A NOVEL ELECTRICITY DEMAND MANAGEMENT SCHEME VIA MULTIAGENT COOPERATIVES IN THE SMART GRID

CHARILAOS AKASIADIS

Master of Science Thesis

Intelligent Systems Laboratory
School of Electronic and Computer Engineering
Technical University of Crete

December 2013 – version 1.1

Η ΙΣΧΥΣ ΕΝ ΤΗ ΕΝΩΣΕΙ

Dedicated to the oppressed citizens of Greece.

ABSTRACT

In this work, we present a directly applicable scheme for power consumption shifting, and the effective flattening of the electricity consumption curve corresponding to some future date (e.g., the day ahead). It is a pro-active scheme, rather than a last-minute peak trimming one; and it can employ the services of either individual or cooperating consumer agents alike. Agents participating in the scheme, however, are motivated to form cooperatives, in order to reduce their electricity bills via lower group prices granted for sizable consumption shifting from high to low demand time intervals. The scheme takes into account individual costs, and uses a strictly proper scoring rule to reward contributors according to efficiency. Cooperative members, in particular, can attain variable reduced electricity price rates, given their different load-shifting capabilities. This allows even agents with initially forbidding shifting costs to participate in the scheme, and is achieved by a weakly budget-balanced, truthful reward sharing mechanism. We provide four variants of this approach, and evaluate them experimentally. One major problem arising in this domain is assessing the participating agents' uncertainty, and correctly predicting their future behavior. Thus, in this work we adopt two stochastic filtering techniques, the Unscented Kalman Filter and the Histogram Filter, and use them to effectively monitor the trustworthiness of agent statements regarding their final actions. Interestingly, our UKF filter is equipped with a Gaussian Process regression model. We incorporate these techniques within our demand management scheme. Our simulation results confirm that these techniques provide tangible benefits regarding enhanced consumption reduction performance, and increased financial gains for the cooperative.

ABSTRACT (GREEK)

Σε αυτήν την εργασία παρουσιάζουμε ένα άμεσα εφαρμόσιμο μηχανισμό για τη μεταφορά κατανάλωσης ηλεκτρικού ρεύματος στο χρόνο, με σκοπό την αποτελεσματική εξομάλυνση της καμπύλης ζήτησης για κάποια μελλοντική ημερομηνία (π.χ. για την επόμενη μέρα). Σε αντίθεση με μηχανισμούς που αποσκοπούν στη μείωση της ενεργειακής κατανάλωσης αφού έχει εμφανιστεί ένα μέγιστο στην σχετική καμπύλη (περιορισμός κατανάλωσης σε περιόδους αιχμής), ο μηχανισμός μας προωθεί την αποφυγή δημιουργίας τέτοιων μεγίστων -και μπορεί να συμπεριλάβει ανεξάρτητους ή και συνεταιριζόμενους πράκτορες. Ωστόσο, στους συμμετέχοντες δίδεται κίνητρο για να σχηματίζουν συνεταιρισμούς, έτσι ώστε να μειώσουν τους

λογαριασμούς του ηλεκτρικού ρεύματος, μέσω της εξασφάλισης χαμηλότερων ομαδικών τιμών που δίδονται για μεταφορά ουσιώδους μεγέθους κατανάλωσης από χρονικά διαστήματα υψηλής ζήτησης σε άλλα χαμηλότερης. Η προσέγγισή μας λαμβάνει υπ'όψιν τα κόστη του κάθε πράκτορα και χρησιμοποιεί έναν «αυστηρά αρμόζοντα κανόνα βαθμολόγησης» (strictly proper scoring rule) για να ανταμείψει τους συμμετέχοντες σύμφωνα με την αποδοτικότητά τους. Συγκεκριμένα, τα μέλη του συνεταιρισμού μπορούν να εξασφαλίσουν μειωμένες τιμές ηλεκτρικού ρεύματος, οι οποίες είναι ενδεχομένως διαφορετικές για κάθε μέλος (και εξαρτώνται από τις συγκεκριμένες δυνατότητες των μελών). Αυτό επιτρέπει σε πράκτορες με αρχικά απαγορευτικά κόστη μεταφοράς κατανάλωσης να συμμετάσχουν, μέσω ενός μηχανισμού διαμοιρασμού των κερδών, ο οποίος αποτρέπει ψευδείς δηλώσεις να τον επηρεάζουν (truthful) και επιτυγχάνει ασθενή ισολογισμό του κεφαλαίου (weak budget-balancedness). Προτείνουμε τέσσερις παραλλαγές αυτής της προσέγγισης και τις αξιολογούμε πειραματικά. Ένα μεγάλο πρόβλημα σε αυτόν τον τομέα είναι η κατανόηση της αβεβαιότητας των συμμετεχόντων και η ακριβής πρόβλεψη των μελλοντικών ενεργειών τους. Ως εκ τούτου, σε αυτή την εργασία υιοθετούμε δύο τεχνικές «στοχαστικού φιλτραρίσματος», το φίλτρο Unscented Kalman, και το φίλτρο ιστογράμματος και τις χρησιμοποιούμε για να παρακολουθήσουμε αποτελεσματικά την αξιοπιστία των δηλώσεων των πρακτόρων σε σχέση με τις τελικές τους πράξεις. Μια περαιτέρω ενδιαφέρουσα συνεισφορά της εργασίας μας είναι ο εξοπλισμός του UKF φίλτρου με ένα μοντέλο στοχαστικής παλινδρόμησης, τις Γκαουσσιανές διαδικασίες (Gaussian Processes). Χρησιμοποιούμε αυτές τις τεχνικές στον κύριο μηχανισμό μας, ο οποίος προσφέρει υπηρεσίες διαχείρισης της ζήτησης του ηλεκτρικού ρεύματος. Τα αποτελέσματα των προσομοιώσεων επιβεβαιώνουν ότι αυτές οι τεχνικές προσδίδουν απτά πλεονεκτήματα όσον αφορά τον περιορισμό της ζήτησης, αλλά και αυξημένα χρηματικά κέρδη για τον συνεταιρισμό.

PUBLICATIONS

Some ideas and figures have appeared previously in the following publications:

Charilaos Akasiadis and Georgios Chalkiadakis. Agent Cooperatives for Effective Power Consumption Shifting. In Proc. of the 27th AAAI Conference on Artificial Intelligence, (AAAI-2013), pages 1263–1269, 2013.

If we knew what it was we were doing,
it would not be called research, would it?

— Albert Einstein

ACKNOWLEDGMENTS

I would like to express my appreciation to my advisory committee: Georgios Chalkiadakis, Mihail Lagoudakis and Antonios Deligiannakis. Thanks for giving me the opportunity to conduct research at the Intelligent Systems Laboratory of the Technical University of Crete. Special thanks to my supervisor, Georgios Chalkiadakis for all the time he spent to help me become better and for all the patience that this procedure required; it has been an honour to work with you.

Many thanks to my family for their continuous support, as if it had not been for them, this work would not have been completed.

My gratitude to the Intelligent Systems Laboratory and to all my colleagues with whom we have been working together these two years, Manos Orfanoudakis, Kostas Babas, Giannis Skoulakis, Lefteris Chatzilaris, Aggelos Aggelidakis, Stelios Piperakis and Nikos Kofinas. You are the best team!

Special thanks to Alex Georgogiannis for the countless hours we spent studying scientific issues. You are a great source of inspiration. Thanks for treating me as a brother!

Thanks to my friends from the Department of Production Engineering and Management, Giorgos Strofilas, Katerina Tzedaki and Kalliroi Porfiri. Thanks for reminding me that there is life outside of the lab.

Thanks to my music associates, Nikitas Papadoulakis, Nick Tzanakakis, and also to the band "Menace", for composing the soundtrack of this two-year period.

Thanks to my old friends Vaggelis Varvatsoulakis, Giorgos Malandrakis, Giorgos Garofalakis and Alex Eleytheropoulos. Everyone helped in his own way.

Special thanks to my trainer, Vladimir Metreveli for keeping me fit.

Thanks to Konstantina Mpolanaki for her solidarity. It really meant a lot!

Finally, thanks to the dismissed employees of the Greek public radio and television for reminding me through their broadcasts that there can be hope and vision without funding or political support.

CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF ALGORITHMS

SYMBOLS AND NOTATION

t	Half hour time interval
t_h	Peak interval
T_H	Set of peak intervals
t_l	Non-peak interval
T_L	Set of non-peak intervals
p^t	Price for electricity consumption during interval t
p_{high}	High electricity price assigned for consumption during peak intervals
p_{low}	Low electricity price assigned for consumption during non-peak intervals
τ	Demand threshold that denotes price level change
sl	Safety limit
$q_\tau^{t_h}$	Amount of load above τ at a peak interval
$q_{min}^{t_h}$	Minimum load eligible for shifting
$Q_{max}^{t_h}$	Maximum load eligible for shifting
$r_i^{t_h}$	Load reduced by agent i at t_h
$q_i^{t_l}$	Load shifted by agent i at t_l
$p_{group}(q)$	Better price granted for shifted load (function of reduced load during t_hs
$\hat{r}_i^t$	Stated reduction capacity of i during t

$c_i^{t_h \to t_l}$	Shifting cost of agent i for shifting load from t_h to t_l
$CRPS_i$	Continuously ranked probability score of agent i
B_i	Electricity bill of agent i
$\hat{\sigma}_i$	Stated shifted uncertainty of agent i
$\hat{p}_i$	Reservation price of agent i
$\tilde{r}_i$	Cooperative estimate of i's reduction capacity
p_i^{eff}	Effective price of agent i
p_i^{Eff}	Final effective price of agent i after CRPS application and normalization
α_i	Unknown coefficient about agent final actions
$\tilde{\alpha}_i$	Estimated coefficient about agent final actions

Part I

SMART GRID AND DEMAND SIDE MANAGEMENT

In this part we introduce the reader to the key ideas related to this work. The first chapter provides motivation and introduction. The second chapter provides the background notions that are relevant to our problem, and reviews concepts and ideas essential to our novel cooperative multiagent electricity demand shifting scheme.

1 INTRODUCTION

The scientific study of electricity and electric power production has led to rapid advancements in the industrial production process and technology. Ending up to be treated as a commodity, electricity is responsible for the more convenient everyday life that people of today enjoy, to the degree that it has become a necessity. As a result, the electricity grid has become a complex "machine" whose operation must be reliable and effective. A characteristic of electric power that is problematic, however, is that its storage is quite difficult and expensive, so it must be consumed at the very time that it is generated [30]. Another important issue about contemporary electricity generation is that it is mainly produced 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 is possible to harm surrounding areas, and their burning also produces gases, which help exacerbate the so-called "greenhouse effect". As a remedy to these concerns, recent trends propose greener approaches that will help future electricity production become less polluting and more sustainable, introducing the hope for a cleaner planet [39, 17, 19]. Governmental acts, in particular, from all around the world, drive producing sides to shutdown fossil fuel burning facilities and substitute them with new, "cleaner" ones. The turn to renewable sources for electric power generation is desirable, as it leads to lower carbon dioxide emissions. Also, such sources can be constructed in a non-industrial and decentralized manner, allowing the average household to contribute and benefit from its participation to the electricity production process [1, 36]. Despite the positive impacts that come with the use of renewable sources, new difficulties arise regarding electricity production and demand management. This is because weather dependent sources are by definition intermittent and potentially unreliable with respect to the size of their output. Thus, we need to construct more robust infrastructures to deal with the aforementioned issues effectively.

The Smart Grid

The plan for the future power grid, also known as the *Smart Grid*, incorporates interconnectivity and communication of both power producing and consuming sides, in order to facilitate the management of generation and consumption of electricity, at times where the use of renewable electricity sources rises [19]. More specifically, the Smart Grid aspires to use secure, two-way communication technologies to exploit valuable information across every stage, from generators and distribution networks to final consumption. In this way we can achieve clean production, safety, security, reliability, resilience, efficiency and sustainability in the electricity grid [17]. In short, the impact of the applications of Smart Grid technologies can lead to:

- Active consumer participation.
- Exploitation of all generation and storage possibilities.
- Development of new electricity products, services and markets.
- Optimization of the Grid operation.
- Anticipation and response to system disturbances.
- Resistance to disasters.

Demand Balancing

The data recorded by sensors placed at all levels of the Smart Grid, can be used to effectively balance supply and demand. Provided with this information, end users, who are potentially represented by *autonomous agents* [54], are able to *react* to signals, like supply insufficiency indicators for example, and alter their consumption patterns accordingly. To this end, several load control programs have been proposed, where electricity consumers are encouraged to abridge their consuming activities, or shift them to off-peak hours in order to reduce *peak-to-average ratio (PAR)* [30, 11]. Apart from industrial and other large-scale consumers, the participation of residential customers in such schemes is also possible, provided that smart meters or similar electricity management systems are available [45, 50]. Such schemes may involve a specialized intermediary company offering demand-response services which manages consumers who agree, for a cash reward, to step in and contribute to the "trimming down" of the demand curve in the event of an impending critical period. Also, there exist demand reduction scheme types that provide reduced electricity consumption rates to consumers for lowering their consumption over a prolonged time period, while others use dynamic, real-time pricing [2, 4, 6].

