3.35 and thus calculate the p_i^{eff} effective prices, the actual b_i bills that agents pay are once again determined via employing Eq. 3.34. We now present the objective functions $f(\mathbf{p}^{eff})$ that are to be optimized, and form affine problems [36] with optimal solutions.¹⁷

¹⁶For ease of notation, \mathbf{p}^{eff} refers to the vector containing all p_i^{eff} s.

¹⁷By affine problems we refer to convex functions with affine equality constraints. This is fundamental in convex optimization formulations, in order to tackle infeasibility caused by the defined constraints [36].

1: Minimize maximum individual gain loss Namely:

$$\underset{\mathbf{p}^{eff}}{\text{minimize}} \max\{gain(\hat{\mathbf{p}}|\mathbf{p}^{eff}) - gain(\hat{\mathbf{p}}|\mathbf{p}_C)\} \quad (3.36)$$

This case measures the gain differences for each agent induced by the internal pricing change, and tries to minimize the maximum of these differences.

2: Maximize sum of individual gains Namely:

$$\underset{\mathbf{p}^{eff}}{\text{maximize}} \sum gain(\hat{\mathbf{p}}|\mathbf{p}^{eff}) \quad (3.37)$$

This case sums up individual gains of the cooperating agents, and strives to maximize this particular sum.

3: Minimize sum of gain loss Namely:

$$\underset{\mathbf{p}^{eff}}{\text{minimize}} \sum \{gain(\hat{\mathbf{p}}|\mathbf{p}^{eff}) - gain(\hat{\mathbf{p}}|\mathbf{p}_C)\} \quad (3.38)$$

This objective sums up the total gain loss of agents induced by the price balancing, and aims at minimizing this particular sum.

4: Minimize the maximum individual gain loss and the sum of gain loss (sum of objectives 1 and 3) Namely:

$$\underset{\mathbf{p}^{eff}}{\text{minimize}} \max\{gain(\hat{\mathbf{p}}|\mathbf{p}^{eff}) - gain(\hat{\mathbf{p}}|\mathbf{p}_C)\} + \sum \{gain(\hat{\mathbf{p}}|\mathbf{p}^{eff}) - gain(\hat{\mathbf{p}}|\mathbf{p}_C)\} \quad (3.39)$$

This case is a combination of cases 1 and 3, that is it minimizes the sum of the maximum gain loss of agents and also the sum of all agents' gain loss.

5: Minimize price differences Namely:

$$\underset{\mathbf{p}^{eff}}{\text{minimize}} \|\mathbf{p}^{eff} - \mathbf{p}_C\|_1 \quad (3.40)$$

By introducing the ℓ_1 -norm as an objective function, what we actually manage to do is to induce the minimum number of gain transfers required inside the cooperative. As such, this case is expected to produce outputs that are very similar to those of our previously proposed heuristic gain balancing scheme—with one important difference: *criterion 5* does *not* guarantee that the initial ranking of participants *wrt.* gains is maintained. The experimental simulations in Section 3.6.6 actually confirm our expectations regarding this method's behaviour.

Moreover, note that the initial gain ranking is not maintained by *any* of the objective functions proposed here, since this would require the enforcement of additional constraints that would change dynamically according to the number of participants in each shifting coalition. This fact makes the optimization problem formulation more complex, and have negative impacts on the genericness of such approaches.

3.5 Mechanism Properties

The mechanism that we described in the previous sections, although simple, exhibits certain desirable properties. First of all, participation in shifting coalitions is individually rational in expectation, since non-negative gains are guaranteed for every

coalescing agent, by the application of the constraints and the internal price balancing; the latter also allows the inclusion of even agents with high shifting costs, that otherwise could only lose by participating.

Another property that is ascertained is weak budget balancedness. Equation 3.34 in particular, dictates that no loss can be generated by the cooperative, as participants are asked to pay a proportion of the cooperative bill, respective to each one's performance and contribution. Note though, that the proportions might sum up to more than 1, so there are cases when the cooperative generates (small) profit. Also, the cost is divided "fairly" among the individuals, judging by their load quantity, their calculated effective price, and, of course, their deviation between promises and final actions, as measured by the *CRPS*.

Moreover, the use of the *CRPS* scoring rule ensures that the presented mechanism is *truthful* with respect to their shifting capacity-related statements. An agent has to be as accurate as possible regarding shifting capacity and corresponding uncertainty, as otherwise it will suffer a gain loss due to a bad *CRPS* score. The truthfulness of the agent statements regarding their shifting costs is more difficult to formally guarantee. Since the agents operate in a large, dynamic, and open environment, one cannot determine an incentive compatible mechanism in the Bayes-Nash sense, because analysing Bayes-Nash equilibria properties is computationally infeasible in such settings. Indeed, it is next to impossible for a member agent to reason on the unknown capabilities or availability of thousands of other agents, and no common prior determining such properties can be reasonably assumed. So, given this uncertainty, the best that an agent can do is to be truthful regarding its shifting costs: If the agent states inflated shifting costs, it runs the danger of not being selected

for C . Similarly, if the agent states shifting costs lower than its real ones, then it risks suffering a high reduction in expected gain (since the lower these costs are, the higher its p_i^{eff} effective price). In addition, the sheer size and dynamic nature of the problem makes it improbable that a rational consumer would be willing to utilize, on a daily basis, the resources necessary to estimate potentially beneficial fake shifting costs, in order to game the scheme. The internal price balancing mechanism makes this even more complex. In practice, the cooperative could use estimates of industry-dependent shifting cost limits, to fend off any such attempts. Moreover, if an agent is able to and chooses to stick to promised actions, this will allow him to achieve high returns.

Additionally, the proposed mechanism is fair¹⁸ for both the Grid and the contributors, in the following sense: first, the Grid is the one that sets the thresholds, safety limits, and group price functions, according to the expected savings that will come up. Next, at the consumer side, fairness is promoted in various ways:

- (a) our mechanism guarantees that there will be no loss in expectation
- (b) accurate participants achieve larger gains
- and (c) we give contribution opportunity to negative gain agents, when required

Even when internal to the coalition price balancing is applied¹⁹ the mechanism

¹⁸We use the term “fair” a bit loosely, without referring to some formal notion of fairness, such as *max-min fairness* or *the Shapley value* [39]. Note however that *CRPS* payments punish non-accurate agents and “boost” the gains of accurate ones. In this sense, *CRPS* itself could be seen as constituting a formal notion of fairness: agents realize that their final billing will take place using *CRPS*, and that their bill is proportional to what they deliver and to what others have promised to deliver and finally deliver.

¹⁹Recall that internal price balancing only takes place when negative gain agents exist, but this is not always the case. The frequency depends on the specific settings.

is still fair in the case of the heuristic approach.

Last but not least, the pricing schemes and coalition formation methods employed by the mechanism, can be readily used by cooperatives offering electricity demand management services, as they are simple enough and require no legislature changes whatsoever—in sharp contrast to *real-time pricing* approaches [23].

Computational complexity Let $|A|$ be the number of the agents in the cooperative; for each peak interval that is announced, only the $|\mathbb{C}^{t_h}| \leq |A|$ available agents communicate their information. Thus, the determination of potential contributors subject to the grids constraints, is of order $\mathcal{O}(|\mathbb{C}^{t_h}|)$. Next, we rank the $|\mathbb{C}^{t_h}|$ agents by their contribution potential $\tilde{r}_i \hat{p}_i$, paying the cost of an off-the-shelf algorithm such as $\mathcal{O}(|\mathbb{C}^{t_h}| \log(|\mathbb{C}^{t_h}|))$. The formation of the minimum coalition of $|c| \leq |\mathbb{C}^{t_h}|$ agents, as required by the coalition formation methods *CF1* and *CF2* (Section 3.3.1), involves two summations and two conditional checks, i.e., it has $\mathcal{O}(|c|)$ complexity. Now, method *CF1* needs additional $\mathcal{O}(|C| - |c|)$ calculations to expand to the maximal coalition, where $|C|$, $|c| \leq |C| \leq |\mathbb{C}^{t_h}|$, is the final number of agents in the coalition. In case negative gain agents exist and price balancing needs to be initiated, an additional ranking according to expected gain is needed, adding $\mathcal{O}(|C| \log(|C|))$ to the complexity when the heuristic internal price balancing is used. Final calculations regarding p_i^{eff} s and b_i s add a linear increase of $\mathcal{O}(|C|)$. Thus, the complexity for the planning of shifting operations during a single peak interval can be kept just linearithmic to $|\mathbb{C}^{t_h}|$ —that is, it increases less than quadratically with the number of available individuals during a peak interval t_h . However, if the cooperative employs the constrained optimization problem formulations of Section 4.2 for internal

price balancing, then the complexity rises to $\mathcal{O}(|C^3|)$, i.e. the typical complexity of quadratic programming algorithms, which can be used to solve affine problems such as ours [36].

3.6 Experimental Evaluation

In this section we conduct extensive experimental simulations of our mechanism on real consumption patterns. We first present our real world dataset, and explain how we augmented it to also contain information which is still unavailable and cannot be obtained, but is nevertheless required in order to illustrate the mechanism's performance. We then proceed to evaluate the coalition formation methods of Section 3.3.1. Following that, we examine the effects of using variable group prices and employing the *CRPS* scoring rule or not; and, finally, study the differences among our various internal price balancing techniques.

3.6.1 The simulations dataset

To experimentally evaluate our methods, we created a simulations dataset based on real electricity consumption data,²⁰ from Kissamos, a municipality at the greek island of Crete. The dataset contains hourly consumption values for the year 2012, as well as contract types and geographical locations, and a summary of its contents appears in Table 3.2. This dataset was selected because it included various consumer types, containing measurements from residential customers, to public and agricul-

²⁰The dataset was provided by the Hellenic-Public Power Company (PPC, www.dei.gr).

tural consumer types.²¹ Analytical descriptions of the datasets that we used in our simulations can be found in Appendix A. We note that participants in our setting belong to two classes of realistic, highly plausible agent behavior. First, 50% of the participants are mainly confident, and also have a high probability to deliver what they promised (they belong to the *BB* class of agents presented in Appendix A); while 50% of the agents are uncertain predictors, which might or might not follow stated forecasts. The exact definition of these two classes of agents appears in Appendix A. Note that once the necessary (due to the lack of real data) agent shifting costs are calculated, the only additional parameters that one might need to adjust for experimentation are $p_g(\cdot)$, q_{\min} , and Q_{\max} .

Table 3.2: Kissamos 2012: Size and corresponding individual average consumption and bills for each consumption contract type.

Type	Count	Avg Daily Individ. Cons. (<i>kWh</i>)	Avg Daily/ Type (<i>kWh</i>)	% Of total	Avg. Bill Cost (€) over 100 days
Residential	5889	7.294	42956.721	35.410	61.66
Commercial	1381	25.080	34636.032	28.550	213.20
Agricultural	271	111.473	30209.372	24.901	933.17
Municipal	295	13.776	4063.979	3.349	113.30
Public	68	76.361	5192.588	4.280	645.36
Industrial	38	107.257	4075.785	3.359	903.74
Public Law	12	14.921	179.053	0.147	120.13
Total	7954	-	121313.532	100	-

²¹For further validation we also run our simulations using an additional dataset that contained industrial customers only from India.

Key parameters of the shifting scheme

We assume a threshold τ for our model, to the 93% of the maximum demand value among all time intervals, and is fixed for all of them—though, it could also be variable across time intervals. The safety limit is 97.5% of τ . We report that in our simulations there are on average 6.2 peak intervals and 14.9 non-peak intervals per day. The p_l, p_h price levels correspond to the day/night retail prices provided by PPC, the greek public power company, i.e. $p_l = 0.0785 \text{ €/kWh}$ and $p_h = 0.094 \text{ €/kWh}$. The p_g group price rate ranges from $p_g^{max} = 0.05625$ to $p_g^{min} = 0.0214 \text{ €/kWh}$, depending on the reduction size q :

$$p_g(q) = \frac{0.0214 - 0.05625}{(Q_{max}^{th})^\kappa - (q_{min}^{th})^\kappa} (q^\kappa - (q_{min}^{th})^\kappa) + 0.05625 \quad (3.41)$$

with q ranging from Q_{max}^{th} , that is the amount of load above τ , to some minimum q_{min}^{th} , which we set to $0.3Q_{max}^{th}$. To also account for non-linear group pricing functions, the reduction size q , as well as Q_{max}^{th} and q_{min}^{th} , are raised to a power of κ . By assigning different values to κ , we can achieve various slopes of the pricing function. For the rest of the experiments, the value assigned to κ is 1, unless otherwise stated.

3.6.2 Evaluation of the proposed coalition formation methods

In our first set of experiments we compare the performance of the different coalition formation methods described in Section 3.3.2.²² Recall that *CF1* maximizes contri-

²²The internal price balancing method used in this and all subsequent experimental subsections up to 3.6.6, is our heuristic price balancing technique.

Table 3.3: Average results over a 100 days simulation period.

		<i>CF1</i>	<i>CF2</i>	<i>CF3</i>	<i>Random</i>
Exp. Coop. Gain (€/day)	μ	52.15	4.31	62.66	7.22
	σ	22.18	2.09	16.75	7.94
Actual Coop. Gain (€/day)	μ	33.16	-6.35	38.35	4.59
	σ	14.02	3.33	10.41	5.22
Total Gain (100 days)	Σ	3316.03	-635.63	3835.61	459.59
Coop. “Surplus” (€/day)	μ	0.18	1.13	0.43	0.09
Expected Reduction (<i>kWh</i> /day)	μ	1153.975	355.924	1341.691	216.499
	σ	496.659	155.061	372.472	235.063
Final Reduction (<i>kWh</i> /day)	μ	966.400	301.963	1125.482	190.791
	σ	410.780	132.343	308.096	206.762
Peak (Demand $\geq \tau$) Trimmed (%)	μ	71.31	22.42	83.33	15.36
	σ	23.83	8.23	3.35	17.17
Avg. Reducing Coalition Size	μ	191.11	17.98	239.22	337.14
Gain per participant (€/day)	μ	0.17	-0.52	0.16	0.01

bution potential and expected capacities, *CF2* maximizes contribution potential but meets minimum capacity requirements, *CF3* maximizes expected gain and capacities, and *Random* selects the contributors at random. Since ours is the first approach to large scale coordinated demand shifting, there exist no benchmark methods to compare ours to.

Table 3.3 shows the average results for the four methods—with each one of them applied on the same input values with the others per day, over a 100 days simulation horizon. We can observe that the most successful coalitions with respect to the amount of actual gain are formed by employing *CF1* and *CF3*. The *CF3* formation method, in particular, generates the highest amounts of expected and actual final

gains, apparently due to its ability to rank agents with respect to expected gain. In addition, *CF3* achieves larger consumption reduction (shifting) amounts, leading to higher percentage of peak load trimmed. This is due to the fact that acting coalitions sizes are on average larger when *CF3* is in use, for the reasons explained in Section 3.3.2; a fact that also means that their members errors with respect to shifting abilities are “cancelling out” better, and performance improves. Moreover, *CF3* appears to more robust overall when compared to *CF1*, since *CF3* results come with a variance that is consistently lower.

By contrast, the other two methods, *CF2* and *Random*, do not do as well. *CF2* does not seem to be effective, neither with respect to shifted quantities, nor with respect to gain. This is because *CF2* forms coalitions that appear to secure the minimum required reduction amount; however, typically the agents are not capable of delivering their promises, and as such, the coalition’s final actual gain is negative. *Random*, on the other hand, does manage to generate positive gain, but judging from the actual reduction amount (which is the minimum of all the methods), it fails to select effective and efficient participants. That was to be expected, since when *Random* is employed the contributors are selected in an entirely random fashion.

Table 3.4 contains additional information regarding the 100-day simulation results. Specifically, it presents the most active participants from each consumption category, when using the different coalition formation methods.

We see there that the most active participants, participating in the scheme dozens of times per month, gain approximately 0.03 euros per *kWh* shifted, as illustrated in Table 3.5. This performance leads to electricity bills that are reduced up to approximately 20% for the most active residential participants when *CF1* or *CF3* is

Table 3.4: Shifting performance of the most active participants per consumer class and coalition formation method, over 100 days.

Method	ID	Consumer Class	Participations	Total load shifted (kWh)	Total gain (€)	Gain/ kWh (€)
<i>CF1</i>	288	Ind.	368	5987.004	174.46	0.029
	372	Pub.	337	1832.405	43.86	0.024
	1955	Comm.	292	934.892	25.26	0.027
	7528	Res.	237	363.520	12.32	0.033
<i>CF2</i>	288	Ind.	326	5700.062	-169.68	-0.029
	372	Pub.	218	1329.302	-31.81	-0.023
	1955	Comm.	147	599.539	-15.32	-0.025
	7528	Res.	83	188.214	-3.37	-0.017
<i>CF3</i>	288	Ind.	398	6602.649	129.19	0.019
	372	Pub.	353	1841.574	47.76	0.025
	1474	Comm.	317	1072.027	25.51	0.023
	7528	Res.	280	414.244	14.00	0.033
<i>Random</i>	5287	Res.	44	5.648	0.07	0.014
	1941	Comm.	41	7.312	0.38	0.052
	349	Pub.	38	0.273	0.03	0.125
	293	Ind.	35	3.577	0.22	0.061

used, *wrt.* the average bill of the respective category. Even participants that are much less active can expect to make substantial gains from scheme participation: Table 3.6 depicts the gains of agents that participate in the scheme at least 15 times in a month. As illustrated there, these agents can reduce their monthly bill by 1.4% – 3.5% on average. Notice that these rates are comparable to discount rates commonly used in consumer rewards programs, and definitely much higher than bank deposit accounts interest rates currently (2016) in effect in most countries in Europe and North America.

Table 3.5: Monthly financial gains for the four most active participants.

Class	Avg. monthly bill (€)	Gain/ <i>kWh</i> (€)		Total gain/month (€)		Bill reduction %	
		<i>CF1</i>	<i>CF3</i>	<i>CF1</i>	<i>CF3</i>	<i>CF1</i>	<i>CF3</i>
Industrial	301.24	0.029	0.019	58.15	43.06	19.3%	14.29%
Public	215.12	0.024	0.025	14.62	15.92	6.7%	7.4%
Commercial	71.06	0.027	0.023	8.42	8.50	11.8%	11.9%
Residential	20.55	0.033	0.033	4.10	4.66	19.9%	22.67%

Table 3.6: Average gains for participants with 15 or more participations per month.

Class	Avg. monthly bill (€)	# of agents	Total load shifted (<i>kWh</i>)	Gain/ <i>kWh</i> (€)	Total gain (€)	Bill reduction %
<i>CF1</i>						
Residential	20.55	559	16.776	0.037	0.62	3.0%
Commercial	71.06	109	42.635	0.031	1.34	1.8%
Public	215.12	11	114.004	0.027	3.13	1.4%
Industrial	301.24	25	173.025	0.031	5.37	1.7%
<i>CF3</i>						
Residential	20.55	559	19.828	0.036	0.72	3.5%
Commercial	71.06	109	47.507	0.032	1.54	2.1%
Public	215.12	11	119.567	0.029	3.51	1.6%
Industrial	301.24	25	190.272	0.026	5.12	1.7%

In the case of *CF2*, the numerical results of Table 3.4 indicate once again that it is not an appropriate method for coalition formation, as even the most active consumers lose from participation. Finally, the use of *Random* does not grant losses to the most active participants, however the gain is very small for every consumer type, indicating that its use does not provide adequate incentives for the shifting of electricity demand.

To further validate our scheme we also conducted simulation experiments on an additional dataset that included only industrial customers from India, which was derived after the statistical analysis of 36 small and medium scale actual industrial facilities. This dataset has been previously used also in [96]. We can see that the results are similar (see Table 3.7).

Table 3.7: Average results over a 100 days simulation period for the dataset with only industrial consumers.

		<i>CF1</i>	<i>CF2</i>	<i>CF3</i>
Exp. Coop. Gain (€/day)	μ	2094.72	415.73	2139.38
	σ	895.56	-273.89	933.24
Actual Coop. Gain (€/day)	μ	348.03	77.30	361.66
	σ			
Coop. "Surplus" (€/day)	μ	6.20	14.81	6.13
Expected Reduction (<i>kWh</i> /day)	μ	32856.45	12868.40	32919.06
Final Reduction (<i>kWh</i> /day)	μ	24454.32	9471.053	24539.96
	σ	8083.3	2556.5	8058.2
Peak (Demand $\geq \tau$) Trimmed (%)	μ	98.61	42.51	98.63
	σ	0.75	5.86	0.75
Avg. Reducing Coalition Size	μ	47.70	15.49	49.06
Gain per participant (€/day)	μ	18.77	-6.44	19.02

In conclusion, employing *CF3* appears to be the most profitable, effective and

robust of our proposed coalition formation techniques. Also, *CF1* is another effective demand-reducing coalition formation method, which could also be considered as an alternative. For the rest of the experiments, however, *CF3* is our coalition formation method of choice, unless otherwise stated.

3.6.3 Assessing the effect of different group pricing slopes

We now apply different κ values to the $p_g(q)$ pricing function of Eq. 3.41, and observe their impact on the total cooperative gain and the percentage of peak-load trimmed. Changing κ from negative to positive values results to different non-linear forms, and to concave and convex curves respectively, as illustrated in Fig. 3-3. For each κ value we used the *CF3* coalition formation method, and the *Conservative* approach for estimating the reduction capabilities of the participants. The results over 100 days, are presented in Table 3.8.

Table 3.8: Results from an 100 days simulation for different values of κ in the group pricing function.

		$\kappa = -4$	$\kappa = -1$	$\kappa = 1$	$\kappa = 4$	$\kappa = 10$
Peak (Demand $\geq \tau$) Trimmed (%)	μ	82.86	81.65	81.36	81.47	81.50
	σ	3.17	4.18	3.57	3.98	3.23
Total Gain (100 days)	Σ	4951.71	4641.15	3835.61	2714.37	1397.49

As we can see, for $\kappa = -4$, where $p_g(q)$ reaches low values faster, the total cooperative gain grows larger when compared to the other four cases. Also, the average percentage of peak-demand trimmed is slightly higher, with lower standard deviation. This means that this particular $p_g(q)$ form is the most incentivizing for the

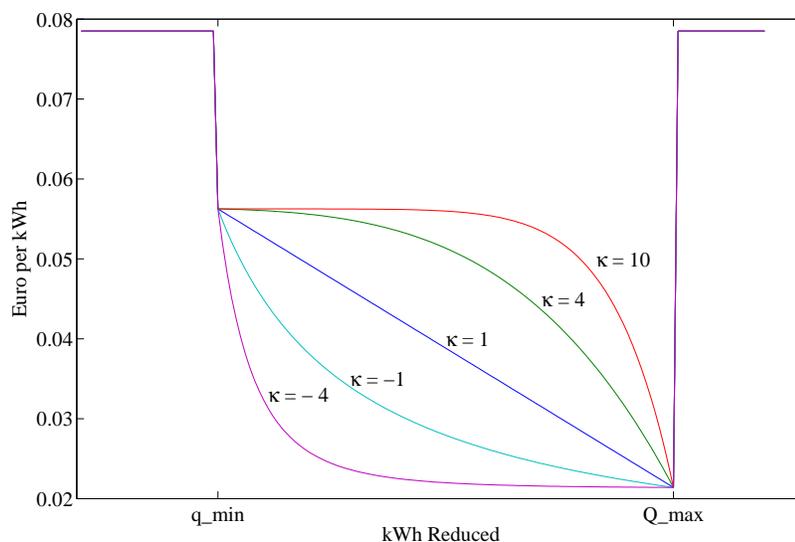


Figure 3-3: Forms of p_g for different values of κ .

consumer side, and respectively demands larger discounts from the Grid's side. As the κ value increases, the total cooperative gain become gradually smaller, without significant changes in the percentage of the trimmed peak-load. In the rest of our experiments below, we adopt the $p_g(q)$ given by $\kappa = 1$ as the middle ground solution.

3.6.4 Coalition size vs. group price range

Overall, it is clear that in order for shifting to take place, the Grid must grant a p_g range that provides enough gain to the agents, in order to overcome the underlying individual shifting costs. Here we study the dynamics associated with this p_g range selection. Specifically, we examine the average reducing coalition size formed at each t_h , given variable p_g prices granted for collective consumption shifting. To do so, we simultaneously add an offset to both the p_g^{max} and group p_g^{min} values produced by

Eq. (3.41), with the offset ranging from -0.015, to +0.025 of their initial values; then, following formation, we observe the average number of agents in reducing coalitions for each peak interval.

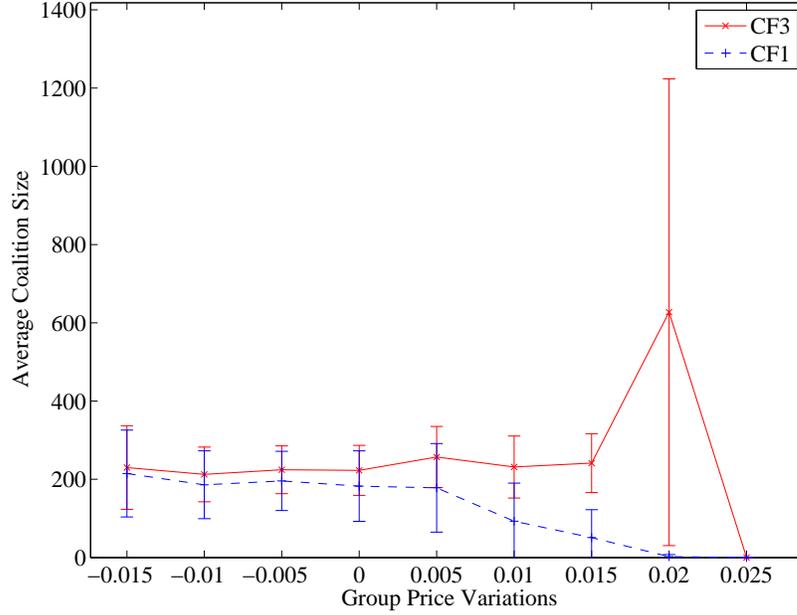


Figure 3-4: Average coalition size vs. p_g increase for the coalition formation methods *CF1* and *CF3*.

Figure 3-4 demonstrates this concept, where average coalition sizes over 50 simulation days are plotted against group price range variations. If the *CF1* method is applied, it is obvious that as p_g increases to get closer to p_{low} , fewer agents decide to contribute—and, subsequently, less consumption is finally shifted. In the case of *CF3*, however, we observe that the average coalition size is in general more stable; moreover, for a $p_g^{max} = 0.05625 + 0.015 = 0.07125$ and above (i.e., for p_g^{max} values close to p_l), the *CF3* mean and standard deviation of the average coalition

size increases, in contrast to *CF1*—and both methods reach the value of zero for the maximum offset added (i.e., when p_g^{max} exceeds p_l and p_g^{min} is also high, it is no longer possible to form profitable coalitions.

This difference in the behaviour is explained as follows. First, recall that *CF1* ranks contributors according to their contribution potential ξ_i , not accounting for the expected monetary gains to be created. This is acceptable in settings where agents have low shifting costs, and the formation of shifting coalitions is relatively easy, i.e. there is abundance of shifting capacities for the offered prices. On the other hand, when the better price p_g is high, it becomes harder to come up with a coalition of contributors offering large shifting quantities at such high prices. Thus, it becomes more probable to create infeasible coalitions that violate one of the constraints of Eqs. (3.25)—(3.27) during the *CF1* process. Hence, the average coalition size gradually drops to zero.

In the case of *CF3*, however, contributors are ranked according to their expected gain, which is a very good indicator of the coalition’s feasibility potential. This method is able to guarantee the coalition’s feasibility, even if a larger number of participants is required in order to achieve it. This is in fact illustrated in Figure 3-4: when few agents in the population can profit directly from the granted p_g price, the required shifting amount is more difficult to gather, and more agents have to join in the coalition to do so.

3.6.5 Assessing the CRPS effect

In this set of experiments, we gradually increase the relative errors of the agents and measure the drop in their final gains after applying the *CRPS* scoring rule.

Specifically, the relative error of each agent was progressively increased from 0 to 1 by 0.1 over 11 complete runs of 20 simulation days each. This naturally leads to a higher (i.e., worse) *CRPS* score for the individual agents, and thus a bad *CRPS* score for their corresponding reducing coalitions. Figure 3-5 plots together the *CRPS* score and losses in gain,²³ for the average individual. We observe that as an individual agent's *CRPS* score gets worse, its gain losses increase. In the vertical axis, we present the relative difference d between the final bill paid b_i (calculated *after* the *CRPS* rule has already been applied) and the $p_i^{eff}q_i$ expected to be paid at the time of the formation of the coalition (i.e., *before* applying *CRPS*):

$$d = \frac{b_i - p_i^{eff}q_i}{p_i^{eff}q_i} \quad (3.42)$$

Notice that this relative difference might be much larger than 2 for a “highly inaccurate” agent. This is because, firstly, if the difference between the statements and the final actions is very high, applying the *CRPS* can double the price (cf. Eq. (3.15)). Secondly, rising individual *CRPS* scores correspond to reduced cooperative effectiveness, and thus the final effective group price (and cooperative bill) is higher than anticipated, resulting to lower-than-expected gains.

It is therefore clear that *CRPS* can induce substantial “penalties” on erroneous agents and coalitions. Thus, employing *CRPS* provides definite incentives for the agents to produce accurate statements regarding their shifting capacities.

²³Note that this does not mean the average gain is negative, just that it decreases as a result of *CRPS*-generated penalties.

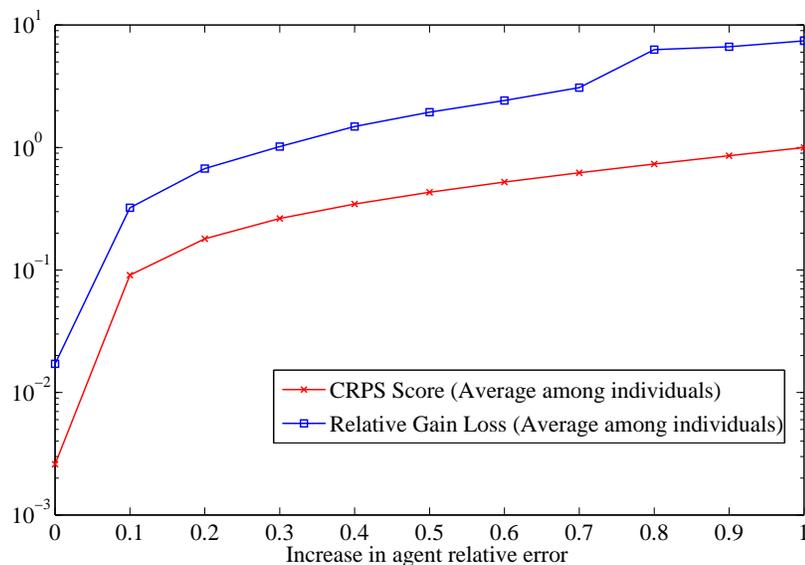


Figure 3-5: Average losses in gain (increase in bill) due to *CRPS*, as induced by increasing participant inaccuracies, across all intervals.

CRPS penalizing only cooperative performance In this part we also test a case where the *CRPS* is applied for producing the *overall* cooperative bill (Eq. (3.31)) only, but *not* for determining the individual bills of the participants. That is, instead of using Eq. (3.34), the final bill of the individuals b_i is calculated by

$$b_i = \frac{p_i^{\text{eff}} q_i}{(\sum_{j \in C} p_j^{\text{eff}} q_j)} B_C \quad (3.43)$$

where B_C is given by Eq. (3.31).

Table 3.9 summarises the differences between skipping internal penalization (via using Eq. (3.43)) and not. The first thing to observe is that the “cooperative surplus” is non-existent when *CRPS* is not applied internally, as expected since the

exact amount of the cooperative bill is now split among the participants (charged via Eq. (3.43)). Then, to better illustrate the desired effect of the internal *CRPS* penalization, we measure the difference between expected and final gain for agents with *CRPS* score higher than 0.2 (which we can consider as inaccurate participants) and for those with score lower than 0.2 (accurate participants). As shown in the middle column of Table 3.9, when there is no internal penalty applied to cooperative members, there exists a difference in final gains and in “losses *wrt.* expected gains” between the highly erroneous and less erroneous agents. However, when applying internal individual penalization, these “losses” for the erroneous participants rise, while for the non-erroneous ones the average final gain exceeds the value of the expected gain. Note that the expected gain is given by Eq. (3.30), that does not preclude the possibility that initial gain estimates are pessimistic. Thus, making more gains than expected when being accurate, in the presence of highly inaccurate participants is not overly surprising. This exacerbates the differences in gain transfers between these two agent classes: it really pays to be truthful and accurate. In a nutshell, applying the *CRPS* “internally”, clearly incentivizes the individual agents to be accurate and deliver what they promised.

3.6.6 Experimenting with different internal price balancing techniques

In this subsection, we compare the performance of the various optimization criteria of Section 3.4 with respect to resulting price assignments and resulting gain per participant. We employ the *cvx* toolkit of *Matlab* to obtain the solutions of the

Table 3.9: Average differences between expected and final gain per contribution, with and without internal *CRPS* penalization.

All surplus and gain values average over participations	Without internal penalty	With internal penalty
Cooperative “surplus” (€/day)	0	0.10
Exp. gain per participation for agents with <i>CRPS</i> > 0.2	0.03132	0.03132
Final gain per participation for agents with <i>CRPS</i> > 0.2	0.02703	0.02646
Difference	0.00429	0.00486
Exp. gain per participation for agents with <i>CRPS</i> ≤ 0.2	0.03076	0.03076
Final gain per participation for agents with <i>CRPS</i> ≤ 0.2	0.03060	0.03095
Difference	0.00016	-0.00019

various mathematical programming methods of Section 3.4.2. All different pricing methods operate on the same input data. In this way, the difference in behaviour occurring by using the various criteria is clearly demonstrated, as shown in Figures 3-6 and 3-7. The x -axis shows the agent IDs coalescing to shift demand at a t_h , in this case from 1 to 68. Note that agent IDs are ranked in *descending order* with respect to the gains granted by considering p_C as the price paid.

Figure 3-6 in particular, depicts the participant \hat{p}_i s and p_i^{eff} s, where the latter are calculated with the various methods that we explained earlier. As we observe, certain pairs of criteria perform similarly: specifically, criteria 2 (maximize sum of individual gains) and 3 (minimize sum of gain losses), criteria 1 (minimize max individual gain loss) and 4 (sum of criterion 1 and criterion 3), and our proposed heuristic balancing and criterion 5 (minimize price differences).

The actual effects of variable pricing can be seen in Fig. 3-7, where we present the gains for each agent, when assigning the calculated \mathbf{p}^{eff} ; also, the gain before we

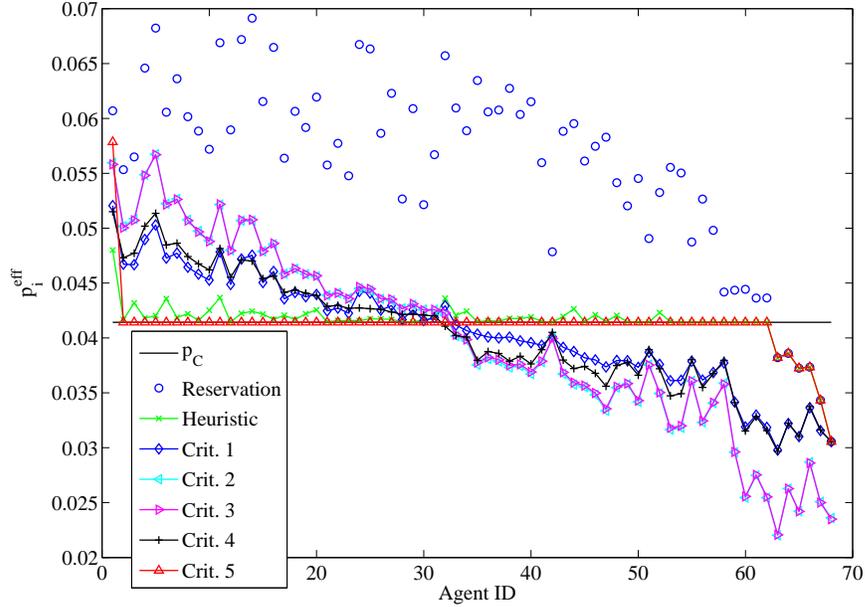


Figure 3-6: Prices (€/kWh shifted) assigned to each individual for a sample peak interval as a result of employing our pricing methods.

perform internal price balancing is shown with the yellow curve. It is clear from the figure that optimisation methods 2 and 3 tend to dispense the amount of gain equally (within some tolerance levels) among the contributors. This can be considered as a violation of the mechanism's fairness, as participants end up enjoying the same amount of gain regardless of the value of their individual contribution.