As in a Smart Grid environment electricity is generated and consumed in a purely distributed fashion, it is more efficient to group users in larger entities. Thus, recent work in the multiagent systems community has put forward the notion of Virtual Power Plants (VPPs). These correspond to coalitions of electricity producers, consumers, or even *prosumers*[1], who cooperate in order to meet market demands, mimic the reliability characteristics of traditional power plants, and deal efficiently with the issues that accrue [4, 11]. Generally, a multiagent systems approach can be considered more natural for these kind of paradigms, as it can model various individual participant types with their alternating preferences and uncertainties, and realistically simulate a distributed Smart Grid scenario [8, 54].

An Effective Power Consumption Shifting Scheme

In this work, we propose a simple and directly applicable power

1 A prosumer is an entity that can both produce and consume energy.

consumption shifting scheme for the Smart Grid. Our scheme motivates self-interested business units, represented by autonomous agents that potentially form coalitions, to shift power consumption from peak intervals to others with lower demand, in order to receive lower electricity price rates for their contribution. In more detail, in our setting the independent system operator (ISO) [23] —for instance, the national 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 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. To promote efficiency, we use a *strictly proper scoring rule CRPS*, proposed in the mechanism design literature [18], which incentivizes agents to report their predicted shifting capabilities as accurately as possible.

Now, to avoid overwhelming computational complexity and its related costs, it is clearly in the interests of the Grid to interact with the smallest number of "shifting" business units possible [7]. Thus, 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 small consumer alone. At the same time, individuals might not be capable of sufficient reduction due to high shifting costs. As a result, they are motivated to *join forces in a cooperative*, and attempt to coordinate their actions in order to reach the expected reduction levels and make their participation in the scheme worthwhile. This is similar to group purchasing in e-marketplaces, where some coalition 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 [28, 56]. Inspired by work in that domain, we devise a 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.

The proposed scheme aims to shift power consumption to non-peak intervals and flatten the consumption curve on a *day-ahead basis*. Our shifting scheme requires no legislature change and can be applied directly, including residential, commercial and industrial customers. We do not assume or propose the use of oscillating hourly electricity prices like real-time pricing (RTP) [30] does, but provide participating consumers with privileges in the form of better electricity prices, as explained above. Our approach can thus be considered as a simple and immediately applicable alternative to the use of real-time pricing. As we detail in Section 2.1, RTP is a highly controversial demand side management mechanism—in terms of perceived efficiency, but also politically and ethically. In contrast, our scheme incentivises *truthful and accurate* reduction capacity statements by participating agents. In this way, it is much less prone to manipulation than schemes employing simplistic reduced "flat" consumption tariffs. Moreover, the scheme encourages agent cooperation in consumer cooperatives. We equip such cooperatives with an *individually rational*, *incentive compatible*, and *budget-balanced* [46] reward-sharing mechanism to be used for awarding members' reduction efforts. Briefly, the overall scheme works as follows:

1. The Grid announces peak and non-peak time intervals[2] 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 in this thesis.

3. A *clearing process* takes place, determining the accepted agent bids.

4. During the next day[3], the agreed consumption shifting activities take place.

Incentivising Truthful Statements

In order for agent cooperatives to be functional, efficient, and profitable, they need to take business decisions regarding which members to include in service-providing coalitions. 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 power Grid, trust between selfish agents is not implied, and must be guaranteed. *Mechanism design* and related approaches —in this case, the CRPS rule incorporated in our mechanism— attempt to build trust among agents via providing them with the incentives to truthfully and accurately report their intended future actions, along with their corresponding uncertainty regarding those actions [46, 7, 26, 43]. Unfortunately, even if participating agents are perfectly truthful regarding their abilities and corresponding uncertainty, their reports and

2 We must note that time intervals can be of any size. In this thesis, we consider them to be 48 half hour intervals per day, that is tha standard division in the energy domain [7].

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

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 purpose, one could 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 [21, 35]. Although such approaches have promising results, they cannot immediately predict the actual behavior of a specific agent, which might be motivated by private knowledge or business concerns, neither do they account for errors due to equipment malfunction. In contrast, we propose the application of generic prediction methods, which are nevertheless able to adapt to a specific agent's behavior.

Stochastic Filtering to Monitor Trustworthiness

To this end, we use *stochastic filtering methods* [55] to keep track of the parameters 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)* [49] and the *Unscented Kalman* filter (UKF) [22] to predict the future actual actions of agents participating in cooperatives offering electricity demand management services.

GP-UKF and HF

The *UKF* method is a non-linear extension of the celebrated *Kalman Filter (KF)* recursive state estimation algorithm [49]. The *UKF* parameters in our work here are, interestingly, obtained via a *Gaussian Process (GP)* regression on past observed data. This is in the spirit of recent approaches that equip *UKF* with *GP* prediction models, to achieve improved estimation performance in other domains [3, 24]. Now, the *UKF* technique assumes that final actions are functions of member-stated uncertainty, and performs accordingly. *HF*, on the other hand, ignores agent-stated uncertainty, and takes into account past performance observations only.

These two methods 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 capacities; 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.

Both 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 filtering 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. In a nutshell, our results indicate that *both* filtering techniques examined are strong candidates for monitoring the trustworthiness of selfish agents in the Smart Grid, with GP-UKF's behavior appearing to be more solid. 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.

Summary and Contributions

To the best of our knowledge, this is the first work to propose a specific *protocol* and *mechanism* for achieving large-scale electricity demand *shifting* without the use of "intermediary third parties" or real-time pricing. Also, it is the first work to use stochastic filtering and regression methods for assessing the performance of autonomous, economically-minded agents participating in Smart Grid cooperatives or other such entities. In addition, this is the first time that Gaussian Processes are combined with UKF in this domain, with very promising results. Our experimental results provide a testimony to the benefits arising from the formation of agent cooperatives in the Smart Grid.

In more detail, this is the first work that presents a complete scheme for *electricity consumption shifting*. Our scheme is *directly applicable* and requires no legislature change whatsoever. Now, in order for the scheme to effectively flatten demand, we promote collective consumer actions via *agent cooperatives*. This kind of approach is more natural to use for the diverse and dynamic Smart Grid environment. Moreover, our scheme possesses certain desirable properties. First of all, it performs *proactive* —ahead of time— demand curve flattening. This results to more stable reduction actions planning. Second, it is *truthful*, by employing a proper scoring rule (CRPS) that ranks agents according to their forecasting precision, and fines them accordingly. Third, it is *individually rational*, as it proposes a lower group price for shifted load, that makes it profitable for agents to participate. Moreover, we developed a *weakly budget balanced* gain transfer scheme, that allows agents with initially forbidding shifting costs, not to suffer monetary losses from participation. Finally, stochastic filtering techniques are used for the first time to *monitor agent final shifting actions uncertainty*, and their application can achieve even more effective electricity demand curve flattening. Such methods are able to adapt to specific agent behavior and can be used to tackle agent inaccuracies, and, as result, make the cooperative even more stable and viable. Given all these, the proposed mechanism can be considered *efficient in the face of the uncertainty*, that governs agent statements and final actions in the scenario of collective electricity consumption shifting.

Chapters Organization

The rest of this thesis is structured as follows. Chapter 2 presents some concepts of past related work. In Chapter 3 the consumption shifting problem is explained. In Chapter 4 we discuss the way consumer cooperatives determine their bids for participating in the shifting process, and the details of the reward sharing mechanism. Chapter 5 shows the two proposed filtering methods to monitor agent uncertainty. Chapter 6 lists our experimental simulations and results. Chapter 7 provides conclusions and outlines future work.

2 BACKGROUND AND RELATED WORK

In this chapter we provide a background for the problem tackled on this thesis. We discuss *demand side management* and *peak-trimming* schemes proposed in the literature or used in the real world. Next we present related work on game theory and mechanism design and introduce the reader to important notions from these fields. One field of study within game theory is the formation of coalitions, or in the case of the Smart Grid, *Virtual Power Plant*. Briefly, VPPs are groupings of distributed energy resources which can collectively be viewed as one simple power plant. As we will be explaining, in our work we propose the formation of VPPs of consumers. Given that our work also incorporates stochastic filtering techniques to monitor consumer uncertainties in order to guarantee stability and effectiveness, a review of contemporary systems using such techniques is also presented.

2.1 PEAK TRIMMING AND DEMAND-SIDE MANAGEMENT

Much research has been made regarding *demand side management* services [4, 47]. Such approaches allow the end users alter their consumption patterns in response to demand levels i.e. reduce consumption at times when generation costs are high [2]. This concept is similar to peak-trimming [27], where a company might call in and ask to reduce consumption during peak periods. On such occasions though, reducing might not be enough, or reduction services might come in too late. Thus, a more realistic approach is to *proactively* determine demand management schemes ahead of time, as in schemes we discuss in this subsection, and the one we propose in our work.

Kota, Chalkiadakis, Robu, Rogers and Jennings [26], were actually the first to propose a demand side management scheme involving electricity consumer cooperatives. Specifically, they proposed the formation of consumer Cooperatives for Demand Side Management (CDSMs), with the aim of "selling back" to the electricity market the amount of load that was not consumed due to proactive reduction measures. 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 *negawatts* in the form of reduction services. Though inspirational, their approach requires a legislature change in order to be applied in real life. This is because, 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, the scheme proposed in this thesis is simpler and easier to apply in practice. In our approach, agents simply run the danger of being granted less profit for their actions than originally promised. Importantly, they are also almost in complete control of that risk, and quite capable of minimizing it, due to the fact that they are guided a priori (and have agreed) to the time slots where they can actually *shift* consumption to.

Several simple reduction schemes that promise reduced flat electricity rates for lower consumption levels over prolonged periods of time, such as critical peak pricing programs, and are already in place in the real world [4]. 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.

Real-Time Pricing

Economists have in recent years been advocating the use of "dynamic", *real-time pricing (RTP)* schemes as a means to avoid market inefficiencies and the aforementioned shortcomings of existing demand reduction schemes, and trim the unwanted peaks [6]. In the MIT introduction to the Smart Grid [30], authors define RTP 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 ("day-ahead pricing") or an hour ahead ("hour-ahead pricing") 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.

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, and perhaps more importantly, 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 [32]. In contrast, our scheme 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. Generally, ahead of time curve flattening is more desirable than peak trimming, at the last minute.

2.2 GAME THEORY AND MECHANISM DESIGN

Due to the economic nature of the control and optimization problems faced by rational entities operating in the Smart Grid, *game theory* [34] approaches are highly appropriate. Such approaches model such problems as games and study the strategies that each player should adopt. Cooperative game theory, and the problem of cooperating agent coalition formation in particular, are very important in the domain of the Smart Grid. For example, Vitelingum *et al.* [52], propose specific strategies for the adaptive management of distributed micro-storage energy devices. Contreras *et al.* [57, 9] provide a *Bilateral Shapley Value* negotiation scheme for forming coalitions of new collaborators in the power transmission network. In [58], a multiagent system modeling of electricity traders in the new "free" electricity market is presented. Authors propose a game theoretic approach for forming coalitions for multilateral trades between Smart Grid entities.

In another work, that of [20], a distributed load management scheme with dynamic pricing strategies is proposed. The problem is treated as a network congestion game where pure Nash equilibrium solutions can be found in a finite number of steps.

In [31], authors try to minimize peak to average ratio, and energy costs. They present a convex formulation of the problem and propose a distributed algorithm that produces solutions calculated as best response on other users consumption patterns. We chose not to follow a pure game theoretic approach, because it requires that every "player" retains a specific and fixed strategy. This cannot be realistically assumed to hold—let alone guaranteed—in any large, open multiagent environment.

Another field of game theory is *mechanism design*, that is also sometimes called as reverse game theory. Here, not only the player strategies are studied, but the structure of the game itself. More specifically, the rules of the game are set in a way, such that the strategies that players choose to follow, lead to desired outcomes of the game. One central goal in mechanism design is *incentive compatibility* [46]. Namely, mechanisms are designed so that agents are incentivized to truthfully reveal their private preferences or information. In our case, the desired outcome is to shift peak load to non peak intervals, without suffering monetary loss. To guarantee this outcome, agents must *truthfully* interact with the Grid.

The work of Kota, Chalkiadakis, Robu, Rogers and Jennings [26] mentioned earlier, for instance, comes complete with certain incentive compatibility guarantees. However, unlike what we do in this thesis, no guidelines whatsoever as to where to shift consumption to are provided by that model, and deals agreed there involve reduction promises only.

Another field of game theory is *mechanism design*, that is also sometimes called reverse game theory. Here, not only the player strategies are studied, but the structure of the game itself. More specifically, the rules of the game are set in a way, such that the strategies that players choose to follow, lead to desired outcomes of the game. In other words, the objective is to design protocols of strategic agent interactions with certain desirable properties; and these properties are often ensured by the mechanism making the promise of certain *payments* (e.g., in the form of monetary trasnsfers) to the participants. One central goal in mechanism design is *incentive compatibility* [46]. Namely, mechanisms are designed so that agents are incentivized to truthfully reveal their private preferences or information. In our case, the desired outcome is to shift peak load to non peak intervals, without suffering monetary loss. To guarantee this outcome, agents must truthfully interact with the Grid. Another desirable property for mechanisms is that of *budget-balancedness*. Budget balancedness means that the mechanism awards the participants the same amount of money it collects from them [46]; *Weak* budget-balancedness, on the other hand, means that the mechanism might actually be able to make a profit (as rewards collected by the participants exceed the amount of payments made to them). In either case, (weak or strong) budget-balancedness implies that there is no need for external transfers into the mechanism. As we detail later, the payment mechanism we designed for our consumers cooperatives is weakly budget-balanced (as cooperative members pay the cooperative a sum that at least matches the electriciity bill charged to the cooperative for its energy consumption).

Other work in mechanism design includes [51], where authors propose the management of the trading of electricity between homes and micro-grids via a market-based mechanism and trading strategies for the Smart Grid. Robu *et al.* in [44] provide two versions on a truthful allocation mechanism for a dynamic population's hybrid electric vehicle charging.

Recent work in the energy production side has proposed the use of *Continuously Ranked Probability Score (CRPS)* scoring rule for evaluating power production or consumption predictions of agents participating in cooperatives in the Smart Grid [43]. 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. In our approach, instead of production, CRPS is used for consumption shifting effort ranking. While this technique provides the agents with strong incentives to stay truthful (and, indeed, provides theoretical guarantees for statement truthfulness), it does not guarantee agent statement accuracy. For this reason, we employ stochastic filtering techniques to monitor agent trustworthiness, as explained later.

2.3 VIRTUAL POWER PLANTS

A concept related to CDSMs is that of *Virtual Power Plants* [13, 12]. VPPs are amalgamations of electricity producers and consumers acting as "power plants" attempting to counter the effects of peak-time consumption before it becomes hazardous for the stability of the Grid [4]. Initially, the concept of VPP was developed to enhance the visibility and control of distributed energy resources to system operators and other market actors, by providing an appropriate interface between these system components [37]. But since the Smart Grid mainly consists of distributed resources and a large number of consumers that perform according to their own beliefs and convenience, large groupings in VPPs make the analysis and interactions easier. The collective action of VPP members makes it possible for the VPP to participate in "critical peak pricing programs", that is members might be rewarded with better consumption rates for reducing their energy demand over some period [4].

In [25], authors propose a decentralized system architecture as a mean of balancing supply and demand in clusters of distributed energy resources. Intelligent distributed coordination is achieved by organizing the DER-agents into a logical tree, assigning them roles and prescribing strategies to use in their interactions.

The work of [29] proposes the use of multiagent coalition formation strategies and market based techniques for the creation of VPPs. Their work though, does not provide mechanism design guarantees, such as individual rationality or incentive compatibility. These concepts are addressed and guaranteed in our work.

The beneficial nature of cooperative producer VPPs is clearly demonstrated by Chalkiadakis *et al.* in [7], where an extended analysis outlines the benefits arising from distributed energy resources coalescing to profitably sell energy to the Smart Grid. That paper is the first to explore the game theoretic view of Virtual Power Plant formation. They provide a mechanism that can be considered as an alternative to feed-in tariffs and promotes reliability of supply. They model the formation of the virtual power plant as a coalitional game and calculate payments that lie in the set of the core of this game. These cooperative VPPs lead to better utilization of distributed energy resources as they take into account every agent that represents them. Similarly, our work employs a better price, awarded for shifting *sizable* amounts of load, that is agents are better off forming large cooperatives and shifting "less costly" peak load to other non peak intervals.

The *Cooperatives for Demand Side Management* proposed in [26], is another form of Virtual Power Plants. Similarly to [26], our approach advocates the creation of cooperative VPPs at the consumer side to provide demand management services, and proposes a specific method to incentivize the *shifting* —instead of only trimming— of the peak electricity consumption loads to promote the flattening of the global electricity demand curve ahead of time.

2.4 ADAPTIVE SYSTEMS AND STOCHASTIC FILTERING TECHNIQUES

Stochastic Filtering Theory uses probability tools to estimate unobservable stochastic processes [55]. Such methods are employed in this work to enhance the accuracy and effectiveness of our mechanism. This is achieved by monitoring past consumer statements and actions, and based on these, cooperative "central beliefs" about future actions are generated, that do not depend on agent statements, as those might be inaccurate. There exist many approaches on adapting systems and corresponding attempts to tackle system related uncertainty. A detailed review of the types of uncertainty influencing the operation of an adaptive system, complete with techniques for uncertainty representation, can be found in [15]. In our case, we are concerned about the uncertainty inherent in electricity shifting statements and actions.

One concrete example of a self-adaptive system is provided by the "RESIST" framework [10]. RESIST is a *situated software system* that monitors various information sources and reconfigures itself proactively. In that work, a *Hidden Markov Model (HMM)* is trained and solved. Contextual parameters are also used in order to make the system adaptive to environmental changes. HMMs could be applied to our setting. However, we chose not to incorporate them, as we are not trying to estimate each agent's exact state (although that could be done in a higher level analysis), but instead, we wish to predict agent future performance based on past actions. A self-adaptive system that is closer to our requirements is the FeatUre-oriented Self-adaptatION framework [14]. The system operates using interconnected features, whose importance is constantly monitored and reconfigured accordingly.

Now, stochastic filtering techniques, in particular Kalman filters and their variants, have been widely used in a variety of domains; and can of course be incorporated in self-adaptive systems such as the above. In general, Kalman filters are algorithms that make use of historically observed series of noisy measurements and generate more precise estimates about future observations. In particular, they estimate the internal states of *linear dynamic systems*. In cases where the dynamic system is non-linear, extensions of the Kalman filter, like the Extended Kalman filter can be used. The Extended Kalman filter linearizes the non-linear function around the current estimate, using multivariate Taylor series expansions. In what follows, we mention only a couple of examples that are most relevant to our work.

In [16], authors use an *Extended Kalman Filter* whose parameters are given by a Particle Swarm Optimization algorithm, to compute the synaptic weights of a neural network. This neural network is then used to predict wind turbine production; however, EKF is prone to errors due to inaccurate approximations. The approach in [24], on the other hand, employs *Gaussian Process* [42] regression for learning motion and observation models by some training measurements. Resulting GP parameter values are fed in to an *Unscented Kalman Filter*, in order to perform tracking of an autonomous micro-blimp. The approach successfully tackles precision problems that appear from the combination of noisy observations and uncertainty in the model. GPs have also been used recently by [5] to forecast electricity demand, and the predictions are tested in the electricity market simulation of PowerTAC. In our proposed scheme, we devise a similar combination of a GP model with a UKF filter, and employ it as one of our two proposed filtering techniques.

In our work, we enhance our generic framework for demand shifting, via the adoption of two distinct filtering techniques, for predicting the power consumption shifting efforts of participating agents. This can be viewed as a type of system adaptation. When an agent shifts some load, its actions are monitored and a corresponding model is induced. Then, instead of simply taking individual forecasts into account, the learned model is used to better predict agent and cooperative action quality, and improve monetary benefits and general performance.

2.5 DISCUSSION

Admittedly, the penetration of Smart Grid technology to the electricity network has already begun. Provided that the designs and architectures about to be adopted keep up with the philosophies of "clean energy" and "sustainability", the expectation of a "greener" future can be considered realistic. Of course, in order to make this possible, one must persuade end-users to actively participate in the effort. Thus, the mechanisms and infrastructures used must be transparent, reliable and fair, and also preserve user privacy[1].

To this end, in this thesis we explore a *cooperative* approach utilizing *multiagent systems* technologies. As we detail in the upcoming Chapters, we incorporate related work notions, such as cooperatives

1 Actually, there are already disputes between corporations running large renewable sources power plants and occupants of the exploited areas about the actual side effects that industrialized plants create [33].

of electricity consumers and strictly proper scoring rules to make possible to shift peak load demand to non-peak time intervals taking into account personal preferences. In addition, the adoption of stochastic filtering techniques that monitor past agent performance and can estimate actual agent uncertainty regarding demand shifting operations, further guarantee the stability and the effective performance of the proposed scheme, assuming that agents uncertainty follows specific underlying models that are inferred by past observations. We now proceed to present our approach in detail.

Part II

THE PROPOSED MECHANISM

This part presents the overall proposed mechanism. First, we provide the pricing and specific consumption shifting constraints that must hold in order for the problem to be feasible. Also, the "bidding process" is explained with respect to single business units. Since it is hard for single units to be effective in their efforts, the formation of cooperatives of consumers are proposed in Chapter 4. This is more realistic, and constitutes a basic feature of our work. In addition, a gain balancing algorithm that sets variable effective prices to each member of the cooperative, as well as different shifting coalition formation methods are provided. In the final chapter of this part, we discuss the uncertainties that govern agent statements and actual final actions further, and propose two stochastic filtering techniques to monitor individual agent performance and predict such uncertainties. It is the first time that these techniques are applied in this domain.

3

A DEMAND-SIDE MANAGEMENT SCHEME FOR EFFECTIVE ELECTRICITY CONSUMPTION SHIFTING

Grid Side

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. Such prices are often set by individual utility companies, or, in many cases, by a nationwide independent system operator (ISO), managing the electricity grid. We hereby term such an authority as "the Grid" for convenience. In our model, we also assume that there exist exactly two different price levels $p_{high} > p_{low}$. These, however characterize *each specific time interval* t, based on a demand threshold τ under which electricity generation costs are lower :

$$p^t = \begin{cases} p_{high}, & \text{if } Demand^t \geqslant \tau \\ p_{low}, & \text{if } Demand^t < \tau \end{cases} \tag{1}$$

The intervals during which $p^t = p_{high}$ 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 available in case of high uncertainty or an emergency, and minimize the risk that high-cost generators would have to be turned on. That is, the Grid would ideally want to reduce consumption by $Q^{t_h}_{max} \geqslant q^{t_h}_{\tau}$, where:

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

Intuitively, $q^{t_h}_{\tau}$ 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

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

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 $q_i^{t_l}$.

3.1 CONSTRAINTS

For a safe demand reduction to take place, the following constraints must hold: First,

$$\sum_i r_i^{t_h} \geqslant q_\tau^{t_h} \tag{2}$$

that is, the amount of load reduced must be higher than the minimum needed at t_h. Second,

$$\sum_{t_l} q_i^{t_l} \leqslant \sum_{t_h} r_i^{t_h}, \forall i \tag{3}$$

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). Moreover,

$$\sum_i \sum_{t_l} q_i^{t_l} \leqslant Q_{max}^{t_h}, \forall t_h \in T_H \tag{4}$$

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; and finally

$$\sum_i q_i^{t_l} \leqslant q_{sl}^{t_l}, \forall t_l \in T_L \tag{5}$$

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. 1, 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.

3.2 A BETTER, "GROUP" PRICE, FOR CONSUMPTION SHIFTING

Our scheme allows sizable load consumption from peak to non-peak intervals where an *even lower, "group" price* $p_{group} < p_{low}$ is granted,

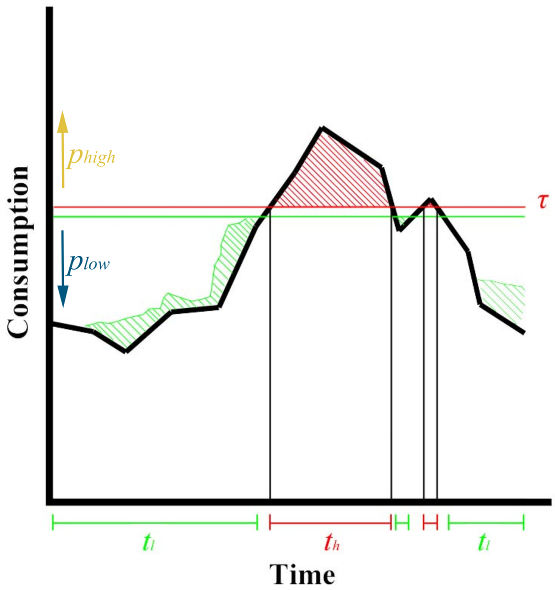

Figure 1: Scheme objective representation. Portions of peak-load, are shifted to lower demand intervals.

and which is a function of the actual load reduction X in a way that for larger load portions, the price becomes better. We term this price as p_{group} because such reduction will likely be possible only by groups of agents. This price is awarded if *the actual quantity of the load shifted* from t_h exceeds some minimum value $q^{t_h}_{min}$, set by the Grid given its knowledge of $q^{t_h}_{\tau}$ (e.g., it could be $q^{t_h}_{min} = q^{t_h}_{\tau}$). Summarizing, p_{group} is a function of the actual load reduction X in a way that for larger load portions, the price becomes better. Thus, p_{group} is later noted as $p_{group}(X)$. We will come back to this issue in Chapter 4.