By contrast, the heuristic price balancing technique (green curve), maintains the ranking with respect to the gain amount, and so does optimization criterion 5, except for the higher gain participant. The performance of criterion 5 is mainly due to the l_1 norm application, which tries to change values in as few p_i^{eff} s as possible. Notice however that, unlike the heuristic balancing method, this criterion does not guarantee that the original agents ranking *wrt.* gain is maintained after balancing. Criteria 1

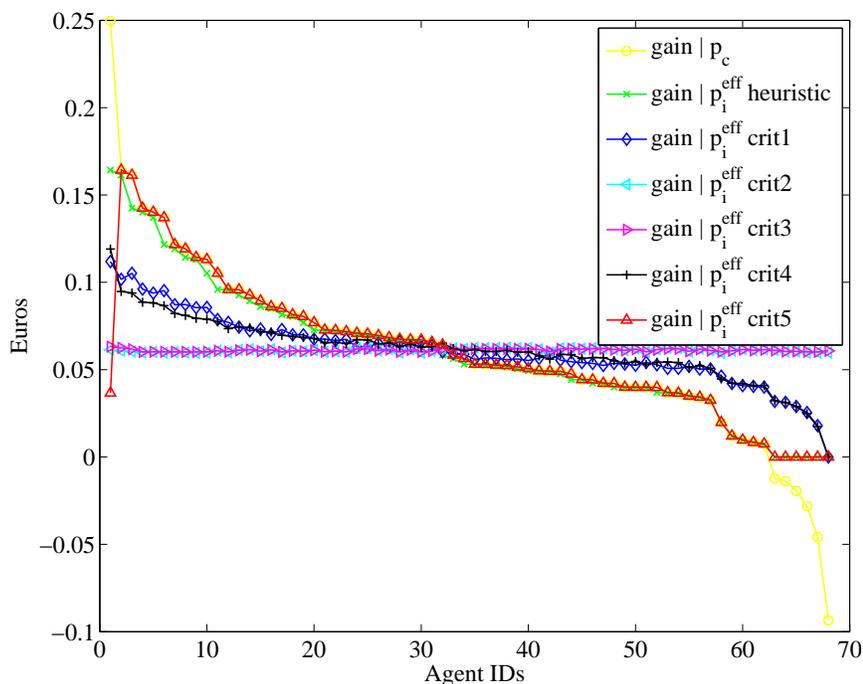


Figure 3-7: Expected gains for each participant for a sample peak interval, when applying different reward sharing approaches.

and 4, on the other hand, appear to maintain this ranking, but are more generous to lower ranking agents (thus, the gain-related “gap” or “social distance” among the agents appears to be closing up when using these methods).

Last but not least, it is important to note that no gain is lost in expectation when adopting any of the proposed internal pricing methods. As we observe in Fig. 3-7, despite the existence of (originally) negative gain agents (the yellow curve for which lies below zero), there are no negative expected gains for any of the participants after internal pricing (regardless of the specific method used).

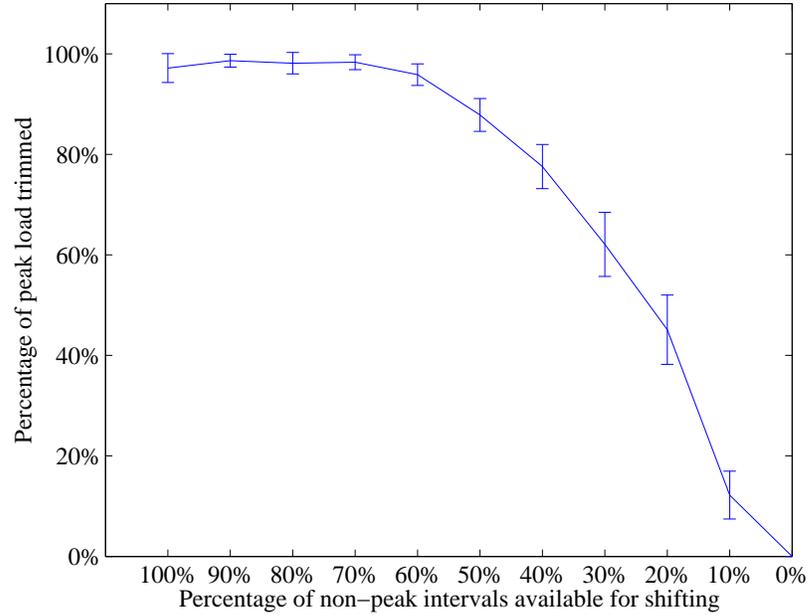


Figure 3-8: Cooperative effectiveness versus contributor availability.

3.6.7 Participant availability

Finally, we examine the relationship between agent availability and cooperative effectiveness. In particular, we ran the simulator for 11 iterations, each lasting for 30 days. Now, in each iteration, we raise the probability of an agent being unavailable for participation during a time interval.

Results are shown in Figure 3-8, where we can observe that the cooperative achieves nearly 100% peak load trimming effectiveness when agents are available for at least 70% of the time. After that point, cooperative effectiveness drops in a nearly linear manner proportionally to the percentage of time intervals during which agents are available. Thus, the higher the participants availability, the greater the effectiveness of the shifting scheme.

3.7 Conclusions

In this chapter, we presented a complete framework for large-scale cooperative electricity consumption shifting, to promote the proactive balancing of the demand curve. Our proposed shifting scheme is directly applicable, and promotes agent efficiency in the face of uncertainty. No additional infrastructure or specialised equipment is required for deployment: off-the-shelf smart metering and transmission equipment can be readily employed, and the computational complexity of the overall mechanism planning the cooperative shifting operations is low. This is achieved via the use of effective coalition formation methods we developed; and via employing the *CRPS* strictly proper scoring rule, which incentivizes truthful and accurate forecasts. Furthermore, our mechanism is equipped with internal pricing schemes that employ gain transfers within a cooperative, to make it worthwhile for individuals to participate in shifting operations and thus guarantee the scheme's effectiveness and profitability. Our mechanism possesses desirable theoretical properties: individual rationality, truthfulness, and (weak) budget balancedness. We ran extensive simulations based on real consumption data, and demonstrated experimentally the effectiveness of our methods. The results of our experiments confirm that our methods could bring tangible benefits to energy cooperatives and other Smart Grid business entities.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 4

Decentralized Cooperative Demand-side Management for Prosumers

In the previous chapter, we showed how it is possible for an electricity consumer to reduce her expenses by rescheduling her energy usage to the most profitable intervals, by setting higher electricity price values for buying energy during intervals of high demand, and lower values during intervals of low demand [7]. This is a task that becomes even more important (and challenging) when it comes to *electricity prosumers*. As prosumers both produce and consume energy [21, 184], they can take advantage of fluctuations in prices, and generate even more profit [111].

However, increased participation to DSM schemes often leads to herding effects. As such, the estimated consumption curve could significantly change, both endangering the Grid's stability, and leading to substantially different economic

outcomes [195]. For this reason, the formation of consumer cooperatives or virtual power plants has been proposed [7, 23, 96, 195], an approach which, however, requires a centralized entity to serve as the cooperative manager. To overcome both herding effects and the need for cooperative manager determination, in this thesis we champion the use of a purpose-designed cryptocurrency protocol for distributed prosumer cooperative coordination. As discussed in Chapter 2, cryptocurrencies and blockchain-oriented algorithms run distributedly, and are transparent. Additionally, they use encryption methods, which guarantee that the transactions are secure, and that no third-parties need to take part in the exchanges [62]. A first generation cryptocurrency protocol has already been used in a setting with electricity prosumers—and, to the best of our knowledge it is the only Smart Grid related work, and is called NRGcoin [122]. Although it incentivizes demand and production balancing, that protocol does not promote large-scale cooperative consumption shifting. In our work, we envisage a next-generation, special-purpose cryptocurrency software, which is executed by each cooperative member in a decentralized fashion, and is used for coordinating electricity consumption shifting actions and the sharing of the rewards.

Here, for the first time, we show how we can combine cryptocurrency with mechanism design for cooperatives formation, to achieve large-scale coordinated shifting of electricity prosumers consumption. The cooperative shifting activities result to increased prosumer profits from electricity trading. Using a cryptocurrency protocol, prosumers autonomously create a *virtual* wholesale mediator between the end-users and the Grid. The protocol takes into account prosumer shifting capacity statements, and distributes personalized rewards given the final collective profits achieved, and the cooperative’s profits sharing policy of choice. The coins awarded represent shares

on the total cooperative profits.

Summarizing, the work described in this chapter contributes to the state of the art as follows.

- First, we model *prosumers* in a market setting with variable prices, and present a *distributed consumption shifting approach for prosumer cooperatives*, which guarantees monetary gains to the participants.
- We apply a *novel cryptocurrency model* for the coordination and management of the cooperative shifting actions. In the proposed model, the rewards from prosumer participation are determined in a personalized manner, in the form of newly mined coins.
- We examine different coin mining methods, and champion one that evaluates prosumers via a scoring rule [67] assessing the difference between promised and final actions. To the best of our knowledge this is the first time that cryptocurrency mining and scoring rules are combined into one method. By penalizing inaccuracy, this method incentivizes prosumers to provide truthful promises.
- We propose specific formation techniques, which select members for participation in cooperative actions.

The contributions of this chapter and the combined scientific fields are also shown in Figure 4-1.

Our approach can be applied in conjunction with any existing regulations or pricing schemes. We evaluate our scheme experimentally on a large dataset that ex-

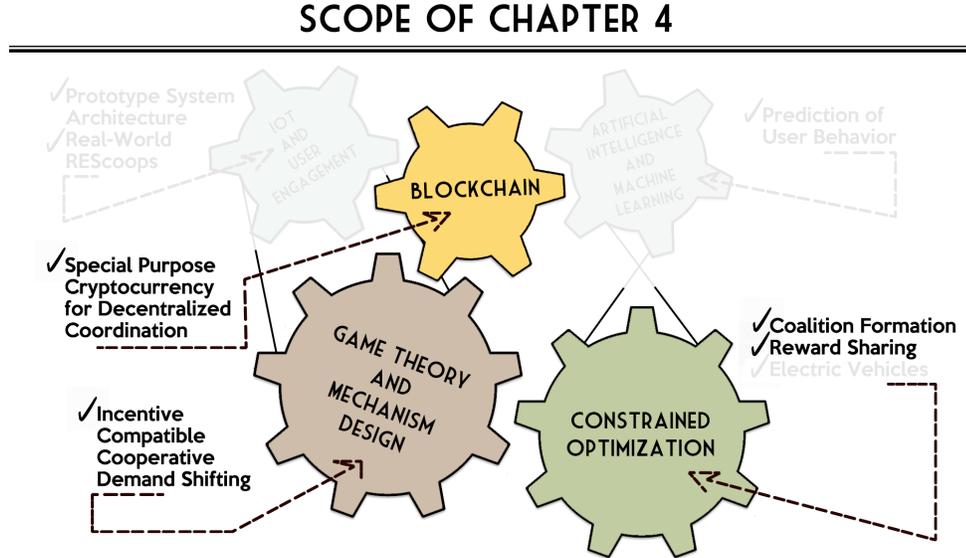


Figure 4-1: Overview of the scientific fields, and our contributions in Chapter 4.

tends over a one-year period, and which is based on real consumption and renewable production data. Simulation results confirm that adopting our mechanism leads to increased profits for the cooperative participants, stabler variable electricity prices, and achieves lower Peak-to-Average Ratio (PAR) values for the difference between electricity supply and demand. Especially when using a scoring rules-based reward redistribution method, accuracy is explicitly incentivized with increased gains for the accurate participants.

This chapter is structured as follows: In Section 4.1 we present the system setting and the individual prosumer financial decisions model. Section 4.2 presents the cooperative model, the cryptocurrency protocol and three different approaches for personalized reward sharing, as well as methods for contributor selection for cooperative actions. Section 4.3 presents the experimental results, and, finally, Section 4.4

concludes this chapter.

Parts of the research described in this chapter appeared originally in [9].

4.1 A Generic Prosumer Consumption Shifting Model

We assume a setting encompassing of *prosumers*, which both import and export energy from and to the Grid; and a *prosumer cooperative*, which is a large coalition of prosumers trading energy as a unique entity. The Grid is regulated by the *distributed system operator (DSO)*, responsible for the transmission of energy and its pricing.

Actors need to take decisions regarding trading at some day-ahead electricity market, or consuming electricity at specific 1 to T intervals during the course of a day (the day-ahead).¹ For ease of notation, we slightly differentiate variables from those presented in Chapter 3, and move the time interval index t to be the subscript of the variable that denotes energy amounts. Each actor i is characterized by the amount of electricity (kWh) imported $q_{i,t}^-$, and the amount exported $q_{i,t}^+$, during the time interval t . The aggregate demand and supply levels for each time interval are given by

$$Q_t^- = \sum_i q_{i,t}^-$$

and

$$Q_t^+ = \sum_i q_{i,t}^+$$

in kWh.

¹A decision theoretic optimization approach for a single prosumer operating in such a setting was proposed by [21, 22]. However, they did not include cooperative electricity trading, nor dealt with consumption shifting.

We assume reliable renewable production and demand forecasting techniques that can achieve high precision—lower than 2% mean absolute percentage error [81, 187]. Predictions are noted as $\tilde{q}_{i,t}^-$, and $\tilde{q}_{i,t}^+$ kWh, for imports and exports respectively. The predicted demand and supply for the planning horizon are noted as \tilde{Q}_t^- for the total imports, and \tilde{Q}_t^+ for the total predicted exports.

4.1.1 Promoting Demand-Side Management

Many methods have been proposed for modelling individual consumption profiles [68, 146, 194]. In our work, we examine the rescheduling of *shiftable loads*, which are those loads that it is possible to shift to in later or earlier time intervals, with minimum impact on the consumer’s well being, e.g., battery charging, water-heaters, washing-machines, etc. Now, to promote demand side management operations, prosumers should be offered better prices to counterbalance the associated shifting costs. Following existing dynamic pricing mechanisms, which promote the balancing of demand and renewable energy supply [122], we assume that *billing functions* are in place (by the DSO) for selling $B_t^{sell}()$, and buying $B_t^{buy}()$ energy to/from the Grid, each with specific properties. First, they are functions of the quantity of energy produced $q_{i,t}^+$ and consumed $q_{i,t}^-$, respectively. Next, and in order to satisfy the supply and demand balancing requirements [182], both also need to be functions of aggregate supply, Q_t^+ , and demand, Q_t^- . Specifically, $B_t^{sell}()$ must take maximum values for fixed $q_{i,t}^-$ s and $q_{i,t}^+$ s, during intervals when $Q_t^+ = Q_t^-$. This incentivizes prosumers to produce exactly the quantity that is required for consumption (since their income is then maximized). Intuitively, it is to the DSO’s interest that prosumers decide to sell when $Q_t^+ = Q_t^-$, since this defers the need to import or export energy.

Assumption 1 ([122]). The pricing for selling energy to the Grid during a time interval t , is a function of the sold quantity, the aggregate quantity produced, and the aggregate quantity consumed during that interval, $B_t^{sell}(q_t^+, Q_t^+, Q_t^-)$; and for fixed q_t^+ , it is maximized as:

$$Q_t^+ - Q_t^- \rightarrow 0$$

Note that, assuming the electricity production of prosumers originates mainly from wind turbines and photovoltaic panels, the quantity produced q_t^+ cannot be easily controlled [182]. Moreover, the selling prices are also functions of aggregate demand Q_t^- , which we later optimize by shifting consumption tasks in a large-scale cooperative manner.

The $B_t^{buy}()$ on the other hand, should produce lower prices with higher renewable production excess, prompting prosumers to buy energy from the Grid (and perhaps store it for future use). Intuitively, it is more efficient to consume the cheap renewable energy produced locally, than import from some external balancing market where prices are in general far worse [182]. This is because exporting or importing electricity involves additional expenditures, e.g. transmission lines, electricity resellers, etc. By contrast, $B_t^{buy}()$ produces higher values as renewable energy supply decreases.

Assumption 2 ([122]). The pricing for buying energy from the Grid during a time interval t , is a function of the acquired quantity, the aggregate quantity produced, and the aggregate quantity consumed during that interval, $B_t^{buy}(q_t^-, Q_t^+, Q_t^-)$; and for fixed q_t^- , it is minimized as:

$$Q_t^+ - Q_t^- \rightarrow +\infty$$

4.1.2 Shifting to profitable time intervals

Given this model, a prosumer can control the quantity consumed during time intervals by shifting consumption tasks. We now characterize each time interval as *peak* or *non-peak*. Peak intervals t_h are those intervals during which reducing consumption can be considered profitable for the prosumer. Specifically, due to Assumptions 1 and 2, this happens when aggregate demand is higher than supply:

Definition 4.1.1 (Peak intervals t_h). Consider a non-negative threshold τ . A time interval t is considered to be a peak interval, t_h , if:

$$\tilde{Q}_t^+ - \tilde{Q}_t^- < \tau$$

Non-peak intervals t_l are those intervals during which, increasing consumption levels up to the reduced amount of energy that was decreased during t_h , results to lower expenses due to a reduced buying price. Specifically, due to Assumptions 1 and 2, this happens when demand is lower than supply:

Definition 4.1.2 (Non-peak intervals t_l). Consider a non-negative threshold λ . A time interval t is considered to be a non-peak interval, t_l , if:

$$\tilde{Q}_t^+ - \tilde{Q}_t^- > \lambda$$

Intuitively, variables τ and λ correspond to load difference thresholds that allow profitable shifting actions. Their values can be based on the statistics of \tilde{Q}_t^+ and \tilde{Q}_t^- , according to each actor's business goals.

Now, we assume that each prosumer can alter her baseline demand value q_i^- . More specifically, during peak intervals prosumers can reduce down to $q_{i,t_h}^- - \hat{r}_i^{t_h}$; while for the non-peak intervals consumption can be increased up to $q_{i,t_l}^- + \hat{r}_i^{t_h}$ where $\hat{r}_i^{t_h}$ is the *stated reduction capacity* of each actor i . Also, as in [10], there is a *shifting cost* $c_i^{t_h \rightarrow t_l}$ associated with shifting from a peak to a non-peak interval. The *actual reduction capacity* $r_i^{t_h}$, refers to the load that is reduced during a t_h , and is shifted to some other, non-peak interval t_l . These values can be obtained by using appropriate smart metering equipment.

We now discuss the price differences induced when a prosumer shifts $\hat{r}_i^{t_h}$ from a peak interval to a non-peak interval. When prosumers decrease their consumption during peak intervals, they accrue gains from the lower buy prices, and the higher sell prices. Namely, the *estimated profit* by the induced price variations for reducing at t_h is given by:

$$\begin{aligned} profit_i^{t_h}(\hat{r}_i^{t_h}) &= B_{t_h}^{sell}(\tilde{q}_{i,t_h}^+, \tilde{Q}_{t_h}^+, (\tilde{Q}_{t_h}^- - \hat{r}_i^{t_h})) - B_{t_h}^{sell}(\tilde{q}_{i,t_h}^+, \tilde{Q}_{t_h}^+, \tilde{Q}_{t_h}^-) \\ &+ B_{t_h}^{buy}(\tilde{q}_{i,t_h}^-, \tilde{Q}_{t_h}^+, \tilde{Q}_{t_h}^-) - B_{t_h}^{buy}((\tilde{q}_{i,t_h}^- - \hat{r}_i^{t_h}), \tilde{Q}_{t_h}^+, (\tilde{Q}_{t_h}^- - \hat{r}_i^{t_h})) \end{aligned} \quad (4.1)$$

The result from subtracting the second term in Eq. (4.1) from the first, indicates the profit from the price differences for selling energy; selling during a peak interval t_h , with lowered aggregate demand, $(\tilde{Q}_{t_h}^- - \hat{r}_i^{t_h})$, grants better prices than with the initial demand, $\tilde{Q}_{t_h}^-$ (cf. Assumption 1 & Definition 4.1.1 above). Now, the last two terms

give the difference in the bill that the prosumer will *pay* for consumption during t_h . Thus, to calculate this quantity, we subtract the billing payed by the prosumer for the reduced consumption $(\tilde{q}_{i,t_h}^- - \hat{r}_i^{t_h})$ from the initial estimated bill $B_{t_h}^{buy}(\tilde{q}_{i,t_h}^-, \tilde{Q}_{t_h}^+, \tilde{Q}_{t_h}^-)$.

Similarly, the *estimated loss* generated by increasing consumption during non-peak intervals t_l is given by:

$$\begin{aligned} loss_i^{t_l}(\hat{r}_i^{t_h}) &= B_{t_l}^{sell}(\tilde{q}_{i,t_l}^+, \tilde{Q}_{t_l}^+, \tilde{Q}_{t_l}^-) - B_{t_l}^{sell}(\tilde{q}_{i,t_l}^+, \tilde{Q}_{t_l}^+, (\tilde{Q}_{t_l}^- + \hat{r}_i^{t_h})) \\ &+ B_{t_l}^{buy}((\tilde{q}_{i,t_l}^- + \hat{r}_i^{t_h}), \tilde{Q}_{t_l}^+, (\tilde{Q}_{t_l}^- + \hat{r}_i^{t_h})) - B_{t_l}^{buy}(\tilde{q}_{i,t_l}^-, \tilde{Q}_{t_l}^+, \tilde{Q}_{t_l}^-) \end{aligned} \quad (4.2)$$

To calculate the *estimated gain* for an actor i , for shifting from a t_h to a t_l , we subtract the estimated loss at t_l and the shifting costs $c^{t_h \rightarrow t_l}$ per kWh from the estimated profit at t_h :

$$g_i^{t_h \rightarrow t_l}(\hat{r}_i^{t_h}) = profit_i^{t_h}(\hat{r}_i^{t_h}) - loss_i^{t_l}(\hat{r}_i^{t_h}) - \hat{r}_i^{t_h} c_i^{t_h \rightarrow t_l} \quad (4.3)$$

Definition 4.1.3 (Eligible interval pairs). Eligible shifting interval pairs for a prosumer i are those (t_h, t_l) pairs for which the gain associated with the shifting is positive, i.e.:

$$g_i^{t_h \rightarrow t_l}(r_i^{t_h}) > 0$$

where $r_i^{t_h}$ is the actual quantity of the shifted consumption.

Summarizing, the strategy for individual consumption rescheduling is to find those shifting interval pairs for which the estimated gain is maximized, and shift accordingly.

4.1.3 Shifting without coordination

When optimizing individually, each agent does not take into account other agent rescheduling actions, and considers their consumption to be the baseline. Thus, the optimizer can exhaustively calculate the $g_i^{t_h \rightarrow t_l}(\hat{r}_i^{t_h})$ values for each shifting interval pair, for a stated reduction capacity $\hat{r}_i^{t_h}$. Then, rescheduling takes place (e.g., shifting to the most profitable ones). However, without coordination or constraint enforcements, and since every prosumer optimizes individually, herding effects take place, resulting to substantially different prices during the intervals with the lowest/highest prices, than those anticipated by the prosumers. These “unexpected” price fluctuations are not in favor of the prosumer, as the estimated gains can end up turning to losses:

Lemma 3. If every participant is rational, and billing follows Assumptions 1 and 2, optimizing the rescheduling of consumption individually does not guarantee positive final gains.

Intuitively, Lemma 3 states that non-coordinated shifting actions in such settings cannot be expected to always lead to monetary gains for the participants, and cooperation is essential. It is straightforward to show this, considering the fact that estimated gains are calculated based on the values $\tilde{Q}_{t_h}^- - \hat{r}_i^{t_h}$ and $\tilde{Q}_{t_l}^- + \hat{r}_i^{t_h}$, which are used in the first and last term of Eq. (4.1), and the second and third term of Eq. (4.2), respectively. However, since participants are rational, every one acts the same manner, resulting to substantially different values finally realized, i.e. final $B^{sell}()$, $B^{buy}()$ prices are calculated using $\tilde{Q}_{t_h}^- - (\hat{r}_i^{t_h} + \sum_{j \in C \setminus i} \hat{r}_j^{t_h})$ and $\tilde{Q}_{t_l}^- + (\hat{r}_i^{t_h} + \sum_{j \in C \setminus i} \hat{r}_j^{t_h})$, resulting to lower $B^{buy}()$ and higher $B^{sell}()$. Moreover, if the total shifting capacity is

not constrained, the conditions in Def. 4.1.1 and 4.1.2 can stop holding, rendering the intervals ineligible for profitable shifting.

4.2 Distributed Shifting and Reward Sharing

Now, cooperatives can be key for the effective coordination of consumption shifting actions [7]. Here we describe the workings of *prosumer* cooperatives, allowing members to both sell and buy energy as a single entity. We assume that cooperative members share common estimates regarding the total production and consumption per interval $\tilde{Q}_{t_h}^+$ and $\tilde{Q}_{t_h}^-$ (obtained, e.g., via the summation of communicated individual estimates). Also, participants execute a novel cryptocurrency mechanism, allowing for distributed management, transparency, and personalized rewards. An overview of the network structure and information exchange is shown in Fig. 4-2. Power and information flow from the substations to the residences, and the residences are connected with additional information links among each other. The substations execute an external protocol (e.g. NRG coin [122]), which runs independently from the cooperative blockchain protocol (termed as COOPcoin). As in any blockchain-oriented algorithm, agent participants share the same information, for example individual agent statements $\hat{r}_i^{t_h}$, $\hat{\sigma}_i$, pricing functions $B_t^{sell}()$ and $B_t^{buy}()$, and aggregate estimates \tilde{Q}_t^+ , \tilde{Q}_t^- , etc., thus, they perform the same calculations. Shared information is encrypted due to privacy concerns and it is assumed that direct mapping between a given cryptocurrency peer and its physical location cannot be achieved. Note that DSO substations do not have to actively participate in the cooperative blockchain scheme, since this is built on top of, and independently from any third party protocol

that the DSO might be adopting. The cooperative mechanism awards contributors with new coins, according to specific participation performance measures, as we explain later. For the scheme to work, each individual i must announce only two values

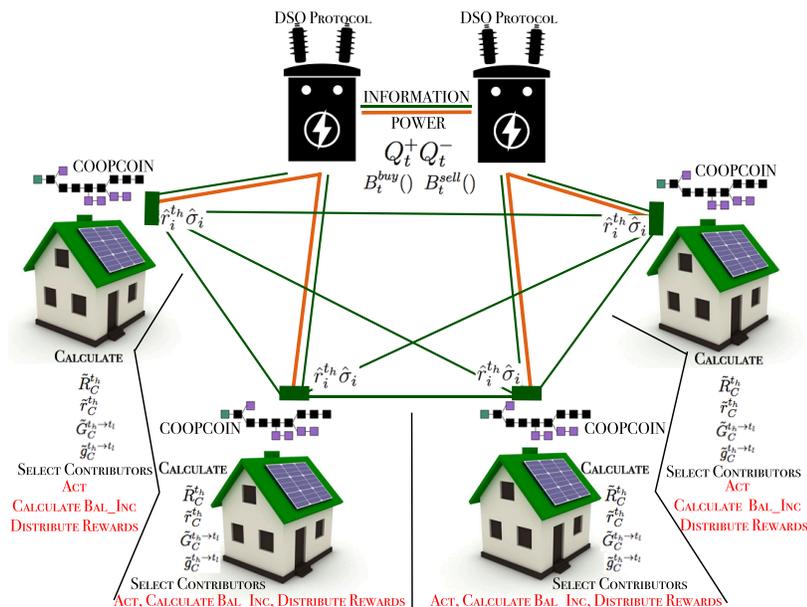


Figure 4-2: Prosumer network and interactions.

for each shifting interval pair: (a) her reduction capacity, \hat{r}_i^{th} ; and (b) her confidence $\hat{\sigma}_i$ for meeting her reduction promises. The confidence represents the variance of a normal distribution (with mean value 0) over the error between the stated and the final action. This is in line with past approaches [7, 155].

An *optimistic estimate* of the cooperative shifting capacity is then collectively

calculated as:

$$\tilde{R}_C^{t_h} = \sum_{i \in C} \hat{r}_i^{t_h} \quad (4.4)$$

and a *pessimistic estimate*, by:

$$\tilde{r}_C^{t_h} = \sum_{i \in C} (1 - \hat{\sigma}_i) \hat{r}_i^{t_h} \quad (4.5)$$

Then, the cooperative determines the shifting interval pairs (t_h, t_l) , as well as the target shifting capacity that will lead to increased profits. In order to guarantee profits, the target shifting capacity is the maximum r^{*,t_h} values to be rescheduled such that Assumptions 4.1.1, 4.1.2 continue to hold. That is, for each shifting interval pair (t_h, t_l) :

$$\text{maximize } r^{*,t_h} \quad (4.6)$$

$$\text{subject to } \tilde{Q}_{t_h}^+ - (\tilde{Q}_{t_h}^- - r^{*,t_h}) < \tau \quad (4.7)$$

$$\tilde{Q}_{t_l}^+ - (\tilde{Q}_{t_l}^- + r^{*,t_h}) > \lambda \quad (4.8)$$

Next, the estimated by the members cooperative gains (minimum and maximum) are calculated, given the total expected consumption and production values of the cooperative for each time interval, $\tilde{q}_{C,t}^- = \sum_{i \in C} \tilde{q}_{i,t}^-$, $\tilde{q}_{C,t}^+ = \sum_{i \in C} \tilde{q}_{i,t}^+$, and the estimates $\tilde{R}_C^{t_h}$, and $\tilde{r}_C^{t_h}$:

$$\tilde{G}_C^{t_h \rightarrow t_l} = \text{profit}_C^{t_h}(\tilde{R}_C^{t_h}) - \text{loss}_C^{t_l}(\tilde{R}_C^{t_h}) \quad (4.9)$$

$$\tilde{g}_C^{t_h \rightarrow t_l} = \text{profit}_C^{t_h}(\tilde{r}_C^{t_h}) - \text{loss}_C^{t_l}(\tilde{r}_C^{t_h}) \quad (4.10)$$

To continue, the estimated gains per kWh are calculated:

$$\tilde{G}_{C,kWh}^{t_h \rightarrow t_l} = \frac{\tilde{G}_C^{t_h \rightarrow t_l}}{\tilde{R}_C^{t_h}}, \tilde{g}_{C,kWh}^{t_h \rightarrow t_l} = \frac{\tilde{g}_C^{t_h \rightarrow t_l}}{\tilde{r}_C^{t_h}} \quad (4.11)$$

For shifting interval pairs $\{t_h, t_l\}$ that

$$\tilde{g}_{C,kWh}^{t_h \rightarrow t_l} > 0 \quad (4.12)$$

holds, the shifting procedure is expected to be profitable, and the shifting interval pair along with its $\tilde{G}_C^{t_h \rightarrow t_l}$ and $\tilde{g}_C^{t_h \rightarrow t_l}$ values are announced to the members.

- If $r^{*,t_h} \geq \tilde{r}_C^{t_h}$, every available contributor can participate in the shifting operations.
- However, in case $r^{*,t_h} < \tilde{R}_C^{t_h}$, the constraint of Eq. (4.6) does not yet hold and gains are not certain, so some agents must be excluded from action.

We will examine different approaches for this in the following section.

Finally, if the cooperative actually reduces $r_C^{t_h} \leq r^{*,t_h}$ given actual Q_t^- & Q_t^+ , the final actual cooperative gain is:

$$g_C^{t_h \rightarrow t_l}(r_C^{t_h}) = profit_C^{t_h}(r_C^{t_h}) - loss_C^{t_l}(r_C^{t_h}) \quad (4.13)$$

Of course, in order for the final gain levels to be inside the estimated range, two conditions must hold. First, the statements $\hat{r}_i^{t_h}$, $\hat{\sigma}_i$, and the predictions \tilde{Q}_t^- & \tilde{Q}_t^+ must be accurate. As mentioned earlier, the accuracy of \tilde{Q}_t^- & \tilde{Q}_t^+ can be ensured by known methods [81, 187]. We examine how we can achieve the accuracy of $\hat{r}_i^{t_h}$, $\hat{\sigma}_i$

in the next section. Second, the cooperative must be sizeable, meaning that it is the only actor that can induce significant price changes by consumption rescheduling. This helps overcome the problems raised by Lemma 3.

Definition 4.2.1 (Sizeable cooperative). A cooperative C is sizeable, if its pessimistic reduction capacity estimate is much greater than the sum of external parties capacity, i.e. when

$$\sum_{i \in C} (1 - \hat{\sigma}_i) \hat{r}_i^{t_h} \gg \sum_{j \notin C} \hat{r}_j^{t_h}, \forall t_h \quad (4.14)$$

Remark 1. The cooperatives that are formed are sizeable, because, due to Lemma 3, every rational agent avoids optimizing individually, thus seeks to cooperate.

Lemma 4. If statements $\hat{r}_i^{t_h}$, $\hat{\sigma}_i$, are accurate, and the cooperative C is sizeable, then, the cooperative's shifting suggestions include only eligible shifting interval pairs for C , in other words C will have:

$$g_C^{t_h \rightarrow t_l} > 0$$

Proof Since the cooperative has accurate knowledge of the total shifting capacity range $\tilde{R}_C^{t_h}$, and $\tilde{r}_C^{t_h}$, and it is sizeable, the $\tilde{g}_C^{t_h \rightarrow t_l}$ and $\tilde{G}_C^{t_h \rightarrow t_l}$ estimates are more accurate than others calculated based on partial knowledge, thus the following holds:

$$\tilde{g}_C^{t_h \rightarrow t_l} < g_C^{t_h \rightarrow t_l} < \tilde{G}_C^{t_h \rightarrow t_l}$$

Now, due to the enforcement of the constraints from Eq. 4.6 and 4.12, for the suggested interval pairs, the minimum gain estimate per kWh is positive, $\tilde{g}_C^{t_h \rightarrow t_l} > 0$. Thus the shifting interval pairs suggested by the cooperative are eligible. \square

In conjunction with Remark 1, this Lemma is important for the following reason. While the calculations above do not take the individual shifting costs into account, cooperative members must weigh the expected gain per kWh (if these are accurate) with their own shifting costs $c_i^{t_h \rightarrow t_l}$ and decide whether they will finally contribute or not. Now, Lemma 4 shows that the cooperative can take advantage of the predicted price differences and create profit by rescheduling consumption. Moreover, since C is sizeable, no other actor can significantly affect prices so that the cooperative does not meet its goals. Therefore, the $\tilde{g}_C^{t_h \rightarrow t_l}$ are accurate (assuming the $\hat{r}_i^{t_h}$, $\hat{\sigma}_i^{t_h}$ are too); and then individuals can safely weigh these against own shifting costs to decide participation. The overall process can be achieved as shown in Alg. 2. The

Algorithm 2 Coordinated shifting for a (t_h, t_l) interval pair

Input: $\tilde{Q}_{t_h}^+, \tilde{Q}_{t_h}^-, \tilde{Q}_{t_l}^+, \tilde{Q}_{t_l}^-, \{\tilde{q}_{i,t}^-\}_C, \{\tilde{q}_{i,t}^+\}_C$
 1: Determine and announce $\tilde{G}_{C,kWh}^{t_h \rightarrow t_l}, \tilde{g}_{C,kWh}^{t_h \rightarrow t_l}$
 2: Receive agent bids $\{\hat{r}_i^{t_h}\}_C, \{\hat{\sigma}_i^{t_h}\}_C$,
 3: Check constraints and select agents
 4: Wait for shifting actions realization, $\{q_{i,t}^-\}_C, \{q_{i,t}^+\}_C$
 5: Distribute revenues to contributors

complexity for solving the algorithm's first step, i.e. finding the peak and non-peak intervals, and respective loads and gains for the daily planning horizon, is a function of the number of time intervals. For example, if the cooperative adopted a constrained optimization approach, it would be $\mathcal{O}(t^3)$,² where t is the number of time intervals. Next comes the selection of the actual contributors during each peak interval, that of Line 3. The duration of this procedure depends on the selection method that each cooperative adopts. The most expensive step of the selection methods we present in

²Since the typical complexity for solvers is cubic.

Section 4.2.3 below, is that of ranking, whose complexity is $\mathcal{O}(n^2t)$ in the worst case, i.e. $\mathcal{O}(n^2)$ for sorting [137], times t time intervals.

4.2.1 Cooperative balance increase

As already discussed (Eq. (4.13)), prosumers generate gain from the price differences for both buying and selling electricity. However, the gain part from buying is immediately awarded to each prosumer in the form of reduced bills, and cannot be redistributed among the members easily. Better sell prices, on the other hand, result to larger income for the cooperative, and this profit can be concentrated into a collective account. The achieved cooperative balance increase by each collective shifting operation is given by:

$$\begin{aligned} bal_inc(r_C^{t_h}) &= B_{t_h}^{sell}(q_{C,t_h}^+, \tilde{Q}_{t_h}^+, (\tilde{Q}_{t_h}^- - r_C^{t_h})) - B_{t_h}^{sell}(q_{C,t_h}^+, \tilde{Q}_{t_h}^+, \tilde{Q}_{t_h}^-) \\ &\quad - B_{t_l}^{sell}(q_{C,t_l}^+, \tilde{Q}_{t_l}^+, \tilde{Q}_{t_l}^-) + B_{t_l}^{sell}(q_{C,t_l}^+, \tilde{Q}_{t_l}^+, (\tilde{Q}_{t_l}^- + r_C^{t_h})) \end{aligned} \quad (4.15)$$

This equation is derived from Eq. (4.1) and (4.2) after removing the parts that include $B_{t_h}^{buy}$, and represents the achieved gain from sell prices alone. Assuming that the initial balance of the cooperative is zero, the cooperative balance over the time horizon of shifting operations is simply the sum of the per time step balance increases:

$$b_{COOP} = \sum_{t_h} bal_inc(r_C^{t_h}) \quad (4.16)$$

However, since each participant contributes to the increase differently, the distribution of rewards must be different as well. A straightforward procedure for this

redistribution is to generate and award new coins, which, nevertheless, are returned to the prosumers based on each one's behavior. For this reason we propose a cryptocurrency protocol that is used simultaneously, to both coordinate, and reward prosumers.

4.2.2 COOPcoin for prosumer cooperatives

To achieve effective rescheduling of prosumers consumption, and reward back members according to their behavior, we propose the employment of a specialized cryptocurrency algorithm, designed for the coordination of prosumer cooperative actions. In existing cryptocurrency schemes, the same protocol is used by all members of the community, and each member executes a program that is linked with a distributed database, called the *blockchain* [25, 94].

The program performs certain calculations—e.g., in our case, consumption and production measurements, gain calculations, and so on, which implement Alg. 2. The individual results are then compared with those of other members, and, if validated—i.e. compared and matched, are written to each user's database that is, added to the blockchain and stored as history. If validation fails for a member, the adopted result is the one calculated by most members. This distributed execution approach removes the need for cooperative managers.