3.3 AN EFFICIENT CONSUMPTION SHIFTING SCHEME

An agent i that wishes to participate in the consumption shifting scheme, is characterized by:

(a) Its stated reduction capacity $\hat{r}^t_i$, namely the amount of load that it is willing to curtail (e.g., by shifting) at a time interval t,

(b) Its shifting costs $c_i^{t_h \to t_l}$, that is the cost that occurs if consumption of a unit of energy is shifted from (a peak interval) t_h to (a non-peak) t_l.

Given the above, the exact shifting protocol we propose is as follows.

Bidding Process

Every day, the Grid announces the forecasted for the day-ahead peak intervals T_H and the most preferable non-peak intervals T_L; and also announces the (quantity-depended) price rates it awards for consumption in the T_L intervals. It then waits for shifting proposals by business units. Each business unit (consisting of a single consumer or more), can interact with the Grid and state its overall load reduction capability during announced T_H intervals, and a number of intervals $t_l \in T_L$, to which it is willing to shift consumption to. This procedure is called *bidding* and, once again it refers to shifting actions pledged for the day-ahead. Note that for each $t \in \{T_H \cup T_L\}$ there can be more than one bidders. Furthermore, bidders can pledge to shift some load from one high consumption t_h interval to several low consumption t_l ones. Moreover, the protocol requires that each bidder reports its confidence (or, degree of uncertainty) regarding its ability to shift $\hat{r}_i^t$ from interval t, in the form of a normal distribution describing its expected *relative error* regarding its reduction forecast. The latter is required to promote efficiency, as we explain below. Algorithm 1 summarizes the bidding process.

Algorithm 1 Bidding Process

Grid generates sets T_H and T_L
Grid announces $p_{group}(X)$ for a reduction of load quantity X in any $t \in T_H$
Business units state:
  {sets of $t_h \in T_H$ and $t_l \in T_l$ intervals of interest;
  shifting (load reduction) capabilities;
  corresponding shifting uncertainty for each t_h;
  and the amount of load they will be moving to each t_l}
Grid checks for constraint violations
if Constraints are met **then**
  The Grid announces the acceptance of the bid to the bidder
  Business units reduce and shift
  Grid awards resulting actions, given their efficiency after the application of the CRPS rule, and announced $p_{group}(X)$
else
  Deny Bid
end if

3.4 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)* [18], which has also been recently used in [43] to incentivize renewable energy-dependent electricity *producers* to accurately state their estimated output when participating in a cooperative. 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 the actual occurrence itself. The rule is strictly proper if it incentivises forecasters to state their true beliefs P *only*, and it does so by maximizing expected reward only when $\hat{P} = P$. Use of CRPS allows us to directly evaluate probabilistic forecasts, and the score is given by:

$$CRPS(\mathcal{N}(\mu, \sigma^2), x) = \sigma[\frac{1}{\sqrt{\pi}} - 2\phi\left(\frac{x-\mu}{\sigma}\right) - \frac{x-\mu}{\sigma}\left(2\Phi\left(\frac{x-\mu}{\sigma}\right) - 1\right)] \tag{6}$$

In our setting, $\mathcal{N}(\mu, \sigma^2)$ is the uncertainty stated over the *expected relative errors*[2] 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 is considered to be zero mean, $\mathcal{N}(0, \hat{\sigma}^2)$, so from now on will be simply noted as $\hat{\sigma}$. Given this notation, an agent i whose bid to shift some load from t_h to t_l is accepted, receives a reduced electricity bill B_i given its actual contribution $r_i^{t_h}$ and its final consumption at t_l, $q_i^{t_l}$ as follows:

$$B_i = (1 + CRPS_i) q_i^{t_l} p_{group}(r_i^{t_h}) \tag{7}$$

Note that $r_i^{t_h}$ is the actual amount shifted from t_h to t_l, which determines the electricity price p_{group} for i at t_l; while $q_i^{t_l}$ is the quantity shifted from t_h to t_l and consumed there. Note that it can be $q_i^{t_l} < r_i^{t_h}$ since an agent can shift $r_i^{t_h}$ to multiple t_ls.

3.5 AGENT INCENTIVES AND DECISION ANALYSIS

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 is associated with a cost $c_i^{t_h \to t_l}$ for the agent. The *gain* that an agent would have for shifting r_i to t_l given t_l's lower price p_{low}, would be equal to

$$gain(i|p_{low}) = r_i(p_{high} - p_{low} - c_i^{t_h \to t_l}) \tag{8}$$

2 The 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.

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,

$$p_{low} + c^{t_h \rightarrow t_l} > p_{high} \quad (9)$$

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_{group} for consumption of $r_i \geqslant q_{\text{min}}$ at t_l s.t.

$$p_{group} + c^{t_h \rightarrow t_l} \leqslant p_{high} \quad (10)$$

then the agent will be incentivized to perform the shift, as his perceived $\mathit{gain}(i|p_{group})$ would now be non-negative for him. This is what is offered by the scheme presented in Section 3.3. Moreover, the use of the CRPS rule guarantees that agents have an incentive to make accurate predictions.

3.6 SCHEME SUMMARY

Summarizing, the Grid announces, for the day ahead, the peak intervals t_h that need consumption reduction by an amount of load X, and the non-peak intervals t_l to which shifting is acceptable. The Grid determines and announces a better price rate $p_{group}(X)$ to offer for the consumption of load X at (any) t_l instead of t_h. This price is awarded if *the quantity of the load shifted* from t_h exceeds some minimum value $q^{t_h}_{\text{min}}$, which is estimated and set by the Grid given its knowledge of the $q^{t_h}_{\tau}$ introduced earlier. In general, we expect $q^{t_h}_{\text{min}}$ to be close to $q^{t_h}_{\tau}$, as the Grid would not be interested in rewarding very small-scale reduction. For simplicity, in our model we assume *a unique* q_{min} *value* used across all time intervals. Consumers then make their bid collectively or alone, and state their reduction capacity $\hat{r}^{t_h}_i$, corresponding uncertainty, and the intervals that they are willing to shift to, along with the shifting costs $c^{t_h \rightarrow t_l}_i$.

We now proceed to explain what happens when such an agent is actually a member of a reducing agents team.

4 AGENT COOPERATIVES FOR EFFECTIVE POWER CONSUMPTION SHIFTING

Chapter 3 presented a shifting protocol which incentivises agents to shift some of their consumption from peak intervals to other non-peak ones, and also to be truthful at stating their shifting capabilities. In the general case, it is very rare even for large industrial consumers to have reduction capacity larger or equal to q_{min}. 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. 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 to a cooperative manager agent, along with its reduction capabilities, its confidence $\mathcal{N}(\mu_i, \hat{\sigma}_i^2)$ on actually reducing that stated amount at t_h, and the set of t_l intervals that it 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 of some of them might be so high that do 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. Any such shifting bid is composed by four parts:

i. t_h, the high cost interval to reduce consumption from.

ii. $\hat{r}_C$, the amount C pledges to reduce at t_h.

iii. The set of low cost intervals t_l to move consumption to, along with the set of corresponding quantities that will be moved to each t_l and the shifting costs for those intervals.

iv. An estimate of its $\mathcal{N}(\mu_C, \hat{\sigma}^2_{C,t_h})$ joint relative error on predicted $\hat{r}_C$.

The bid is determined so that the collective expected gain from the shifting operation is non-negative. We provide the details of how this is ensured below. Assuming that C was selected and reduced by $r_C^{t_h}$ at a peak interval t_h, the bill B_C charged to the cooperative for consuming $q_C^{t_l}$ at a non-peak interval t_l is given by Eq. 7 (substituting C for i):

$$B_C = (1 + CRPS_C)\, q_C^{t_l} p_{group}(r_C^{t_h}) \tag{11}$$

Gain Transfers

Now, even if the *collective* expected gain from the bid is positive, it is not certain that all individuals in C have a positive expected gain as well. Nevertheless, with positive collective expected gain, the possibility of internal *gain transfers* is raised, allowing non-negative (expected) gain for all participants. These transfers have to be performed in such a way so that the budget-balancedness of any cooperative bid is ensured, at least in the weak sense. We now describe exactly how the cooperative determines its bid at t_h.

4.1 COOPERATIVE BIDDING PROCESS

Since the Grid-awarded group rate depends on quantity reduced, we (originally) assume that the cooperative attempts to select a subset C with maximal reduction capacity (we modify this assumption in subsequent algorithm variants). We now present an algorithm that achieves this, while ensuring that C and each one of its members has a non-negative gain, and that budget-balancedness is ensured. In what follows, we drop time indices where these are clearly implied.

Reservation Price

To begin, let:

$$\hat{p}_i = (p_{high} - c_i^{t_h \to t_l}) \tag{12}$$

be agent i's (implicitly stated) *reservation price*, that is, the highest price that i is willing to pay for moving from t_h to t_l (in order to not suffer a loss). The algorithm then proceeds as follows.

First, for every i, we check whether $\hat{p}_i \leqslant 0$. 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). 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.

In this stage, an estimate regarding actual user final reduction, $\tilde{r}_i$, needs to be calculated. Details about the different estimating techniques of $\tilde{r}_i$ are provided in Section 5.1. As a general approach, the algorithm can set:

$$\tilde{r}_i = \hat{r}_i - \hat{\sigma}_i \hat{r}_i \tag{13}$$

for all agents in $\mathbb{C}^{t_h}$, that is, the cooperative makes a *conservative* estimate of an agent's expected performance, given its stated uncertainty. Generally, the $\hat{\alpha} = 1 - \hat{\sigma}$ estimate about future reduction in the general case is different than the actual $\alpha = 1 - \sigma$ of the actual final reduction action. This factor can be considered as a random variable that follows an unknown distribution, thus modeling the uncertainty regarding the statements and final actions. We elaborate on two methods that can help us obtain a more trusted index about α later, in Chapter 5.

Contribution Potential

Now, we can define the *agent contribution potential*, that is the product of agents actual reduction estimate $\tilde{r}_i$ and reservation price $\hat{p}_i$,

$\tilde{r}_i\hat{p}_i$. The algorithm then ranks the agents by $\tilde{r}_i\hat{p}_i$ in decreasing order. Then, starting from the agent with the highest contribution potential 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 both of the following conditions are met for the *maximum possible* group of agents C:

$$\sum_{i \in C} \tilde{r}_i \hat{p}_i \geqslant \tilde{r}_C p_C \tag{4.1.i}$$

$$\tilde{r}_C \geqslant q_{min} \tag{4.1.ii}$$

where q_{min} is the minimum quantity admitting a "group price", and

$$p_C = p_{group}(\tilde{r}_C)$$

is the price rate offered by the Grid for reduction $\tilde{r}_C$, with $\tilde{r}_C = \sum_{i \in C} \tilde{r}_i$.

Shifting Gain

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

$$gain(j|p_C) = \tilde{r}_j(\hat{p}_j - p_C) \tag{14}$$

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 and individual effective price of its members have to be dynamically determined by the cooperative, so that individual rationality is ensured.

Now, if *all* agents in $\mathbb{C}^{th}$ are inserted in C and $\tilde{r}_C$ is still lower than q_{min}, the problem is infeasible and we stop. Likewise, if all agents are in C and

$$\sum_{i \in C} \tilde{r}_i \hat{p}_i - \tilde{r}_C p_C < 0 \tag{15}$$

holds, 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[1]. This means that there is *at least* one agent j in C with positive expected gain, given p_C. That is,

$$gain(j|p_C) = \tilde{r}_j(p_{high} - c_j - p_C) = \tilde{r}_j\hat{p}_j - \tilde{r}_j p_C > 0 \tag{16}$$

1 That is, after a subset C has met $\tilde{r}_C \geqslant q_{min}$, we kept adding agents to C until by adding some k we constructed a C' for which $\sum_{i \in C'} \tilde{r}_i \hat{p}_i - \tilde{r}'_C p'_C < 0$, in which case k is removed from C'.

if not, then no agent has a positive $gain$, and thus

$$\sum_{i \in C} \tilde{r}_i \hat{p}_i - \tilde{r}_i p_C \leqslant 0 \tag{17}$$

leading to

$$\sum_{i \in C} \tilde{r}_i \hat{p}_i \leqslant \tilde{r}_C p_C \tag{18}$$

contradicting condition (4.1.i) above. Having met the two conditions for (maximal) C, also means that agents in C are *collectively* willing to pay a total amount for moving their r_i consumptions to t_l, which is greater than what their group will be asked to pay for, given the offer p_C for $\tilde{r}_C$.

Thus we have ended up with the maximal C so that (4.1.i) and (4.1.ii) hold, and which contains some agents with positive and some with negative gain given p_C, and which we can now use to implement a gain transfer scheme so that all individual agents in C end up with non-negative gain themselves.

4.2 SETTING VARIABLE EFFECTIVE PRICES

Effective Price

At this point the cooperative pre-assigns 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 C will receive an actual price rate that will be dependent on its CRPS score.

Thus, the cooperative initially sets $p_i^{\mathit{eff}} = p_C, \forall i \in C$, given the price p_C expected, and proceeds to rank in decreasing order agents in C according to their *expected gain*, that is:

$$gain(i | p_i^{\mathit{eff}} = p_C) = \tilde{r}_i(\hat{p}_i - p_C) \tag{19}$$

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. If negativities exist, 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. Then, 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 $g_j \geqslant 0$). The value of p_i^{eff} is calculated by:

$$p_i^{\mathit{eff}} = \frac{\tilde{r}_i \hat{p}_i - g_j}{\tilde{r}_i} \tag{20}$$

with g_j being the target gain, i.e. the gain of the agent below.

Then we do the same for the second top agent, until its gain reaches that of the third. We continue this way until all requested gain is transferred, or one's gain reaches zero. If the latter happens, we move to the top again and repeat. The procedure is described by Algorithm 2.

Algorithm 2 Variable Effective Prices

```
Input: p_group(r_C), r̃_i, c_i ∀agent_i ∈ C
Output: p_i^eff ∀agent_i ∈ C
  Compute reservation prices p̂_i and gain g_i = gain_i, ∀i ∈ C
  Sort agents by gain in decreasing order
  if Negative expected gain agents exist (i.e., the set G⁻ is non empty) then
     Count total negative gain L = ∑ g_j, ∀j ∈ G⁻
     Assign agents ∀j ∈ G⁻, p_j^eff = p̂_j
     donation := 0
     while donation < L do
        for all Positive gain agents do
           if donation < L then
              if agent_i is the last positive gain agent in sorted gain list then
                 if donation + gain_i ⩽ L then
                    donation = donation + gain_i
                    p_i^eff = p̂_i
                 else
                    donation = L
                    Assign p_i^eff s.t. only the remaining gain needed is transfered
                 end if
              else
                 if donation + (gain_i − gain_{i+1}) ⩽ L then
                    donation = donation + (gain_i − gain_{i+1})
                    Assign p_i^eff s.t. the amount of i's gain is equal to that of i + 1's
                 else
                    donation = L
                    Assign p_i^eff s.t. only the remaining gain needed is transfered
                 end if
              end if
           else
              Assign p_i^eff = p_group(r_C)
           end if
        end for
     end while
  else
     Assign p_i^eff = p_group(r_C), ∀i ∈ C
  end if
  return p_i^eff, ∀i ∈ C
```

Member Bill

The p_i^{eff} prices thus determined represent internally pre-agreed prices 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:

$$b_i = \frac{(1 + CRPS_i)p_i^{eff} q_i}{(\sum_{j \in C \setminus \{i\}}(1 + \mathrm{CRPS}_j)p_j^{eff} q_j) + p_i^{eff} q_i} B_C \tag{21}$$

Strict propriety is ensured by this rule, as it is an affine transformation of a member's $CRPS_i$ score; and 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*.

4.3 ALTERNATIVE COALITION FORMATION TECHNIQUES

The scheme presented aims to achieve the lowest possible group price, through the addition of as many agents as possible into the reducing set of agents in any given t_h, as long as budget balancedness and individual rationality are respected. Budget balancedness means that all the requested utility is provided by the participating agents themselves. Individual rationality dictates that every single agent suffers no loss from participation in the cooperative shifting process. Now, though this goal is clearly efficient for the Grid (since it apparently promotes the maximum possible reduction at any t_h), it is not necessarily efficient for the reducing coalition at t_h. That is, it does not necessarily maximize the sum of the members expected *gains*: since agents with potentially high costs keep being added until it is possible for the coalition to sustain them through "gain transfers", there might exist different reducer sets with higher overall gain.

The bid determination mechanism proposed in Section 4.1 above can thus be summarized as *Method 1: Rank agents by potential and maximize expected capacity*. We now proceed to provide a more detailed example of its use and then present four variants of that approach. We note that all variants include the "internal" reward transfer phase described in Section 4.2.

4.3.1 *Method 1: Rank by potential, maximize capacity*

Consider an agent i characterized by his shifting capabilities at t_h, $\hat{r}_i^{t_h}$ and $\hat{\sigma}_i$, and his reservation price at t_l, $\hat{p}_i^{t_l}$. These include information for shifting costs.[2] Then, a measure of each agents contribution potential is the product of the estimated reduction capacity and reservation price: $\tilde{r}_i^{t_h} \cdot \hat{p}_i^{t_l}$ We can rank all available agents at t_h by this measure, in decreasing order.

A shifting coalition is valid if its overall capacity, noted by $\tilde{r}_C$ is between certain limits that the Grid designates for each day or even interval, that is $q_{min}^{t_h} \leqslant \tilde{r}_C \leqslant Q_{max}^{t_h}$. In this method we maximize the coalition capacity and we add agents in the coalition until $Q_{max}^{t_h}$ is exceeded provided that the problem is feasible.

2 Recall that $\hat{p}_i^{t_l} = p^{t_h} - cost_i^{t_h \to t_l}$, where p^{t_h} is the price originally charged by the utility company at the interval t_h. Agents with low costs are comfortable by paying higher prices (still lower than p^{t_h}), whereas high cost agents demand very low reservation prices.

To give an example, consider 8 agents that are available for consumption shifting at a specific time interval, with their features as shown in Table 1. *Example 1*

Table 1: Example 1 participants.

ID	$\tilde{r}_i$	$\hat{p}_i$	$\tilde{r}_i \cdot \hat{p}_i$
a	50	0.03	1.5
b	200	0.04	8
c	30	0.038	1.14
d	55	0.05	2.75
e	6	0.067	0.402
f	75	0.032	2.4
g	60	0.001	0.06
h	4	-1.5	-6

Suppose that in the specific interval, $q_{min} = 380$ and $Q_{max} = 480$. We sort agents by their $\tilde{r}_i \cdot \hat{p}_i$ value and we add them into the coalition until q_{min} is reached. Then, we continue adding until either Q_{max} is exceeded, or until $\sum \tilde{r}_i \cdot \hat{p}_i \geqslant \tilde{r}_C \cdot p_{group}(\tilde{r}_C)$ does not hold.

Table 2: Example 1 coalition formation.

ID	$\tilde{r}_i$	$\hat{p}_i$	$\tilde{r}_i \cdot \hat{p}_i$	$\sum \tilde{r}_i \cdot \hat{p}_i$	p_C	$\tilde{r}_C$	$\tilde{r}_C \cdot p_g(\tilde{r}_C)$
b	200	0.04	*8*	8	inf	200	-
d	55	0.05	*2.75*	10.75	inf	255	-
f	75	0.032	*2.4*	13.15	inf	330	-
a	50	0.03	*1.5*	14.65	0.05625	380	**21.375**
c	30	0.038	*1.14*	15.79	0.045795	410	**18.77595**
e	6	0.067	*0.402*	16.192	0.043704	416	**18.180864**
g	60	0.001	*0.06*	**16.252**	0.022794	476	10.849944
h	4	-1.5	-6	10.252	0.0214	480	**10.272**

The process of forming the reducing coalition is summarized in Table 2. Here we can see that until the q_{min} amount is reached, a better price can not be awarded, thus corresponding values are set to inf. Now, note that although adding agent a makes $\tilde{r}_C = q_{min}$, the constraint $\sum \tilde{r}_i \cdot \hat{p}_i \geqslant \tilde{r}_C \cdot p_{group}(\tilde{r}_C)$ is not satisfied and thus, we continue adding agents until agent g is included and the constraint is satisfied. Since we seek to maximize capacity, we try to include agent h also, in order to achieve $\tilde{r}_C = Q_{max}$. Because agent h probably has high shifting costs, it requires to be paid by the cooperative, in order to contribute without suffering a loss. On the other hand, the cooperative is in position to only pay up to the difference $\sum \tilde{r}_i \cdot \hat{p}_i - \tilde{r}_C \cdot p_{group}(\tilde{r}_C)$ and suffer no loss, but this quantity is less than agent h asks, and thus $\sum \tilde{r}_i \cdot \hat{p}_i < \tilde{r}_C \cdot p_{group}(\tilde{r}_C)$; namely the constraint is violated and so, agent h is excluded from the team.

4.3.2 *Method 2: Rank by potential, minimize capacity*

This method is the same as the aforementioned one, with the difference that we stop adding agents in C the moment when the $q_{\min}$ requirement is met. That is, it forms reducing teams, such that $\tilde{r}_C \simeq q_{\min}$, and thus, the team is rewarded the highest p_{group} possible.

4.3.3 *Method 3: Rank by potential, maximize capacity, exclude agents with negative expected gain*

This method is the same as Method 1 (4.3.1), but once $q_{\min}$ is met, an agent in the ranked list is added in the coalition only if its expected gain is non-negative with respect to p_{group} at the moment of its entry. More specifically, we retain contribution potential ranking and come up to the maximal coalition, with respect to reduction capacity, as in Method 1. The difference is that agents are added in the coalition only if they are granted non negative expected gain. When an agent enters the coalition, it's expected gain is checked with respect to the better price granted to the coalition at that particular moment. If the gain is negative then the agent is excluded and the search for contributors continues in the ranked list to the agent below. Here we must note that it is possible for an agent that could be favored by the final price $p_{group}(\tilde{r}_C)$ to be excluded from the coalition, because his contribution potential was checked earlier in the process, when $\tilde{r}_C < Q_{max}$ and for $p_{group}(\tilde{r}_C) > p_{group}(Q_{max})$ he would have suffered gain loss and consequently was rejected.

Example 2 Consider an example where agent i is characterized by $\tilde{r}_i = 500$, $\hat{p}_i = 0.04$ and $\tilde{r}_i \cdot \hat{p}_i = 20$ and an agent $i+1$ with $\tilde{r}_{i+1} = 50$, $\hat{p}_{i+1} = 0.38$ and $\tilde{r}_{i+1} \cdot \hat{p}_{i+1} = 19$. Suppose that, without loss of generality, q_{min} is just reached by adding agent $i-1$ in the coalition and we proceed to check i. The expected gain of i is: $gain(i|p_{group}(\tilde{r}_C)) = \tilde{r}_i(\hat{p}_i - p_{group}(q_{min})) = 500(0.04 - 0.5625) = -8.125$ that is negative and therefore agent i is excluded from the coalition. If we proceed to $i+1$, by following the same process we observe that $gain(i+1|p_{group}(\tilde{r}_C)) = 16.1875$, positive, thus agent $i+1$ is selected. Later on, when Q_{max} is reached, the better price offered is much lower, $p_{group}(q_{max}) = 0.0214$. If we check agent i's expected gain at this moment we observe that $gain(i|p_{group}(\tilde{r}_C)) = 9.3$ that is positive. Furthermore, i ends up being more valuable contributor than $i+1$, as $\tilde{r}_i \cdot p_i^{eff} = 10.7$ and $\tilde{r}_{i+1} \cdot p_{i+1}^{eff} = 1.07$. Thus, this method does not select agents based on their real contribution impact.

4.3.4 *Method 4: Rank by expected gain, maximize capacity*

This method ranks prospective contributors by their expected gain wrt. p_{group} offered at the moment they are checked for entering C. However, whenever an agent enters C, p_{group} changes, and so does the expected gain of every C member. Thus, it has to be recalculated with every new entry, which increases the complexity. In this method instead of ranking probable contributors by their $\tilde{r}_i \cdot \hat{p}_i$ we rank them by their expected gain with respect to the better price offered at the moment they are checked for entering the coalition.

The issue that rises is that for every new agent entering the coalition, the price offered changes, the expected gain for every agent changes too and thus it has to be recalculated with every new entry. If there are n possible contributors and k of them are included to form the minimal coalition with $\tilde{r}_C = q_{min}$, then the added complexity is $\sum_{i=1}^{n-k} i = \frac{(n-k)(n-k+1)}{2}$ more calculations, order of $\mathcal{O}((n-k)^2)$.

4.3.5 *Method 5: Rank by potential, maximize growth rate of expected gain.*

Provided again ranking by $\tilde{r}_i \cdot \hat{p}_i$, an additional concept is to monitor the *average growth rate of coalescing agents' gain* and stop adding when it drops under a certain level. This method is not included in our simulations, for reasons we will be explaining later, in Section 6.2.