Note that the distributed nature of such an algorithm is guaranteed with the use of existing cryptocurrency protocols. Such protocols offer many desirable features, e.g. distributed consensus, transaction transparency, and anonymized data sharing [141]. Particularly, although data are shared among all participants freely, they are encrypted, and only the issuer and trusted parties can actually recover actual

information, and link data with real persons. Increased privacy, transparency, and the ability to operate democratically without a “manager”, are important for cooperatives [1]. Moreover, by ensuring these properties via the use of cryptocurrency, our approach can naturally extend to virtual power plants [23, 155] (where the trust among the constituting entities is even lower).

Our cryptocurrency scheme is specifically designed for prosumer cooperatives, and is called *COOPcoin*. The proposed cryptocurrency protocol “mines” coins according to a small number of measurements and calculations regarding the shifting behaviors, and *does not* require computationally intensive operations like other existing cryptocurrency algorithms, e.g. Bitcoin [94]. Here, “mining”³ is performed collectively, i.e. utility is generated by the better electricity rates as a result of collective shifting, and in place of the Bitcoin’s “proof-of-work” concept [94], we use what we term “proof-of-physical-action”: in order to get rewarded with COOPcoins, certain actions (i.e., electricity consumption shifting actions) must take place in the real world. For the sharing of the rewards, the protocol generates COOPcoins based on the collectively achieved profit and uses these to distribute that profit to the participants. The actual number of COOPcoins returned to each prosumer is determined based on their shifting behavior.

More specifically, the number of COOPcoins awarded depends on two terms: the first, $bal_inc(r_C^{t_h})$, is the actual balance increase due to the shifting operations, given by Eq. (4.15); and the second one, s_k , is a scaling factor used for the personalized

³As explained in Chapter 2, since cryptocurrency is not issued by a central authority, the process that creates new coins is performed by end-users and it is called *mining*. According to this procedure, users check if the data of the available transactions are valid, i.e. signatures are genuine, amounts in transactions are correct, etc., and are given newly created coins as a reward [134].

rewarding.

$$b_{i,t_h}^+ = \text{bal_inc}(r_C^{t_h}) \cdot s_k^{i,t_h} \quad (4.17)$$

Now, the value of s_k that actually scales each participant's share over the balance, depends on the reward sharing policies of each cooperative. We examine three approaches:

The *proportionate to estimated reduction (PROPest)* approach distributes back the balance increase to participants in a proportionate manner, according to their capabilities stated prior to the rescheduling actuation.

$$s_{PROPest}^{i,t_h} = \frac{(1 - \hat{\sigma}_i) \hat{r}_i^{t_h}}{\sum_{i \in C} (1 - \hat{\sigma}_i) \hat{r}_i^{t_h}} \quad (4.18)$$

The *proportionate to actual reduction (PROPact)* approach rewards according to the achieved individual reduction.

$$s_{PROPact}^{i,t_h} = \frac{r_i^{t_h}}{\sum_{i \in C} r_i^{t_h}} \quad (4.19)$$

The *accurate (ACCU)* approach uses the *normalized* CRPS scoring rule, to assess the absolute relative error ϵ_i between promised $\hat{r}_i^{t_h}$, and actual $r_i^{t_h}$ performance, with $\hat{\sigma}_i$:

$$s_{ACCU}^{i,t_h} = \frac{1 - CRPS(\mathcal{N}(0, \hat{\sigma}_i), \epsilon_i)}{\sum_{j \in C \setminus \{i\}} (1 - CRPS(\mathcal{N}(0, \hat{\sigma}_j), \epsilon_j)) + 1} \quad (4.20)$$

As discussed earlier in this thesis, CRPS has been used in the past [7, 155] to

directly evaluate probabilistic forecasts, and the score is given by:

$$CRPS(\mathcal{N}(\mu, \sigma^2), x) = \sigma \left[\frac{1}{\sqrt{\pi}} - 2\phi\left(\frac{x - \mu}{\sigma}\right) - \frac{x - \mu}{\sigma} \left(2\Phi\left(\frac{x - \mu}{\sigma}\right) - 1 \right) \right] \quad (4.21)$$

In our setting, $\mathcal{N}(\mu, \sigma^2)$ is the uncertainty stated over the expected *absolute relative errors* regarding the reduction capacity, as reported by an agent; while x is the actually observed error, ϕ the PDF, and Φ the CDF of a standard Gaussian variable. The mean μ and variance σ^2 of the predictive distribution can be estimated by each agent through private knowledge of its consumption requirements and business needs. Here, it reduces the prosumer COOPcoin reward when her actual performance is not inside the stated confidence range. It is used in negative orientation and is normalized, so that perfect forecasts generate a value of zero, while the worst ones produce a value of 1. This incentivizes participants to be accurate.

Theorem 1. When using ACCU for electricity prosumers cooperative reward sharing, participants are incentivized to be accurate regarding their statements.

Proof A scoring rule is a function $S(P, Q)$ that assesses the distance between a predictive distribution P and an actual distribution Q . When the rule is *strictly proper*, then $S(Q, Q) \geq S(P, Q)$, with the equality holding iff $P = Q$, i.e. the value is maximized for exact forecasts [67]. Since CRPS is used here in negative orientation, and is normalized, i.e. $CRPS \in [0, 1]$, we have that

$$CRPS(Q, Q) \leq CRPS(Q, P)$$

with the equality holding if and only if $Q = P$. Also, because any affine combination of a strictly proper scoring rule is also strictly proper [123], we exclude agent's i CRPS from the denominator of Eq. (4.20), guaranteeing that s_{ACCU}^i is also strictly proper. Now, due to CRPS placement in Eq. (4.20) for fixed $r_i^{t_h}, r_C^{t_h}$, the share from the positive balance increase (Lemma 4) for the participant i is maximized when CRPS=0, leading to

$$s_{ACCU}^{i,t_h}(Q, Q) \geq s_{ACCU}^{i,t_h}(Q, P)$$

with the equality holding iff $Q = P$. Thus, the reward for i is maximized when the forecast $\hat{\sigma}_i$ is accurate. \square

Note that, to maintain strict propriety, i is excluded from the denominator, leading to a small surplus of gain not being directly awarded to the participants in the form of COOPcoins. This *weak budget balancedness* does not affect the other properties of our approach, and the surplus can be returned to the actors in various ways (e.g., via the purchase of new equipment, or as bonus to new members).⁴

4.2.3 Selection of contributors

As pointed out earlier, it is probable that shifting capacity is larger than the maximum eligible for profitable actions. In such cases, the cooperative must select only a subset from the available participants in C to include in shifting operations. The actual method used for the selection can vary among cooperatives, according to their

⁴ Alternatively, considering that COOPcoins represent shares, if no surplus redistribution actions are performed, the result is an increase to the exchange rate between the COOPcoin and the "external" currency used to pay the cooperative, benefiting this way everyone with COOPcoins in their possession.

business plans and policies. In any case, it is to the actors' best interest to form reliable cooperatives in order to both achieve gains, and contribute to Grid stability. Here, we examine three selection methods. The first one, is the same as the *Random* method of Section 3.3.1. The other two, rank selection candidates as the rest of the Section's 3.3.1 methods do, however, they differentiate in the measure that take into account to rank them.

- The *Random selection* method picks contributors uniformly until the required $\tilde{R}_C^{t_h}$ is covered for each t_h .
- The *Reduction Capacity selection* method sorts contributors *wrt.* reduction capacity for each shifting interval pair, in an ascending order. Then, it starts including from the one with the lowest value, until $\tilde{R}_C^{t_h}$ is covered for each t_h . The objective here is to include as many contributors as possible.
- The *Engagement selection* method ranks contributors *wrt.* their wealth in COOPcoins. Then, starting from the richer one, it continues with the rest, until $\tilde{R}_C^{t_h}$ is covered for each t_h . The intuition is to include active and valuable members, since the wealth in COOPcoins does not only indicate participation frequency, but overall effectiveness as well.

Following selection, the accepted contributors are called for action, and rewards are dispensed after the actions occur. As a final note, the computational complexity of the proposed DSM mechanism is similar to that of Chapter 3, as presented in Section 3.5.

4.3 Experimental Evaluation

In this section we present the dataset and the results from our simulations. First, we report the origin of our dataset, and describe its augmentations to account for missing values. Next, we show the different impacts of individual and cooperative shifting actions, as well as selection methods comparison results; the stability of our proposed scheme is illustrated with a sensitivity test and, finally, we show how different COOPcoin reward sharing techniques incentivize statement accuracy.

4.3.1 Simulations setting

Our simulations employ a dataset based on consumption data from Kissamos, a district of Crete, Greece, and renewable production data from Galicia,⁵ Spain, both in 2012. More thorough descriptions of the dataset are given in Appendix A, The consumption data represent hourly demand from different contract categories, and include seasonal variabilities.⁶ In particular, the load profiles come from *residential*, *commercial*, *industrial*, *agricultural*, *public*, and *municipal* customers. However, due to the nature of agricultural and municipal demand profiles (i.e. mainly pumps, street lighting, etc.), which are tasks that cannot be shifted in time, these two categories do not participate in the prosumer cooperative of our simulation. In total, there are 7,376 prosumers in our setting.

The production data come from real wind generators and solar power plants, and have been scaled and divided, so as to represent the production of each prosumer.

⁵<http://www.sotaventogalicia.com>

⁶The production and consumption values in a Matlab file format can be obtained from <http://intellix2.intelligence.tuc.gr/~akasiadi/ProsumerCoop/>.

Prosumers are equipped either with both wind generators and solar panels, or with a single type of generation only. The average daily prosumer electricity production in our setting is 7.68kWh. The numerical results presented in this section are averages over 10 yearly iterations—that is, averages over 10 simulation runs, with each simulation run encompassing 344 days in 2012.⁷

External Variable Pricing. Now, as the external currency to our mechanism, we assume that the NRGcoin protocol is adopted by the market and thus $B_t^{sell}()$ and $B_t^{buy}()$ are calculated based on the formulations shown in [122]:

$$B_t^{sell}(q_t^+, Q_t^+, Q_t^-) = (0.1 \cdot q_t^+) + \frac{0.2 \cdot q_t^+}{e^{\left(\frac{Q_t^+ - Q_t^-}{Q_t^-}\right)^2}} \quad (4.22)$$

and

$$B_t^{buy}(q_t^-, Q_t^+, Q_t^-) = \frac{(0.65 \cdot Q_t^-) \cdot q_t^-}{Q_t^- + Q_t^+} \quad (4.23)$$

Both billing functions satisfy our assumptions regarding variable pricing, as explained in Section 4.1.1. The parameter values 0.1,0.2,0.65 are set arbitrarily, so that Eq. (4.22) and Eq. (4.23) to produce reasonable results. An illustrative example of the B_t^{buy} and B_t^{sell} values during the eleventh and twelfth weeks of the simulation, are shown in Fig. 4-3.

Shifting Behaviours. Unfortunately, no shifting costs and capacities were available in the dataset, and we are not aware of any datasets including such values.

⁷ Note that since the simulated time horizon extends to a year, all seasonal variabilities and additional uncertainties (e.g. individual agent availability, individual reduction capacities for each time interval) are sufficiently taken into account.

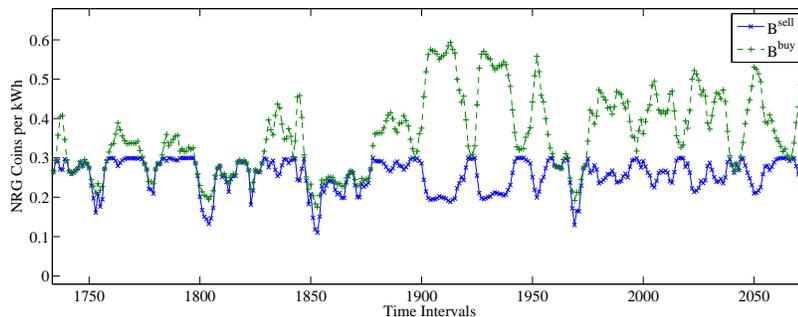


Figure 4-3: Time-series of B_t^{buy} and B_t^{sell} used in the simulation as the Grid’s pricing mechanism.

Thus, we assumed shifting capacities that were on average 35% of the demand, varying among categories (e.g. higher for industrial prosumers, and lower for residential ones).

Shifting costs increase inversely proportionally to the prosumer baseline demand, meaning that shifting to an interval which typically has less demand, induces increased comfort loss. In this setting, the shifting costs result to an average value of 3.9 profitable interval pairs/day, for each prosumer.

Moreover, participants are divided into two different accuracy classes that describe the relationship between agent confidence statements $\hat{\sigma}_i$, and final realized shifting actions. The first one contains the *accurate* predictors; this describes the realistic case where agents are mainly confident about their statements, and also have a high probability to deliver what they promised. The second one, corresponds to the *inaccurate* predictors, where prosumers might or might not follow stated forecasts. For *accurate* predictors, the confidence statements and the parameters for calculating the absolute relative error ϵ_i , are sampled from $\mathcal{B}(1, 5)$ and $\mathcal{B}(4, 2)$ respectively; note that the actual r_i^{th} is calculated by the product of a sample α_i^{th} and the stated

shifting capacity \hat{r}_i^{th} , i.e. $r_i^{th} = (1 - \alpha_i^{th})\hat{r}_i^{th}$. For *inaccurate* agents, confidence statements $\hat{\sigma}_i$ are sampled from a “wider” Gaussian $\mathcal{N}(0.5, 0.15)$, and the α_i^{th} parameter is set to $1 - \hat{\sigma}_i$. About 50% of the participants in our setting belong to the accurate class, with the rest being inaccurate agents (since agents are assigned to a specific class with 50% probability).

Parameters τ and λ . To calculate the τ, λ parameters for each day, we first determined the average kWh difference between supply and demand, $Q_t^+ - Q_t^-$, across the 24 values. Then, τ is placed at the 75% of the distance between the average and minimum difference value, while λ at the 25% of the distance between the average and the maximum difference value. Nevertheless, these particular values are application specific, and other algorithms can be used for their calculation as well, according to each cooperative’s capabilities and business goals. We now proceed to describe the numerical results from our experiments.

4.3.2 Individual vs. cooperative action

We first compare two different scenarios, one with the prosumers shifting according to individually optimized plans, and a second where they shift according to the cooperative suggestions. Contributors are selected randomly, and everyone is accurate with respect to their promises and final actions, i.e. no specific “accuracy” classes are used for this set of experiments. Table 4.1 shows the difference in the total bills of the prosumers, and the average across all year daily peak-to average ratio (PAR) values, for the total demand and supply difference, buying, and selling prices.

First, we observe that the “collective bill” when prosumers cooperate drops by

Table 4.1: Performance of individual and cooperative action.

Method	Total bill difference	Avg. PAR $ Q_t^+ - Q_t^- $	Avg. PAR B^{sell}	Avg. PAR B^{buy}
Initial	-	1.96	1.28	1.21
Individ.	-2.4%	4.75	1.38	1.19
Coop.	-4.4%	1.86	1.24	1.19

a factor of 2 (-4.4% vs. -2.4%) compared to its reduction when they do not. Additionally, the cooperative approach outperforms individual optimization in terms of peak-to-average ratio (PAR) values (average across 344 days) for the $|Q_t^+ - Q_t^-|$ and B^{sell} columns. Lowering the PAR of $|Q_t^+ - Q_t^-|$ means that demand and supply difference is flattened, thus less electricity is exchanged to the balancing market, and consumption of locally produced electricity is promoted. By contrast, the increase of the PAR for the individual approach shows the scale of the herding effects that take place. Furthermore, reduction in the PAR value of the selling price when cooperating, means smaller fluctuations, a fact that allows for more realistic planning. Lastly, both cooperative and individual optimization leads to buying prices that are quite stable.

4.3.3 Evaluating contributor selection methods

Henceforth, we assume that each prosumer belongs to the two different accuracy classes introduced earlier. In this setting, we first evaluate the three different contributor selection approaches we put forward, namely *Engagement*, *Reduction Capacity* and *Random*. Table 4.2 presents the (average) total cooperative gains and balance for 2012, for each of the three proposed contributor selection methods, in NRG coins.

Table 4.2: Total cooperative gains and balance in 2012 for each selection method (NRG coins).

Measure	<i>Engagement</i>	<i>Reduction Capacity</i>	<i>Random</i>
Total Coop. Gains	133805.51	123813.43	80714.62
Total Coop. Balance	68116.49	58134.89	15029.42

The prior COOPcoin wealth distribution required by the *Engagement selection*, is determined by conducting simulations using *Reduction Capacity selection* and *ACCU* reward sharing. We can observe that *Engagement* achieves the highest cooperative gains, and, consequently, the highest cooperative balance. *Reduction Capacity* also performs well *wrt.* gains in NRG coins. Lastly, *Random* has the worst outcome, which, nevertheless, is still profitable.

4.3.4 Shifting capacity sensitivity test

In this set of experiments we gradually change the shifting capacity of every individual, from -50%, to +50% of their initial, and examine the impacts to the average individual gains and average coalition sizes during 2012. Note that during all experiments we enforced two constraints: (a) shifting capacity cannot exceed the hourly demand, and (b) shifting capacity cannot be negative. Results are presented in Figure 4-4. As we observe, the average size of the shifting coalitions drops when the shifting capacity of the prosumers increases, for all selection methods. This is natural, since increase in the shifting capacity helps meeting cooperative shifting requirements with fewer members. Regarding the individual gains of participants for 2012, we can see that *Engagement* and *Reduction Capacity* selection methods are not significantly affected by the difference in the shifting capacity of the individuals. Also, differ-

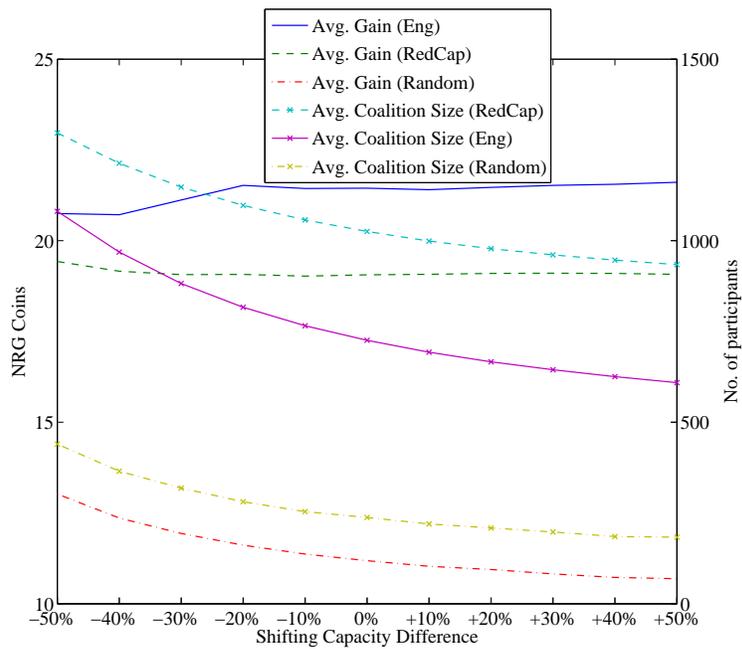


Figure 4-4: Average individual gains during 2012, and average shifting coalition sizes vs. shifting capacity changes for all three selection methods.

ences in the shifting capacity do not induce changes in the selection methods relative ranking; for all values, *Random* consistently ranks last, while *Engagement* produces higher gains than the *Reduction Capacity*.

The increase in average individual gain for the *Engagement* selection is due to the fact that “better” performing agents are selected continuously, resulting to the gain being shared by fewer agents. Lastly, when contributors are selected using the *Random* method, the average individual gain decreases for higher shifting capacity percentages. This happens because shifting operations are overtaken by fewer members, who, however, are not examined wrt. their truthfulness. Thus, it is highly probable that the final cooperative shifting actions deviate from those promised, leading to less overall gain.

Figure 4-5 presents the *total cooperative balance* in 2012, for different shifting capacity percentages. As we can see, the balance of the cooperative *decreases* for all methods as prosumers’ shifting capacity increases. This can be interpreted if we consider the average coalition size values from Fig. 4-4: when the number of acting prosumers decreases, each individual inaccuracy regarding the shifting operation has a larger impact on the cooperative performance, leading to reduced cooperative gains. When actors come in large numbers, their individual errors have less impact, since they cancel out by others to the opposite direction. Regardless of the observed balance reduction, we can see that *Engagement* selection method consistently produces the highest balance, with *Reduction Capacity* following, and *Random* producing the least, irrespective of the way the shifting capacity changes.

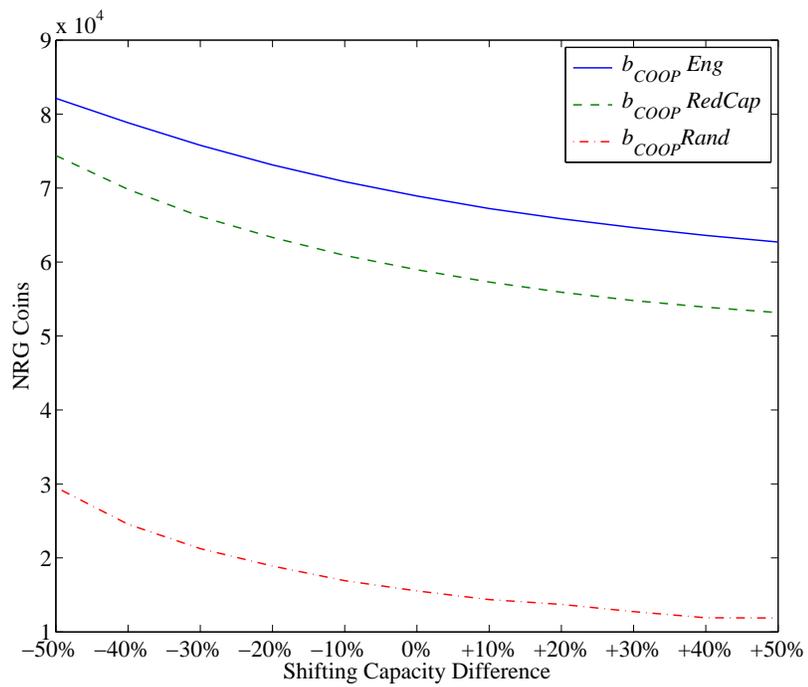


Figure 4-5: Cooperative balance in 2012 vs. shifting capacity changes for all three selection methods.

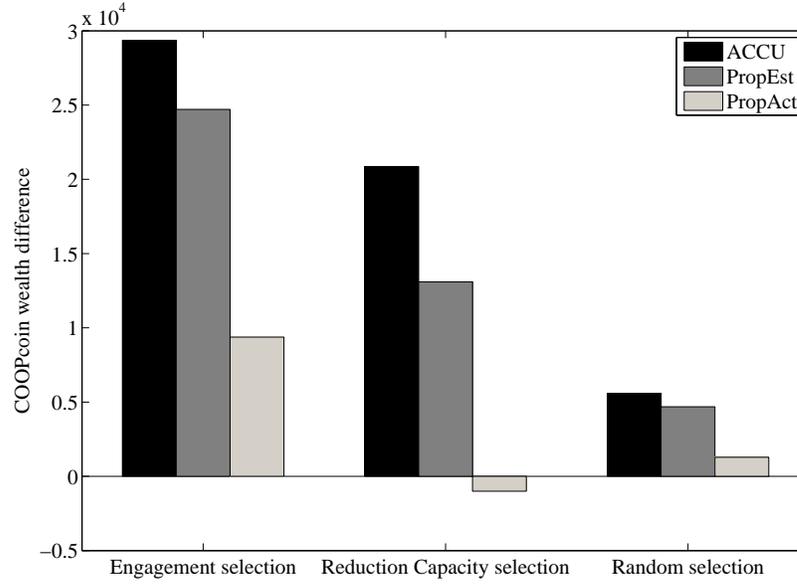


Figure 4-6: COOPcoin total wealth difference between accurate and inaccurate actors after 344 days.

4.3.5 Reward sharing methods evaluation

Finally, we use the three reward sharing methods with each selection method, in order to find the one incentivizing accurate statements the most. Figure 4-6, presents the *difference* in the total COOPcoin wealth between accurate and inaccurate actors when using different selection and reward sharing approaches. We observe that, for every selection method, this difference is higher when the *ACCU* approach is used for rewarding. Also, as expected, when using *ACCU* redistribution combined with *Engagement selection*, the difference in COOPcoin wealth between accurate and inaccurate participants reaches its highest levels. Interestingly, when using *ACCU*, accurate participants are rewarded more, even when the selection criterion does not distinguish between the two classes (i.e., when using *Reduction Capacity* and *Random*

selection). For all these reasons, *ACCU* is clearly the most effective in promoting statements accuracy.

4.4 Conclusions

In this chapter, we presented a cooperative prosumer consumption shifting scheme that employs a distributed cryptocurrency mechanism for the coordination of members' actions. This special purpose cryptocurrency is combined with scoring rules to incentivize electricity prosumers to offer decentralized and coordinated DSM services, and also allows for personalized pricing. We proposed contributor selection and reward sharing methods which incentivize truthfulness, guarantee increased profits from the trading of electricity, and help flatten electricity demand curve. By simulating a Smart Grid setting with prosumers that was based on real-world data, we tested the proposed reward sharing and participant selection methods. Results show that the application of our method leads to more stable electricity prices, more balanced local demand and renewable supply, as well as increased gains for DSM participants. When incorporating the method that rewards accurate participants the most, and also use a selection method based on user engagement, the prosumer cooperative has increased profits. Though illustrated in the domain of prosumer cooperatives, our approach immediately applies to the broader domain of prosumer virtual power plants as well.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 5

Incorporating EVs and Energy Storage Systems

In this chapter we deal with the G2V / V2G (grid-to-vehicle / vehicle-to-grid) problem in modern “smart” grids and microgrids¹. In V2G, battery-equipped EVs communicate and strike deals with the electricity grid to either lower their power demands or return power back to the network when there’s a peak in the request for power [147]. G2V is V2G’s reverse problem, in which EVs must connect and draw power from the grid without overloading it. In such cases, *Electric Vehicles (EV)* are connected to the Grid for charging. However, due to the large number of vehicles expected to occur in the near future, the additional load amount that the Grid is called to cover is expected to be large enough to create disturbances in the electricity supply. For this reason, additional techniques and mechanisms need to be developed,

¹A microgrid is similar to a Smart Grid, but smaller in size and scale e.g., a small community of interconnected prosumers [23].

which will properly manage supply and demand for electricity specifically for electric vehicles whose needs are slightly different from conventional demand [79].

Here, we propose an extension of the DSM schemes presented in the previous chapters, concerning consumers and prosumers, to also incorporate the augmented capabilities of EVs, that is the storage of electricity in the vehicle's battery. The various energy cooperatives of consumers or prosumers, along with additional renewable generators that are locally present, are here grouped under the term *Microgrid*.

In particular, the extended model that incorporates EVs in our setting is shown in Figure 5-1. We assume the existence of a microgrid that includes consumers (buildings and other facilities), renewable energy sources, and EVS, which are linked to each other via an authority, e.g the DSO or the *Microgrid Regulator*. The regulator may additionally import, or export electricity from and to an external power market, e.g. the Smart grid, where, however, buying and selling prices are worse due to the use of additional equipment, taxes, etc. [110]. The Microgrid regulator executes a specific algorithm for V2G / G2V, whose operation we explain later in the chapter. Arrows indicate electricity flow; note that for the cases of the Smart Grid and EVs, the flow can be bidirectional. This algorithm constitutes our main contribution in this chapter, and possesses several desirable characteristics which we analyze below. It should be noted that the algorithm provides the possibility of importing electricity from the vehicles themselves to share with the rest of the network, *whenever circumstances allow*, i.e.:

- When the price for importing electricity from the smart grid are high enough during a period of time.

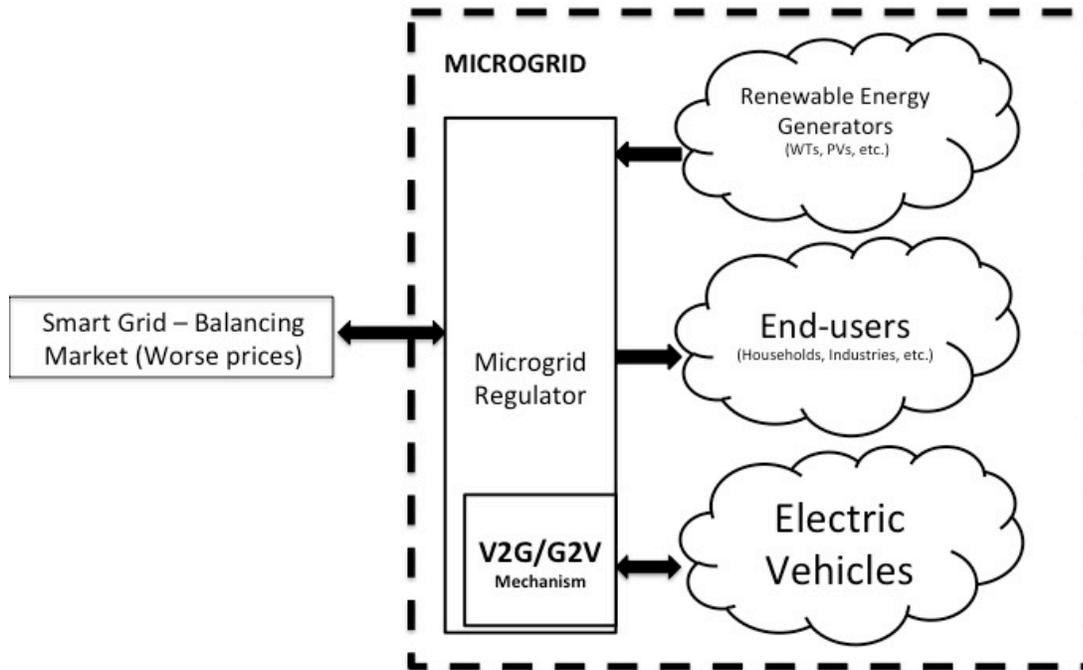


Figure 5-1: Participant types and interactions of the assumed microgrid with V2G/G2V services.

- There are connected vehicles that are already charged, with an expected disconnection time several hours later than that time.

At the heart of the V2G / G2V algorithm there is a process that estimates the most appropriate time intervals for charging and discharging the battery of the vehicle according to

- (a) short-term electricity generation and demand forecasts;

- (b) input and output power values from Smart grid; and
- (c) the preferences of electric car owners, the expected time they will be connected to the grid, and their tolerance to power the grid with energy stored in their batteries.

Therefore, whenever a vehicle is connected to the grid for either charging or discharging, it reports the expected disconnection time, a value representing the target *State-of-Charge SOC* value until the expected time, and a *V2G cost* value for discharging an energy unit from the battery in the interim period. The V2G cost value represents the willingness of the owner of each vehicle to offer stored energy to the grid, and it represents the reservation price that was introduced in Chapter 3. If the network needs energy and it is drained from electric vehicle batteries, the owners are rewarded according to the stated V2G cost of discharge. The algorithm applies a series of constraints, which result in the V2G / G2V mechanism choosing the cheapest source between the Smart Grid and the EVs to import electricity from.

The proposed V2G / G2V algorithm has the following objectives:

1. *Increase the reliability of the electricity grid.* This is achieved by constantly checking demand and production balance, taking into account forecasts, and additional information on the availability of the vehicle batteries that are provided by their owners. Scheduling is performed on time intervals-ahead basis, overcoming this way the drawbacks of planning as last minute action.
2. *Reduce consumption costs.* The Microgrid regulator draws energy from the most advantageous sources at any moment, based on the current purchase/sale

prices from each source (Smart Grid, vehicle batteries), thus ensuring the lowest cost for consumers.

3. *Creating new entrepreneurship models.* Individual electric vehicle owners, but also facilities such as parkings, companies with large fleets of electric vehicles, can benefit from the network by offering stored energy from their vehicle's batteries to the grid for a fee.
4. *Personalized rewarding.* We extend the state of the art by proposing an algorithm that rewards each vehicle for providing stored electricity to the microgrid during times of renewable production shortage, in a personalized manner. Furthermore, this is performed in a non-intrusive way, and, as such, it allows independent decision making for the EV owners side.
5. *Incorporation of large numbers of EVs.* Calculations are simple and increase in EV numbers and participation is not forbidding for the algorithm to work.
6. *Protection of the environment.* The smart grid is made up of a large number of heterogeneous energy sources, possibly including conventional generators, which harm the environment. The use of stored energy in vehicle batteries implies that imports from the smart grid will be less, so the operation of polluting sources is limited. Additionally, by having EVs charging when locally generated renewable energy is available, we promote self-consumption and the better utilization of renewable energy sources.

The scope of the work presented in this chapter is shown in Figure 5-2.

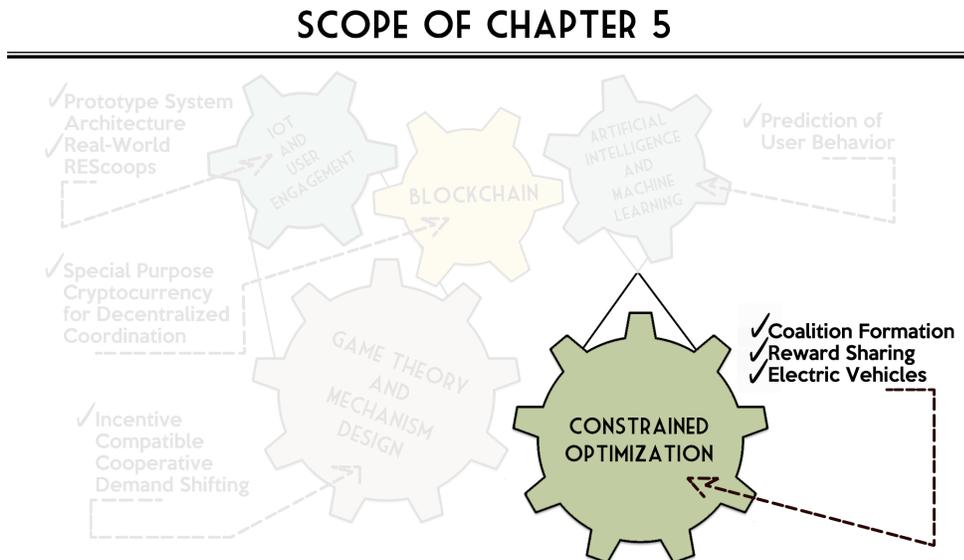


Figure 5-2: Overview of the scientific fields, and our contributions in Chapter 5.

As a final note, several smart charging schemes have been applied in the real world, e.g. the case of Pacific Gas and Electric company.² Since 2015, the company created incentives for EV vehicle owners to participate in Demand Side Management (DSM) programmes. During the first half of related pilots, approximately 20,000kWh were shifted as a result of EV DSM. This amount is enough to power the electricity for two homes over one year, thus avoiding costly and carbon-intensive electricity generation.

The remainder of this chapter is structured as follows: Section 5.1 discusses existing approaches to the EV charging problem. Section 5.2 introduces the V2G/G2V algorithm and the additional constraints that should be taken into account when addressing this problem. In Section 5.3 we present the simulation results, which are

²<https://www.engerati.com/article/pacific-gas-and-electric-company-electric-vehicles-grid>

based on real-world data. Finally, Section 5.4 concludes this chapter.

Parts of the research ideas described in this chapter appeared originally in [11].

5.1 Related Work

The work of [130] presents an extensive review of Smart Grid systems that incorporate EVs with renewable energy generators and V2G capabilities. The authors also outline the imminent market changes that will lead to new business opportunities for EV owners, either individually, or collectively. It is a fact that as the number of available electric vehicles increases, demand for power also increases. Now, because driving is closely linked to people's daily routine, the demand for electricity in the day often co-ordinates, resulting in a rapid increase in specific time periods during the day, reaching levels that can not be covered by the available production, the *herding effect*. This problem is not new, and several approaches have been proposed in the past to efficiently schedule the charging and discharging of individual vehicles: we review some of them below.

Several recent approaches deal with the formation of EV coalitions, to tackle the aspects of DSM that relate to the *V2G* and *G2V* problems (see, e.g., [54, 85]). The existing approaches, however, typically exhibit the following characteristics: (a) they either attempt to form coalitions with respect to a single criterion, or employ lengthy negotiation protocols in order to capture various coalitional requirements while respecting the constraints of individuals; and (b) they attempt to form *optimal* coalitions or *coalition structures* (i.e., optimal *partitions* of the agents space), to maximize the social utility derived by the formation of coalitions. The inherent

hardness of this *optimal coalition structure generation (CSG)* problem [39], however, and the fact that negotiation protocols can be time consuming, severely restricts the practicality and scalability of such algorithms: they can handle at most a few hundred EVs. In reality though, hundreds of thousands of EVs connect to the Grid, and could potentially offer their services; and, if the aim is to balance electricity demand at real time, any such service should be offered almost instantaneously.

For instance, [54] attempts to tackle the optimal CSG problem, and relies on a heavy agents negotiation protocol to form coalitions among EVs that sell power to the Grid. Though it is empirically shown to produce solutions that are close ($\sim 93\%$) to the optimal, this is when tested in scenarios with *a few dozens* of agents only. Moreover, there is only a single formation criterion—namely, the *physical distance* among the EVs. In the real world, however, it is imperative that *a multitude of criteria* is taken into account—such as capacity, discharge power, charging and discharging efficiency, and perceived reliability [85]. In the works of [47] and [48], EVs coalition formation that is able to examine multiple criteria for V2G is provided via the use of hypergraphs. This coalition formation algorithm is shown to scale effectively as the number of EV rises. In our approach, we also examine the G2V possibility.

Now, in many DSM settings, agents might be restricted to interact with only a subset of other agents in the environment, due to limited resources, or existing legal and physical barriers. In such settings, the environment can be seen to possess some structure that forbids the formation of certain coalitions. This can be captured by an undirected graph providing a path connecting any two agents that can belong to the same coalition [39].

In that spirit, the approach of Vinyals *et al.* [198] constitutes an attempt to solve the CSG problem over *social networks* of electricity consumers (such as EVs). Forming *virtual energy consumer (VEC)* coalitions, it attempts to flatten the demand in order to get better prices in what could be a G2V arrangement. By solving the CSG, it finds the best VEC coalition for every consumer on the market; and guarantees a *stable* [39] payoff distribution outcome (i.e., one that removes any monetary incentives for the consumers to leave their coalitions). Unfortunately, this solution is shown to work on social graphs with only a handful of agents. In our approach for forming DSM coalitions, we do not attempt to form optimal coalitions, but rather efficient and effective ones that can form in a limited time. In contrast to existing DSM schemes incorporating MD, our scheme can handle thousands of agents.

Lopez *et al.* [109] present an optimization approach, which takes power distribution factors into account, and by using auctions it manages to prevent demand congestion. Agents calculate their bids based on bid and ask prices, and consider that charging / discharging can be done at different rates for each period of time. The difference with our own method is that uncertainty is not modelled at all. By contrast, we consider uncertainty over drivers' preferences, and in order to model it we ask drivers for information about when each vehicle will be connected.

In [183], a method for real-time charging is proposed. This is achieved by controlling the levels of active and inactive demand using AIMD algorithms (Additive Increase Multiplicative Decrease) so that a balance is achieved in the network. This approach is applied to various scenarios through simulations, and the results show that it performs well, but there is no mention of the cost of energy, a factor that we consider to be the main concern. A further difference is that in our approach, the

scheduling is done in advance, one day earlier, instead of as means of “last minute action”.

The work of [57] presents two V2G schemes that are in effect in the real world, specifically in Brazil. In their analysis, the authors include both battery damage and photovoltaic production. However, the authors make some very strong considerations, such as that all vehicles are connected the same hours in the day, and that charging is only between 00:00 and 06:00. In our approach, each driver-electric vehicle owner decides independently on when to connect and disconnect her vehicle, thus rendering the mechanism non-intrusive with regard to everyday life and habits of drivers.

Another approach for optimizing the current flow in systems including electric vehicles based on mixed integer linear programming (MILP), is presented in [76]. According to this method, the production of wind turbines is forecasted and compared with the storage of electric vehicle fleets. It then optimizes the flow of electricity by controlling constraints related to the distribution of electricity but also to the resulting greenhouse gas emissions. Although both concern in advance the next day’s scheduling of the current flow, the main difference of this work from ours is that *static* fleets of EVs are being considered, which permanently offer their energy in the network, without examining cases of dynamic arrival / departure of vehicles.

In an approach more similar to ours, Karfopoulos and Hatziargyriou [87] propose a distributed cooperative mechanism that achieves the transfer of demand to periods of lower total demand (valley-filling). This approach takes into account restrictions on current electricity prices, availability and capabilities of electric vehicle batteries, and forecasts for power generation from distributed sources. However, their approach

does not analyze at all the financial gains or incentives of EV owners resulting from the procedure. Incentives (e.g. monetary, social, etc.) are very important in mechanisms that involve the engagement of users, which, in a way, are asked to change their behavior and their everyday habits [128].

In [88], authors propose an approach for transferring electricity demand from EV charging by pre-planning. Neural networks and fuzzy logic methods are used to predict the demand and production levels of each future time. The problem of smoothing the demand curve is modelled as a 0-1 knapsack on the Microgrid side, and on the other hand, a gain-cost function is optimized for each EV. In the approach of [89], a stochastic mechanism based on “Unscented Transformation (UT)” is proposed. Here, uncertainties about wind turbine generation, as well as the active demand for vehicle fleets, are modelled, and cost functions are optimized, subject to limitations to the capacity of each battery and their decay factors. In that work also, financial incentives are not taken into account, but the importance of modelling uncertainty is highlighted.

Now, [106] also examines uncertainty, albeit regarding the arrival / departure of vehicles and models it as a non-stationary Markov chain. The best policy is made up of two thresholds: when the battery has more load than the upper threshold, then the vehicle delivers energy, and when lower than the lower threshold, it charges; at intermediate levels, it does nothing. In our approach, because of restrictions on renewable production levels, and because of the ranking of each agent based on its bid price, we do not have to search for an optimal response policy. In addition, using a dashboard, each driver can provide data regarding the time of departure, clearing this way out the uncertainties, hence predicting the state of each vehicle is

not necessary.

In the work of [103], the electricity trade is being considered, in which the participating companies also offer V2G capabilities. In their analysis, the cost of supply of stored energy is examined in conventional facilities, and their equivalence with EVs offering their battery energy is investigated. The main difference with our approach is that the EV side has to offer a unique (collective) selling price, whereas in our case, each owner reports his or her own, thus removing the need for an additional mediator.

The work of [150] proposes the use of special pricing for V2G participants to promote user participation. These are calculated by the DSO. According to our method, participants are rewarded for the supply of electricity, but based on their own pricing. In cases where those are more advantageous than the Smart Grid's power import price, they are included in the mechanism, otherwise they are ignored. This removes the responsibility of calculating special pricing from the DSO-side, and simply rewards participants according to their preferences.

A good example of a Mechanism Design approach for G2V is provided by [154], which details a *dominant-strategy incentive compatible (DSIC)* online mechanism for charging EVs. A greedy pre-allocation policy is tuned to guarantee that deviating from reported values leads to the same, or worst outcomes for the individuals. This is achieved by *cancelling* electricity units' allocations; or by *discharging* any over-allocated units before vehicle departure.

Another approach based on game theory is given by [204]. Here, coalitions are created from e-vehicles that work together to offer V2G services. Each agent (vehicle) calculates the best response on whether to charge or discharge, given market prices.

On the basis of these calculations, it either participates in a coalition or decides to withdraw from it, taking into account the expected profits from the participation of all members. By contrast, our approach creates a single coalition, including those who offer bargains. This way the communication burden and the number of messages exchanged are kept at a very low level. In [205] a group-selling mechanism is proposed by individual electric vehicles that cooperate following a two-stage auction process. Initially, group offers are registered, and at a second level, the network selects the most profitable ones. The mechanism turns out to be ex-post budget balanced, individually rational, and promotes truthfulness. This work, however, makes no mention of levels of renewable energy production, a measure which is key in our case.

To summarize, our contribution to the state-of-the-art is a day-ahead V2G/G2V mechanism for microgrids that incorporate renewable generators as a primary energy source. The V2G/G2V mechanism can incorporate thousands of EVs and carefully assess individual EV preferences to better utilize renewable production, by satisfying grid-related and EV-related constraints, and, what is more, reward EVs for offering their stored energy to the microgrid according to personalized reservation prices. Also, the algorithm allows participating EVs to connect and disconnect from the chargers at different time intervals, constituting this way a non-intrusive approach for the EV drivers.

5.2 The V2G/G2V Extension

The main steps followed by our V2G/G2V algorithm are three. First, it categorizes each time interval (in our case, one hour) as peak or non-peak, namely t_h and t_l , similarly as in the previous chapters. During peak intervals, demand levels of buildings and EVs exceed the levels of renewable energy production, and the microgrid regulator is therefore required to buy electricity from an external source, termed as the Smart Grid. An additional import alternative, is the charge of EV batteries that can be offered by V2G services. During non-peak, more renewable energy is generated than the demand, so the charging of EV batteries could be shifted then to store the renewable energy excess.

The second step is to rank the peak intervals on the basis of the deficit, and the non-peak on the basis of the excess. This step is imperative for the mechanism to determine which time intervals are the most problematic, and then address them with the corresponding priority.

The last step, and the most complicated one, is to choose the appropriate groups of EVs which will either (a) delay their charging, or (b) offer their stored energy to the microgrid, i.e. provide V2G services.

Starting with the time intervals, we first address the one with the highest imbalances. In the first stage, it is checked whether the EVs themselves generate the shortage, and if so, their demand is transferred to subsequent non-peak time intervals, as long as they remain connected to the grid until then. This check can be performed, e.g., by taking into account demand forecasts from EV chargers, and historical input preferences of participating drivers and their EV parameters (e.g.

charge/discharge rate, intervals of EV connections, etc.). The satisfaction of this constraint is ensured by examining the *connection indicator matrix*, M :

$$M = \begin{bmatrix} ci_{1,1} & ci_{1,2} & \dots & ci_{1,24} \\ ci_{2,1} & ci_{2,2} & \dots & ci_{2,24} \\ \vdots & \vdots & \ddots & \vdots \\ ci_{N,1} & ci_{N,2} & \dots & ci_{N,24} \end{bmatrix}, ci_{i,t} \in [0, 1]$$

where $ci_{i,t}$ indicates the expected connection status of the EV during the day i.e. one value for each of the 24 hourly intervals, for each EV. Intervals to move demand to are selected in priority, according to the selling price of electricity; the lower the price, the better for the microgrid regulator.

In the second stage, if there exist connected vehicles with their batteries charged, and demand for vehicles does not itself create a deficit, vehicles with better offers than the Smart Grid's import prices can give the required energy to the microgrid. For each candidate EV for participation, all relevant limitations are considered, such as (a) charging up to the required limits to disconnection; (b) doing so in non-peak time intervals to avoid further imbalances; and (c) the reservation price $\hat{p}_{i,t}$ set by the owner is lower than that of the Smart Grid's current import price. The operation of the algorithm is described by the flow diagram of Figure 5-3. More formally, the algorithm requires the following input:

1. Estimated microgrid renewable production levels from generator j , $\tilde{q}_{j,t}^+$
2. Estimated demand levels from building or infrastructure k , $\tilde{q}_{k,t}^-$
3. Prices for importing electricity from the Smart Grid, for each interval, p_t^{SG}

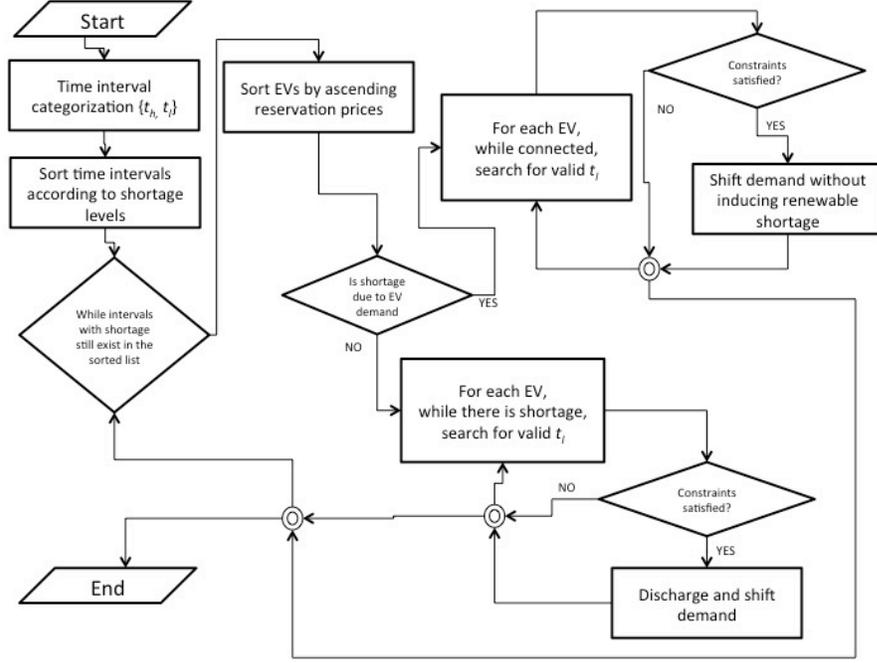


Figure 5-3: Flowchart of algorithm for V2G/G2V services.

And for each EV:

- (a) The connection indicators of the EV, $cl_{i,t}$
- (b) Estimated consumption induced by charging EV batteries, $\tilde{E}_{i,t}^-$
- (c) State of Charge of battery upon connection, \overline{SOC}_i
- (d) Target State of Charge upon disconnection, SOC_i
- (e) Reservation price, $\hat{p}_{i,t}$
- (f) Battery capacity CAP_i
- (g) Maximum charge/discharge rate DR_i

A key measure for the algorithm, is the *imbalance between local production and demand* during a time interval t , γ_t :

$$\gamma_t = \sum_k \tilde{q}_{k,t}^- + \sum_i \tilde{E}_{i,t}^- - \sum_j \tilde{q}_{j,t}^+ \quad (5.1)$$

As a general constraint for the V2G/G2V process, $\tilde{E}_{i,t}^-$ values cannot exceed the maximum charge/discharge rate of DR_i , since it signifies an upper threshold for battery charging demand:

$$\tilde{E}_{i,t}^- < DR_i \quad \forall i, t \quad (5.2)$$

To categorize the time intervals to t_h and t_l , we subtract the sum of every EV's and building's demand from the sum of renewable generation during a time interval t , and classify according to the following rule:

$$\begin{cases} t \in t_l, & \text{if } \gamma_t < 0 \\ t \in t_h, & \text{if } \gamma_t \geq 0 \end{cases} \quad (5.3)$$

Next, we sort time intervals according to γ_t , and for each t_h , sort the EVs according to their reservation prices $\hat{p}_{i,t}$. In cases where G2V can be employed to overcome imbalances—i.e. the demand of EVs themselves induce the imbalance, then for each connected vehicle from the sorted list we search for non peak intervals to shift charging to. The additional constraint that must be examined is to find a non peak interval t_l before the time of expected disconnection according to the connection indicators in M . The next, and the final step for the G2V part, is to check if battery charge shifting to that particular t_l does not change its categorization to be a t_h , i.e. for an

EV i ,

$$\gamma_t + (SOC_i - \overline{SOC}_i) * CAP_i \leq 0 \quad (5.4)$$

continues to hold. If these constraints hold, the rescheduling is validated, and actuated.

Now, if imbalances are not induced due to EV demand, *battery discharging* (V2G) is required to allow the microgrid to balance demand. Here, apart from the demand for the recharging of the battery before disconnection, similar to the G2V constraint check process above, an additional constraint must be examined, that is EV's reservation price $\hat{p}_{i,t}$, and the Smart Grid import price during that particular t_h , p_t^{SG} :

$$\hat{p}_{i,t} < p_t^{SG} \quad (5.5)$$

indicating that the price at which the EV provides the energy from its battery to the microgrid, is lower than the price asked for Smart Grid imports. If the vehicle finally participates in V2G, the driver is rewarded back according to the stated reservation price, and to the amount of energy that was granted to the microgrid from the EV battery. The above procedure iterates for every vehicle and every peak time interval that has been identified by the microgrid regulator. Note that, for microgrid cases with really large EV fleets (e.g. tens of millions), fast multi-criteria coalition formation methods can be incorporated, e.g. those proposed in the work of [48].

In the following section we present the numerical results of the V2G/G2V extension application, that indicate less Smart Grid imports and increased EV profits when applied on real-world data.

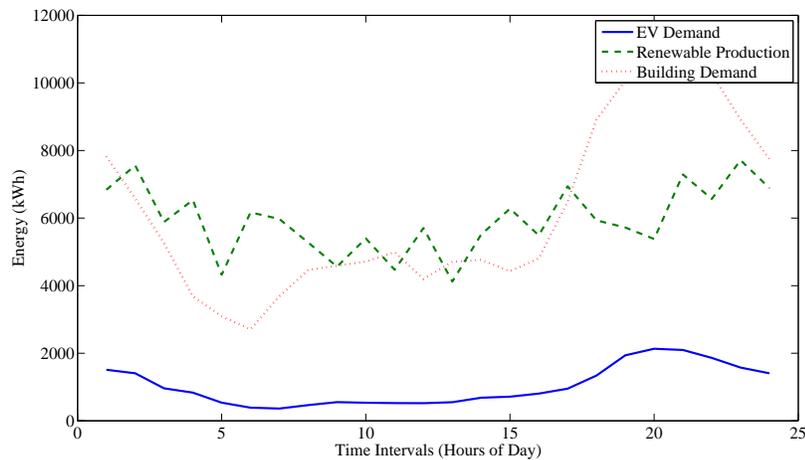


Figure 5-4: Hourly aggregate values for a sample day that is produced from the simulator.

5.3 Experimental Evaluation

The real-world data that constitute the basis of our dataset comes from two different sources. One is a renewable energy park in Spain that includes photovoltaics and wind turbines,³ the one used in Chapter 4 as well. Consumer data, electric vehicles and buildings demand come from a UK-funded customer-led network revolution (CLNR) project.⁴ Both datasets are thoroughly described in Appendix A. An example of demand and production values generated by our simulator for one day is shown in Figure 5-4. As shown in the graph, at some time intervals, renewable energy production is more than the demand of buildings and EVs, while in others less. Also, on this particular day, demand is high at about 17:00 in the afternoon. According to the original dataset, the cost of microgrid’s consumption for this day

³<http://www.sotaventogalicia.com/en>

⁴<http://www.networkrevolution.co.uk/resources/project-data/>

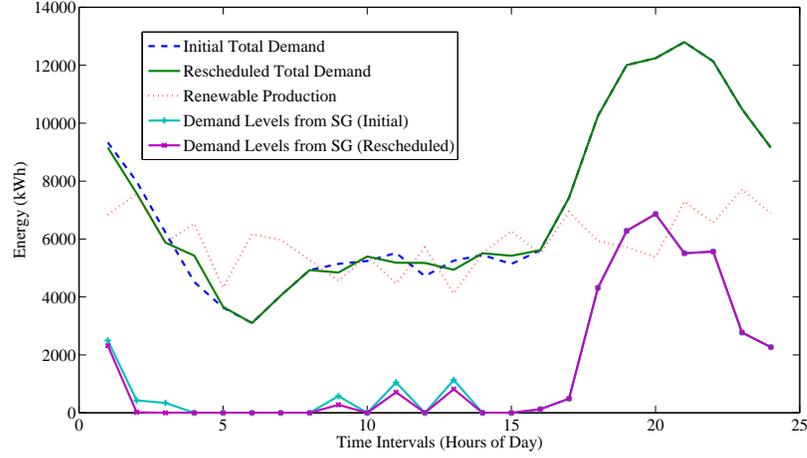


Figure 5-5: Hourly aggregate values for a sample day that is produced from the simulator after V2G/G2V is applied.

is shown to be €4,941.52. The number of EV candidates for participation that were included in the simulation was 1000 vehicles.

Next, we apply the proposed V2G/G2V mechanism, and the charging and discharging of specific vehicles is rescheduled. This has the effect of reducing demand at some peak time intervals, and increasing to some non-peak. The reprogrammed consumption profile is such that the production of renewable energy is better utilized. The result of the algorithm using the data of the sample day of Figure 5-4 is shown in Figure 5-5. As we can see, during peak intervals, demand decreases and is shifted to non-peak intervals. Importantly, the increase in non-peak is not enough so as demand to exceed the levels of production. Also, the size of the aggregate demand without V2G/G2V applied, is equal to the size of the rescheduled demand as in the results from the V2G/G2V application, meaning that no vehicle has remained with a charge that is lower than the levels desired, despite the involvement of some of

them in the mechanism.

In addition, the cost of microgrid's consumption dropped from €4,941.52 to €4,868.34, resulting in a profit of €73.18, or 1.4% of the initial cost. It is also worth mentioning that part of this amount was given back to the end-users as a payment for the electricity they supplied by discharging their batteries.

5.3.1 Simulation over a one-year time horizon

In this part, we run the daily simulations for 365 days, and observe the average indices over the year. Figure 5-6 shows the average initial consumption, reprogrammed consumption curves and renewable generation levels. As we can see, consumption curves generally vary, depending on whether the initial consumption is higher or lower than the total output. Exceptions are the time intervals 18-24, where due to the scheduling horizon length (24 hours) and the particular set of data, it is not possible to find suitable later non-peak intervals to transfer the consumption to. This issue can be solved easily, either by increasing the scheduling horizon, or by executing the algorithm twice a day, in partially overlapping periods of time. In any case, however, in the "final" time horizon, this phenomenon will still be displayed. A more detailed view of the original and reprogrammed vehicle consumption is given in Figure 5-7. Regarding the microgrid's total energy costs i.e., the required funds for importing energy from the Smart Grid during the 365 days of V2G/G2V application, it decreased from €1,951,125 to €1,932,206, resulting in a profit of €18,919 in total, a decrease of 0.9%. The average profit per day is €51.83.

Turning now to the gains of vehicle owners in their total, with average values per participation, the results are shown in Table 5.1: As we can see, EV owners that

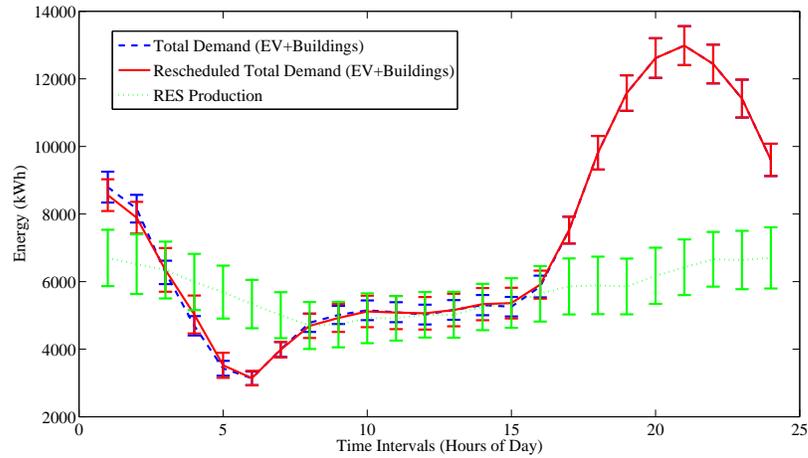


Figure 5-6: Daily mean and standard deviation values from the 365 days simulation.

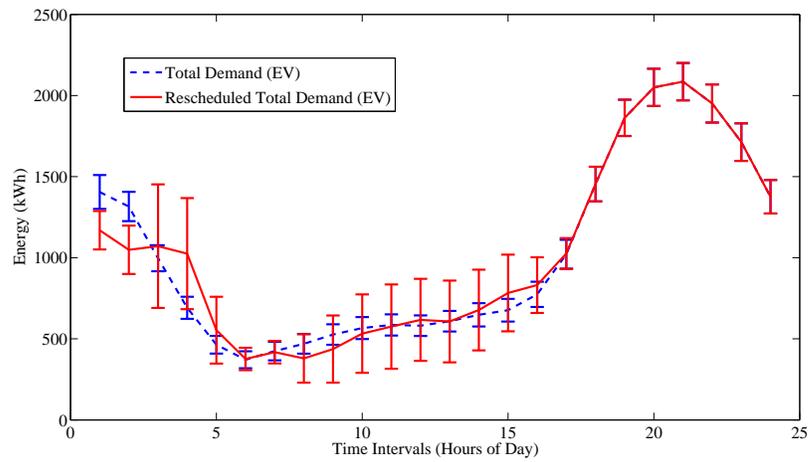


Figure 5-7: Daily mean and standard deviation from EV demand values of the 365 days simulation.

Table 5.1: V2G participations—energy from batteries only.

No. of participations (daily avg.)	44588 (122.15)
Avg. reservation price of accepted offers in euros	0.082
Avg. rescheduled load per participation in kWh	6.89
Avg. profit per participation in euros (percentage of daily cost)	0.56 (0.01%)
Total profit for 365 days in euros	25.302,07
Total cost of shifted demand in euros	3.073,15
Actual profit for EVs in euros	22.228,91
Profit percentage wrt. initial cost	1.29%

Table 5.2: Imports and exports of energy.

	Initial levels	Levels with V2G/G2V	Diff. (%)
Avg daily imports from SG (kWh)	42.706,38	42.116,73	-1.38
Avg daily microgrid renewable excess (kWh)	6.942,17	6.352,53	-8.49
Yearly energy imported from SG (kWh)	15.587.830	15.372.610	-1.38
Yearly daily microgrid renewable excess (kWh)	2.533.894	2.318.673	-8.49

participated in V2G, managed to achieve significant monetary gains, by rescheduling only 6.89 kWh per participation on average.

Next, regarding results of the imports and exports of electricity, we present the values before, and after the implementation of the V2G / G2V algorithm; Figure 5-8 shows the input levels (time intervals for which the curve is greater than zero) and the output (time intervals for which the curve lies below zero). Note that after the implementation of V2G / G2V, the difference is closer to zero, reflecting the balancing between consumption and production levels of renewable. This results in a more efficient absorption of renewable energy, and in reductions of imports from external networks. Cumulative results for the 365 days simulation are shown in Table 5.2.

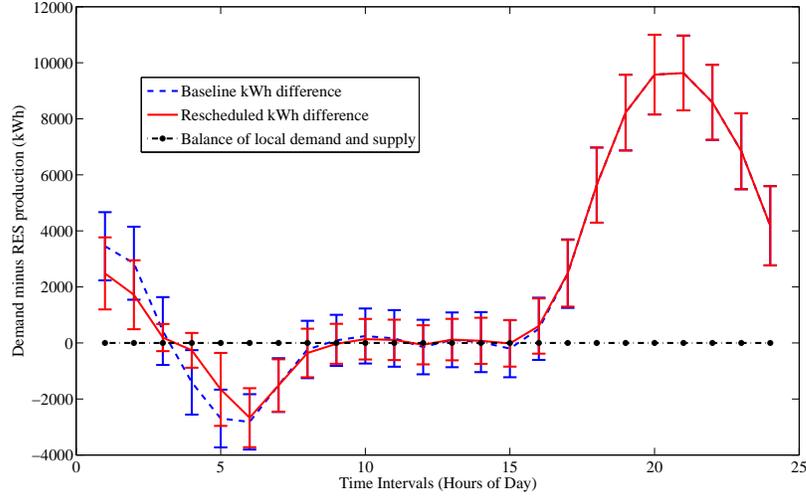


Figure 5-8: Typical daily imbalance curve during the year.

Finally, we note that the average time for the algorithm to run for one day is 0.41 seconds.

5.3.2 EV reservation prices sensitivity test

In this simulation setting, we add a factor (termed as reservation price offset) to the cost of each vehicle owner so that we can observe the earnings of the microgrid based on this change. Specifically, in subsequent iterations, we add to each agent's reservation price ⁵ (each electric vehicle, every hour of the day) a value starting at -0.06 and reaching up to +0.025, with a 0.005 increment per iteration. The results are shown in Figure 5-9. As a matter of fact, the lower the owner costs, the greater the cost savings for the microgrid regulator's perspective. However, if reservation prices

⁵Recall that an agent's reservation price is the price the EV owner charges for providing one kilowatt-hour of energy from the EV's battery to the microgrid.

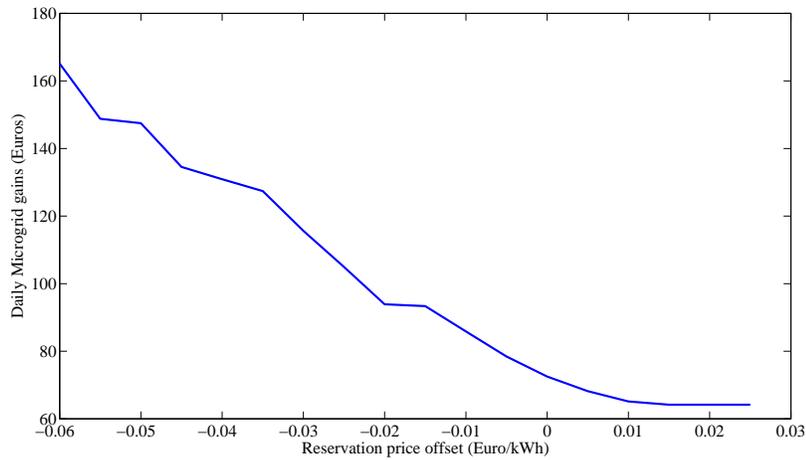


Figure 5-9: Daily microgrid gain vs. EV reservation prices offset.

rise significantly i.e., as we move to the right on the x-axis, then Smart Grid input is more advantageous, and none of the EV's batteries are selected for discharging. Moreover, we observe that although reservation prices may significantly increase after some point, a profit still exists. This (much smaller) profit is generated by the application of the G2V part of the mechanism, i.e. shifting demand of EVs at later time intervals when renewable generation is in abundance.

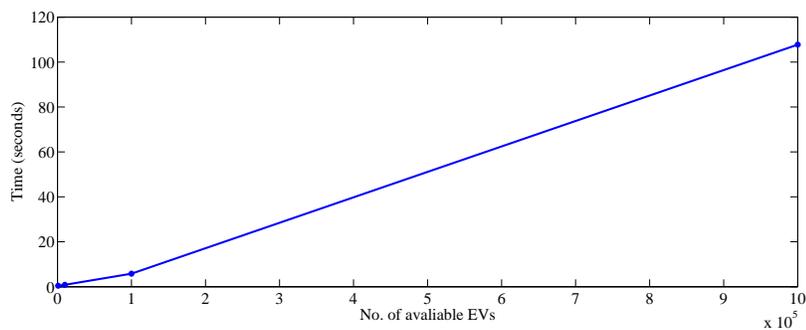


Figure 5-10: Average time for the calculation of daily consumption rescheduling as the number of incorporated EVs rises.

Finally, Figure 5-10 shows the average time in seconds for running the proposed V2G/G2V algorithm. We can see that the algorithm scales well (less than two minutes) for rescheduling the daily consumption patterns even when one million EVs are present in the microgrid.

5.3.3 Performance comparison

Here, we compare the performance of the proposed *V2G/G2V* algorithm with an already existing charging/discharging optimization method found in literature, that of [87]. This method, termed as *V2G/G2V-OPT*, formulates the problem using mathematical programming, and assumes an iterative procedure where each EV optimizes a day-ahead its charging/discharging curve based on a virtual price signal. Note that this virtual price is not related to the actual individual costs of EV owners, and is based on the difference between the locally produced renewable energy and the imports from the Smart Grid. As renewable excess increases among time intervals, the virtual price gets lower, guiding this way EVs to discharge during intervals with higher virtual prices, and to charge when these are lower. In this experimental setting, due to the fact that according to *V2G/G2V-OPT* each EV solves an optimization problem—one per iteration until the algorithm converges, and that our simulations were executed on a single computer, we drop the number of EVs to 30, 50, and 100, respectively, in order to allow the simulations to complete within a reasonable time frame. Results are shown in Table 5.3.

As we can see, both approaches manage to reduce microgrid energy imports from the Smart Grid and to increase the utilization of locally generated renewable energy, performing better as participating EVs increase in numbers. We can also see that the

Table 5.3: Shifting performance of the most active participants per consumer class and coalition formation method, over 100 days.

Measure	Num. of EVs	$V2G/G2V$	$V2G/G2V-OPT$
Average reduction in microgrid daily imports from Smart Grid (%)	30	0.05	0.46
	50	0.16	1.00
	100	0.37	1.86
Average reduction in microgrid daily renewable energy excess (%)	30	0.07	0.72
	50	0.33	1.99
	100	0.74	3.75
Average microgrid daily gains from rescheduling (€)	30	1.02	-890.00
	50	1.01	-1411.00
	100	2.38	-2848.88
Average execution time per EV (sec)		$7.2 \cdot 10^{-5}$	5.07

$V2G/G2V-OPT$ method is more effective with respect to these two measures. However, this performance improvement comes with an increased cost for the microgrid side, as result from the rescheduling. This is natural, since $V2G/G2V-OPT$ does not account for EV reservation prices at all. On the contrary, our $V2G/G2V$ algorithm effectively constraints negative microgrid gains, and guarantees that there will be profit from the rescheduling operations, both for the EV side, and as concerns the microgrid regulator as well. Additionally, $V2G/G2V$ is significantly faster than its $V2G/G2V-OPT$ alternative, indicating that it can incorporate really large numbers of EVs without requiring strong computational power from the EV side. Note that, although centralized by formulation, $V2G/G2V$ can be executed in a decentralized fashion as well, using e.g. a blockchain-oriented protocol as proposed in Chapter 4.

5.4 Conclusions

In this chapter, we studied the incorporation of a V2G / G2V mechanism for EVs in DSM, which helps to balance demand and production from renewable energy sources in day-ahead planning in smart electricity networks. The algorithm asks drivers to state the vehicle's connection / disconnection schedule, as well as a personal energy reservation price for discharging batteries and providing power to the grid. After checking relevant limitations, such as the duration of each vehicle's electrical connection, the cost of discharging, the energy costs from external imports, and the next day's demand and production forecasts, vehicle batteries either delay their charging, or offer part of their stored energy to the grid. This results to increased use of the microgrid's renewable resources, and to reduction in imports from external energy sources. Upon participation, the owners of EVs are rewarded, without impacting the comfort of using the vehicle. Results from an experimental simulation process based on real-world datasets, showed the benefits of using V2G/G2V mechanisms with respect to both monetary and environmental concerns, since renewable self-production is better utilized.

Chapter 6

Machine Learning for Enhanced Cooperative Performance

In order for agent cooperatives to be functional, efficient, and profitable, they need to take business decisions regarding which members to include in consumption shifting coalitions, from a very large set of available end-users. These decisions naturally depend on the abilities (e.g., electricity production or consumption reduction capacities) of individual agents. These abilities need to be either monitored by some central cooperative-managing agent, or need to be truthfully and accurately communicated to it. However, it is clear that in the large and dynamically changing scene of the Smart Grid, trust between selfish agents is not implied, and must be guaranteed. As extensively mentioned in this dissertation already, mechanism design and related approaches attempt to build trust among cooperative members, via providing them with the incentives to truthfully and accurately report their intended future actions, along with their corresponding uncertainty regarding those actions [7, 40, 96, 155].

Unfortunately, even if participating agents are perfectly truthful regarding their abilities and corresponding uncertainty, their reports and estimates can still be highly inaccurate. This can be due to, for example, communication problems, malfunctioning equipment, or prejudiced beliefs and private assumptions—e.g., a truthful reporting agent might be overly pessimistic or optimistic.

As a result, monitoring the performance of individuals and correctly predicting their future contributing potential is of utmost importance to a cooperative or an organization relying on the services of selfish, distributed, autonomous agents. To this end, several approaches try to explicitly estimate agent electricity consumption and production amounts, by incorporating prediction models that rely on agent geographical location and weather forecasts, or the processing of macroeconomic data [84, 136]. Although their results are promising, such methods cannot immediately predict the actual behaviour of a specific agent, which might be motivated by private knowledge or business concerns, neither do they account for errors due to equipment malfunction. By contrast, this chapter proposes the application of generic prediction methods, which are nevertheless able to adapt to a specific agent’s behaviour regarding the promised and final consumption shifting actions. These techniques could also be used to assess the reliability of the stated preferences and constraints of the EV drivers that were used as input in the coalition formation algorithm of the previous chapter (though we did not employ them in that context of this thesis).

In more detail, the techniques presented in this chapter are incorporated within the effective mechanism of Chapter 3, that promotes the formation of agent cooperatives for power consumption shifting. Taking that specific model for granted, we now showcase the use of *machine learning methods* to keep track of the pa-

rameters that best describe agent behaviour, and effectively estimate actual future agent performance. These techniques are able to not only fit the dynamics of the processes governing agent performance, but can also imbibe the potential errors of electricity metering or information transmission devices. In particular, we adopt the *Histogram Filter (HF)* [189], a method that estimates the underlying distribution of agent behavior using past occurrences, and three regression techniques, i.e. *Gaussian Processes for regression (GP)* [144], *k-Nearest Neighbours Regression (k-NN)*, and *Kernel Regression (KR)* [63, 75], which receive as input agent stated estimates regarding their future behavior, and regress using past observed occurrences to provide a prediction. These predictions represent the future actual actions of agents that participate by large numbers in cooperatives and offer electricity demand management services. Broad user involvement in electricity demand shifting is crucial, because during peak times the amounts of load required to be shifted or curtailed by the power grid is in general quite larger than the total demand of a single consumer. Thus, in order to gather enough shifting capacity, the cooperative must consider as many candidates for participation as possible.

The methods that we adopt here are very generic, and have wide areas of application. Their employment in the power consumption shifting domain ensures that member agents can be ranked by the cooperative according to their perceived consumption shifting capabilities; and thus untruthful or inaccurate agent statements regarding their capacity and corresponding uncertainty will not be able to jeopardize the stability and effectiveness of the overall mechanism governing the cooperative business decisions (e.g., which agents to select for consumption shifting at a given point in time). This is key for the economic viability of any such cooperative.

To the best of our knowledge, ours is the first work to use filtering and regression methods for assessing the performance of autonomous, economically-minded agents participating in Smart Grid cooperatives or other such entities. Our simulations demonstrate that employing any of the methods tested leads to improved performance in the consumption shifting domain, when compared to the “baseline” mechanism of Chapter 3 that makes no use of such prediction tools. Specifically, when using these enhancements, the cooperative achieves higher overall electricity consumption reductions; and enjoys financial rewards that are higher than those generated by the baseline algorithm.