4.4 ALGORITHM PROPERTIES

The reward transfer scheme and the overall cooperative bid determination algorithm presented above have several desirable properties. It is worth noting that these properties hold for all team formation methods presented in 4.3. First, *Individual rationality* is ascertained for all agents in C, as they all have non-negative expected gain from participation. In addition, the member final effective prices p_i^{eff} are set in such way so that overall group value is shared among agents (with some paying more than group price p_C and some less), resulting to $B_C = \sum_i p_i^{Eff} q_i$ (q_i being the members' actual contribution and p_i^{Eff} the final effective price i has to pay, after applying CRPS and normalizing). Considering the total cooperative consumption at non-peak intervals q_C, that will be,

Individual rationality

Budget Balancedness

$$p_i^{Eff} = \frac{(1 + CRPS_i)p_i^{eff}}{(\sum_{j \in C \setminus \{i\}}(1 + CRPS_j)p_j^{eff} q_j) + p_i^{eff} q_i}(1 + CRPS)q_C p_C \quad (22)$$

This means that the (weak) budget-balancedness of the mechanism is ensured.

Truthfulness

Moreover, the transfer scheme presented is *truthful*. Of course, since the agents operate in a large, open environment, one cannot determine an incentive compatible mechanism in the Bayes-Nash sense, since analysing Bayes-Nash equilibria properties is computationally infeasible in this setting. Indeed, it is next to impossible for a member agent to reason on the capabilities or availability of other agents, of which he is unaware of, and no common prior determining such properties can be reasonably assumed. Given this uncertainty, however, the best that an agent can do is *to be truthful* regarding its shifting costs and capacity (and corresponding confidence in reduction). If the agent states inflated shifting costs, it runs the danger of not being selected for C while otherwise it would be. Similarly, if the agent states shifting costs lower than its real ones, then the agent risks being among the ones to suffer the highest reduction in expected gain. This is because the lower these costs are, the higher effective price the agent will be asked to pay via the transfer scheme described above. Finally, it is clearly to the interest of the agent to be as accurate as possible regarding its shifting capacity and corresponding uncertainty, since otherwise it will suffer a gain loss due to a bad CRPS score. Thus, the transfer scheme and overall algorithm are incentive compatible for members, given the dynamic and open nature of this large multiagent environment.

Low Computational Complexity

Last but not least, the computational cost of the bid determination process (including the process of setting effective prices) is quite reasonable. Specifically, it is proportional to the cost of sorting at most $|\mathbb{C}^{t_h}|$ agents in every t_h of interest twice (once when they are ranked according to $\hat{r}_i\hat{p}_i$, and once when they are ranked according to perceived gain). Of course, if the number of members interested in reducing at some t_h is in the dozens of thousands, this can become problematic. However, the cooperative *(a)* does have time in its disposal to execute the algorithm, as bids are determined a day in advance, and *(b)* it can impose constraints to curtail the number of members considered, if this is necessitated by the cost. Summarizing, applying a fast sorting algorithm results to a complexity of $\mathcal{O}(n \log n)$ for methods 1, 2 and 3. Methods 4 and 5 require the additional computational cost of the expected gain calculation, that is $\mathcal{O}((n-k)^2)$.

5 STOCHASTIC FILTERING METHODS TO MONITOR AGENT BEHAVIOR

As discussed previously in this thesis, it is very important for the cooperative to be accurate in its predictions. For cooperative actions to be profitable and effective, the predictions must be highly accurate, that is a challenging task given private agent preferences and underlying uncertainty. In Chapters 3 and 4 we provided a protocol that incentivizes agent accuracy in the face of uncertainty and private preferences. A complementary approach would be to incorporate "reputation-based" models of trust used in MAS settings, see e.g. [40]. However, trust among multiple agents in such a large and dynamic setting is not easy to maintain. Moreover, even if all agents are indeed truthful, exogenous factors, like equipment malfunctions can still put the stability of the cooperative in danger. For these reasons, it is of utmost importance that agent actions are somehow monitored and beliefs about individual agent performance are maintained. To this purpose, we will be incorporating into our scheme two stochastic filtering techniques that have already been successfully applied in other domains.

5.1 MONITORING AGENT TRUSTWORTHINESS

As established in previous Chapters, agent statements greatly affect cooperative decisions, and, if inaccurate, endanger the scheme's stability and effectiveness. Specifically, the cooperative prediction method presented in Section 4.1, assumes that agent i's estimated reduction $\tilde{r}_i$ will be given as follows, after factorization of Equation (13):

$$\tilde{r}_i = (1 - \hat{\sigma}_i)\hat{r}_i \tag{23}$$

where $\hat{r}_i$ is the stated reduction capacity of agent i and $\hat{\sigma}_i$ its stated uncertainty. That is, the cooperative (e.g. cooperative manager agent) bases its decisions as to whom to include in shifting coalitions on the estimated $\hat{r}_i$ quantity, which is entirely based on information provided by the agent i itself. However, we can do better than that. Instead of using $(1 - \hat{\sigma}_i)\hat{r}_i$, the cooperative could adopt a *trusted index* r^*_{i,t_h}—one not stated explicitly by i, but which nevertheless reveals the distribution best describing future agent actions. This index can then be used as $\tilde{r}_i^{t_h}$ to calculate a more accurate *contribution potential* for i. We now describe a process by which to acquire this r^*_{i,t_h} index.

To begin, observe that the actual *relative error* of agent i's performance is:

$$e = \frac{\hat{r}_i^{t_h} - r_i^{t_h}}{\hat{r}_i^{t_h}} \tag{24}$$

where $r_i^{t_h}$ is the actual amount of load that will be reduced, and which can be, in general, assumed to be provided by a transformation of the stated $\hat{r}_i^{t_h}$ amount:

$$r_i^{t_h} = \alpha_i \cdot \hat{r}_i^{t_h} \tag{25}$$

with α_i corresponding to a random variable following some unknown probability distribution. Had we known this α_i quantity, we would be able to calculate $r_i^{t_h}$ and use that instead of $\tilde{r}_i^{t_h}$. The α_i quantity is not known, however. Given this, an objective of our work is to build models for agent performances by approximating the distributions that α_is follow. We can then sample such a distribution to obtain $\tilde{\alpha}_i$, an α_i estimate. Next, we can then use this estimate to obtain our trusted index r^*_{i,t_h} to replace $\tilde{r}_i^{t_h}$:

$$r^*_{i,t_h} = \tilde{\alpha}_i \cdot \hat{r}_i^{t_h} \tag{26}$$

As a result, more accurate predictions about individual agent and cooperative shifting abilities can be obtained.

We can assume that given all underlying uncertainty, an individual agent's final behavior most likely corresponds to a complex, non-linear function of its past behavior. As explained in Section 2.4, one can use stochastic filtering in order to attempt to approximate such a function. We chose to test two filtering approaches that are expected to fit such a function well:

(a) A non-linear *Kalman filter* approach, the *Unscented Kalman Filter (UKF)*, combined with *Gaussian Processes (GP)*

(b) The *Histogram Filter (HF)*, a non-parametric filtering technique.

Both methods require an adequate amount of historical data collected, in order to form an elementary model that can be sampled. For this reason, the conservative ranking method can be employed initially, and get replaced by one of the proposed methods once the required data is available. We now proceed to describe these methods and their application to our setting in detail.

5.2 THE HISTOGRAM FILTER

The first stochastic filtering method we examine is the Histogram Filter (HF). Histogram filters decompose a continuous state space to a finite set of areas or bins:

$$\mathrm{dom}(X) = \mathbf{x_1} \cup \mathbf{x_2} \cup .. \cup \mathbf{x_K}$$

These bins are a partition of the initial space

$$\forall j \neq k, \mathbf{x_j} \cap \mathbf{x_k} = \emptyset$$

and

$$\bigcup_k \mathbf{x_k} = \mathrm{dom}(X)$$

The HF uses a histogram to map a probability p_k to each of the bins $\mathbf{x_k}$. The value of each p_k depends on the frequency of the observations in the range of bin k.

With this approach, agent forecasts $\hat{\sigma}$ 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 it's range is selected. The final estimate $\tilde{\alpha}_i$ is another sample from a uniform distribution normalized to have range equal to that of the bin obtained.

5.3 THE UNSCENTED KALMAN FILTER

Past work has shown that the classic *KF* algorithm is limited to systems with linear transition and observation models; while the *Extended Kalman Filter (EKF)* can handle non-linearities, but not in an optimal manner [49]. The *UKF*, first introduced in [22], uses the so-called *unscented transform* to obtain a better estimate than the *EKF* when dealing with highly non-linear models [53]—such as those describing electricity consumption shifting capabilities.

In order to apply Kalman Filtering we need to define our model in an appropriate form. Here, forecasts[1] of $\hat{\sigma}$ are considered to be depended on past forecasts and linked to actual values of α governing actual agent behavior as follows:

$$\hat{\sigma}_\tau = A\hat{\sigma}_{\tau-1} + w_\tau \tag{27a}$$

$$\alpha_\tau = H(\hat{\sigma}_\tau) + u_\tau \tag{27b}$$

that is, the forecast $\hat{\sigma}_\tau$ depends on its past value transformed by A, a state transition matrix, plus noise w_τ. The real α_τ is the sum of a function of the forecast $H(\hat{\sigma}_\tau)$ and some noise u_τ. We need no initial assumption about the H function as it is approximated by the GP and can be of any form. A brief presentation of the Unscented Kalman Filter follows.

The Unscented Transform

Let $x \in \mathbb{R}^L$ be a Gaussian random variable with mean $\bar{\mathbf{x}}$ and co-

1 To ease notation, we henceforth drop the agent index i from σ and α, when this is implied.

variance $\mathbf{P_x}$ that is propagated through a nonlinear function $\mathbf{y} = g(\mathbf{x})$. A matrix $\mathcal{X}$ can be constructed that contains $2L + 1$ *sigma vectors* $\mathcal{X}_j$ and their corresponding weights W_j, via the *unscented transform* procedure:

$$\begin{aligned}
\mathcal{X}_0 &= \bar{\mathbf{x}} \\
\mathcal{X}_j &= \bar{\mathbf{x}} + (\sqrt{(L+\lambda)\mathbf{P_x}})_j, \qquad j = 1, ..., L \\
\mathcal{X}_j &= \bar{\mathbf{x}} - (\sqrt{(L+\lambda)\mathbf{P_x}})_{j-L}, \quad j = L+1, ..., 2L
\end{aligned} \tag{28}$$

$$\begin{aligned}
W_0^{(m)} &= \frac{\lambda}{L+\lambda} \\
W_0^{(c)} &= \frac{\lambda}{L+\lambda} + (1 - \eta^2 + \xi) \\
W_j^{(m)} &= W_j^{(c)} = \frac{1}{\{2(L+\lambda)\}}, \quad j = 1, ..., 2L
\end{aligned} \tag{29}$$

with $\lambda = \eta^2(L+\kappa) - L$ being the scaling parameter. The η and κ parameters can be tuned to change the spread of the *sigma vectors* around the random variable, ξ incorporates prior knowledge of the distribution of $\mathbf{x}$, and $(\sqrt{(L+\lambda)\mathbf{P_x}})_j$ is the jth row of the matrix square root. Then, the *sigma vectors* are propagated via the nonlinear function:

$$\mathcal{Y}_j = g(\mathcal{X}_j), \ j = 0, ..., 2L \tag{30}$$

Finally, the mean and covariance of $\mathbf{y}$ are approximated by a weighted sample mean and covariance of the posterior *sigma vectors*:

$$\bar{\mathbf{y}} \approx \sum_{j=0}^{2L} W_j^{(m)} \mathcal{Y}_j \tag{31}$$

$$\mathbf{P_y} \approx \sum_{j=0}^{2L} W_j^{(c)} \{\mathcal{Y}_j - \bar{\mathbf{y}}\}\{\mathcal{Y}_j - \bar{\mathbf{y}}\}^T \tag{32}$$

UKF equations as presented in [53] follow:

$$\begin{aligned}
\hat{x}_0 &= E[x_0] \\
\mathbf{P_0} &= E[(x_0 - \hat{x}_0)(x_0 - \hat{x}_0)^T] \\
\hat{x}_0^\alpha &= E[x^\alpha] \\
\mathbf{P}_0^\alpha &= E[(x_0^\alpha - \hat{x}_0^\alpha)(x_0^\alpha - \hat{x}_0^\alpha)^T]
\end{aligned}$$

For $k \in \{1, ..., +\infty\}$, calculate sigma points:

$$\mathcal{X}_{k-1}^\alpha = [\hat{x}_{k-1}^\alpha \ \hat{x}_{k-1}^\alpha \pm \sqrt{(L+\lambda)\mathbf{P}_{k-1}^\alpha}]$$

Time update:

$$\mathcal{X}^x_{k|k-1} = \mathbf{F}[\mathcal{X}^x_{k-1}, \mathcal{X}^v_{k-1}]$$

$$\hat{x}^-_k = \sum_{i=0}^{2L} W_i^{(m)} \mathcal{X}^x_{i,k|k-1}$$

$$\mathbf{P}^-_k = \sum_{i=0}^{2L} W_i^{(c)} [\mathcal{X}^x_{i,k|k-1} - \hat{x}^-_k][\mathcal{X}^x_{i,k|k-1} - \hat{x}^-_k]^T$$

$$\mathcal{Y}_{k|k-1} = \mathbf{H}[\mathcal{X}^x_{k|k-1}, \mathcal{X}^n_{k-1}]$$

$$\hat{y}^-_k = \sum_{i=0}^{2L} W_i^{(m)} \mathcal{Y}_{i,k|k-1}$$

Measurement update equations:

$$\mathbf{P}_{\tilde{\mathbf{y}}_k \tilde{\mathbf{y}}_k} = \sum_{i=0}^{2L} W_i^{(c)} [\mathcal{Y}_{i,k|k-1} - \hat{y}^-_k][\mathcal{Y}_{i,k|k-1} - \hat{y}^-_k]^T$$

$$\mathbf{P}_{\mathbf{x}_k \mathbf{y}_k} = \sum_{i=0}^{2L} W_i^{(c)} [\mathcal{X}_{i,k|k-1} - \hat{x}^-_k][\mathcal{Y}_{i,k|k-1} - \hat{y}^-_k]^T$$

$$\mathcal{K} = \mathbf{P}_{\mathbf{x}_k \mathbf{y}_k} \mathbf{P}^{-1}_{\tilde{\mathbf{y}}_k \tilde{\mathbf{y}}_k}$$

$$\hat{x}_k = \hat{x}^-_k + \mathcal{K}(y_k - \hat{y}^-_k)$$

$$\mathbf{P}_k = \mathbf{P}^-_k - \mathcal{K} \mathbf{P}_{\tilde{\mathbf{y}}_k \tilde{\mathbf{y}}_k} \mathcal{K}^T$$

with, $x^\alpha = [x^T v^T n^T]^T$, $\mathcal{X}^\alpha = [(\mathcal{X}^x)^T (\mathcal{X}^v)^T (\mathcal{X}^n)^T]^T$, λ is a composite scaling parameter, L the dimension of augmented state, $\mathbf{P_v}$ the process noise covariance, $\mathbf{P_n}$ the measurement noise covariance and finally, W_i the weights.