All methods appear to be able to provide reasonable predictions regarding the actual performance of individual agents, given the agents’ stated intended actions and related uncertainty. As such, the efficiency of these machine learning methods is not restricted in the demand management and peak-trimming domains, but they can arguably be readily employed to monitor the trustworthiness of electricity producers’ statements regarding their intended actions. Additionally, the aggregation of the individual agent values by the cooperative has certain advantages. Apart from the fact that the Grid has to deal with a single entity instead of hundreds, or thousands (individual agent numbers), we also showcase the high prediction accuracy in the aggregate cooperative forecasts, despite the presence of inaccuracies in the individual agent predictions. In a nutshell, our results indicate that the techniques examined in this chapter are strong candidates for monitoring and predicting the trustworthiness of selfish agents in the Smart Grid. We thus believe they deserve to be further evaluated in this direction, since they can bring tangible benefits to business entities operating in this domain. Figure 6-1 summarizes the combined scientific domain and

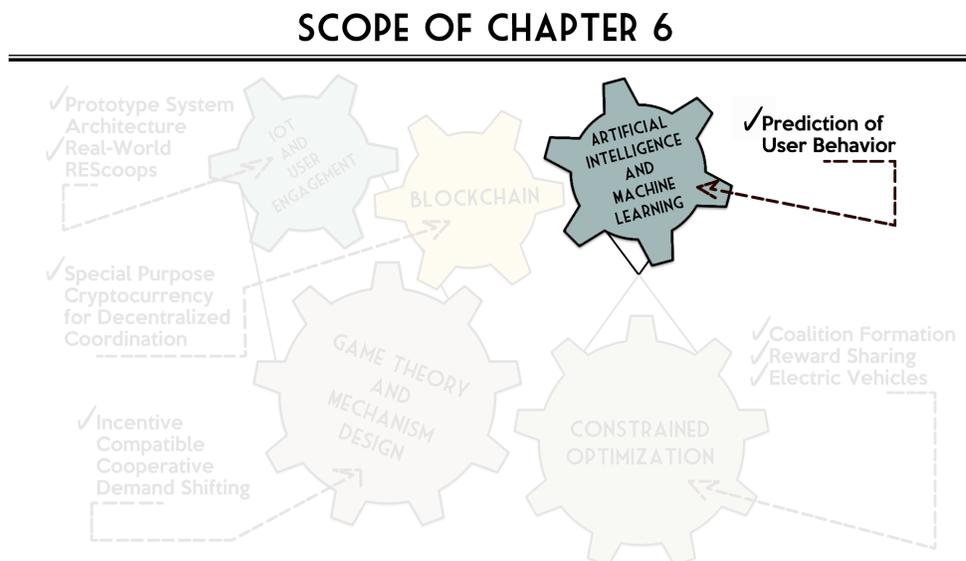


Figure 6-1: Overview of the scientific fields, and our contribution in Chapter 6.

contribution.

The rest of the chapter is structured as follows. Section 6.1 describes past related work. Section 6.2 describes the adopted methods and their application to our problem in detail. Section 6.3 presents our simulations on a large, real-world dataset, and the experimental results. Finally, Section 6.4 concludes.

Parts of the research described in this chapter appeared originally in [8] and in [13].

6.1 Related Work

Here, we briefly review some work on machine learning methods that are used to make predictions in Smart Grid settings. To begin, there is a number of past studies

incorporating methods for forecasting the electricity demand in power networks. For example, in [151], authors employ Support Vector Regression to forecast building consumption in an urban energy management setting. Artificial Neural Networks (ANN) are reported by [206] to also be appropriate for forecasting electricity demand, however they require fine-tuning procedures, which might not be easy to perform in some cases. The work of [30] presents an algorithm to predict the usage of household appliances, and therefore the demand of the household in a real-world setting. The authors of [174] investigate the impact of measurement aggregation in the prediction accuracy of two demand forecasting methods, the Holt-Winters, and the Seasonal Naïve. In our work, we predict the *future shifting behavior* of agents that form cooperatives and participate in a DSM scheme. Agents interact with the Grid collectively, thus aggregation of future statements regarding shifting actions comes natural; and by also adding machine learning approaches for prediction we come up with results of high precision. As our experimental results show, although individual forecasts might not be very accurate and precise, the aggregate load finally shifted by a cooperative is very close to the expected value.

Approaches for forecasting the production of electricity have also been examined in the past. For example, the work of [170] combines machine learning techniques and weather forecasts to predict the output of solar power generators. Authors in [188] also predict the output of photovoltaic (PV) panels using an ANN with input the current and past power measurements, as well as the air and panel temperatures. The approach of [37] uses autoregressive time-series to model the probability density function of a PV's production levels. Machine learning techniques have been used in the realistic Smart Grid market simulator 'powerTAC' [90] as well. Specifically, the

work of [34] incorporates Gaussian Processes (GP) to predict electricity demand; and the approach of [46] uses ANNs, Decision Trees, and Linear Regression to predict prices in the wholesale market. Supervised machine learning has been applied for predicting the levels of user comfort when altering the usage patterns of devices, in order to seamlessly achieve more energy efficient operation [104]. A comparison of the performance of three approaches is presented, a Support Vector Machine, an ANN, and a Naive Bayes, and results show that even with small numbers of training samples, high overall accuracy can be achieved.

A system for predicting the output of renewable energy resources is proposed in the work of [136]. Authors utilize non-linear approximation components to offer a low-cost web-based tool for forecasting the production levels of wind turbines and photovoltaic generators.

The work of [44] deals with the problem of estimating the actual electricity demand, during DR procedures. Estimating the actual demand during DR is not so straightforward as estimating the baseline demand, due to lack of historical data, and due to the concept drift that is implied, i.e. two different consumption strategies, during baseline and during DR. So, authors examine ways to effectively estimate these values, by comparing some simple averaging of historical measures methods. The performance produces low error levels, however, authors argue for an ensemble prediction method that adapts accordingly.

In our work here, we enhance the scheme presented in Chapter 3, by proposing four distinct prediction techniques that predict the power consumption shifting efforts of participating agents. When an agent shifts some load, its prior statements and final actions are monitored, and a corresponding model is induced. Then, instead of

only taking individual statements into account, the learned model is used to better predict individual agent and cooperative action quality, and thus, improve monetary benefits and general performance.

6.2 Method Description

To begin, we denote the *actual* amount of load reduced by r_i^{th} . In general, it can be assumed to be provided by a transformation of the stated \hat{r}_i^{th} amount:

$$r_i^{th} = \alpha_i \hat{r}_i^{th} \quad (6.1)$$

with the (observed) ‘accuracy factor’ α_i corresponding to a random variable characterizing the accuracy of the statement regarding the promised shifting amount. This variable follows some unknown probability distribution.

Let $\hat{\sigma}_i$ denote the stated variance of agent’s i relative error between stated and final actions;¹ the objective of our work in this section is to find a regression function f such that $f(\hat{\sigma}_i)$ is a ‘good approximation’ of the ‘accuracy factor’ α_i , i.e., $\tilde{\alpha}_i = f(\hat{\sigma}_i)$ and α_i are close in some sense. We can then use this estimate to obtain our trusted index r_{i,t_h}^* as follows:

$$r_{i,t_h}^* = \tilde{\alpha}_i \hat{r}_i^{th}$$

Given all underlying uncertainty, an individual agent’s final behaviour most likely corresponds to a complex, non-linear function of its past behaviour. This is translated

¹Recall from the previous section that $\hat{\sigma}_i$ models the uncertainty regarding the reduction capacity of each agent i .

into a highly non-linear (unknown) function f in our regression model $\tilde{\alpha}_i = f(\hat{\sigma}_i)$. We choose to test four different approaches that are able to provide an estimate \tilde{a}_i , given $\hat{\sigma}_i$ and a set of T past observations $\mathcal{D} = \{(\hat{\sigma}_i^{-T}, \alpha_i^{-T}), \dots, (\hat{\sigma}_i^{-1}, \alpha_i^{-1})\}$: (i) the *Histogram Filter (HF)* [189], (ii) an approach utilizing a *Gaussian Process (GP)* [144], (iii) one using *k-Nearest Neighbours (k-NN)* regression [63], and (iv) one using *Kernel Regression (KR)* [75]. These methods are chosen among others due to their ability to fit any relationship, either linear or non-linear, between the variables of the model. Additionally, *k-NN* and *KR* are (weakly) universally consistent methods; in words, as the size of the dataset increases, the precision of the estimates increases as well [75].

Here, we must note that these methods require an adequate amount of historical data collected, in order to obtain enough information to create reasonable predictions. To this purpose, the *conservative estimate* can be employed at first, and then get replaced by one of the proposed methods once enough data is available.

The *GP* uses historically observed pairs of agent statements and final actions values to fit normal distributions for the underlying random variables for each statement value. The *HF* ignores user stated uncertainty over the performance and takes into account only past observations for its predictions. On the other hand, the rest of the chosen methods assume that final actions are functions of user stated uncertainty, and perform accordingly. All four methods are generic, with very wide areas of application. Moreover, not only are they able to fit the dynamics of the processes governing agent performance, but can also imbibe the potential errors of electricity metering or information transmission devices. Here we employ them to enhance the performance of our energy consumption shifting mechanism. By so doing, more accurate agent rankings can be obtained, and inaccurate statements will not be able

to jeopardize the stability and effectiveness of the mechanism. We now describe methods that we examine, and their application to our setting in detail.

6.2.1 Histogram filter

Assume that for each of the agents, the accuracy factor α_i is bounded below and above by L_i and U_i respectively, i.e., $\alpha_i^t \in [L_i, U_i]$, for all $t \in \{1, 2, \dots\}$. Histogram filters partition the interval $[L_i, U_i]$ into a set of K disjoint bins $[l_k^i, u_k^i]$ such that:

$$[L_i, U_i] = \cup_{k=1}^K [l_k^i, u_k^i]$$

The *HF* uses a histogram to map a probability p_k to each of the K bins $[l_k^i, u_k^i]$. The value of each p_k depends on the frequency of the observations in the range of bin k .

With this approach, agent forecasts $\hat{\sigma}_i$ are completely ignored and only past observations of α_i are taken into account. Every time an agent participates in a consumption shifting coalition, its actions are monitored and stored. A histogram is calculated over the set of available observations. Then, according to each bin's height, a colored roulette wheel is constructed that can be sampled to obtain the most probable ranges of α_i , i.e. the more frequent values appear in a bin the more probable its range is selected. So, we sample the corresponding roulette wheel, and come up with a specific bin. The final $\tilde{\alpha}_i$ estimate is another sample from a uniform distribution normalized to have range equal to that of the bin obtained, i.e.:

$$\tilde{\alpha}_i \sim \mathcal{U}(l_k^i, u_k^i)$$

The advantage of *HF* is that it requires no prior knowledge about the form of the distribution that α_i follows, and adapts effectively to all kinds of non-linearities [189]. On the other hand, it needs a number of measurements before it starts working accurately and performance might be unacceptable in initial stages with no actual measurements. Another drawback is that if the distribution changes over time, the length of a history window must be re-set, in order to get rid of expired measurements interference. Also, it does not take into account the observed error standard deviation $\hat{\sigma}_i$ that is, the agent-stated confidence.

6.2.2 Gaussian process filter

Here we describe the application of Gaussian processes for probabilistic regression, to construct a regressor that is used for monitoring and prediction purposes.² Given a set of training samples, $\mathcal{D} = \{(\mathbf{x}_j, y_j), j = 1, \dots, n\}$ (\mathbf{x}_j inputs and y_j noisy outputs) we need to predict the value of y_* for an input vector \mathbf{x}_* . We assume the following model:

$$y_j = f(\mathbf{x}_j) + \epsilon_j, \text{ where } \epsilon_j \sim \mathcal{N}(0, \sigma_{noise}^2)$$

with σ_{noise}^2 the variance noise.

GP regression is a Bayesian approach that assumes a priori that function values follow: $p(\mathbf{f}|\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n) = \mathcal{N}(\mathbf{0}, K)$ where $\mathbf{f} = [f_1, f_2, \dots, f_n]^T$ is the vector of latent function values, $f_j = f(\mathbf{x}_j)$ and K is the covariance matrix that is computed by

²Note that in a previous short paper [8], we also explored the use of *GP-UKF*, an *unscented Kalman filter combined with Gaussian process regression*. However, in order to exploit the full power of that technique, in reality one needs to have access to a realistic model of the stochastic dependencies among the past $\hat{\sigma}_i$ agent statements. Without such a model, one cannot observe significant differences between using *GP-UKF* and *GP* alone.

a covariance function $K_{jk} = k(\mathbf{x}_j, \mathbf{x}_k)$. The joint *GP* prior and the independent likelihood are both Gaussian with mean and variance as follows:

$$GP_\mu(\mathbf{x}_*, \mathcal{D}) = K_{*,f}(K_{f,f} + \sigma_{noise}^2 I)^{-1} \mathbf{y} \quad (6.2a)$$

$$GP_\sigma(\mathbf{x}_*, \mathcal{D}) = K_{*,*} - K_{*,f}(K_{f,f} + \sigma_{noise}^2 I)^{-1} K_{f,*} \quad (6.2b)$$

GPs also require value assignments to the vector $\theta = [\mathbf{W} \ \sigma_f \ \sigma_{noise}]$ that contains the hyperparameters, with \mathbf{W} holding the distance measure of each input in its diagonal, σ_f being the variance of the input and σ_{noise} the variance of the process noise. We can find the optimal values for θ by maximizing the log likelihood:

$$\theta_{max} = \arg \max_{\theta} \{\log(p(\mathbf{y}|\mathbf{X}, \theta))\} \quad (6.3)$$

In our application setting, \mathbf{x}_* is the reported agent confidence $\hat{\sigma}_i$ for the interval under examination, and $\mathcal{D} = \{(\hat{\sigma}_i^{-T}, \alpha_i^{-T}), \dots, (\hat{\sigma}_i^{-1}, \alpha_i^{-1})\}$ the T pairs of past agent values. Finally, the estimate of future agent behavior can be calculated by:

$$\tilde{\alpha}_i = GP_\mu(\hat{\sigma}_i, \mathcal{D}) + u_\tau \quad (6.4)$$

with noise u_τ following $\mathcal{N}(0, GP_\sigma(\hat{\sigma}_i, \mathcal{D}))$. Note that this approach takes into account agent statements regarding their uncertainty, $\hat{\sigma}_i$, thus incorporates more prior information than the *HF* approach. Also, the matrix inversion due to the term $(K_{f,f} + \sigma_{noise}^2 I)^{-1}$ results to a cubic computational complexity in order to obtain the predictions.

6.2.3 k -nearest neighbours regression

Nearest Neighbor methods are among the simplest of all machine learning algorithms. The idea is to memorize the training set and then to predict the value on any new input instance on the basis of the values of its closest neighbors in the training set [63].

More formally, let $\mathcal{D} = \{(\mathbf{x}_j, y_j), j = 1, \dots, n\}$ be a sequence of training examples. For each x , let $\pi_1(x), \dots, \pi_n(x)$ be a reordering of $\{x_1, \dots, x_n\}$ according to their distance from x , $d(x, x_i)$; d is any distance(metric) function. That is, for all $i < n$,

$$d(x, x_{\pi_i(x)}) \leq d(x, x_{\pi_{i+1}(x)})$$

For a number k , the k -NN rule for regression is defined as follows:

- Input: a training set $\mathcal{D} = \{(\mathbf{x}_j, y_j), j = 1, \dots, n\}$ and a query point x
- Output: return $f_{k\text{-NN}}(x) = \frac{1}{k} \sum_{l=1}^k y_{\pi_l(x)}$

As previously for the GP , here, the query point x is the reported agents' confidence $\hat{\sigma}_i$ for the interval under examination, and $\mathcal{D} = \{(\hat{\sigma}_i^{-T}, \alpha_i^{-T}), \dots, (\hat{\sigma}_i^{-1}, \alpha_i^{-1})\}$ consists of the T pairs of past agent values. Let $\pi_1(\hat{\sigma}_i), \dots, \pi_T(\hat{\sigma}_i)$ be a reordering of $\{\hat{\sigma}_i^{-T}, \dots, \hat{\sigma}_i^{-1}\}$ according to their distance $d(\hat{\sigma}_i, \hat{\sigma}_i^{-t}) = |\hat{\sigma}_i - \sigma_i^{-t}|$ from the target point $\hat{\sigma}_i$.

The estimate $f_{k\text{-NN}}(\hat{\sigma}_i)$ is calculated as the average of the k nearest values of $\alpha_i^{-\pi_l(\hat{\sigma}_i)}$ as

$$\tilde{\alpha}_i = \frac{1}{k} \sum_{l=1}^k \alpha_i^{-\pi_l(x)} \quad (6.5)$$

The main computational load in the calculation of $f_{k\text{-NN}}(\hat{\sigma}_i)$ is the calculation of distances between the query point $\hat{\sigma}_i$ and all samples $\hat{\sigma}_i^{t-}, t = \{1, \dots, T\}$ in the training set \mathcal{D} . We discuss the choice of k in the experimental section. We just mention that for the k -NN to perform well in practice the fraction $\frac{k}{T}$ should tend to 0. In words, $\frac{k}{T} \rightarrow 0$ implies that the samples in the training set and the size of the neighborhood are such, that local changes in the distribution of $\hat{\sigma}_i$'s can be effectively captured by k -NN.

The graph of the function $f_{k\text{-NN}}(\cdot)$ is “bumpy”, since $f_{k\text{-NN}}(x)$ is discontinuous in x . To see this, consider the simple case of a dataset \mathcal{D} with n pairs (x_j, y_j) of scalar values. As we move from a query point x from left to right, the k -nearest neighborhood remains constant, until a point x_i to the right of x becomes closer than the furthest point $x_{i'}$ in the neighborhood to the left of x , at which time x_i replaces $x_{i'}$. The average in $\frac{1}{k} \sum_{l=1}^k y_{\pi_l(x)}$ changes in a discrete way, leading to discontinuous $f_{k\text{-NN}}(x)$.

6.2.4 Kernel regression

Now, the discontinuity of $f_{k\text{-NN}}(x)$ can be tackled. Rather than give equal weights to all points in the neighborhood of point x , we can assign weights that die off smoothly with distance from x . This is the approach taken in the so-called Nadaraya-Watson (NW) kernel weighted regression, presented in [75]. Specifically, continuing with the same set up as in k -NN regression, i.e., a set of training samples \mathcal{D} and a query point

$x \in \mathbb{R}^d$, in kernel regression we consider the weighted average of the form

$$f_{\text{NW}}(x) = \frac{\sum_{i=1}^n K_h\left(\frac{x-x_i}{h}\right)y_i}{\sum_{i=1}^n K_h\left(\frac{x-x_i}{h}\right)} \quad (6.6)$$

where $K_h : \mathbb{R}^d \rightarrow \mathbb{R}$ is the kernel function, and $h > 0$ a bandwidth.

There are a number of details that one has to attend to in practice when using the kernel regression approach:

- The bandwidth h has to be determined. Large h implies lower *variance* (average over more observations), but higher *bias* (we essentially assume the true function to be constant within the window).
- The observation weights induced by the kernel; there are many choices for $K_h(\cdot)$, some of them having bounded support (for example the Epanechnikov kernel or the tri-cude function) and others not (for example the Gaussian kernel) [63].

As in k -NN, the query point in this section x is the reported agents' confidence $\hat{\sigma}_i$ for the interval under examination, and $\mathcal{D} = \{(\hat{\sigma}_i^{-T}, \alpha_i^{-T}), \dots, (\hat{\sigma}_i^{-1}, \alpha_i^{-1})\}$ consists of the T pairs of past agent values. The estimate $f_{\text{NW}}(\hat{\sigma}_i)$ is calculated as:

$$f_{\text{NW}}(\hat{\sigma}_i) = \frac{\sum_{t=1}^T K_h\left(\frac{\hat{\sigma}_i - \hat{\sigma}_i^{-t}}{h}\right)\alpha_i^{-t}}{\sum_{t=1}^T K_h\left(\frac{\hat{\sigma}_i - \hat{\sigma}_i^{-t}}{h}\right)} \quad (6.7)$$

The main computational cost in kernel regression is the calculation of distances between the query point $\hat{\sigma}_i$ and all samples $\hat{\sigma}_i^{-t}$, $t \in \{1, \dots, T\}$ in the training set \mathcal{D} .

6.3 Experimental Evaluation

In this section we conduct extensive simulations of our scheme on the real-world consumption patterns of the dataset used in Section 3.6.1. We compare the results of the collective consumption shifting procedure, when employing the proposed techniques for monitoring and prediction.

Agent statements and final shifting actions

We need a model to describe how *(i)* the agent statements on their uncertainty regarding shifting capacities at t_h occur; and *(ii)* their actual, final shifting actions occur. To this end, we define four agent classes; the first one, denoted as *UU*, describes the most uncertain setting where participants have equal probability to state *any* $\hat{\sigma}_i$ value, and also equal probability to deliver *any* α_i , that is both variables follow a uniform distribution, and are independent. The second one, denoted as *BB*, describes the realistic case where agents are mainly *confident* about their statements, and also have a high probability to deliver what they promised. In the *BB* class, the stated error standard deviation $\hat{\sigma}_i$ and the observed α_i accuracy factor follow two *Beta* distributions, $\mathcal{B}(1, 5)$ and $\mathcal{B}(4, 2)$ respectively. We choose *Beta* distributions primarily because they are very good at representing and updating probabilistic beliefs regarding potential behaviours, as also manifested by their widespread use for simulating behaviours and uncertainties related to real world scenarios (see, e.g., the work of [96]). The use of the particular two aforementioned *Betas* corresponds to error statements having a low mean of $= \frac{1}{6}$ (i.e., to agents stating high certainty), whereas the final observed accuracy factor is closer to 1 (implying increased actual

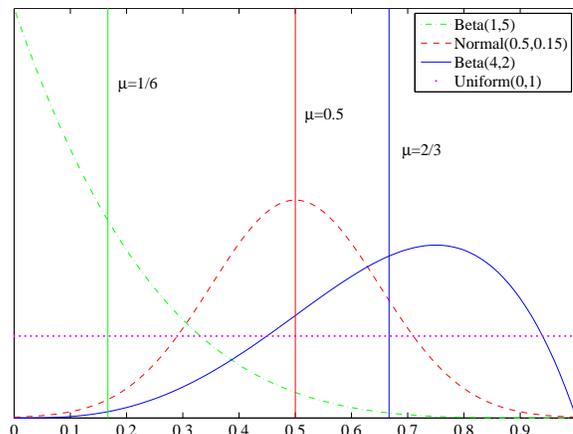


Figure 6-2: Probability density functions of agent behaviors.

performance effectiveness).

The third agent class is the GG , where consumers might or might not follow stated forecasts, so α_i and $\hat{\sigma}_i$, although independent, both follow the same Gaussian—i.e., the $\mathcal{N}(0.5, 0.15)$ distribution.

Finally, the last agent behavior class that we consider is the $1-G$, which is the only case where agent statements and final actions are *not* independent. In particular, agent statements $\hat{\sigma}_i$ are sampled from a Gaussian distribution, i.e. $\mathcal{N}(0.5, 0.15)$, however, the α_i s are calculated by subtracting the $\hat{\sigma}_i$ from 1. In this way, we model a specific underlying relationship between agent statements and final actions. The distributions that we use in our evaluation are depicted in Fig. 6-2. Table 6.1 summarizes the parameters and distributions associated with every class of agents behaviour.

We believe that BB and GG sufficiently capture two realistic, highly plausible scenarios of agent behaviour. Of course, the underlying models of agents behaviour

Table 6.1: Behavioural classes of the scheme participants used in the simulation.

Class	Parameter	Distribution
<i>UU</i>	$\hat{\sigma}_i$	$\mathcal{U}(0, 1)$
	α_i	$\mathcal{U}(0, 1)$
<i>BB</i>	$\hat{\sigma}_i$	$\mathcal{B}(1, 5)$
	α_i	$\mathcal{B}(4, 2)$
<i>GG</i>	$\hat{\sigma}_i$	$\mathcal{N}(0.5, 0.15)$
	α_i	$\mathcal{N}(0.5, 0.15)$
<i>1-G</i>	$\hat{\sigma}_i$	$\mathcal{N}(0.5, 0.15)$
	α_i	$1 - \mathcal{N}(0.5, 0.15)$

could follow any other distribution as well, this is why we also incorporate the *UU* and *1-G* agent behavior.³ However, these behavioural classes are less interesting since their behaviour is easier to predict; and, moreover, they are likely in practice, as (i) in realistic settings, errors do occur, while (ii) if predictions are highly inaccurate, the agents would most probably be acting upon them already, to avoid penalties. In the following, we simulate five scenarios where we change the agent behavior classes for each case.

6.3.1 Evaluation of monitoring techniques

To begin, we provide some details regarding the implementation of the simulations, which were conducted in Matlab, using a 3GHz CPU, and 4GB of RAM. The *HF* approach is tailor made for our needs, using standard Matlab methods. For the *GP* implementation, the GPML Toolbox [149] was used, with a composite mean function (sum of linear and constant), isotropic Matern covariance, and Gaussian likelihood.

³For instance, in [8], the classes of *accurate* and *inaccurate* predictors were also introduced.

The k - NN method was a custom implementation using the *euclidean distance*, and $k = \sqrt{n}$, where $n = 100$ the number of past samples. Finally, the KR method is custom made as well, using $h = 1$ as parameter value, and the exponential kernel.

To test-evaluate the performance of the HF and GP techniques, we first applied them on a single agent of the BB class, trained over 1000 past value couples that were generated using the same distributions as in the simulation. Figures 6-3 and 6-4 show the outcomes for the HF and GP , respectively. We can observe that although the HF does not take into account agent $\hat{\sigma}_i$ statements, it fits well to the real underlying distribution.⁴

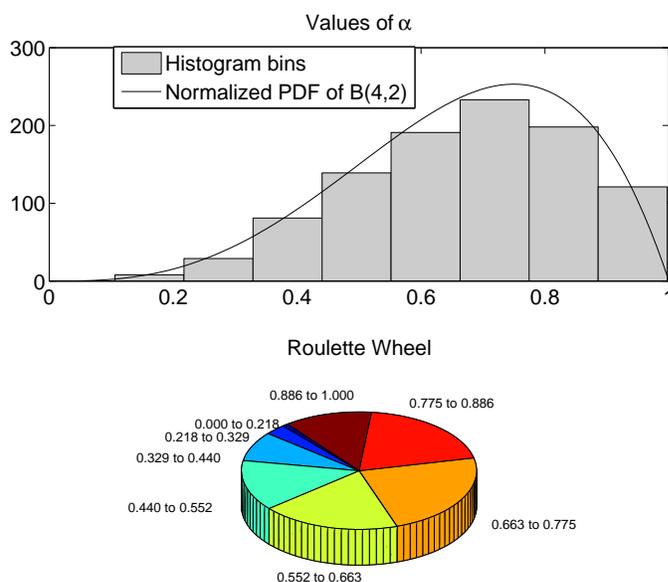
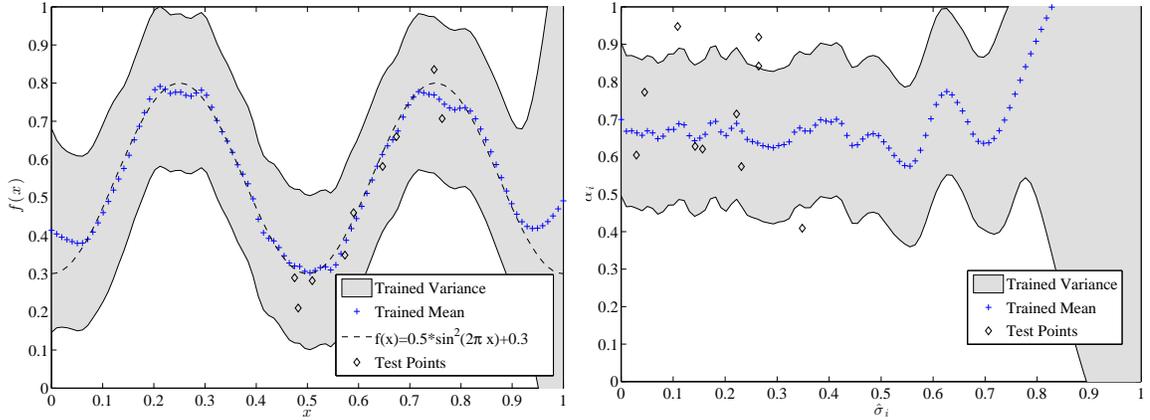


Figure 6-3: Histogram with 8 bins and corresponding roulette wheel.

The GP , on the contrary, does take agent commitments into account. As we can

⁴Recall that in the BB setting, the $\hat{\sigma}_i, \alpha_i$ values are statistically independent.

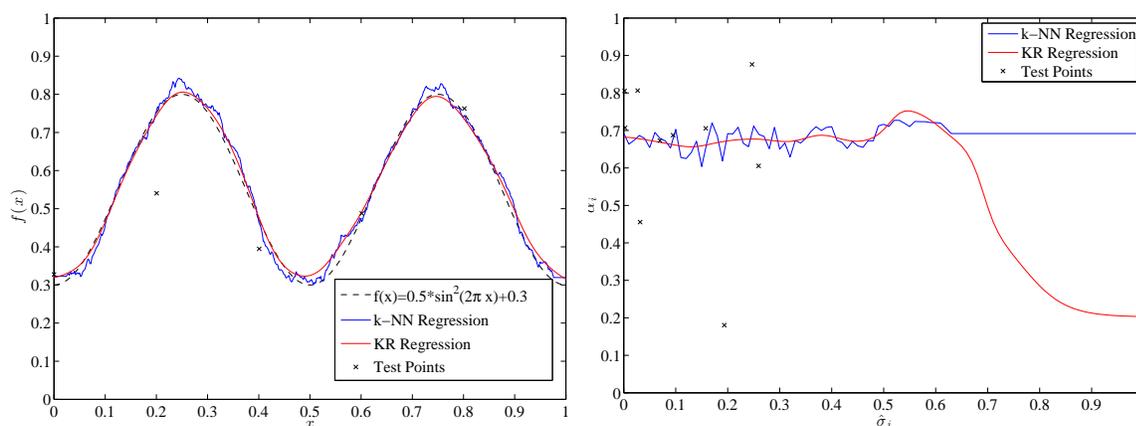


(a) Illustrating the *GP* capabilities: fitting an arbitrary non-linear function. (b) Fitting the *BB* behavioural class of our simulation scenario.

Figure 6-4: *GP* fits. Illustration of *GP* training for two non-linear cases.

observe in Figure 6-4a, the *GP* fit an arbitrary noisy non-linear scenario, proving that non-linearities can be handled. Fig. 6-4b presents the trained variances and means for a typical agent participating in the cooperative shifting process for the setting where $\hat{\sigma}$ and α_i follow $\mathcal{B}(1, 5)$ and $\mathcal{B}(4, 2)$ respectively. In this case, there is no function of a specific known form for the Gaussian process to approximate, as the points are both random variables following different distributions. Despite that, *GP* has converged to some relationship between input and output values. We can infer that this estimated complex function is meaningful, by the fact that most ‘test points’ fall within the shaded area representing the *GP* output variance; the ‘test points’ plotted in this case are random $\hat{\sigma}_i, \alpha_i$ values, sampled by the $\mathcal{B}(1, 5)$ and $\mathcal{B}(4, 2)$ respectively. Thus, *GP* is apparently able to produce meaningful predictions, even when the relationships between variables are governed by some highly complex function. We can observe that the main mass of the observations is in the upper

left quadrant ($E[\hat{\sigma}_i] \approx 0.16$ and $E[\alpha_i] \approx 0.66$) and that the trained means are close to that area. Note that because $\mathcal{B}(1, 5)$ gives very low to zero probability for $\hat{\sigma}$ values between 0.7 and 1, the number of corresponding training points is very low, so uncertainty in that region is very high. This is not an issue though, as GP is not likely to be asked to provide predictions in that range.



(a) Illustrating the k - NN and KR capabilities: fitting an arbitrary non-linear function. (b) Fitting the BB behavioural class of our simulation scenario.

Figure 6-5: k - NN and KR fits. Illustration of the regression functions for two non-linear cases, with $k = \sqrt{n}$ for the k - NN and $h = 1$ for the KR .

Similarly, we test the two remaining methods, k - NN , and KR . As we can see in Figure 6-5, both methods are able to converge to a regression model that is quite close to the real values of the non-linear function, and of the distribution generating α_i s respectively, quite close to what the GP did. Also, here, we illustrate the ‘bumpiness’ effect of the k - NN approach, and the smoother outcome of the KR , as discussed in Sections 6.2.3, and 6.2.4. It is clear for the case of Figure 6-5a that KR tends to fit the function better, while the k - NN tends to over-fit the given data.

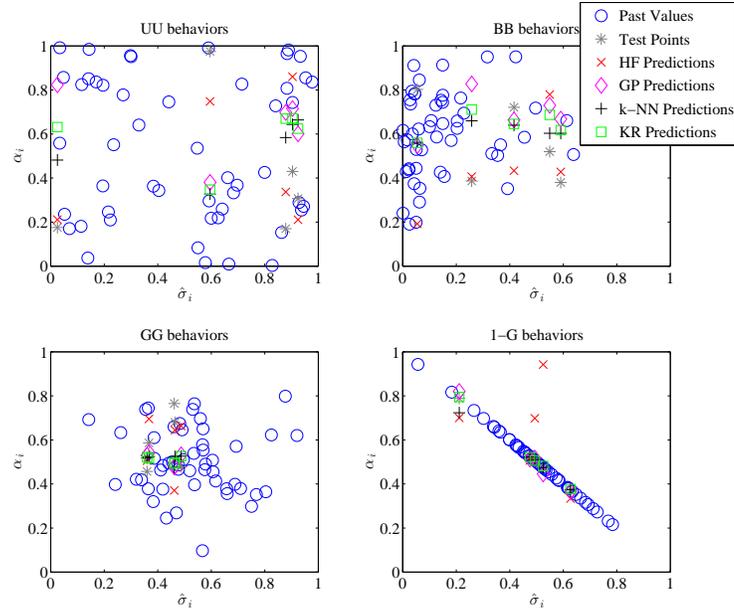


Figure 6-6: *HF*, *GP*, *k-NN*, and *KR* prediction examples for the various agent behavior scenarios.

To continue with the methods evaluation, we conduct some initial experiments to provide further intuition regarding each method's performance. Initially, we sample 50 $(\hat{\sigma}_i, \alpha_i)$ pairs for each of the behavior scenarios of Table 6.1 and consider these to be the past observations. Next, we employ the four proposed techniques one by one, and use them to obtain our predictions. Results are shown in Fig. 6-6. As we can see, for the first behavioral class (*UU*), the predictions $\tilde{\alpha}_i$ of our methods fall quite far from the test points (asterisks). This is expected, however, since there is no actual model that correlates past values of α_i and $\hat{\sigma}_i$ —both variables are sampled from a uniform distribution. Additionally, as we will be explaining in the following section, the individual inaccuracies have less impact when the values are aggregated,

e.g. when a large set of agents shift their electricity consumption in a coordinated fashion.

Regarding the two following behavioral classes BB , and GG , there is no direct connection between agent statements and final actions, but each are sampled from specific distributions, i.e. $\mathcal{B}(1, 5)$, $\mathcal{B}(4, 2)$, and $\mathcal{N}(0.5, 0.15)$. Thus, the main mass of the historical values, as well as the predictions, are concentrated around the mean values of the distributions. Lastly, when tested on the fourth behavioral class $I-G$, all methods apart from the HF achieve very high precision. Recall that the HF does not take into account $\hat{\sigma}_i$, hence, since in this case $\hat{\sigma}_i$ and α_i are *not* independent, valuable information is not incorporated, and the predictions are off the actual occurrences. Regarding the rest of the methods, they all take as input agent statements $\hat{\sigma}_i$, and are able to “fit” the underlying model quite well.

We now proceed to examine how these prediction techniques can be used to enhance the performance of an agent cooperative that offers electricity consumption shifting services. The shifting of consumption is such, that can make the aggregate demand curve smoother, or alter the consumption levels during specific time intervals, so that renewable generation levels are matched more effectively.

6.3.2 Experiments with agents of the same behavior

In this part, we run one 100-days simulation per agent behavioral class, and report the performance of the various methods when the individual predictions are aggregated in the cooperative electricity consumption shifting scenario. For training, we used 100 $(\hat{\sigma}_i, \alpha_i)$ couples generated from the corresponding distributions; these were considered to be the historical values. We compare the results to those accrued when using

“conservative” estimates for performance prediction. The *Conservative*, *HF*, *GP*, *k-NN*, and *KR* estimates are calculated as described in Section 6.2.

The average final gain of the cooperative from the shifting operations, as well as the accuracy of each of the methods between the expected and actually shifted load, are shown in Table 6.2. As we can see, the performance of *GP*, *k-NN*, and *KR*

Table 6.2: Average daily gain (€) and accuracy in aggregate amount of *kWh* shifted (%), from 100 days simulations, one for each agent behavior class.

	<i>Conserv.</i>		<i>HF</i>		<i>GP</i>		<i>k-NN</i>		<i>KR</i>	
	€	%	€	%	€	%	€	%	€	%
<i>UU</i>	37.37	80.25	37.00	79.82	57.84	97.71	57.97	97.97	58.64	98.86
<i>BB</i>	32.36	78.46	54.92	95.10	59.59	99.76	59.34	99.58	59.39	99.60
<i>GG</i>	55.63	94.34	52.57	94.15	60.37	98.91	60.78	99.34	60.81	99.41
<i>1-G</i>	47.51	89.46	53.68	94.67	58.82	99.65	58.21	99.38	59.23	99.79

methods is consistently very good, reaching accuracy higher than 97%, for all agent behavioral classes. Also, the *HF* approach also performs well—accuracy higher than 94%, except for the case of *UU* agent behaviors, where it drops significantly. Overall, all methods perform much better than the *Conservative* estimation approach, for all agent behaviors, apart from the *GG* case, where *Conservative* also achieves high precision.

The average daily gain of the cooperative is strongly related with the accuracy of the performance prediction methods, as it can be clearly seen from the numerical results; the more accurate the predictions regarding the aggregate performance, the higher the gain that the cooperative receives. Note that each simulated day generates different agent demand values and variable shifting capabilities, according to each

agent's contract type. This is why, e.g. *KR* achieves higher gain for the *GG* case than for *BB* and *1-G*, despite the fact that it achieved slightly lower precision. Also, here we validate the result from the work of [174], i.e. when aggregating the potentially not accurate individual forecasts to a single cooperative estimate, the accuracy of the aggregate prediction is very high. This happens because the differences in the individual predictions and the actual realized actions are canceled out among each other; while some agents might over-perform *wrt.* their initial statements, it is highly probable that some others will under-perform, helping reduce the error gap induced by the pessimistic participants.

6.3.3 Trial with mixed agent behavior classes

In the last experimental setting, we assume all aforementioned agent behavioral classes to be present in the same 100-days scenario, which is simulated 10 times. Specifically, agents are assigned one of the four behavioral classes with equal probability at the beginning of each of the ten trials of 100 days, and are governed by this particular behavior throughout each of the simulations.⁵ The numerical results are presented in Table 6.3, where we present the means and standard deviations of the daily final cooperative gains that are achieved, the expected and actual reduction during the peak intervals, the accuracy of each method and the percentage of the trimmed peak demand, as well as the daily average reducing coalition sizes, and the elapsed time for the generation of the predictions.

All four methods generate actual cooperative gain values that are quite higher

⁵Note that although the distributions that describe the actions of an individual remain the same for every day of each simulation, the actual values change, since every day they are resampled.

than the value achieved when using the *Conservative* approach. Regarding the daily average *expected* size of the reduced load, we observe that all evaluated methods provide almost identical results. However, this is not the case when we examine the *final* size of the load actually reduced by the cooperative. Here, it is obvious that *GP*, *k-NN*, and *KR* achieve final reduction values that are very close to the expected levels, reaching accuracy of 99% and above. The same hold for the percentage of the peak demand trimmed, i.e. the quantity of load that is above the threshold τ . We can see that the incorporation of the regression methods for the prediction of agent actions helps trim almost the whole amount of peak demand, i.e. more than 97%, in contrast to the *Conservative* and *HF* approach, which trim 84% and 91% respectively.

Another important performance indicator is the required time in seconds needed for the predictions to be generated. The numerical results indicate that the fastest method, with the significant difference of the order of hundreds of times faster than *GP*, is the *k-NN* approach. The *HF* and *KR* are also quite fast, especially when compared to the *GP* approach. This happens, as discussed in Section 6.2, because *GP* requires the inversion of a large matrix, which is a quite complex and time consuming operation. However, since the scheme that we adopt is based on day-ahead planning and is not real-time, there is enough time to execute *GP* to generate models for all of the 7954 agents that participate in the setting.

Finally, we can observe that *k-NN* employs the largest number of participating agents on average, on a daily basis. In the cooperative electricity consumption shifting setting this is quite important, as we need participants to come in large numbers, since this has positive impacts in the stability of the cooperative's performance. Also,

Table 6.3: Average results from an 100 days simulation scenario including all agent accuracy classes.

		<i>Conserv.</i>	<i>HF</i>	<i>GP</i>	<i>k-NN</i>	<i>KR</i>
Actual Coop. Gain (€/day)	μ	39.77	49.11	57.18	57.41	57.61
	σ	12.95	14.99	15.73	15.88	15.78
Expected Reduction (<i>kWh</i> /day)	μ	1332.291	1333.079	1334.746	1334.605	1334.708
	σ	373.471	373.321	372.935	373.212	373.127
Final Reduction (<i>kWh</i> /day)	μ	1113.266	1219.697	1322.723	1327.406	1331.412
	σ	326.037	354.558	375.141	378.918	376.991
Accuracy (%)	μ	83.56	91.49	99.09	99.46	99.75
Peak Demand Trimmed (%)	μ	82.54	90.28	97.34	97.41	97.76
	σ	4.82	4.62	2.66	2.77	2.53
Avg. Reducing Coalition Size	μ	296.05	357.38	382.19	389.09	385.79
	σ	90.23	111.28	114.45	116.42	115.15
Estimation Time (sec)	μ	-	1.56	444.03	0.006	0.44

the aggregation of a large number of participants leads to better accuracy and more effective demand shifting operations.

Summarizing, the evaluation demonstrates that *KR* is potentially the most effective of the prediction techniques examined by all aspects, with both *GP*, and *k-NN* producing cooperative results that are very close to each other, regardless of the behavior model by which each individual agent is described. This is best illustrated by the better performance with respect to prediction accuracy, and the nearly perfect effectiveness in terms of peak load trimmed. Recall that for the *HF*, agent forecasts are not taken into account, so potentially important information is ignored. Intuitively, *GP*, *k-NN*, and *KR* can effectively learn and adapt to the underlying model that relates agent forecasts and final actions, thus enabling the cooperative to choose reducing coalitions that often deliver what they promised. *HF*, however, exhibits a strong performance also. Thus, in a nutshell, our results indicate that

all techniques examined are strong candidates for predicting the accuracy of selfish agent statements in Smart Grid consumer cooperatives.

6.4 Conclusions and Future Work

In this chapter we presented four different methods for monitoring and predicting agent actions in a power consumption shifting scheme, a Histogram Filter, a Gaussian Processes regression method, a k -Nearest Neighbours, and a Kernel Regression approach, to recognize possible underlying relationships between agent forecasts and final actions in a cooperative electricity consumption shifting scheme. The methods outperform the formerly used method *wrt.* prediction accuracy—and financial gains generated. Our techniques are generic and can be integrated within different types of systems (e.g., they can be used for monitoring the accuracy of electricity production statements). Also, the computation time is reasonable even when the number of participants comes in the order of thousands. This is the first time these methods are applied in the electricity consumption shifting domain; and the potential value of this work to any real-world enterprise operating in the Smart Grid can be very high. Similar techniques can also be used to monitor the behavior of EV drivers, managing this way to predict the time intervals of EV connection/disconnection to chargers more accurately. To apply the proposed scheme in the real world, it is imperative that smart metering equipment is available for each of the participating agents, so that to be able to measure consumption and reductions with precision. Moreover, the presence of devices or applications with intuitive interfaces, is key for the effective performance monitoring at the user side.

Chapter 7

Supporting RES Cooperatives and Aiding the End Users

In the previous chapters, we described the game theoretic and machine learning tools behind DSM mechanisms that can be used to engage end-users to participate in coordinated electricity consumption shifting actions, in order to better utilize renewable energy sources, and effectively balance demand and supply. In the current chapter, we take all the aforementioned tools for granted, and present a system architecture for DSM services based on the Internet of Things (IoT) paradigm—a concept that is expected to drive the next industrial revolution in the following years. Additionally, we design a *serious game* [152] and implement two different versions of Graphical User Interfaces (GUI), which are employed to test the effectiveness of DSM incentive types and data visualization techniques on human users. In the proposed serious game, players aim to maximize their economic and social revenues by accepting or not DSM actions that their home energy management agent suggests,

after accounting for the pricing mechanism and the user preferences. Finally, we devise a methodology for statistically analyzing the effectiveness of DSM schemes and energy efficiency (EE) interventions in general. Since interconnected devices and smart meters produce large quantities of data, it is desirable to propose ways for effectively manipulating them. Also, statistical analysis can be used to evaluate implemented DSM and EE interventions. The work presented in this chapter, along with the methods of the previous chapters, integrates our approach to offer readily applicable and complete large-scale DSM solutions, whose incorporation in modern electricity grids can benefit multiple stakeholders.

In most cases, IoT approaches follow a service-based architecture [15]. The provided services can be categorized into three distinctive yet interdependent types, which all can be combined to deliver other, more complex applications:

- i services that capture properties of the physical world and provide raw or slightly processed measurements (sensing services);
- ii services that process the acquired measurements and provide the inferred results (processing services);
- iii services that enable certain actions, based on these results (actuating services).

Thus, IoT applications involve distributed sensor networks at various scales, e.g. over the human body, small indoor areas (such as a house), or even world-scale outdoor areas. In any case, they are interconnected with other distributed services, possibly located on a cloud infrastructure. Their purpose is to process the measurements gathered and autonomously extract results, i.e. without the need of human

intervention. Upon processing, some actuation elements may be triggered as well. Such services are offered to a variety of users and can be combined to provide other, more complex, services and applications, via IoT services marketplaces. To achieve this, they must be easily discoverable and interconnectable by developers using the IoT ecosystem, and, more importantly, they should be able to be deployed and executed on a cloud infrastructure [98].

In the energy sector in particular, during the recent years, utility companies and DSOs around the world, and specifically in most European and North American countries, have started replacing old consumption meters, with new, smarter ones. These smart meters allow the real-time monitoring of energy consumption and provide the opportunity to use energy more efficiently and when it is available from RES. Furthermore, there is a huge proliferation of mobile devices, which are connected to the internet – moving towards the Internet of Things (IoT). Even household devices will be connected in the next few years, e.g. the fridge will report online and in real time its status and its energy consumption [164]. Thus, Internet of Things involves technologies that will enable the development of large-scale and distributed complex services, such as DSM mechanisms that is our main concern in this thesis. These services can deliver any game-theoretic or decision theoretic approach for managing energy consumption.

In the scope of a past research project,¹ we designed a specialized Graphical User Interface (GUI), that complements the work of this thesis, i.e. explicitly provides the means for the DSM mechanism to interact with the end-user that participates in DSM. The GUI was presented to real users in a form of a serious game, where

¹<http://iot.synaisthisi.iit.demokritos.gr>

groups of energy consumers were asked to plan (a day-ahead) the rescheduling of their consumption, and receive rewards, either social, or economic/monetary. The goal of this set of trials was to test whether specific incentive types and data visualization techniques are more suitable to convince the end-user to participate into DSM.

Additionally, in this chapter, we describe a data analysis procedure to evaluate the actual performance of such services, which can be used by real world energy cooperatives. The results from the trials and the analysis show that indeed some approaches are more effective than others in engaging the end-user and in achieving demand reduction. The work described in this chapter completes the multiagent DSM approach that is presented in this thesis, and can be used as a guideline for the development and application of large scale, real-world DSM solutions. The contributions of this chapter and the combined scientific fields are also shown in Figure 7-1.

This chapter is structured as follows: In Section 7.1 we review past work regarding decision theoretic optimization approaches for electricity demand rescheduling. Being modular and IoT-based, our architecture approach can incorporate any method as the optimization module. Section 7.2 presents the system architecture for delivering large-scale DSM services. Section 7.3 thoroughly describes the proposed GUI and some results from the real-world trials. Section 7.4 describes the methodology for analysis of EE measures effectiveness. Finally, Section 7.5 concludes this chapter.

Parts of the research described in this chapter appeared originally in [12], in [50], and in [14].

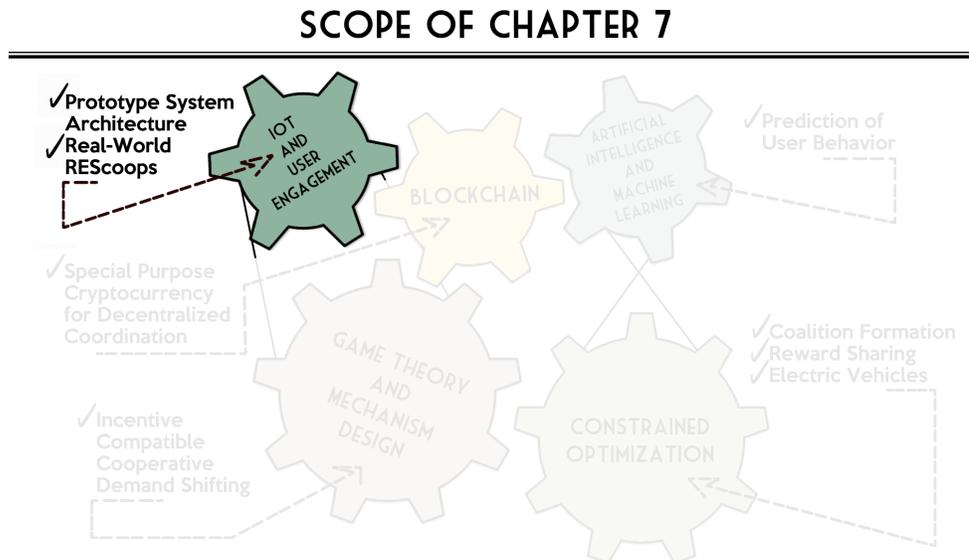


Figure 7-1: Overview of the scientific fields, and our contributions in Chapter 7.

7.1 Related Work

We now provide a brief review of a few methods that rely on decision theoretic optimization for managing the energy consumption in residences and other energy consuming buildings. The high-level architecture we describe in the following section is such, that allows the incorporation of any energy consumption management method to serve as the agent optimizer. To begin, Logenthiran *et al.* [108] present a heuristic evolutionary algorithm to implement a generalized, day-ahead DSM strategy that optimizes the consumption schedule of shiftable devices according to specific objectives (economic, environmental, etc.). Then, when a device is enabled, it either turns on if this occurs within its operation timeframe dictated by the optimized schedule, or simply does not turn on and notifies the user about the time of its future

availability.

Then, [107] formulates an appliances scheduling problem that accounts for energy costs and consumer preferences in a residential smart grid, and solves it using a distributed subgradient method. The objective is to reduce the demand's *Peak-to-Average (PAR)* ratio. Intuitively, a reduced PAR corresponds to a reduction of peak demand, contributing to system efficiency and reliability. In [91] a generic DSM model is introduced, which also aims to minimize PAR, costs, and the *waiting time* of electrical appliances from households, industries, and commercial buildings. Here, a multi-objective fitness function is optimized by a *genetic algorithm*. To optimize for multiple users, however, every single user must share electricity consumption-related information with all others. This raises privacy concerns, and induces security risks for the buildings' residents.

In the work of [56], a distributed multi-agent optimization is proposed, for the automatic rescheduling of domestic loads, using ant colony system optimization. Although metaheuristics are appropriate for solving the highly complex problem of domestic load rescheduling, this approach does not deal with the problem of reward sharing of the cooperative profits, nor for the end-user related flexibility (shifting) costs, both big issues that have impact on user engagement, which is the key for effective DSM schemes. Also, very important Grid constraints are not taken into account and their consideration is left as future work. The work of [176] explicitly takes into account user preferences regarding the time intervals of shifting, and proposes a genetic algorithm to minimize domestic electricity bills. User preferences are modelled as soft and hard constraints, however, here too, few are mentioned regarding Grid's incentives and economics. More importantly, this approach does not

promote cooperative action, which, as we have stressed in the previous chapters of this thesis, is crucial for DSM schemes to be effective and to have positive impacts.

In general, automatized appliance operation schedules prescribed by optimization methods like the ones above, restrains the consumer's ability to directly control consumption tasks. By contrast, MD approaches such as ours aim to *guide* consumers towards efficient consumption schedules—while allowing them to choose themselves whether they will contribute in DSM or not, explicitly presenting them with participation-associated gains and losses.

Another work, that of [112] studies ways of providing the incentives for reliable demand response. Participants choose to participate in a two-period mechanism, by initially expressing their intentions—that is, costs and delivery probability values, and then are rewarded, or penalized, according to their final delivery. In [17], matters regarding autonomous agents and their interactions with humans are studied. More specifically, authors design and test an agent for electricity tariff switching, with the ultimate goal to choose the best energy tariff, taking into account user and electricity production uncertainties, as well as user comfortability and incentives regarding the interaction with the agent. Authors consider two tariffs, one for standard power generation which is constant, and a second for wind generation that varies with time. In order to choose a particular tariff, the agent must predict the user consumption and the tariff from wind power, which are both subject to uncertainties. Thus, the “Tariffagent” monitors and predicts electricity usage, predicts the energy output of wind powered plants, and provides the necessary interfaces in order to interact with the user in a comfortable, trustful, and non-intrusive way.

In [78] a study on the acceptance of monitoring technologies in contemporary

urban places is presented. The questionnaire that was answered by 127 humans of a wide age range, had a medical monitoring context. In short, participants were negative against surveillance technologies (microphone, camera, position), with the least worse being the positioning system, and age is not related to acceptance, neither is gender. On the contrary, health is related to acceptance, with the *ill people being more negative* to monitoring technologies. Correlating with the place of surveillance, participants were more positive of being monitored in public places, rather than in private.

In our work, we monitor participants power consumption shifting actions. Since the residential sector is expected to have large impact on demand shifting schemes, it is crucial to study the acceptance of monitoring techniques on such context. Indeed, there is a need for such schemes to not be intrusive on privacy issues and thus have a deterrent effect, as we must maintain participation at high levels in order for the scheme to be effective.

The works of [59, 116, 101] study ways of providing incentives to users through gamification techniques and reward mechanisms. Rodden *et. al* [157] study consumer responses when faced against future energy infrastructures. Making one step further, we build a smart EMS paradigm that is ready to be deployed in real residences, and present it to human subjects. In our work, we provide different incentives for achieving the same goal —to obtain a “greener” consumption profile. We test these two games with real human users and compare the incentive categories with respect to their effectiveness, i.e. making users willing to participate. In [68] authors address the reaction of smart houses when introduced to time-based electricity tariffs. They develop and validate a model for smart households that are billed according to time-

based prices. The simulation model consist of a generation component and a scheduling one. Regarding the smart scheduling part, first authors identify and classify the appliances that can be shifted. More specifically, washing machines, dishwashers and dryers are considered of semi-automatic control —as consumer interaction is needed, e.g. the dishwasher must be loaded. Another class is fully-automatically controlled loads, such as storage water and space heating, freezers, refrigerators, etc. Finally, non-shiftable loads include stove, lighting, electronics, among others, and their use requires energy consumption at the very time it is requested. The best time for starting an appliance is determined by the solution of an optimization problem. However, as authors themselves state, this scheme is vulnerable to avalanche/herding effects. This particular household modelling is also used in the PowerTAC competition.

Now, the most appropriate architecture type for our proposed system to follow, is an IoT-based architecture [132]. According to this approach, each simpler module of a complex applications is virtualized as a cloud service, and specific protocols that allow the dynamic interconnection of such services are available [92, 72]. Examples include publish/subscribe protocols [28], ubiquitous data access methods [124], service composition tools [180], etc. IoT-based architectures have already been proposed in the literature for industrial environments [191], smart cities [64], healthcare [69], and safety and security settings [86], however, IoT-oriented DSM approaches is still an open field for research [61, 167]. Regarding the energy management domain, the work of [166] proposes an HVAC² optimization approach that is based on IoT for services interconnection, and considers both variable energy prices and user preferences and comfort. Another work, that of [201], presents guidelines for designing

²Heating-Ventilation and Air Condition

energy monitoring and consumption management systems for intelligent buildings incorporating IoT technologies. Our approach, apart from the individual building energy management and consumption optimization modules, we also assume an external DSM mechanism module that is used to drive collective consumption shifting and aggregate the actions of agent cooperatives.

7.2 System Architecture for DSM services in an IoT Ecosystem

The proposed system architecture is shown in Figure 7-2. We assume that there is a DSM mechanism service implemented on an Internet-of-Things platform (e.g., the ones proposed in the previous chapters of this dissertation) that can take into account forecasts of electricity production and consumption levels a day-ahead. Along with the predictions, the mechanism module also collects data regarding the consumption profiles of the individual consumers, and then sends feedback and rewards to the agents in the house that represent them. Typical IoT platforms provide effective ways to manage the information flow between multiple services, regardless of their physical position and communication protocol. Executed on a machine inside the house, a device controller also exists that is able to turn on or off electrical appliances. The controller sends data about the consumption of these appliances to a monitoring module, that in turn stores the data into a storage infrastructure. Additionally, a GUI module exists that runs in parallel with the agent (indicated as demand scheduler/optimizer) and is used to communicate user preferences and

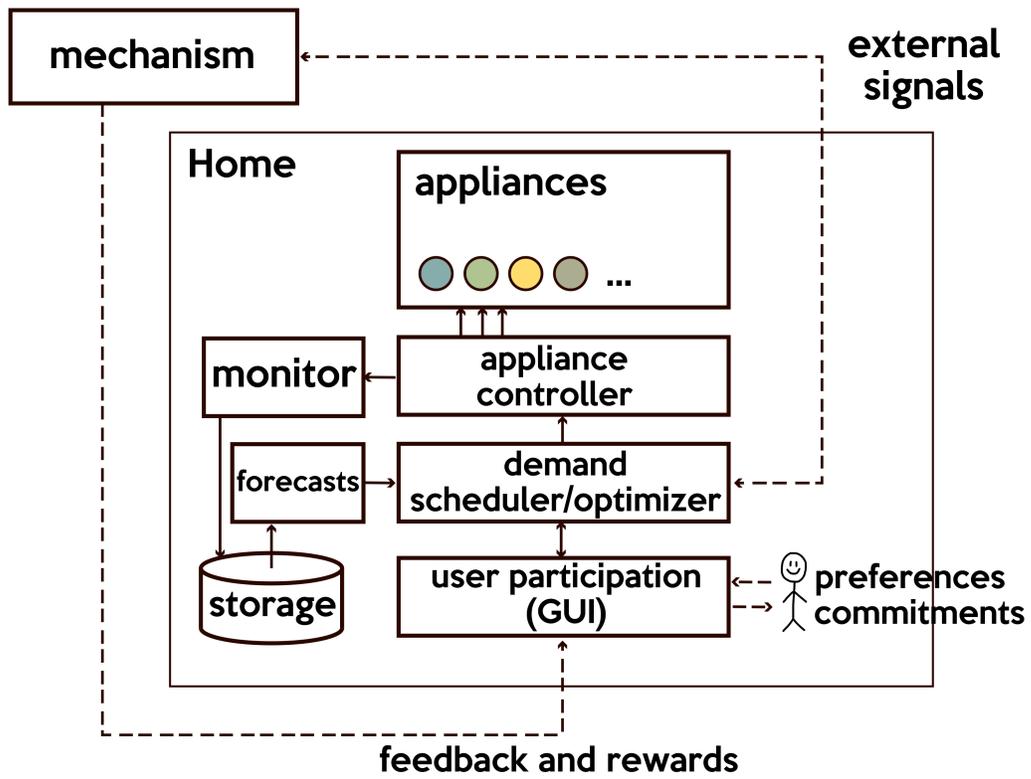


Figure 7-2: IoT based architecture for large-scale DSM mechanisms.

commitments. The optimizer is responsible for handling the external signals, user preferences, and forecasts regarding future home consumption. We assume, as it is common in IoT approaches [77], that the different services (modules) communicate with publish/subscribe protocols,³ allowing this way their dynamic interconnection. By incorporating this kind of modular IoT architecture we can easily replace each component with another that has similar functionality though implemented differently, giving the opportunity to the solution provider to choose appropriate procedures and methods, according to each use-case. This fact overcomes the limitations found on other, application-specific approaches.

7.3 A Serious Game for Demand Shifting

Using the proposed architecture, we developed a serious gaming approach that was used to test whether energy monitoring GUIs can effectively incentivize users to become more aware energy consumers, and to encourage them to use green energy. Specifically, we developed two versions of an Energy Saving Game, that is used to present to the player the same GUI used by the autonomous agent managing the consumption in the residences that would participate in DSM. The data used during these games and the offered interaction are the same: the difference is in the visualization and incentives. One version, termed as the *economic game*, is based on economic information and links the user's consumption data with pricing schemes which offer lower prices at times when abundant "green" energy is available, visualizing how the proposed suggestions affect the energy bill and highlighting the

³See, e.g., the MQTT protocol (<http://mqtt.org>).

consequent savings. The other version, termed as the *social game*, relies on social features and emphasizes the competition with other users, rewarding the player with respective amounts of virtual, “green” coins for each DSM suggestion realized.

The general interface used for both games is composed of several individual linked visualizations designed to inform, incentivize, and facilitate communication between the user and the agent module. Within each GUI version we include a number of coordinated views along with aspects of both pragmatic and aesthetic information visualization to ensure that:

- (a) we are informing users
- (b) we persuade them to follow incentivized actions
- (c) we allow them to effectively and intuitively make selections and provide feedback to the system

Though some visualizations in the social and economic interfaces are the same, there are small context-driven differences within each module, which highlight the most relevant information to each game, as we detail later on.

For participating in the contest, first, human users assumed that they were located inside their house, and that at some time of the day (a configuration that can be personalized for each case), the agent consulted them for the residence’s future (e.g. a day-ahead) consumption rescheduling possibilities. Participants should consider their own preferences and decide whether they would shift specific tasks at different hours of the day, where renewable production is higher, i.e. accept or reject the agent’s proposals for a period of 10 days for each GUI incentive type. Finally, they



(a) Social interface.

(b) Economic interface.

Figure 7-3: The two different interfaces.

were also asked to fill in a questionnaire regarding their experience with the energy saving GUI. The two GUI versions of the prototype design are shown in Figure 7-3.

7.3.1 Game controls

As every video game, ours has specific game controls i.e., methods to interact with the player. The interactive components are the “All”, “Agents temperature preference”, and “Suggestions” boxes in the top left corner. These serve as the primary mode for user input and contain both aesthetic and pragmatic features. The first, the suggestions box in the upper left corner, is the most detailed of the three components and presents the proposed shifted tasks, the original time the user has been predicted to complete the tasks, the suggested shifted times to complete the tasks, the duration of the tasks, the watts used by the tasks, and, in the case of the economic game, the old and new prices resulting from accepted shifts. In addition, for ease of use, we give the option to always accept particular tasks. The “Always Accept” box minimizes user interaction and makes the process less invasive while the pictographs allow the

user to more quickly process information. In addition, the visual coins serve as a motivating feature used to sway the user toward a “green” scheduling decision. The second, the “Agent’s temperature preference” box, contains a slider allowing the user to switch the heating (in the form of an average target temperature) between the agent proposed temperature (marked with a red arrow) and the user’s preferred temperature (marked with a black stickman). The third of these, the “All button” in the very left corner, gives the user the option to quickly select all suggestions if he or she desires, and the “I’m done” button ends the input session and delivers the final results to the multi-agent module. Combined, these three visual components provide information in an aesthetic and persuasive manner.

7.3.2 Game feedback

The next important visual, the bar chart/line graph in the lower left corner, contains both pragmatic and aesthetic properties designed to inform the user while also persuading him to make “green” or “economic” decisions. In addition, it serves as a complementary visualization to the first three input mechanisms where the user can better understand and explore the consequences of choosing a particular schedule. More specifically, the bar graphs show the average (the aggregated consumption divided by the number of households) of the predicted consumption for the next 24 hours, the green bars correspond to the average of consumed green energy and the gray bars correspond to the conventionally produced, non-renewable energy. The yellow line projected over these bars shows the predicted user’s consumption for the next 24 hours, before taking into account any shifting or heating reduction, and the purple line shows the predicted consumption after the changes caused by accepted

suggestions. Furthermore, each point is marked with a different color to indicate the quality of the user's consumption at each hour. The *green* color indicates a good consumption pattern, *black* indicates an average consumption pattern and *red* indicates a poor consumption pattern. This visualization is pragmatic in the sense that it clearly communicates important technical information and aesthetic in the sense that the colors and design of the layout are chosen specifically to encourage the user to follow a green rather than dirty scheme. Also, this visualization supports the user in his decision making process, as the purple line showing predicted consumption is dynamically updated every time the user selects different combinations of accepted suggestions. Thus, the user can test a variety of combinations before making his final selection.

The remaining components on the right provide summaries that give the user a better idea of where he stands within both the current round and entire game. In the upper-middle part of the window a stylized man carrying a money bag provides information about the effectiveness of accepted suggestions in the form of the number of earned green coins in the social game and money saved for the next day in the economic game. The box under the man with the coin describes the differences before and after shifts as well as the number of green coins or money he has saved. The "Demand Scheduler" box on the right provides a friendly greeting and a record of the day of the year, and the "Leader Board" in the social game and "Economic Board" in the economic game each give the user a better idea of how she compares to her "competitors". Lastly, the red timer at the top provides an indication of the amount of time that remains.

Summarizing, each of the visualizations work together to inform users, persuade

them to follow incentivized actions, and allow them to effectively make selections and understand their actions. The games were implemented as web applications, based on Java Servlets and using JavaScript, JQuery and Jqplot to obtain the necessary dynamism during the interaction—such as the changes in the consumption levels by accepting suggestions.

7.3.3 The game with economic incentives

Figure 7-3b shows the interface for the economic version of the game. As mentioned before, the economic game focuses on the amount of money that can be saved by following the proposed suggestions, rather than focusing on the amount of green energy used. Consequently, there are several small differences in the sub-visualizations between the two though. For example, the suggestion table provides the old price and the new price, i.e. the amount of money corresponding to the job in the original time and the amount of money due in the case of shifting, in addition to the other categories shown in the social interface. Furthermore, the bar of coins is proportional to the difference between the old price and new price rather than the amount of “green” energy used and the money bag shows the total amount of money saved by accepting suggestions related to the heating and shiftable appliances. Next to the graph, the user can visualize the before-shifting total price, the after-shifting total price and a projection of savings for the year that would come up after always accepting this suggestion. In the right column of the page, under the demand scheduler information, projections for the year are shown, including the cost of non-shiftable appliances, the total projected cost for the year, the total projected cost after shifting (including heat shifting) and the total projected savings.

7.3.4 The game with social incentives

In the social game, the emphasis is placed on the comparison and the “green” competition among friends rather than monetary savings. A gain in terms of green coins is associated with each of the suggestions by the mechanism. Thus, in this interface type there is no information about the pricing and the money icons are substituted by green coin icons. To engage the user and motivate him or her to be “greener”, a *leader board* is shown in the right column of the interface. This way, the user is able to compare her consumption with that of her friends’, using the earned green coins as the means for this comparison. Furthermore, the money bag contains the total coins that can be earned with accepted suggestions, including both heating and shiftable appliances, whilst the box beside the graph gives this information only for the shiftable ones.

7.3.5 Results from trials with real users

In this set of experiments we offered the opportunity to 50 actual humans to play the energy saving game. The population included ages from 25 to 55, and participants education levels ranged, from undergraduate students to post doctoral researchers. Divided into ten groups of five persons, participants were asked to play each one of the two games, one after the other, for a simulation time of 15 days⁴. Each participant was assigned to a specific simulated house, and was asked to act as if the agent was controlling and suggesting the schedules of his own electrical appliances. After the end of the two consecutive games that would allow us to assess the effectiveness of

⁴Note that each consecutive group played the two games in different order.

the two different incentive types, they were also asked to fill in a questionnaire, in order to share with us their experience and provide some valuable feedback.

The first indicator measured from these experiments is the suggestion acceptance ratio, and its value is 69.1%, indicating that the agents proposals are accepted in their majority by the subjects. More specifically, for the social game the measure was 69.8%, whereas for the economic one it was 68.4%, namely the different type of game did not affect incentives significantly. In addition, the total hourly average consumption was initially 16758.99Wh, where the total consumption after the acceptance of agents heating optimization suggestions was 15172.92Wh for the social and 15441.71Wh for the economic. The total consumption during a typical day is shown in Figure 7-4. We can see that by participating in the scheme via the GUI, aggregate consumption can be effectively reduced and rescheduled. Judging from the figure, we can observe a slight differentiation between the two incentives effectiveness, with the social performing slightly better. Also, the amount of consumption shifted or curtailed indicates that our approach follows the right path.

For every game round, and both game types, the *most accepted* agent suggestions were those about water heating and washing machines, including also clothes dryers and dishwashers. On the other hand, the *least accepted* suggestions regarded cooking appliances. In particular, Table 7.1 holds the acceptance ratios for the shifting of appliances' consumption in each of the two games.

The experiment with the real users helped answer specific questions regarding the effectiveness of the proposed system, and we now proceed to analyze the results from the questionnaire. Note that the ranking was made in the scale of 1 to 5.

Table 7.1: Acceptance ratio over shifting different appliances per type of game.

Appliance	GameType	Acceptance Ratio
Clothes Dryer	Economic	77.2%
	Social	82.1%
Cooker	Economic	58.7%
	Social	58.0%
Dish Washer	Economic	70.6%
	Social	81.0%
Oven	Economic	53.2%
	Social	53.8%
Washing Machine	Economic	79.4%
	Social	74.4%
WaterHeater	Economic	71.3%
	Social	69.4%

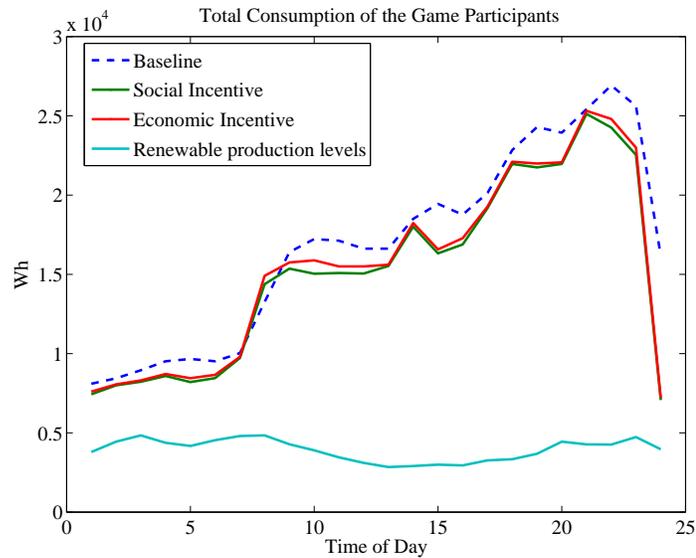


Figure 7-4: Typical day of the game rounds.

How good was the energy saving system in general? The feedback regarding the ease of use of the system indicated that in general the system was perceived as easy, with an average rating of 3.78.

Were the agents optimized suggestions of any use? The set of questions regarding the usefulness of the agents suggestions gave a mean of 3.56, indicating that the suggestions were quite useful to the participants.

Is the social aspect of the game incentivizing? The feedback of the players showed that the social aspect of the game can provide adequate incentives, with an average rating of 3.29.

Is the economic aspect of the game incentivizing? The rating regarding the economic part of the game gave an average ranking of 4.18, indicating that the economic incentives are most preferable for DSM operations.

Which of the above two gamified aspects induced more demand rescheduling actions? Judging from the participants answers, although most seem to be willing to participate in DSM and energy saving operations, there is a small preference towards economic incentives, rather social, “green” ones. Note that this slightly contradicts the observed incentive effectiveness that was shown by the game results to be better for the social type.

We can conclude that both incentive types proved to be quite effective in inducing DSM participation; green energy usability and system stability are effectively promoted. Although human feedback showed a preference towards the economic type,

the social type proved by the numerical results to be slightly more effective.

7.4 Creating Methodologies for Analyzing Data from Real-World RES Cooperatives

The design of a data analysis methodology, in order to conduct a statistical analysis of energy-related data, is a complex process that requires significant technical expertise. Such analyses could be used, for example, to get a better understanding on the broader impact of DSM schemes, but also other energy efficiency interventions, and their influence in the behavior and the energy savings of end users. Also, since available data from IoT devices and smart meters may be highly heterogeneous and large in size, a fact that makes the management and analysis harder, the definition of data collection and analysis methodologies becomes necessary. Then, based on the results of such processes, policy advises can be given; best practices for energy efficiency can be highlighted; and new business models that accommodate energy savings can be created. In the scope of an ongoing EU, H2020 project,⁵ we collaborated with real world cooperatives to test the performance of specific energy efficiency measures and techniques that were applied to actual energy consumers to reduce their demand, in both the residential and commercial sectors. For the purposes of this research project we devised a specific methodology for the gathering and the analysis of real consumption data, and for the assessment regarding the effectiveness of each energy efficiency measure.

Now, there are many adjustments required in order to reach the optimum bal-

⁵<http://rescoop.eu>

ance between reliability, accuracy and significance. In this section, we outline the methodological approach.

The proposed methodology of analysis consists of the following steps:

1. Identification of the current state of the data from the REScoops under examination.
2. Consideration of a common data format and inventory-related issues.
3. Receiving the datasets
4. Completion of the datasets missing values from additional third party sources or by using regression techniques.
5. Analyze the development of the average monthly and yearly consumption for all cooperatives and separate all possible variables that might influence the decrease in consumption (climate, change in habits, etc.).
6. Validate the decrease in average consumption levels.
7. Check whether there is a decrease in usage of conventionally produced energy (from non-renewable resources).

In order to validate the decrease in average consumption levels, we must first divide the sample population into different groups, based on whether specific energy efficiency interventions have been applied or not, and other factors of similarity (e.g. contract type, number of residents, etc.). Then, we obtain the kernel density estimates of the underlying distributions, and check if the mean and variance values

of those distributions differ significantly. To determine whether the difference is significant, the *p-value* and analysis of variance (ANOVA) methods [55] were used.

The proposed methodology was applied on data originating from 7 REScoops located in different European countries. Briefly, results suggest:

- i Specific energy efficiency interventions can be effective in driving reductions in end-users energy consumption.
- ii Joining a REScoop leads to more than 20% reductions in energy demand.
- iii Subscribing to consumption monitoring and savings suggestions software platforms results to approximately 35% consumption reduction.
- iv REScoop members significantly contribute to energy conservation and to the reduction of harmful gases emissions (a projection given analysis results estimates these savings as approximately 1,500 tons of CO_2 per month).

For more information regarding the REScoop Plus project, and a thorough analysis of the results, the reader can refer to [12], [50], and [142].

7.5 Conclusions

In this chapter we presented an IoT service-oriented system architecture that can be used by implementations of large-scale DSM that incorporate large numbers of end-users, which have control over their electricity consumption. We also developed and tested two GUI versions with real users, one that promoted economic incentives, and one with more social-oriented ones. Results from questionnaires suggest that

such GUIs can be effective in promoting renewable energy usage to the consumers. Finally, we devised a statistical analysis methodology to assess the effectiveness of various energy efficiency measures that are applied all around Europe by renewable energy sources cooperatives and their members. The positive environmental impact of such energy efficiency interventions is shown in many cases.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 8

Conclusions

In this dissertation we presented a multi-level approach for delivering effective large-scale cooperative Demand-Side Management (DSM) solutions in real-world Smart Grid environments. Balancing demand and supply, although imperative for the stable operation of the electricity grid, is a quite difficult task in cases where the output of renewable generators is intermittent and influenced by factors that cannot be directly controlled, for example, weather conditions that drive solar and wind generators' output levels. In such cases, DSM techniques i.e., the altering of energy consumption patterns in time, can be applied so that renewable energy is efficiently utilized.