5.4 GAUSSIAN PROCESSES

For the UKF to be effective, we first need to obtain some information about the underlying model, that is, the function through which the forecasts are turned into final actions. For this purpose we incorporate *Gaussian Processes* (GP) [38], a powerful tool that can be used for regression and classification without a parametric model assumption.

In this work we employ Gaussian processes for probabilistic regression: for a set of training samples, $\mathcal{D} = \{(\mathbf{x}_j, y_j), j = 1, ..., n\}$ ($\mathbf{x}_j$ inputs and y_j noisy outputs) we need to predict the distribution of the noisy output at some test locations $\mathbf{x}_*$. We assume the following model:

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

with σ^2_{noise} the variance noise. GP regression is a Bayesian approach that assumes a priori that function values follow: $p(\mathbf{f}|\mathbf{x}_1, \mathbf{x}_2, ..., \mathbf{x}_n) =$

$\mathcal{N}(\mathbf{0}, K)$ where $\mathbf{f} = [f_1, f_2, ..., f_n]^T$ is the vector of latent function values, $f_j = f(\mathbf{x}_j)$ and K is the covariance matrix that is computed by a covariance function $K_{jk} = k(\mathbf{x}_j, \mathbf{x}_k)$.

In order to proceed to the inference, we must combine the joint GP prior $\mathbf{f}$ and $\mathbf{f}_*$ obtained by the test values with the likelihood $p(\mathbf{y}|\mathbf{f})$, via Bayes rule. The result is the joint posterior:

$$p(\mathbf{f}, \mathbf{f}_*|\mathbf{y}) = \frac{p(\mathbf{f}, \mathbf{f}_*)p(\mathbf{y}|\mathbf{f})}{p(\mathbf{y})} \tag{33}$$

To produce the posterior for our predictions, we marginalize out the unwanted training set variables:

$$p(\mathbf{f}_*, \mathbf{y}) = \int p(\mathbf{f}, \mathbf{f}_*|\mathbf{y})d\mathbf{f} = \frac{1}{p(\mathbf{y})} \int p(\mathbf{y}|\mathbf{f})p(\mathbf{f}, \mathbf{f}_*)d\mathbf{f} \tag{34}$$

where

$$p(\mathbf{f}, \mathbf{f}_*) = \mathcal{N}(\mathbf{0}, \begin{bmatrix} K_{\mathbf{f},\mathbf{f}} & K_{*,\mathbf{f}} \\ K_{\mathbf{f},*} & K_{*,*} \end{bmatrix}), \text{ and } p(\mathbf{y}|\mathbf{f}) = \mathcal{N}(\mathbf{f}, \sigma^2_{noise}I).$$

The joint GP prior and the independent likelihood are both Gaussian with mean and variance as follows:

$$GP_\mu(\mathbf{x}_*, \mathcal{D}) = K_{*,f}(K_{f,f} + \sigma^2_{noise}I)^{-1}\mathbf{y} \tag{35a}$$

$$GP_\sigma(\mathbf{x}_*, \mathcal{D}) = K_{*,*} - K_{*,f}(K_{f,f} + \sigma^2_{noise}I)^{-1}K_{f,*} \tag{35b}$$

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:

$$\theta_{max} = \arg\max_\theta \{\log(p(\mathbf{y}|\mathbf{X}, \theta))\} \tag{36}$$

5.5 COMBINING GP WITH UKF

The output of the GP can be used in conjunction with the UKF in order to generate more accurate predictions regarding agents' α_is. When an agent states an uncertainty forecast $\hat{\sigma}_{\tau-1}$, it is propagated through the dynamic model and the expected mean GP_μ and variance GP_σ of the corresponding α_i are calculated via the GP. These two values are used to transform Eq.(27a) and (27b) into:

Dynamic Problem Formulation

$$\hat{\sigma}_\tau = A\hat{\sigma}_{\tau-1} + w_\tau \tag{37a}$$

$$\alpha_\tau = GP_\mu(\hat{\sigma}_\tau, \mathcal{D}) + u_\tau \tag{37b}$$

with noise u_τ following $\mathcal{N}(0, GP_\sigma(\hat{\sigma}_\tau, \mathcal{D}))$. Since in our setting a realistic transition model is hard to define given a lack of related real data regarding the progression of uncertainty statements over time, state transition matrix A is set to the identity matrix, and w_τ is assumed to follow $\mathcal{N}(0, 1)$. The method we described here, GP-UKF, is then used over this updated dynamic model to obtain the predictions $\tilde{\alpha}_i$.

Here we must note that it is possible to use the GP alone, sampled at the most recent agent statement. However, such an approach would then not be able to model time varying uncertainty estimates. In contrast, the GP-UKF is able to incorporate even time varying agent uncertainty estimate statements, should such statements become available.

5.6 METHODS EVALUATION

We provide a detailed evaluation of our methods in Chapter 6. Our simulations demonstrate that employing any of the two filters tested leads to improved performance in the consumption shifting domain, when compared to a "baseline" mechanism that makes no use of such performance monitoring tools. Specifically, when using these enhancements, the cooperative achieves a higher overall electricity consumption reduction; and enjoys financial rewards that are higher than those generated by the baseline algorithm.

Each filtering method's exact performance depends on the uncertainty characterizing a particular simulation setting. We observed, however, that *HF*'s performance is stable, and the method is able to forecast the future performance of an agent quite closely. At the same time, combining *UKF* with *GP* regression appears to lead to increased prediction accuracy, and higher economic gains overall.

Part III

SIMULATIONS

In this part we present the setting and results of the experimental evaluation of our mechanism. First we describe the dataset we used and its various initializations. Next we compare the different coalition formation methods that we have proposed and present the general features of our mechanism. Finally we adopt the stochastic filtering techniques to monitor single agents and agent classes and demonstrate the performance boost that such methods provide.

6 EXPERIMENTAL RESULTS

In this Chapter, the results from the experimental evaluation are presented and discussed in detail. In Section 6.1 we present the dataset that we use to simulate a realistic cooperative shifting scenario. Section 6.2 evaluates our overall scheme and compares the performance of the alternative methods for forming reducing coalitions proposed in Section 4.3, as well as various results that demonstrate the effectiveness of our proposed scheme. Finally, the evaluation of the incorporation of the two stochastic filtering methods is presented in Section 6.3. The experimental results provided in this Chapter demonstrate the effectiveness of our demand shifting scheme, as it delivers what it promised, i.e. providing effective demand-side management services. Also, it is shown that the reducing cooperatives can obtain monetary gains by participation in this scheme. Both earnings and effectiveness are further improved by the incorporation of the stochastic filtering methods we proposed.

6.1 DATASET

Population Modeling

In our experiments, we use simulated data coming from real consumption patterns of 36 industrial customers in India[1]. We were provided with two 36×48 arrays, one holding the mean consumption values for each customer at each time interval and one holding the standard deviation of the means. For the agents' consumption patterns generation we repeatedly sample these normal distributions, until the desirable agent population number is met, that is 4968 agents. In other words, every resulting agent pattern emerged from sampling the 48 normal distributions for each interval of one of the 36 real consumer profiles.

Parameter Initialization

The τ threshold is fixed to 96.5% of the maximum total demand across all time intervals. The safety limit is set to 99% of τ, while $q_{\min}^{t_h}$ to 1% of the total load at t_h; and the p_{high} & p_{low} values are set to the day-night prices specified by the greek utility company P.P.C.[2]. The p_{group} rate (in € / KWh) ranges from $p_{group}^{\max} = 0.05625$ to $p_{group}^{\min} = 0.0214$, depending on reduction size q:

$$p_{group}(q) = \frac{p_{group}^{\min} - p_{group}^{\max}}{Q^{t_h} - q_{\min}^{t_h}} \cdot (q - q_{\min}^{t_h}) + p_{group}^{\max} \tag{38}$$

1 The exact same consumers as in the experiments of [26]; we thank that paper's authors for providing the dataset.

2 http://www.dei.gr

with q ranging from $q^{t_h}_{min}$ to a maximum (t_h-specific) Q^{t_h}.

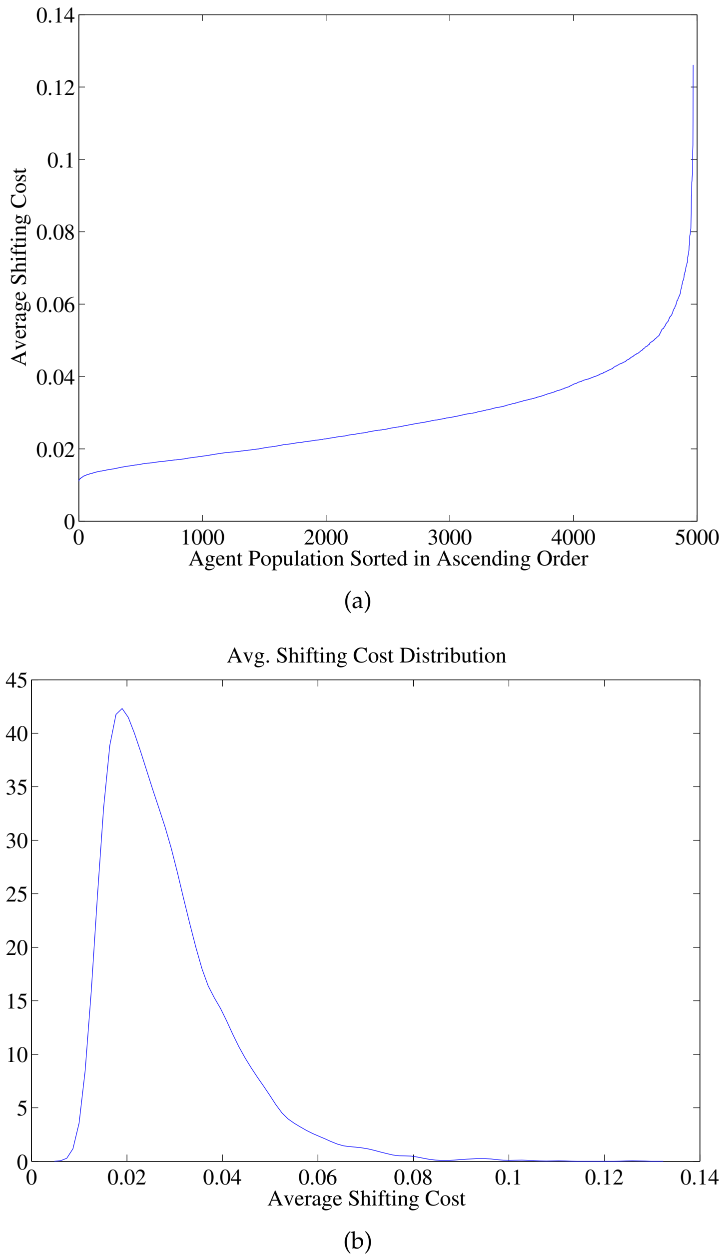

Figure 2: Average Shifting Costs.

Shifting Capabilities

Individual shifting preferences are not provided in the dataset and are generated as follows. A beta distribution[3] ($a = 1, b = 43.444$) is sampled twice for each agent, giving the means (a *higher* mean and a *lower* mean) of two normal distributions ($\sigma = 0.01$) that are next sampled for each interval, resulting to the actual agent *shifting cost*.