The issue with DSM is that, frequently, along with the electrical appliances demand profile, changes are required in the daily routine of human end users as well. For such requests to be accepted though, additional incentives must be provided to the end users, and sophisticated mechanisms must be designed, so that user engagement and energy efficiency is promoted. Our approach advances the state-of-the-art in many ways. The results demonstrate that it can be readily applied in large, real-

world environments; deliver monetary gains to participants; lead to increased Grid stability; and help achieve better utilization of renewable energy.

This final thesis chapter is organized as follows. Section 8.1 reviews each previous chapter, highlighting their most important points. Then, in Section 8.2 we provide possible extensions of our research, and state and discuss open questions related to this thesis.

8.1 Summary

Throughout this dissertation, we introduced the reader to the DSM problem, set the theoretical background and motivation behind our work, and proposed detailed solutions for each problem part. In the beginning, we addressed cooperatives of consumers for large-scale and effective demand shifting. We presented a mechanism for the coordination of profitable shifting actions to promote participation for individually rational agents. We stressed the fact that formulating the problem as a multi-agent system is more natural, since we try to engage a large number of individuals, rather than a centralized and inflexible system of homogenized population, as a large number of past approaches for DSM do. Additionally, our work promotes massive participation from the end-user side, and applies personalized incentives, according to the behavior, capabilities, and performance of each agent. The proposed mechanism allows even the participation of agents, which have increased related shifting costs that would normally forbid them from participating in DSM. Then, we combined a cryptocurrency protocol with coalition formation and reward sharing methods to also promote prosumer cooperatives in DSM. Participation in our mecha-

nism ensures increased prosumer profits from the selling of energy, and reduced costs for consumption. Specific participant selection and reward sharing methods promote engagement and the effective use of renewable energy sources. Additionally, to incorporate electric vehicles in DSM, we proposed a new V2G/G2V algorithm. Electric vehicles are expected to invade the automobile market in the near future, a fact that will most probably disrupt electricity grids. Our EVs-incorporating DSM algorithm can be employed in a small or a larger scale, and, apart from shifting the charging of connected EVs, it can also make use of existing (stored) energy in the vehicles batteries in order to reduce electricity imports from external markets, and effectively manage energy produced by (local) renewable generators. Furthermore, we proposed the application of machine learning methods to monitor and predict DSM participation behavior, delivering this way even more effective DSM mechanisms. Finally, we researched implementation architectures for DSM services, and statistical analysis methodologies for validating the changes (decrease) in consumption from non-renewable, legacy generation systems as a result of energy efficiency interventions.

For each of the proposed approaches, we conducted extensive simulations on datasets originating from the real-world. Numerical results indicate that the incorporation of our methods can lead to significant gains for DSM participants, promote truthfulness and engagement, increase the Grid stability between supply and demand, and make better use of renewable energy generators. What is left, is to deploy such approaches in large-scale real-world settings, and also test the influence of specific factors related to real-world power systems, such as reactive power, losses induced by physical distance, etc.

In more detail, in Chapter 3, we presented a readily applicable framework for large-scale cooperative electricity consumption shifting, to promote the proactive balancing of the demand curve. Our proposed shifting scheme promotes agent efficiency in the face of uncertainty, achieved via the use of effective coalition formation methods we developed. The employment of the *CRPS* strictly proper scoring rule allows the calculation of personalized prices according to the participation efficiency the cooperative and each agent. Furthermore, our mechanism is equipped with internal pricing schemes that employ gain transfers within a cooperative, to make it worthwhile for individuals to participate in shifting operations and thus guarantee the scheme's effectiveness and profitability. The proposed mechanism comes with desirable theoretical properties, specifically individual rationality, truthfulness, and (weak) budget balancedness. Extensive simulations based on real consumption data, and demonstrated experimentally the effectiveness of our methods. The results of our experiments confirm that our methods could bring tangible benefits to energy cooperatives and other Smart Grid business entities. Our model can easily be adapted in the future to also include more social and behavioral aspects of the users. Then, we could study the stability of the formed coalitions in this augmented setting.

Next, in Chapter 4, we presented a cooperative consumption shifting scheme that employs a distributed cryptocurrency mechanism for the decentralized coordination of prosumer shifting actions. This special purpose cryptocurrency is combined with *CRPS* to incentivize electricity prosumers to offer decentralized and coordinated DSM services, and also to allow personalized rewarding. We formulated the problem of coordinated cooperative prosumer consumption shifting, and introduced theorems that guarantee accuracy and effectiveness in the shifting procedure. To validate

reductions as a result of mechanism participation in real-world settings, the measurements from smart meters must be profiled. Also, we proposed and analyzed contributor selection and reward sharing methods which incentivize truthfulness, guarantee increased profits from the trading of electricity, and help flatten electricity demand curve, as well as the dynamic energy prices. Results from simulation experiments show that the application of our method leads to more balanced local demand and renewable supply, as well as increased gains for the prosumers that participate in the DSM scheme.

Continuing, in Chapter 5, we proposed an algorithm for the incorporation of a V2G / G2V mechanism for EVs in DSM. The algorithm requires the drivers to state the vehicle's connection / disconnection schedule, as well as a personal energy reservation price for discharging batteries and providing power to the grid. Then, it checks all relevant limitations, such as the duration of each vehicle's electrical connection, the cost of discharging, the energy costs from external imports, and the next day's demand and production forecasts. As a result, vehicle batteries either delay their charging, or offer part of their stored energy to the grid. This results to increased use of the local renewable resources and to reduction in imports from external energy sources. Results from an experimental simulation process based on real-world datasets, showed the benefits of using V2G/G2V mechanisms with respect to both monetary and environmental concerns. Further extensions could include mathematical programming formulations that can effectively balance both the costs of the DSO manager, and that of the EV drivers as well.

In Chapter 6, we provided an enhancement of the proposed DSM schemes, which incorporates Machine Learning methods to monitor and predict user behavior re-

garding DSM participation. The proposed methods include a Histogram Filter, a Gaussian Processes regression method, a k -Nearest Neighbours, and a Kernel Regression approach, which are used to recognize possible underlying relationships between agent forecasts and final actions in a cooperative electricity consumption shifting scheme. The techniques greatly improve the performance of cooperative shifting actions, with some achieving more than 95% accuracy between the promised and the actual load shifted in total by the participants. Also, the computation time was reasonable, even when the number of participants comes in the order of thousands. To apply the proposed scheme in the real world, it is imperative that smart metering equipment is available for each of the participating agents, so that to be able to measure consumption and reductions with precision. Also, the evolution of the participants behavior after long periods of contributions could also be studied. To this goal, real-world behavioral data should be collected and used, a task that, unfortunately, is quite difficult due to privacy regulations and user unwillingness. The techniques could be employed for predicting the intervals of EV connection and disconnection to chargers as well. However, such predictions can be easily be performed by the drivers according to their daily routines or professional schedules, making the application of such prediction techniques obsolete.

Finally, in Chapter 7 of this dissertation, we discussed additional notions that make our work complete, such as the proposal of a service-oriented architecture for the implementation of large-scale DSM that can incorporate large numbers of end-users, intuitive graphical user interfaces that promote specific types of incentives, and a statistical analysis methodology to validate changes in consumption behavior of real-world REScoop members. We also note that, the software suite that has already

been developed to be used in our experiments, can be offered as a commercial solution to various types of companies for the analysis of energy markets, the identification and the planning of profitable DSM opportunities, consulting services, etc. Our theoretical model is such, that it can be incorporated to describe most electricity markets, that include different types of actors and electricity loads.

8.2 Future Work and Research Directions

The mechanisms investigated in this thesis were shown to be effective through simulations. Regardless, it is imperative to actually deploy them in the real-world. This could be done in collaboration with utility companies or energy cooperatives that are currently operating worldwide in various locations in the globe. Due to the threatened disruption of the electricity grid due to increased RES integration, Grid regulators and experts are proposing new business models and alternative pricing schemes that will deliver the so-called 3Ds; i.e., *decentralization*, *de-carbonization*, and *digitization* [175]. The DSM schemes proposed in Chapters 3 and 4 of this dissertation make effective use of ICT (Information and Communication Technology) tools, promote *decentralization*; the use of renewable energy sources (*de-carbonization*); and, most importantly, offer personalized incentives to drive active participation on the end-user side. Although many approaches have been proposed to deliver energy usage optimization, only a few will finally be accepted for real-world implementations; apart from the economic aspects, adopted in the analysis of Chapters 3 and 4, the social aspects must be also taken into account. This fact generates new research questions e.g. *how can ICT tools be used to increase users' social revenues from DSM*

and *how DSM actions can influence users reputation in the community*. One first step towards this direction was given by the serious game and in the interactive GUI that we designed and tested in Chapter 7.3 (promoting *digitization*). However, this domain though should be researched in more depth. For example, our work could be extended in the directions given by [16, 18, 169], for applications in the real-world that could also incorporate appropriate AI techniques to make user experience easier and more convenient.

Social and economic incentives are also reinforced by open business models [45]. Cooperatives like the ones promoted in this dissertation have the ability to rapidly expand since they connect members directly, and the need for advertising expenses vanishes, which is not the case for conventional businesses where such costs are quite high. For this, however, many legislation restrictions still exist, since electricity (most probably rightfully) is considered in most countries as a public good. Thus, additional regulations will be needed for the implementation of Grid-oriented open businesses.

Now, although incentive compatible mechanisms are a requirement to drive user engagement, a detail that is crucial for both centralized and decentralized approaches and should gather attention in future work, is how to drive *participation accuracy*. The availability of EV batteries and storage equipment, as well as the shifting capabilities of end users, constitute *the flexibility of demand*, which is the actual resource that must be managed and should be accurately modelled for each use case. Mostly, this can be achieved with the incorporation of *non-intrusive* performance and behavior monitoring tools. However, the extent of behavior monitoring data and their storage is a strong counterincentive for participation by residential users. Knowing

if and when each resident is present in the house, or if and when a vehicle is parked at a specific location can be manipulated in unethical ways from non-trusted and malicious third-parties. A potential answer to the privacy concerns is the aggregation of individual values to represent larger clusters of the population in the form of statistical values, making this way impossible to reconstruct the original private data from which these values came from. Also, the feature of anonymization that is introduced by the use of cryptocurrency can also help to the privacy aspects. Then, one could build behavior models using techniques as those presented in Chapter 6, and more accurately predict the resource amounts that should be managed, and produce effective suggestions and usage schedules.

An interesting extension of our work would be to study the stability of the formed coalitions, when confronted with others, of different nature, e.g. following the directions of [82] for coalitions of energy store and renewable generators. It is interesting to assess the impact on the market when such coalitions exist, e.g. the changes in dynamic prices and incentives in general. Then, multi-objective indices should be examined when forming the coalitions, extending the coalition formation methods of Chapters 3, 4, and 5, and also provide near-optimality guarantees. The schemes presented in this dissertation can also be applied in settings that include the *interaction of multiple cooperatives*. In such cases it would be interesting to investigate how contributors would select in which cooperative to participate in, and how cooperatives form their participation incentives. This would most probably highlight the existence of different agent idiosyncrasy types, e.g. ones that are more economically minded, others being more environmental sensitive, and so on. To this end, future work should devise a realistic model defining possible stochastic dependencies (or

“transitions”) among past stated forecasts of the agents based on their types, and operate upon it with Machine Learning techniques, for improved monitoring ability. In addition, methods for automatically clustering and re-clustering consumers with respect to inferred preferences should be developed.

Furthermore, it would be very interesting to investigate how incentivization schemes such as the ones in this thesis alter over time the baseline behavior of participants; This raises the question of where and when baseline and incentivized behaviors converge.

Moreover, we are in the process of designing a new architecture for a multi-agent system that aggregates demand flexibility, which includes specific MD tools to incentivize participants to be truthful. In this setting, agents that represent either small residencies or large companies, do not necessarily trust each other, and, what is more, might not even have the ability to directly interconnect and coordinate. Thus, there is a need for a trusted central manager to effectively aggregate private individual flexibility estimates and use them to offer demand flexibility services. The simulations will be based on a new dataset that originates from a real-world company.

Applications to open problems in different fields, other than Smart Grids

Our work can be also deployed in alternative application areas, where large numbers of end users create demand for limited resources. This, in a way, has been formulated by [199] as Prediction of Use (POU) games, where consumers aggregate their demand in a group buying scheme, and reduce the deviation between stated (forecasted) and actual final actions. Applications can be found in the cloud computing field, since it contains the problem of allocating limited resources to end-users that

have an increasing demand. Computational resources in a cloud infrastructure are fixed, but their utilization varies across time intervals—just as energy demand and productions imbalances do in our case—and the balancing of tasks is required for achieving energy efficient operation and increased quality of service. In this domain, however, the shifting constraints and capabilities are different than in the electricity grid domain, and potentially different incentivization schemes are required. Unlike the everyday activities of users, computational tasks could be paused or postponed with minimum impact in the service quality. What is more challenging, is the *fog computing* domain, where the physical equipment is not owned by a specific company or corporation, but provided by multiple, smaller-scale, individuals or companies. In this setting, a plethora of additional constraints must be taken into account, which do not describe mere computational capabilities of homogeneous equipment, but rather differing usage policies, prices, and service availability. Interestingly, literature from the cloud computing field has already cited our work see, e.g., [178, 207].

Another, even more promising area for extending our research is the allocation of the resources of IoT services. The vast number of services (sensors, processors, actuators) that are interconnected and available for use in many combinations and in complex applications, makes it extremely hard to manage and allocate them to requests. Moreover, since in the IoT setting such services are owned and managed by many and different parties (companies, individuals, etc.), it would be interesting to propose the formation of “IoT services cooperatives” composed of multiple (relatively small) individuals, and employ methods such the ones proposed here to make them profitable and effective. Here again, balancing between demand and supply is required, and the reward redistribution must be based on the quality of service, the

probability of failure, the duration of usage, and so on.

Furthermore, schemes such as the ones we propose in this dissertation can be readily incorporated in *crowdsourcing* and *crowdfunding* applications [162]. Here, cooperatives *offer* funds, goods, and services, instead of creating demand as in the Smart Grid application. In recent approaches crowdfunding is augmented with prediction markets [41]. This can also be supported by our schemes, and includes the problem of selecting the right participants based on their promises and commitments, as well. The incorporation of coalition formation techniques, the CRPS and other strictly proper scoring rules, and our proposed reward sharing methods, could guarantee that agents are incentivized to be accurate and effective in delivering crowdsourced tasks. Also the incorporation of cryptocurrency protocols enables the regulation-free and instant funding of proposals, and provides guarantees against fraud and malpractice.

At the same time, there also exist open philosophical challenges regarding the applications of MAS and AI methods, such as in the artificial social intelligence (ASI) and man-machine collaboration fields [171]. Particularly ASI is considered to be a revolutionary approach that is expected to advance the field of sociology [115]. Artificial Intelligence methods are employed to analyze behavior and predict user reactions in social media and marketing. When it comes to man-machine collaboration, more critical processes are offered, such as search and rescue during natural disasters [148], industrial handling operations [99], car driving [42], and many more. In such applications, human and taskforce abilities are greatly enhanced with the help of advanced computation methods and robotics. The fast formation of coalitions using approaches such as those described in Chapters 3 and 4, also enhanced

by the behavior monitoring and prediction techniques presented in Chapter 6, could be useful for effectively managing emergency situations. Also, personalized rewarding schemes such as those presented in Chapters 3, 4, and 5, could be employed to engage individuals into taking action for the common good e.g., for promoting social services, volunteering, etc. There, participants are rewarded according to how much their individual actions contribute to the common goal. Each collective action has an impact that creates a reward, which varies given the current demand and supply levels of that given service. The collective reward can then be redistributed to the actors according to their marginal contribution impact with respect to that of the others. In more complex cases, forecasts of contribution could be also incorporated, along with scoring rules to rank their accuracy.

On a final note, this dissertation has, broadly, made contribution towards the development of autonomous agents that take decisions to serve goals of humans. However, we should remember that science and technology can be used for unethical purposes as well, with one of the most dangerous possibilities being the design and deployment of autonomous weapons [179]. This purpose, among others, requires that the machine would take life-changing (probably life-taking) decisions for humans, allowing this way for potentially unexpected errors to occur, and removing the responsibility of the human decision maker. Despite the fact that it is nearly impossible to define decision models that are “universally fair”, there would no one to take the blame in cases where that would be required. Thus, it is imperative to debate advanced AI and MAS applications and solutions, targeting to build *trust* regarding technology and providers, given the substantiality of privacy concerns, and other ethical issues. Although it is not currently clear how this can be per-

formed, it could be possible that an ex-post auditing system regarding autonomous agents decisions and actualized operations is employed. The evaluation should be performed by experts, and society itself. Guidelines for developing “ethical” solutions are given in the 23 Asilomar principles by the Future of Life institute.¹ These are categorized into Research Issues, Ethics and Values, and Longer-term Issues, with the leading concepts being transparency, privacy, responsibility, and common good. Future developments and applications should definitely take these principles into consideration, in order to guarantee a better and sustainable future for human societies.

¹<https://futureoflife.org/ai-principles/>

Appendix A

Datasets Used in the Simulations

For the end users consumption dataset used in chapters 3,4, and 6, we were based on real consumption values provided by the Hellenic Public Power Company. Using this data we retrieved the $\bar{D}_{i,x}^d$ average daily consumption of each consumer i of type x , calculated over the entire 2012 year—and then estimated its $D_{i,x}^t$ consumption per each t time-interval. We note that the highest bill amount is paid by agricultural customers—whom, however, we do not include in the simulated shifting cooperative, due to their practically non-existing load-shifting ability, as we will be explaining below. The second highest amount is charged to industrial customers, and the lowest to the residential ones.

For the additional consumption dataset from industries located in India (used in Chapter 3), we were provided directly with the mean and the variances of the hourly consumption for 36 different cases. Using these values we sampled the daily hourly consumption of 4968 agents. The initial values of the 36 industrial sites were the exact same as in the experiments of [96].

A.1 Agent shifting capacities

The participants' *shifting capacities* are not provided by our dataset. Therefore, we derive the values of agent shifting capacities based on their type as follows. First, we used the publicly available *typical* ranges of *shiftable load percentages*,¹ as summarized in Table A.1. Thus, the shifting load percentage at t_h of any particular agent of a particular type, was *sampled* from the uniform distribution corresponding to the ranges typical for its type (appearing in Table A.1). Each participant's shifting capacity at t_h , was then calculated as the product of its type's corresponding percentage and its overall consumption at t_h . The shifting capacities thus estimated, are the ones reported by the agents as their $\hat{r}_i^{t_h}$ amounts.

Table A.1: Consumption contract types and their shifting capacities (as percentages of their total demand).

Type	Percentage of shiftable load
Residential	up to 80% ²
Commercial	up to 27% ³
Agricultural	0%
Municipal	0%
Public	up to 27% ⁴
Industrial	up to 90% ⁵
Public Law Entity	0%

¹Examples of shiftable loads for residential consumers include those associated with water heaters, washing machines, dishwashers, and floor heating [160].

²<http://www.eea.europa.eu/data-and-maps/figures/households-energy-consumption-by-end-uses-4>

³<http://www.eia.gov/analysis/studies/demand/miscelectric/pdf/miscelectric.pdf>

⁴http://www.enerintown.org/download.ashx?f=EnerInTown_Report.pdf

⁵The 90% shifting capacity might be optimistic for many industries; however, it is realistic for certain others (as reported, e.g., in <http://info.ornl.gov/sites/publications/files/Pub45942.pdf>). Moreover, this number corresponds to an estimate derived given observed consumption be-

Note that not all consumer types have loads available for shifting. Agricultural consumption for example, mainly consist of pumps whose consumption could be shiftable only if additional infrastructure is installed. Also, municipal consumption mainly consists of street lighting, which is also a non-shiftable load category too. Finally, public law entities have a very tiny overall contribution. Thus, we choose not to include agricultural, municipal, and public law entities in the shifting scheme. Given this, the total number of the cooperative participants is 7376; i.e., the sum of residential, commercial, industrial, and public buildings consumers.

A.2 Agent statements and final shifting actions in the setting of Chapter 3

We need a model to describe how *(i)* the agent statements on their uncertainty regarding shifting capacities at t_h occur; and *(ii)* their actual, final shifting actions occur. To this end, we define two main agent classes; the first one, denoted as *BB*, describes the realistic case where agents are mainly *confident* about their statements, and also have a high probability to deliver what they promised. In the *BB* class, the stated error standard deviation $\hat{\sigma}$ and the observed α_i accuracy factor follow two *Beta* distributions, $\mathcal{B}(1, 5)$ and $\mathcal{B}(4, 2)$ respectively, which are depicted in Fig. A-1. We choose *Beta* distributions primarily because they very good at representing and updating probabilistic beliefs regarding potential behaviours, as also manifested by their widespread use for simulating behaviours and uncertainties related to real world

haviour in the (very small) industrial dataset used in [7, 96]. We also note that in our dataset the industrial consumption is a very small percentage of the total consumption (see Table 3.2).

scenarios (see, e.g., [43, 83, 96]).

The use of the particular two aforementioned *Betas* corresponds to error statements having a low mean of $= \frac{1}{6}$ (i.e., to agents stating high certainty), whereas the final observed accuracy factor is closer to 1 (implying increased actual performance effectiveness). The second agent class is the *uncertain predictors*, *UP*, where consumers might or might not follow stated forecasts, so α_i and $\hat{\sigma}_i$ both follow the same normal distribution with a slightly raised variance—i.e., the $\mathcal{N}(0.5, 0.15)$ distribution depicted in Fig. A-1. Table A.2 summarizes the parameters and distributions associated with *BB* and *UP* agents behaviour. About 50% of the participants in our setting belong to the *BB* class, with the rest being *UP* agents (since agents are assigned to a specific class with 50% probability). *BB* and *UP* sufficiently capture two realistic, highly plausible scenarios of agent behaviour. Of course, the underlying models of agents behaviour could follow any other distribution as well.⁶

Table A.2: Behavioural classes of the scheme participants.

Class	Parameter	Distribution
<i>BB</i>	$\hat{\sigma}$	$\mathcal{B}(1, 5)$
	α_i	$\mathcal{B}(4, 2)$
<i>UP</i>	$\hat{\sigma}$	$\mathcal{N}(0.5, 0.15)$
	α_i	$\mathcal{N}(0.5, 0.15)$

⁶For instance, in [8], the classes of *accurate* and *inaccurate* predictors were also introduced. However, these behavioural classes are less interesting since their behaviour is easier to predict; and, moreover, they are less likely in practice, since (i) in realistic settings, errors do occur, while (ii) if predictions are highly inaccurate, the agents would most probably be acting upon them already, to avoid penalties.

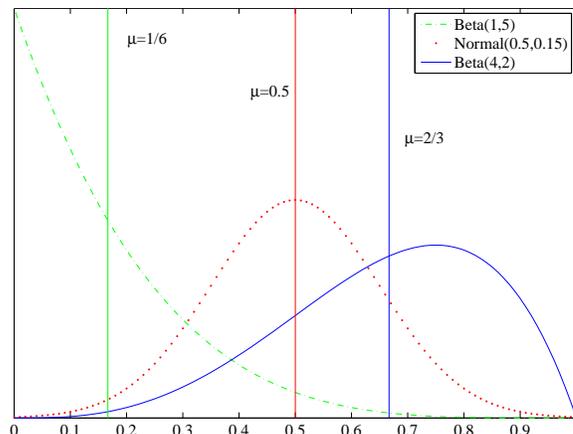


Figure A-1: Probability density functions of agent behaviors.

A.3 Shifting costs

In the absence of actual real world data, we determine the agents *shifting costs* via the use of certain probability distributions. We use Beta distributions for the same reasons outlined above, and due to the fact that their support is in $[0, 1]$: shifting costs are not likely to be negative, and the value 1 is too large to be assumed as a realistic shifting cost—as it would mean that the consumer would require a compensation of 1 €/kWh for shifting, which is not realistic. We now proceed to explain in detail the way we generate the costs for each category that is eligible for the shifting operations.

Residential For this consumer type, shifting costs are initialized according to the following procedure. First, the following must hold, in order for a consumer not to diverge from its baseline consumption without participating in our scheme:

$$\text{cost}_i^{t^h \rightarrow t^l} > p_{high} - p_{low}$$

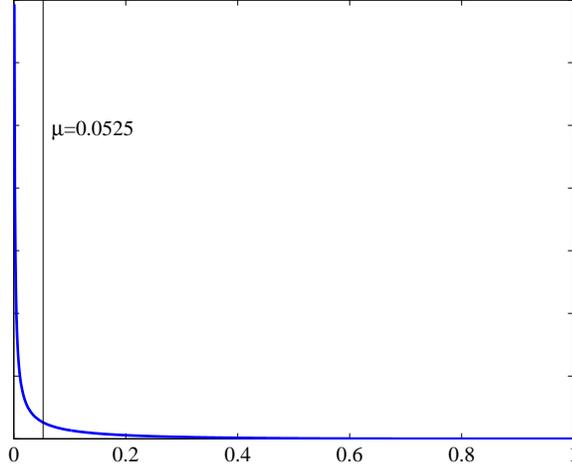


Figure A-2: Probability density function of $\mathcal{B}(0.20865, 3.76572)$.

$$\text{cost}_i^{t^h \rightarrow t^l} > 0.0155 \quad (\text{A.1})$$

Thus, we choose a beta distribution [83], $\mathcal{B}(0.20865, 3.76572)$, which has $\mu \approx 0.0525$ and $\sigma \approx 0.1$. The pdf is shown in Fig. A-2. The mean of this distribution is thus set to a value which is greater than the 0.0155 cost required by Eq. (A.1), and lower (but not much) than the p_g^{max} price—thus ensuring that there exists a substantial number of agents that would be willing to participate in the shifting efforts (assuming they can form coalitions that shift substantial consumption quantities), given the better prices granted by the Grid.⁷

To calculate the *actual shifting costs* for each residential participant, we obtain two samples from the aforementioned *Beta* distribution, a higher and a lower one that constitute the means of two *distinct Beta* distributions, each with $\sigma = 0.01$.

⁷This is a rather plausible scenario, since the Grid would most likely make an effort to offer prices that would be appealing to the participants. Moreover, after sampling the distribution, we add 0.0156 to all values to place them in the correct range—a process that results to a sample mean that is actually higher than p_g^{max} .

The distribution with the higher mean is used for shifts from intervals with baseline agent consumption above the agent’s daily average to others below. Intuitively, it is difficult for an agent that has “abnormally high” consumption at some t_j to shift to other t_l ’s (because the abnormally high consumption most probably means that she “is obliged to” consume at t_j). The *Beta* with the lower mean is sampled for all other shifting operations—that is, for shifting from an interval with lower baseline than the agent’s daily average to higher or lower, and for shifting from an interval with baseline consumption higher than the average to an interval with baseline also higher than the average. Finally, the costs for shifting to a *peak* time interval, from any other interval, are set to $+\infty$ for all agents.

Commercial For the commercial shifting costs generation, we adopt the following process. First, we take three samples from our *Beta*, and randomly assign the individuals in the commercial population to one of three cost types (with equal probability): (a) *low cost* (with costs derived by a *Beta* with the lower sampled value), (b) *medium cost* (with costs derived by a *Beta* whose mean is the centre of the interval between the lower and the higher sampled value), and (c) *high cost* (with costs derived by a *Beta* whose mean is the higher value sample).

Then, we define the *cost scaling factor* s_{cost}^f for an interval t , which is *inversely proportional* to a *consumption scaling factor* s_t^f :

$$s_{cost}^f = 1/s_t^f$$

with s_t^f describing the ratio between the hourly demand of a particular consumer

type, and the mean daily demand for that type (in this case, the commercial one). Intuitively, if a commercial consumer consumes much less than its daily mean at a particular interval t , then it could be the case that its business is actually slow or even closed at t —and thus shifting consumption to t is more problematic (implying a higher cost scaling factor for that interval).

Now, according to the non-peak time interval t that the consumer shifts to, we multiply the cost type-associated *Beta* mean with the scaling factor sf_{cost}^t for t , and then sample a *Beta* whose mean is the resulting number, and has $\sigma = 0.01$.

Public Exactly as the commercial.

Industrial The same as the commercial, but no scaling factors apply; every interval is a sample from a beta distribution with the mean of the corresponding cost type. There are three different types: low, medium, and high cost, exactly as above.

In Figure A-3 we present the actual agent shifting cost model generated from the aforementioned procedures. Specifically, Fig. A-3a is a histogram of the distribution of average-per-agent daily shifting cost values. We can observe that the larger mass is concentrated close to the minimum eligible shifting cost value, but many agents with higher average values also exist. Figure A-3b illustrates the actual shifting cost values sampled for all agents over the course of a typical day, sorted in ascending order with respect to the cost values sampled. Note that the shifting cost values spread smoothly across the $[0, 1]$ interval.

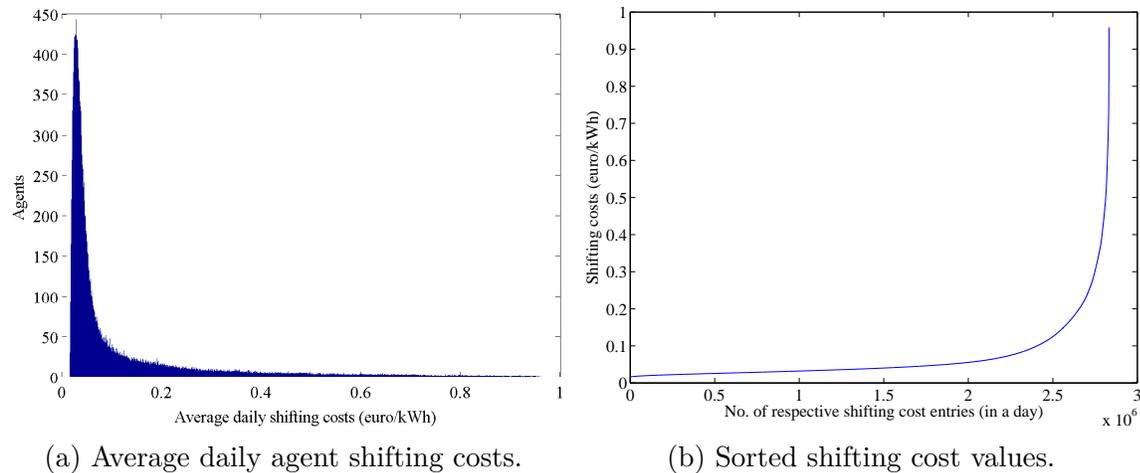


Figure A-3: Model of agent costs.

A.4 Renewable energy sources production levels

For the purposes of the simulations performed in Chapters 4 and 5, we fitted normal distributions on real-world measurements from a renewable energy park located in Spain. The particular park is called Sotavento Galicia, and is equipped with both photovoltaic and wind generators.⁸

To build the models for the renewable generators that were used in our experiments, we calculated the typical hourly values from the measurements collected during the year 2012. Then, by sampling normal distributions described by resulting means and variances, we managed to retrieve realistic hourly renewable production values.

⁸<http://www.sotaventogalicia.com>

A.5 Electric vehicles demand levels

For the purposes of the simulation experiments executed in the scope of Chapter 5, consumption models of electric vehicles were required. The Customer-Led Network Revolution project ⁹ provides hourly and monthly EV consumption measurements that were collected from May of 2013 until July of 2014.

Now, due to the fact that the consumption of each EV is constrained by the charging rate and battery capacity, obtaining samples from normal distributions with the resulting means and variances from the dataset values would not be appropriate. On the contrary, without loss of generality, we assume fixed battery capacities (240 kWh) and charging rates (up to 86.4 kW per hour), and use the measured values to infer each EV's time of connection. During intervals of higher mean EV consumption, the probability of an EV connecting to a charger is higher. Thus, the original dataset's consumption distributions were used to only sample the time of connection of an electric vehicle, and not to indicate the demand levels; these were determined according to the state of charge (SOC) of each battery, and the target SOC value that the driver requested upon disconnection. These resulted by sampling a beta distribution $\mathcal{B}(1, 5)$ and a normal distribution $\mathcal{N}(0.5, 0.1)$, respectively. The duration of a vehicles connection to the charger was sampled by a uniform random distribution, that generates a random integer between 1 and 10 (hours). For the original demand curves, the batteries charge immediately upon connection until the target SOC is reached. When the V2G/G2V algorithm is applied, charging and discharging can be shifted to any time interval that the vehicle has an uninterrupted connection.

⁹<http://www.networkrevolution.co.uk/resources/project-data/>

Bibliography

- [1] European Federation for Renewable Energy Sources Cooperatives. <https://rescoop.eu>. Accessed: 2017-04-15.
- [2] Kyoto protocol to the united nations framework convention on climate change. Technical report, United Nations Treaty Database, December 1997.
- [3] REEEP/UNIDO Training Package: Sustainable Energy Regulation and Policymaking for Africa. Module 14: Demand-side Management, 2008.
- [4] Paris agreement. Technical report, United Nations Treaty Collection, July 2016.
- [5] T. Ackermann, G. Andersson, and L. Söder. Distributed generation: a definition. *Electric Power Systems Research*, 57(3):195–204, 2001.
- [6] S. Aggarwal. *Bitcoin Magnet*. Notion Press, 2017.
- [7] C. Akasiadis and G. Chalkiadakis. Agent Cooperatives for Effective Power Consumption Shifting. In *Proc. of the 27th AAAI Conference on Artificial Intelligence, (AAAI-2013)*, pages 1263–1269, 2013.
- [8] C. Akasiadis and G. Chalkiadakis. Stochastic filtering methods for predicting agent performance in the smart grid. In *ECAI 2014 - 21st European Conference on Artificial Intelligence, 18-22 August 2014, Prague, Czech Republic - Including Prestigious Applications of Intelligent Systems (PAIS 2014)*, pages 1205–1206, 2014.
- [9] C. Akasiadis and G. Chalkiadakis. Decentralized large-scale electricity consumption shifting by prosumer cooperatives. In *ECAI 2016 - 22nd European Conference on Artificial Intelligence*, pages 175–183, 2016.

- [10] C. Akasiadis and G. Chalkiadakis. Cooperative electricity consumption shifting. *Sustainable Energy, Grids and Networks*, 9:38 – 58, 2017.
- [11] C. Akasiadis and G. Chalkiadakis. Mechanism design for demand-side management. *IEEE Intelligent Systems*, 32(1):24–31, Jan 2017.
- [12] C. Akasiadis, G. Chalkiadakis, M. Mamakos, N. Savvakis, T. Tsoutsos, T. Hoppe, and F. Coenen. Analyzing Statistically the Energy Consumption and Production Patterns of European REScoop Members: Results from the H2020 Project REScoop Plus. In *Proc. of the 9th International Exergy, Energy and Environment Symposium (IEEEES - 9)*, May 2017.
- [13] C. Akasiadis and A. Georgogiannis. Predicting agent performance in large-scale electricity demand shifting. *Advances in Building Energy Research*, pages 1–22, in press, 2017.
- [14] C. Akasiadis, K. Panagidi, N. Panagiotou, P. Sernani, A. Morton, I. A. Vetsikas, L. Mavrouli, and K. Goutsias. Incentives for rescheduling residential electricity consumption to promote renewable energy usage. In *2015 SAI Intelligent Systems Conference (IntelliSys)*, pages 328–337, Nov 2015.
- [15] A. Al-Fuqaha, M. Guizani, M. Mohammadi, M. Aledhari, and M. Ayyash. Internet of Things: A survey on enabling technologies, protocols and applications. *IEEE Communications Surveys & Tutorials*, (99), 2015.
- [16] A. Alan, E. Costanza, J. Fischer, S. Ramchurn, T. Rodden, and N.R. Jennings. A field study of human-agent interaction for electricity tariff switching. In *Proceedings of the 13th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-2014)*, pages 965–972, 2014.
- [17] A. Alan, E. Costanza, J. Fischer, S. D. Ramchurn, T. Rodden, and N. R. Jennings. A field study of human-agent interaction for electricity tariff switching. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems (AAMAS-2014)*, pages 965–972, 2014.
- [18] Alper T Alan, Mike Shann, Enrico Costanza, Sarvapali D Ramchurn, and Sven Seuken. It is too hot: An in-situ study of three designs for heating. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pages 5262–5273. ACM, 2016.

- [19] S. Alessandrini, L. Delle Monache, S. Sperati, and G. Cervone. An analog ensemble for short-term probabilistic solar power forecast. *Applied energy*, 157:95–110, 2015.
- [20] S. Aman, Y. Simmhan, and V.K. Prasanna. Energy management systems: state of the art and emerging trends. *Communications Magazine, IEEE*, 51(1):114–119, January 2013.
- [21] A. Angelidakis and G. Chalkiadakis. Factored MDPs for optimal prosumer decision-making. In *Proc. of the 2015 International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2015)*, pages 503–511, 2015.
- [22] A. Angelidakis and G. Chalkiadakis. Factored MDPs for optimal prosumer decision-making in continuous state spaces. In *Proceedings of EUMAS-2015*, pages 91–107, 2015.
- [23] P. Asmus. Microgrids, virtual power plants and our distributed energy future. *The Electricity Journal*, 23(10):72–82, 2010.
- [24] L. Atzori, A. Iera, and G. Morabito. The internet of things: A survey. *Computer Networks*, 54(15):2787 – 2805, 2010.
- [25] A. Back, M. Corallo, L. Dashjr, M. Friedenbach, G. Maxwell, A. Miller, A. Poelstra, J Timon, and P. Wuille. Enabling Blockchain Innovations with Pegged Sidechains. <https://blockstream.com/sidechains.pdf>, 2014. [Online; accessed 9-August-2017].
- [26] D. Bandyopadhyay and J. Sen. Internet of things: Applications and challenges in technology and standardization. *Wireless Personal Communications*, 58(1):49–69, 2011.
- [27] S. Bandyopadhyay, R. Narayanam, R. Kota, Pg Dr M. I. Pg Hj Petra, and Z. Charbiwala. Aggregate demand-based real-time pricing mechanism for the smart grid: a game-theoretic analysis. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI-2015)*, pages 2554–2560. AAAI Press, 2015.

- [28] A. Banks and R. Gupta. MQTT protocol specification. <http://docs.oasis-open.org/mqtt/mqtt/v3.1.1/mqtt-v3.1.1.html>, 2014 (accessed July 22, 2017).
- [29] A. Barbato and A. Capone. Optimization models and methods for demand-side management of residential users: A survey. *Energies*, 7(9):5787–5824, 2014.
- [30] A. Barbato, A. Capone, M. Rodolfi, and D. Tagliaferri. Forecasting the usage of household appliances through power meter sensors for demand management in the smart grid. In *Smart Grid Communications (SmartGridComm), 2011 IEEE Int. Conference*, pages 404–409, Oct 2011.
- [31] S. Behboodi, D. P. Chassin, C. Crawford, and N. Djilali. Renewable resources portfolio optimization in the presence of demand response. *Applied Energy*, 162:139 – 148, 2016.
- [32] F. Bellotti, B. Kapralos, K. Lee, P. Moreno-Ger, and R. Berta. Assessment in and of serious games: An overview. *Advances in Human-Computer Interaction*, 2013.
- [33] C. Bikcora, L. Verheijen, and S. Weiland. Semiparametric density forecasting of electricity load for smart charging of electric vehicles. In *Control Applications (CCA), 2015 IEEE Conference on*, pages 1564–1570. IEEE, 2015.
- [34] M. Blum and M. Riedmiller. Electricity Demand Forecasting using Gaussian Processes. In *Workshops at the Twenty-Seventh AAAI Conference on Artificial Intelligence*, 2013.
- [35] C. Boutilier. Eliciting forecasts from self-interested experts: scoring rules for decision makers. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2012)-Volume 2*, pages 737–744, 2012.
- [36] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge University Press, 2004.
- [37] A. Bracale, P. Caramia, G. Carpinelli, A. R. Di Fazio, and G. Ferruzzi. A bayesian method for short-term probabilistic forecasting of photovoltaic generation in smart grid operation and control. *Energies*, 6(2):733, 2013.

- [38] J. Bushnell, B. F. Hobbs, and F. A. Wolak. When it comes to demand response, is ferc its own worst enemy? *The Electricity Journal*, 22(8):9–18, October 2009.
- [39] G. Chalkiadakis, E. Elkind, and M. Wooldridge. Computational aspects of cooperative game theory. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 5(6):1–168, 2011.
- [40] G. Chalkiadakis, V. Robu, R. Kota, A. Rogers, and N. R. Jennings. Cooperatives of distributed energy resources for efficient virtual power plants. In *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2011)*, pages 787–794, May 2011.
- [41] S. Gujar Chandra, P. and Y. Narahari. Crowdfunding public projects with provision point: A prediction market approach. In *ECAI 2016 - 22nd European Conference on Artificial Intelligence*, pages 778–786, 2016.
- [42] V. Charissis and S. Papanastasiou. Human–machine collaboration through vehicle head up display interface. *Cognition, Technology and Work*, 12:41–50, 03 2010.
- [43] C. Chatfield and G.J. Goodhardt. The beta-binomial model for consumer purchasing behaviour. In *Mathematical Models in Marketing*, volume 132 of *Lecture Notes in Economics and Mathematical Systems*, pages 53–57. Springer Berlin Heidelberg, 1976.
- [44] C. Chelmis, S. Aman, M. R. Saeed, M. Frincu, and V. K. Prasanna. Estimating reduced consumption for dynamic demand response. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI-2015)*, 2015.
- [45] H. Chesbrough. *Open business models: How to thrive in the new innovation landscape*. Harvard Business Press, 2006.
- [46] M. M. P. Chowdhury. Predicting prices in the Power TAC wholesale energy market. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence (AAAI-2016)*, pages 4204–4205.
- [47] F. Christianos and G. Chalkiadakis. Efficient multi-criteria coalition formation using hypergraphs (with application to the V2G problem). In *Multi-Agent Systems and Agreement Technologies - 14th European Conference, EUMAS*

- 2016, and 4th International Conference, AT 2016, Valencia, Spain, December 15-16, 2016, Revised Selected Papers, pages 92–108, 2016.
- [48] F. Christianos and G. Chalkiadakis. Employing hypergraphs for efficient coalition formation with application to the V2G problem. In *ECAI 2016 - 22nd European Conference on Artificial Intelligence, 29 August-2 September 2016, The Hague, The Netherlands - Including Prestigious Applications of Artificial Intelligence (PAIS 2016)*, pages 1604–1605, 2016.
- [49] K. Christidis and M. Devetsikiotis. Blockchains and smart contracts for the internet of things. *IEEE Access*, 4:2292–2303, 2016.
- [50] F. Coenen, T. Hoppe, G. Chalkiadakis, C. Akasiadis, and T. Tsoutsos. Exploring energy saving policy measures by renewable energy supplying cooperatives (REScoops). In *ECEEE 2017 Summer Study on Energy Efficiency*, May-June 2017.
- [51] Litos Strategic Communication. Smart grid: An introduction. *U.S. Department of Energy*, 2008.
- [52] C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou. The complexity of computing a nash equilibrium. *SIAM Journal on Computing*, 39(1):195–259, 2009.
- [53] B. Davito, H. Tai, and R. Uhlener. The smart grid and the promise of demand-side management. *McKinsey*, 2010.
- [54] G. de Oliveira Ramos, J. C. Burguillo, and A. L.C. Bazzan. Dynamic constrained coalition formation among electric vehicles. *Journal of the Brazilian Computer Society*, 20(1):1–15, 2014.
- [55] M.H. DeGroot and M.J. Schervish. *Probability and Statistics*. Pearson custom library. Pearson Education, 2013.
- [56] T. Dethlefs, T. Preisler, and W. Renz. Multi-agent-based distributed optimization for demand-side-management applications. In *Computer Science and Information Systems (FedCSIS), 2014 Federated Conference on*, pages 1489–1496. IEEE, 2014.

- [57] L. Drude, L. Carlos Pereira Junior, and R. R  ther. Photovoltaics (PV) and electric vehicle-to-grid (V2G) strategies for peak demand reduction in urban regions in brazil in a smart grid environment. *Renewable Energy*, 68:443 – 451, 2014.
- [58] A. Ehsanfar and B. Heydari. An incentive-compatible scheme for electricity cooperatives: An axiomatic approach. *IEEE Transactions on Smart Grid*, (99):1–1, 2016.
- [59] T. Erickson, M. Li, Y. Kim, A. Deshpande, S. Sahu, T. Chao, P. Sukaviriya, and M. Naphade. The dubuque electricity portal: Evaluation of a city-scale residential electricity consumption feedback system. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '13, pages 1203–1212, New York, NY, USA, 2013. ACM.
- [60] X. Fang, S. Misra, G. Xue, and D. Yang. Smart grid - the new and improved power grid: A survey. *Communications Surveys Tutorials, IEEE*, (99):1–37, 2011.
- [61] G. Fiorentino and A. Corsi. Internet of things for demand side management. *Journal of Energy and Power Engineering*, 9:500–503, 2015.
- [62] P. Franco. *Understanding Bitcoin: Cryptography, Engineering and Economics*. The Wiley Finance Series. Wiley, 2014.
- [63] J. Friedman, T. Hastie, and R. Tibshirani. *The elements of statistical learning*, volume 1. Springer series in statistics Springer, Berlin, 2001.
- [64] I. Ganchev, Z. Ji, and M. O'Droma. A generic IoT architecture for smart cities. In *25th IET Irish Signals Systems Conference 2014 and 2014 China-Ireland International Conference on Information and Communications Technologies (ISSC 2014/CIICT 2014)*, pages 196–199, June 2014.
- [65] A. Gerossier, R. Girard, G. Kariniotakis, and A. Michiorri. Probabilistic day-ahead forecasting of household electricity demand. In *CIREN 2017-24th International Conference on Electricity Distribution*, page 0625, 2017.

- [66] L. Gkatzikis, I. Koutsopoulos, and T. Salonidis. The role of aggregators in smart grid demand response markets. *IEEE Journal on Selected Areas in Communications*, 31(7):1247–1257, July 2013.
- [67] T. Gneiting and A. E. Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007.
- [68] S. Gottwalt, W. Ketter, C. Block, J. Collins, and C. Weinhardt. Demand side management: A simulation of household behavior under variable prices. *Energy policy*, 39(12):8163–8174, 2011.
- [69] J. Granados, A.-M. Rahmani, P. Nikander, P. Liljeberg, and H. Tenhunen. Towards energy-efficient healthcare: An internet-of-things architecture using intelligent gateways. In *Wireless Mobile Communication and Healthcare (Mobihealth), 2014 EAI 4th International Conference on*, pages 279–282. IEEE, 2014.
- [70] J. Gubbi, R. Buyya, S. Marusic, and M. Palaniswami. Internet of things (IoT): A vision, architectural elements, and future directions. *Future generation computer systems*, 29(7):1645–1660, 2013.
- [71] R. Guerrero-Lemus and J.M. Martínez-Duart. *Renewable Energies and Co2: Cost Analysis, Environmental Impacts and Technological Trends- 2012 Edition*. Lecture Notes in Economics and Mathematical Systems. Springer, 2012.
- [72] D. Guinard, V. Trifa, S. Karnouskos, P. Spiess, and D. Savio. Interacting with the SoA-based internet of things: Discovery, query, selection, and on-demand provisioning of web services. *IEEE transactions on Services Computing*, 3(3):223–235, 2010.
- [73] B. Guo, D. Zhang, Z. Wang, Z. Yu, and X. Zhou. Opportunistic IoT: Exploring the harmonious interaction between human and the internet of things. *Journal of Network and Computer Applications*, 36(6):1531–1539, 2013.
- [74] S. Gyamfi, S. Krumdieck, and T. Urmee. Residential peak electricity demand response—highlights of some behavioural issues. *Renewable and Sustainable Energy Reviews*, 25:71–77, 2013.

- [75] L. Györfi, M. Kohler, A. Krzyzak, and H. Walk. *A distribution-free theory of nonparametric regression*. Springer Science & Business Media, 2006.
- [76] G. Haddadian, N. Khalili, M. Khodayar, and M. Shahiedehpour. Security-constrained power generation scheduling with thermal generating units, variable energy resources, and electric vehicle storage for V2G deployment. *International Journal of Electrical Power and Energy Systems*, 73:498 – 507, 2015.
- [77] A. Hakiri, P. Berthou, A. Gokhale, and S. Abdellatif. Publish/subscribe-enabled software defined networking for efficient and scalable IoT communications. *IEEE communications magazine*, 53(9):48–54, 2015.
- [78] S. Himmel, M. Ziefle, and K. Arning. From living space to urban quarter: acceptance of ICT monitoring solutions in an ageing society. In *International Conference on Human-Computer Interaction*, pages 49–58. Springer, 2013.
- [79] S. S. Hosseini, A. Badri, and M. Parvania. The plug-in electric vehicles for power system applications: The vehicle to grid (V2G) concept. In *2012 IEEE International Energy Conference and Exhibition (ENERGYCON)*, pages 1101–1106, Sept 2012.
- [80] C. Ibars, M. Navarro, and L. Giupponi. Distributed demand management in smart grid with a congestion game. In *SmartGridComm 2010*, pages 495 –500, oct. 2010.
- [81] R. K. Jain, K. M. Smith, P. J. Culligan, and J. E. Taylor. Forecasting energy consumption of multi-family residential buildings using support vector regression: Investigating the impact of temporal and spatial monitoring granularity on performance accuracy. *Applied Energy*, 123:168 – 178, 2014.
- [82] P. Janovsky and S. A. DeLoach. Increasing use of renewable energy by coalition formation of renewable generators and energy stores. 2016.
- [83] N.L. Johnson, S. Kotz, and N. Balakrishnan. *Continuous univariate distributions*. Number v. 2 in Wiley series in probability and mathematical statistics: Applied probability and statistics. Wiley & Sons, 1995.

- [84] J. B. Jorgensen and F. Joutz. Modelling and Forecasting Residential Electricity Consumption in the U.S. Mountain Region. Technical Report 2012-003, The George Washington University, Dept. of Economics, January 2012.
- [85] S. Kamboj, W. Kempton, and K. S. Decker. Deploying power grid-integrated electric vehicles as a multi-agent system. In *The 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2011)- Volume 1*, pages 13–20, 2011.
- [86] B. Kantarci and H. T. Mouftah. Trustworthy sensing for public safety in cloud-centric internet of things. *IEEE Internet of Things Journal*, 1(4):360–368, 2014.
- [87] E. L. Karfopoulos and N. D. Hatziargyriou. Distributed coordination of electric vehicles providing V2G services. *IEEE Transactions on Power Systems*, 31(1):329–338, Jan 2016.
- [88] K. Kaur, A. Dua, A. Jindal, N. Kumar, M. Singh, and A. Vinel. A novel resource reservation scheme for mobile phev in v2g environment using game theoretical approach. *IEEE Transactions on Vehicular Technology*, 64(12):5653–5666, Dec 2015.
- [89] A. Kavousi-Fard, T. Niknam, and M. Fotuhi-Firuzabad. Stochastic reconfiguration and optimal coordination of v2g plug-in electric vehicles considering correlated wind power generation. *IEEE Transactions on Sustainable Energy*, 6(3):822–830, July 2015.
- [90] W. Ketter, J. Collins, and P. Reddy. Power TAC: A competitive economic simulation of the smart grid. *Energy Economics*, 39:262–270, 2013.
- [91] M. A. Khan, N. Javaid, A. Mahmood, Z. A. Khan, and N. Alrajeh. A generic demand-side management model for smart grid. *International Journal of Energy Research*, 39(7):954–964, 2015.
- [92] S. U.; Zaheer R.; Khan Shahid Khan, R.; Khan. Future internet: The internet of things architecture, possible applications and key challenges. In *10th International Conference on Frontiers of Information Technology, 2012*, 2012.
- [93] D. Kirschen and G. Strbac. *Fundamentals of Power System Economics*. Wiley, 2004.

- [94] N. Koblitz and A. J. Menezes. Cryptocash, cryptocurrencies, and cryptocontracts. *Designs, Codes and Cryptography*, 78(1):87–102, 2015.
- [95] R. Kosara. Visualization criticism—the missing link between information visualization and art. In *Information Visualization, 2007. IV'07. 11th International Conference*, pages 631–636. IEEE, 2007.
- [96] R. Kota, G. Chalkiadakis, V. Robu, A. Rogers, and N. R. Jennings. Cooperatives for demand side management. In *The Seventh Conference on Prestigious Applications of Intelligent Systems (PAIS @ ECAI)*, pages 969–974, August 2012.
- [97] G. Kotsis, I. Moschos, C. Corchero, and M. Cruz-Zambrano. Demand aggregator flexibility forecast: Price incentives sensitivity assessment. In *2015 12th International Conference on the European Energy Market (EEM)*, pages 1–5, May 2015.
- [98] F. A. Kraemer and P. Herrmann. Creating internet of things applications from building blocks. *ERCIM news*, (101):19–20, April 2015.
- [99] J. Krueger. *Human-Machine Collaboration*, pages 668–671. Springer Berlin Heidelberg, Berlin, Heidelberg, 2014.
- [100] J. Kumagai. Virtual power plants, real power. *Spectrum, IEEE*, 49(3):13–14, 2012.
- [101] G. E. Lee, Y. Xu, R. S. Brewer, and P. M. Johnson. Makahiki: An open source game engine for energy education and conservation. *Department of Information and Computer Sciences, University of Hawaii, Honolulu, Hawaii*, 96822:11–07, 2012.
- [102] J. H. Lee, M. G. Hancock, and M.-C. Hu. Towards an effective framework for building smart cities: Lessons from seoul and san francisco. *Technological Forecasting and Social Change*, 89:80–99, 2014.
- [103] S. Lefeng, Z. Qian, and P. Yongjian. The reserve trading model considering v2g reverse. *Energy*, 59:50 – 55, 2013.

- [104] B. Li, S. Gangadhar, S. Cheng, and P. K. Verma. Predicting user comfort level using machine learning for smart grid environments. In *Innovative Smart Grid Technologies (ISGT), 2011 IEEE PES*, pages 1–6. IEEE, 2011.
- [105] C. Li, U. Rajan, S. Chawla, and K. Sycara. Mechanisms for coalition formation and cost sharing in an electronic marketplace. In *ICEC '03*, pages 68–77, New York, NY, USA, 2003. ACM.
- [106] H. Liang, B. J. Choi, W. Zhuang, and X. Shen. Optimizing the energy delivery via v2g systems based on stochastic inventory theory. *IEEE Transactions on Smart Grid*, 4(4):2230–2243, Dec 2013.
- [107] Y. Liu, C. Yuen, S. Huang, N. Ul Hassan, X. Wang, and S. Xie. Peak-to-average ratio constrained demand-side management with consumer’s preference in residential smart grid. *IEEE Journal of Selected Topics in Signal Processing*, 8(6):1084–1097, Dec 2014.
- [108] T. Logenthiran, D. Srinivasan, and Tan Zong Shun. Demand side management in smart grid using heuristic optimization. *IEEE Transactions on Smart Grid*, 3(3):1244–1252, Sept 2012.
- [109] M.A. López, S. Martín, J.A. Aguado, and S. de la Torre. V2g strategies for congestion management in microgrids with high penetration of electric vehicles. *Electric Power Systems Research*, 104:28 – 34, 2013.
- [110] H. Lund and E. Münster. Integrated energy systems and local energy markets. *Energy Policy*, 34(10):1152 – 1160, 2006.
- [111] Y. Luo, S. Itaya, S. Nakamura, and P. Davis. Autonomous cooperative energy trading between prosumers for microgrid systems. In *Local Computer Networks Workshops (LCN Workshops), 2014 IEEE 39th Conference on*, pages 693–696, Sept 2014.
- [112] H. Ma, V. Robu, N. Li, and D. C. Parkes. Incentivizing reliability in demand-side response. In *Proc. Of 25th int. Joint conf. On artificial intelligence, IJCAI 2016*, pages 352–358, 2016.
- [113] M. Ma and A. Oikonomou. *Serious Games and Edutainment Applications*. Number v. 2. Springer International Publishing, 2017.

- [114] R. L. Machete. Contrasting probabilistic scoring rules. *Journal of Statistical Planning and Inference*, 143(10):1781–1790, 2013.
- [115] T. Malsch. Naming the unnamable: Socionics or the sociological turn of/to distributed artificial intelligence. *Autonomous agents and multi-agent systems*, 4(3):155–186, 2001.
- [116] B. Marques and K. Nixon. The gamified grid: Possibilities for utilising game-based motivational psychology to empower the smart social grid. In *AFRICON, 2013*, pages 1–5, Sept 2013.
- [117] J. Mattila. The Blockchain Phenomenon – The Disruptive Potential of Distributed Consensus Architectures. ETLA Working Papers 38, The Research Institute of the Finnish Economy, May 2016.
- [118] R. C. Merkle. *A Digital Signature Based on a Conventional Encryption Function*, pages 369–378. Springer Berlin Heidelberg, Berlin, Heidelberg, 1988.
- [119] S. Mhanna, G. Verbic, and A. C. Chapman. Towards a realistic implementation of mechanism design in demand response aggregation. In *Proceedings of the 18th Power Systems Computation Conference*, pages 1–7, 2014.
- [120] D. R. Michael and S. L. Chen. *Serious games: Games that educate, train, and inform*. Muska & Lipman/Premier-Trade, 2005.
- [121] M. Mihaylov. *Decentralized Coordination in Multi-Agent Systems*. PhD thesis, Vrije Universiteit Brussel, 2012.
- [122] M. Mihaylov, S. Jurado, N. Avellana, K. Van Moffaert, I.M. de Abril, and A. Nowe. NRGcoin: Virtual currency for trading of renewable energy in smart grids. In *11th International Conference on the European Energy Market (EEM), 2014*, pages 1–6, May 2014.
- [123] N. Miller, P. Resnick, and R. Zeckhauser. Eliciting informative feedback: The peer-prediction method. *Management Science*, 51(9):1359–1373, 2005.
- [124] Da. Miorandi, S. Sicari, F. De Pellegrini, and I. Chlamtac. Internet of things: Vision, applications and research challenges. *Ad Hoc Networks*, 10(7):1497–1516, 2012.

- [125] MIT authors. Engaging electricity demand. In *MIT Interdisciplinary Study on the future of the electric grid*. MIT, 2011.
- [126] H. Mohsenian-Rad and A. Leon-Garcia. Optimal residential load control with price prediction in real-time electricity pricing environments. *IEEE Transactions on Smart Grid*, 1(2):120–133, sept. 2010.
- [127] H. Mohsenian-Rad, V.W.S. Wong, J. Jatskevich, R. Schober, and A. Leon-Garcia. Autonomous demand-side management based on game-theoretic energy consumption scheduling for the future smart grid. *IEEE Transactions on Smart Grid*, 1(3):320–331, dec. 2010.
- [128] E.F. Moran. *Environmental Social Science: Human - Environment interactions and Sustainability*. Wiley, 2011.
- [129] K. Moslehi and R. Kumar. A reliability perspective of the smart grid. *IEEE Transactions on Smart Grid*, 1(1):57–64, June 2010.
- [130] F. Mwasilu, J. J. Justo, E.-K. Kim, T. D. Do, and J.-W. Jung. Electric vehicles and smart grid interaction: A review on vehicle to grid and renewable energy sources integration. *Renewable and Sustainable Energy Reviews*, 34:501–516, 2014.
- [131] S. Nakamoto. Bitcoin: A peer-to-peer electronic cash system, 2008.
- [132] H. Ning and Z. Wang. Future internet of things architecture: like mankind neural system or social organization framework? *IEEE Communications Letters*, 15(4):461–463, 2011.
- [133] N. Nisan. Introduction to Mechanism Design (for Computer Scientists). In N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*, pages 209–242. Cambridge University Press, 2007.
- [134] K. J. O’Dwyer and D. Malone. Bitcoin mining and its energy footprint. In *Irish Signals Systems Conference 2014 and 2014 China-Ireland International Conference on Information and Communications Technologies (ISSC 2014/CIICT 2014)*. 25th IET, pages 280–285, June 2014.

- [135] P. Oliveira, L. Gomes, T. Pinto, P. Faria, Z. Vale, and H. Morais. Load control timescales simulation in a multi-agent smart grid platform. In *IEEE PES ISGT Europe 2013*, pages 1–5, Oct 2013.
- [136] A. A. Panagopoulos, G. Chalkiadakis, and E. Koutroulis. Predicting the power output of distributed renewable energy resources within a broad geographical region. In *ECAI-2012/PAIS-2012: 20th European Conference on Artificial Intelligence, Prestigious Applications of Intelligent Systems Track*, August 2012.
- [137] C. H. Papadimitriou. Computational complexity. In *Encyclopedia of Computer Science*, pages 260–265. John Wiley and Sons Ltd., Chichester, UK.
- [138] D. C. Parkes. Iterative combinatorial auctions: Achieving economic and computational efficiency. PhD Dissertation, Department of Computer and Information Science, University of Pennsylvania, May 2001.
- [139] O. Parson, S. Ghosh, M. Weal, and A. Rogers. Non-intrusive load monitoring using prior models of general appliance types. In *Twenty-Sixth Conference on Artificial Intelligence (AAAI-12)*, 2012.
- [140] G. Pepermans, J. Driesen, and D. Haeseldonckx. Distributed generation: definition, benefits and issues. Energy, Transport and Environment Working Papers Series ete0308, Katholieke Universiteit Leuven, Centrum voor Economische Studiën, Energy, Transport and Environment, August 2003.
- [141] M. Pilkington. Blockchain technology: Principles and applications. *Research Handbook on Digital Transformations*, 2016.
- [142] REScoop Plus. Data Analysis Report. Deliverable 2.3, H2020-EE-2015-3-MarketUptake, 2017.
- [143] V. Pukelienė and I. Maksvytienė. Economy scale impact on the enterprise competitive advantages. *Engineering Economics*, 57(2), 2015.
- [144] J. Quiñonero Candela and C. E. Rasmussen. A Unifying View of Sparse Approximate Gaussian Process Regression. *Journal of Machine Learning Research*, 6:1939–1959, December 2005.

- [145] S. D. Ramchurn, D. Huynh, and N. R. Jennings. Trust in multi-agent systems. *The Knowledge Engineering Review*, 19(1):1–25, 2004.
- [146] S. D. Ramchurn, P. Vytelingum, A. Rogers, and N. Jennings. Agent-based control for decentralised demand side management in the smart grid. In *The 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2011)*, pages 5–12, 2011.
- [147] S. D. Ramchurn, P. Vytelingum, A. Rogers, and N. R. Jennings. Putting the ‘smarts’ into the smart grid: a grand challenge for artificial intelligence. *Communications of the ACM*, 55(4):86–97, 2012.
- [148] S. D. Ramchurn, F. Wu, J.E. Fischer, S. Reece, W. Jiang, S. J. Roberts, T. Rodden, and N. R. Jennings. Human-agent collaboration for disaster response. *Journal of Autonomous Agents and Multi-Agent Systems*, pages 1–30, 2016.
- [149] C. E. Rasmussen and H. Nickisch. The gpml toolbox. <http://www.gaussianprocess.org/>, January 2013.
- [150] D. B. Richardson. Encouraging vehicle-to-grid (v2g) participation through premium tariff rates. *Journal of Power Sources*, 243:219 – 224, 2013.
- [151] J. K. Rishee, K. M. Smith, P. J. Culligan, and J. E. Taylor. Forecasting energy consumption of multi-family residential buildings using support vector regression: Investigating the impact of temporal and spatial monitoring granularity on performance accuracy. *Applied Energy*, 123:168 – 178, 2014.
- [152] U. Ritterfeld, M. Cody, and P. Vorderer. *Serious games: Mechanisms and effects*. Routledge, 2009.
- [153] V. Robu, G. Chalkiadakis, R. Kota, A. Rogers, and N. R. Jennings. Rewarding cooperative virtual power plant formation using scoring rules. *Energy*, 117:19–28, 2016.
- [154] V. Robu, E. H. Gerding, S. Stein, D. C. Parkes, A. Rogers, and N. R. Jennings. An online mechanism for multi-unit demand and its application to plug-in hybrid electric vehicle charging. *Journal of Artificial Intelligence Research*, 48:175–230, 2013.

- [155] V. Robu, R. Kota, G. Chalkiadakis, A. Rogers, and N. R. Jennings. Cooperative virtual power plant formation using scoring rules. In *Proceedings of the 26th Conference on Artificial Intelligence (AAAI-12)*, pages 370–376, August 2012.
- [156] R. Roche, S. Suryanarayanan, T. M. Hansen, S. Kiliccote, and A. Miraoui. A multi-agent model and strategy for residential demand response coordination. In *2015 IEEE Eindhoven PowerTech*, pages 1–6, June 2015.
- [157] T. A. Rodden, J. E. Fischer, N. Pantidi, K. Bachour, and S. Moran. At home with agents: Exploring attitudes towards future smart energy infrastructures. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1173–1182. ACM, 2013.
- [158] J. Rodríguez-Molina, M. Martínez-Núñez, J.-F. Martínez, and W. Pérez-Aguiar. Business models in the smart grid: challenges, opportunities and proposals for prosumer profitability. *Energies*, 7(9):6142–6171, 2014.
- [159] H. Rose, A. Rogers, and E. H. Gerding. A scoring rule-based mechanism for aggregate demand prediction in the smart grid. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2012)*, pages 661–668, 2012.
- [160] A. Rosin, H. Hõimoja, T. Möller, and M. Lehtla. Residential electricity consumption and loads pattern analysis. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence (IJCAI-2010)*, Electric Power Quality and Supply Reliability Conference (PQ), pages 111–116, 2010.
- [161] H. Saboori, M. Mohammadi, and R. Taghe. Virtual power plant (VPP), definition, concept, components and types. In *2011 Asia-Pacific Power and Energy Engineering Conference*, pages 1–4, March 2011.
- [162] Y. Sakurai, M. Shinoda, S. Oyama, and M. Yokoo. Flexible reward plans for crowdsourced tasks. In *International Conference on Principles and Practice of Multi-Agent Systems*, pages 400–415. Springer, 2015.
- [163] P. Samadi, H. Mohsenian-Rad, R. Schober, and V. W.S. Wong. Advanced demand side management for the future smart grid using mechanism design. *IEEE Transactions on Smart Grid*, 3(3):1170–1180, 2012.

- [164] C. Sarkar, S. N. A. U. Nambi, R. V. Prasad, and A. Rahim. A scalable distributed architecture towards unifying iot applications. In *2014 IEEE World Forum on Internet of Things (WF-IoT)*, pages 508–513, March 2014.
- [165] M. Scherr. Multiple and coordinated views in information visualization. *Trends in Information Visualization*, pages 38–45, 2008.
- [166] J. Serra, D. Pubill, A. Antonopoulos, and C. Verikoukis. Smart hvac control in iot: Energy consumption minimization with user comfort constraints. *The Scientific World Journal*, 2014, 2014.
- [167] K. Shahryari and A. Anvari-Moghaddam. Demand side management using the internet of energy based on fog and cloud computing. In *10th IEEE International Conference on Internet of Things (iThings 2017)*, volume 9, pages 500–503, 2015.
- [168] P. H. Shaikh, N. B. Mohd Nor, P. Nallagownden, I. Elamvazuthi, and T. Ibrahim. A review on optimized control systems for building energy and comfort management of smart sustainable buildings. *Renewable and Sustainable Energy Reviews*, 34(0):409–429, 2014.
- [169] Mike Shann, Alper Alan, Sven Seuken, Enrico Costanza, and Sarvapali D Ramchurn. Save money or feel cozy?: A field experiment evaluation of a smart thermostat that learns heating preferences. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems (AAMAS-2017)*, pages 1008–1016, 2017.
- [170] N. Sharma, P. Sharma, D. Irwin, and P. Shenoy. Predicting solar generation from weather forecasts using machine learning. In *IEEE International Conference on Smart Grid Communications (SmartGridComm), 2011*, pages 528–533, Oct 2011.
- [171] T. B. Sheridan. Human–robot interaction. *Human Factors*, 58(4):525–532, 2016. PMID: 27098262.
- [172] J. Shneidman and D. C. Parkes. Overcoming rational manipulation in mechanism implementations. Technical report.

- [173] Y. Shoham and K. Leyton-Brown. *Multiagent systems: Algorithmic, game-theoretic, and logical foundations*. Cambridge University Press, 2009.
- [174] P. Goncalves Da Silva, D. Ilić, and S. Karnouskos. The impact of smart grid prosumer grouping on forecasting accuracy and its benefits for local electricity market trading. *IEEE Transactions on Smart Grid*, 5(1):402–410, Jan 2014.
- [175] F.P. Sioshansi. *Innovation and Disruption at the Grid’s Edge: How distributed energy resources are disrupting the utility business model*. Elsevier Science, 2017.
- [176] A. Soares, Á. Gomes, C. H. Antunes, and H. Cardoso. Domestic load scheduling using genetic algorithms. In *European Conference on the Applications of Evolutionary Computation*, pages 142–151. Springer, 2013.
- [177] M. Soliman, H. and A. Leon-Garcia. Game-theoretic demand-side management with storage devices for the future smart grid. *IEEE Transactions on Smart Grid*, 5(3):1475–1485, 2014.
- [178] H. Song, J. Zhu, Y. Jiang, and B. Li. A new method for balancing cloud resource. In *Web Information System and Application Conference (WISA), 2015 12th*, pages 95–98. IEEE, 2015.
- [179] R. Sparrow. Killer robots. *Journal of applied philosophy*, 24(1):62–77, 2007.
- [180] P. Stelmach. Service composition scenarios in the internet of things paradigm. In *Technological Innovation for the Internet of Things*, pages 53–60. Springer, 2013.
- [181] S. Stoff, J. Eto, and S. Kito. *DSM shareholder incentives: Current designs and economic theory*. Energy & Environment Division, Lawrence Berkeley Laboratory, University of California, 1995.
- [182] G. Strbac. Demand side management: Benefits and challenges. *Energy Policy*, 36(12):4419 – 4426, 2008. Foresight Sustainable Energy Management and the Built Environment Project.
- [183] S. Stüdli, E. Crisostomi, R. Middleton, and R. Shorten. Optimal real-time distributed V2G and G2V management of electric vehicles. *International Journal of Control*, 87(6):1153–1162, 2014.

- [184] Q. Sun, A. Beach, M. E. Cotterell, Z. Wu, and S. Grijalva. An economic model for distributed energy prosumers. In *System Sciences (HICSS), 2013 46th Hawaii International Conference on*, pages 2103–2112, Jan 2013.
- [185] M. Swan. *Blockchain: Blueprint for a New Economy*. O’Reilly Media, 2015.
- [186] S. B. Taieb, R. Huser, R. J. Hyndman, and M. G. Genton. Forecasting uncertainty in electricity smart meter data by boosting additive quantile regression. *IEEE Transactions on Smart Grid*, 7(5):2448–2455, 2016.
- [187] J. W. Taylor. Triple seasonal methods for short-term electricity demand forecasting. *European Journal of Operational Research*, 204(1):139–152, 2010.
- [188] T. T. Teo, T. Logenthiran, and W. L. Woo. Forecasting of photovoltaic power using extreme learning machine. In *IEEE Innovative Smart Grid Technologies - Asia (ISGT ASIA), 2015*, pages 1–6, Nov 2015.
- [189] S. Thrun, W. Burgard, and D. Fox. *Probabilistic Robotics*. MIT Press, 2005.
- [190] Ngoc Cuong Truong, Tim Baarslag, Gopal Ramchurn, and Long Tran-Thanh. Interactive scheduling of appliance usage in the home. July 2016.
- [191] I. Ungurean, N.-C. Gaitan, and V. G. Gaitan. An iot architecture for things from industrial environment. In *Communications (COMM), 2014 10th International Conference on*, pages 1–4. IEEE, 2014.
- [192] K. Valogianni, W. Ketter, and J. Collins. A multiagent approach to variable-rate electric vehicle charging coordination. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2015)*, AAMAS ’15, pages 1131–1139, 2015.
- [193] S. S. van Dam, C. A. Bakker, and J. D. M. van Hal. Home energy monitors: impact over the medium-term. *Building Research & Information*, 38(5):458–469, 2010.
- [194] A. Veit and H.-A. Jacobsen. Multi-agent device-level modeling framework for demand scheduling. In *IEEE Int. Conf. on Smart Grid Communications, SmartGridComm ’15*, 2015.

- [195] A. Veit, Y. Xu, R. Zheng, N. Chakraborty, and K. Sycara. Demand side energy management via multiagent coordination in consumer cooperatives. *Journal of Artificial Intelligence Research*, pages 885–922, 2014.
- [196] A. Veit, Y. Xu, R. Zheng, N. Chakraborty, and K. P. Sycara. Multiagent coordination for energy consumption scheduling in consumer cooperatives. In *Proc. of the 27th AAAI Conference on Artificial Intelligence, (AAAI-2013)*, pages 1362–1368, 2013.
- [197] E. Viardot, T. Wierenga, and B. Friedrich. The role of cooperatives in overcoming the barriers to adoption of renewable energy. *Energy Policy*, 63:756 – 764, 2013.
- [198] M. Vinyals, F. Bistaffa, A. Farinelli, and A. Rogers. Coalitional energy purchasing in the smart grid. In *Energy Conference and Exhibition (ENERGYCON), 2012 IEEE International*, pages 848–853. IEEE, 2012.
- [199] M. Vinyals, V. Robu, A. Rogers, and N. R. Jennings. Prediction-of-use games: a cooperative game theory approach to sustainable energy tariffs. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*, pages 829–836, 2014.
- [200] P. Vytelingum, T. D. Voice, S. D. Ramchurn, A. Rogers, and N. R. Jennings. Theoretical and practical foundations of large-scale agent-based micro-storage in the smart grid. *Journal of Artificial Intelligence Research*, 42:765–813, 2011.
- [201] C. Wei and Y. Li. Design of energy consumption monitoring and energy-saving management system of intelligent building based on the internet of things. In *2011 International Conference on Electronics, Communications and Control (ICECC)*, pages 3650–3652, Sept 2011.
- [202] M. Wooldridge and N. R. Jennings. Intelligent agents: Theory and practice. *The knowledge engineering review*, 10(2):115–152, 1995.
- [203] J. Yamamoto and K. Sycara. A stable and efficient buyer coalition formation scheme for e-marketplaces. In *Proceedings of the 5th International Conference on Autonomous Agents (AGENTS '01)*, pages 576–583, New York, NY, USA, 2001. ACM.

- [204] R. Yu, J. Ding, W. Zhong, Y. Liu, and S. Xie. Phev charging and discharging cooperation in v2g networks: A coalition game approach. *IEEE Internet of Things Journal*, 1(6):578–589, Dec 2014.
- [205] M. Zeng, S. Leng, S. Maharjan, S. Gjessing, and J. He. An incentivized auction-based group-selling approach for demand response management in V2G systems. *IEEE Transactions on Industrial Informatics*, 11(6):1554–1563, Dec 2015.
- [206] H. T. Zhang, F. Y. Xu, and L. Zhou. Artificial neural network for load forecasting in smart grid. In *2010 International Conference on Machine Learning and Cybernetics*, volume 6, pages 3200–3205, July 2010.
- [207] J. Zhu, H. Song, Y. Jiang, B. Li, and J. Wang. On cloud resources consumption shifting scheme for two different geographic areas. *IEEE Systems Journal*, 2016.
- [208] Y. Zhu. *Measuring Effective Data Visualization*, pages 652–661. Springer Berlin Heidelberg, Berlin, Heidelberg, 2007.