3 Here, we choose a Beta distribution, as we want the costs to have mean values between zero and one. Additionally, such distributions are more appropriate for modelling uncertainties, as they can concentrate probability in a desired range. Despite these facts, other forms can also be used.

Consider the mean consumption of some agent i during a day, $\bar{c}_i$. If the potential shifting operation trims consumption from a time interval t^{above} where the agent's baseload demand is higher than $\bar{c}_i$ to another interval, t^{below}, that baseload demand is lower than $\bar{c}_i$, then the value sampled from the normal distribution with the *higher* mean is assigned to that $cost_i^{t^{above} \to t^{below}}$ for agent i, while the value sampled from the normal distribution with the *lower* mean is assigned to all other transitions, i.e. $t^{above} \to t^{above}$, $t^{below} \to t^{above}$ and $t^{below} \to t^{below}$. The intuition behind this higher cost assignment is that peaks in demand are mainly costly to shift (but there exist some agents that can contribute with lower costs or shifting capacities). Also, because the outcome of the beta distribution is between 0 and 1, we add 0.01151 to all values in order to move cost values to an acceptable cost range[4]. In Fig 2(a), agents are sorted by their average daily shifting costs. Population with average shifting cost below 0.02 is considered *low* cost, above 0.02 and below 0.04 *medium* cost, and above 0.04 *high* cost.

On an average run, 811.86 agents were *high cost*; 2809 were *medium cost*; while 1347.14 were *low cost*. The kernel density estimate of the daily average agent shifting costs is shown in Fig. 2(b).

Reduction capacities are estimated based on the variance of each agent baseline consumption, as this is a good indicator for its demand elasticity[5]. Agent uncertainty stated for bidding, is provided by sampling a beta distribution, with $a = 1, b = 5$ i.e., the great mass of the agent population has low to average uncertainty and actual agent shifting actions are provided by sampling another beta, with $a = 4, b = 2$ modeling the realistic case that at best the agents deliver what they promised, but often fail to do so.

In our experiments, we first assume that all 4968 agents participate in the cooperative. Note that running one simulation day involving all agents takes on average only 1.5 sec on a 3.3 GHz PC (with a further 110 sec for demand curves and shifting costs initialization period before simulation starts).

A typical picture of what happens in a given day is shown in Fig. 3. The diamond pointed curve is the resulting shifted demand, the dashed curve shows the initial demand and the two vertical lines depict threshold τ and safety limit.

As Fig. 3 shows, intervals 19 to 20, 22 to 25, 27 to 29 and 31 to 37 are peak intervals. Consumption that was reduced is shifted to other, non-peak intervals. Notice that at interval 27, an effective reducing coalition could not be formed due to lack of available reducing

4 Costs that are lower than $p_{high} - p_{low} = 0.01151$ are not realistic, because such agents would shift by themselves to lower demand intervals, changing their baseline consumption

5 Elasticity of demand is defined as the degree to which demand for a good or service varies according to its price.

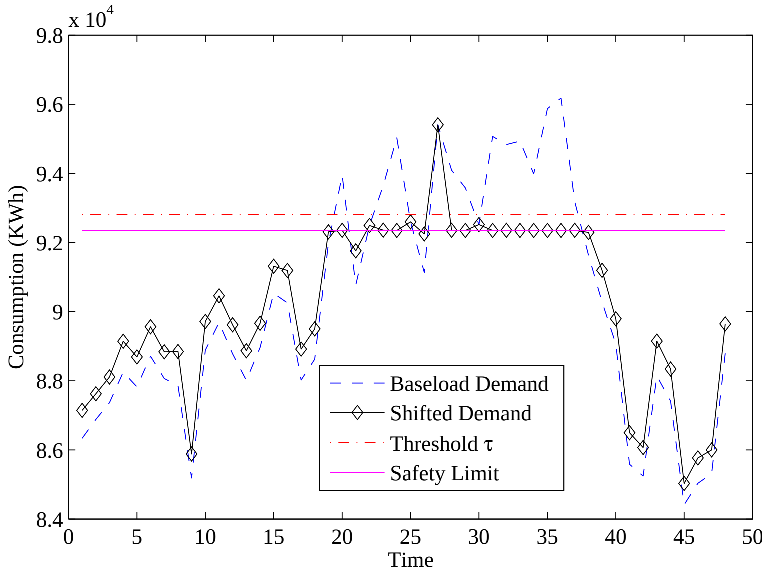

Figure 3: Consumption shifting scheme output.

agents. Despite this infeasibility though, the amount of total peak-load trimmed is very significant.

6.2 METHOD EVALUATION

We continue with the simulation of 100 days, and compare the four methods presented in Section 4.3. Results are shown in Table 3.

Method 2 clearly ranks lower than all others both in terms of co-operative gains and trimming effectiveness. This is because it forms coalitions of a "minimum" size, capable *in expectation* to shift just q_{min}. Thus, due to uncertainty governing actual agent behavior, profits suffer when agent promises fail to materialize (and q_{min} is not reached).

In contrast, *Methods 1, 3* and *4* all trimmed more than 98% of peak load, and have similar performance. When adopting *Method 3* though, it is possible for an agent that could be favoured by the final p_{group} price to be excluded from the coalition, if its contribution potential was checked early-on in the process, when the p_{group} price awarded at that point happened to grant negative expected gain to that agent. *Method 4* results to the highest cooperative gain, and highest consumption reduction, but is the most expensive computationally. Moreover, from the Grid's point of view, it is probably not worth it to hand over an additional 37.28€/day, 4.1% of the amount "paid" to *Method 1*, for a mere 0.35% increase in consumption reduction. Therefore, *Method 1* appears to be the most appropriate for our purposes, as it is comparable to the rest both in terms of cooperative gains and trimming ability; is cheaper for the Grid to use; and allows even agents with initially negative expected gains to participate in the scheme. For those reasons, we chose *Method 1* for further evaluation. After having examined the general behavior of our methods, we then conducted experiments to examine the CRPS effect, the average size of the reducing coalitions when p_{group} perturbs and the cooperative reduction performance subject to different agent participation percentages.

We also simulated the fifth method proposed for coalition formation in Section 4.3, and explored its behavior. Provided ranking by $\tilde{r}_i \hat{p}_i$, the concept is to monitor the *average growth rate of coalescing agents' gain* and stop adding when it drops under a certain level. Suppose that in a coalition C with $gain_C$, we add another agent i, then the gain growth obtained by this addition is $growth = gain_{C \cup i} - gain_C$. Trying to monitor growth values for each peak interval we observe that gain growth is not smooth and thus it is difficult to set a certain limit, as shown in Fig. 4. For this reason, this method was not chosen for further evaluation. *Method 4.3.5*

	Method 1	Method 2	Method 3	Method 4
Expected Cooperative Gain (€/day)	2094.72	415.73	2099.40	2139.38
Actual Cooperative Gain (€ /day)	895.56 (σ=348.03)	-273.89 (σ=77.30)	899.34 (σ=349.36)	933.24 (σ=361.66)
Cooperative "Surplus" (€ /day)	6.20	14.81	6.20	6.13
Expected Reduction (KWh)	32856.45	12868.40	32861.60	32919.06
Final Reduction (KWh)	24454.32 (σ=8083.3)	9471.053 (σ=2556.5)	24461.93 (σ=8086.1)	24539.96 (σ=8058.2)
Peak (Demand $\geqslant \tau$) Trimmed (%)	98.616 (σ=0.75)	42.511 (σ=5.86)	98.619 (σ=0.73)	98.639 (σ=0.75)
Avg. Reducing Coalition Size	47.70598	15.49350	47.82045	49.06499
High Cost Participants (%)	5.609709	4.396607	5.384088	4.151038
Medium Cost Participants (%)	57.78608	56.61016	57.91386	55.56619
Low Cost Participants (%)	36.60421	38.99323	36.70205	40.28277

Table 3: Average Results (100 days simulation); σ denotes standard deviation from average values.

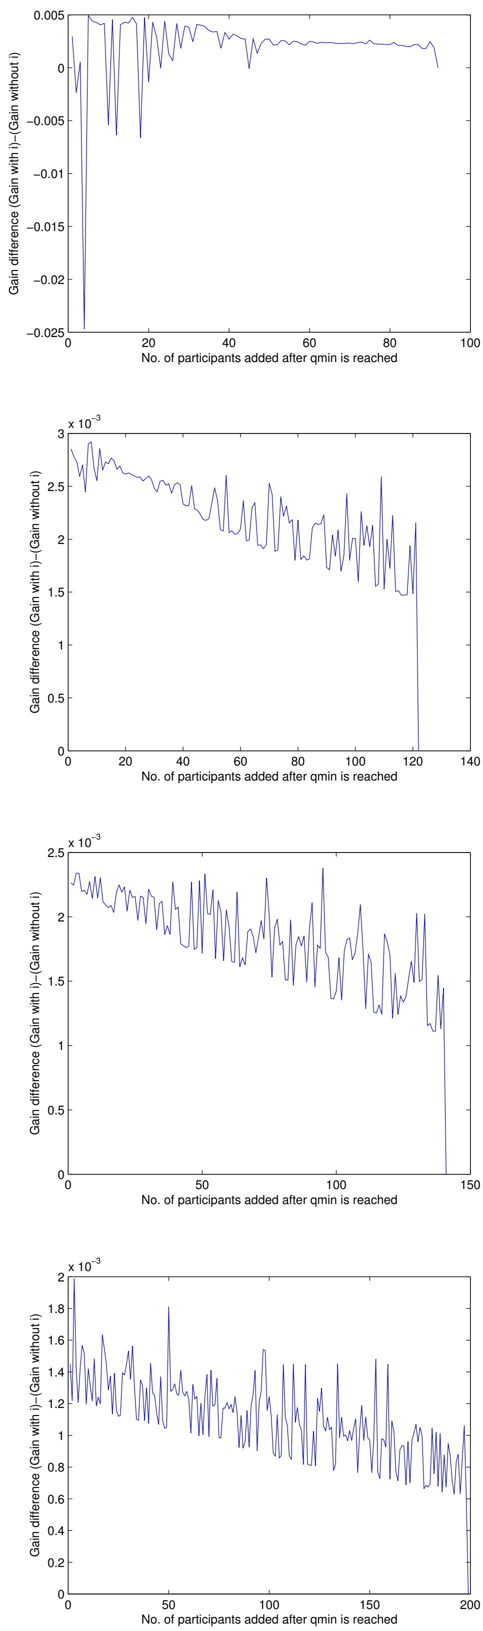

Figure 4: Gain growth for various t_h.

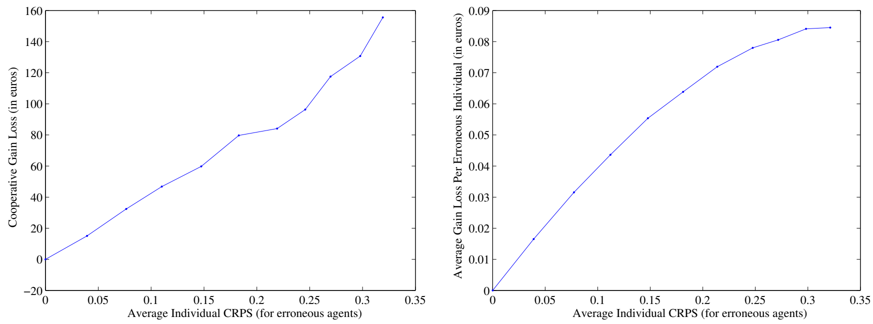

(a) (b)

Figure 5: Daily loss of gain as CRPS rises.

6.2.1 *CRPS Effect*

In this set of experiments, every reducing coalition formed at any t_h is made to contain 10 agents whose relative error is non-zero. The relative error was progressively increased, over 11 complete runs of 50 simulation days each; that is, the relative error was constant during a run lasting for 50 simulation days, and was then increased for later runs. 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. As shown in Figure 5, as the individual agent's CRPS score gets worse, its gain losses increase; and the cooperative as a whole suffers progressively increasing gain losses as well. Thus, CRPS clearly incentivizes the agents to accurately state their reduction capabilities, and deliver what they promised.

6.2.2 *Coalition Size vs Pgroup*

Next, we examine the average reducing coalition size formed at each t_h given different p_{group} prices granted for collective consumption shifting. More specifically, we simultaneously increase the p_{group}^{max} and p_{group}^{min} values produced by Eq. 38 up to $+0.05$ of their initial values, and observe the average number of agents in reducing coalitions for each peak interval. Figure 6 shows this concept, where average coalition sizes over 100 simulation days are plotted against group price range variations. It is obvious that as p_{group} increases to get closer to p_{low}, fewer agents decide to contribute—and, subsequently, less consumption is finally shifted. Thus, in order for shifting to take place, the Grid must grant a p_{group} range that provides enough gain to the agents, given individual shifting costs.

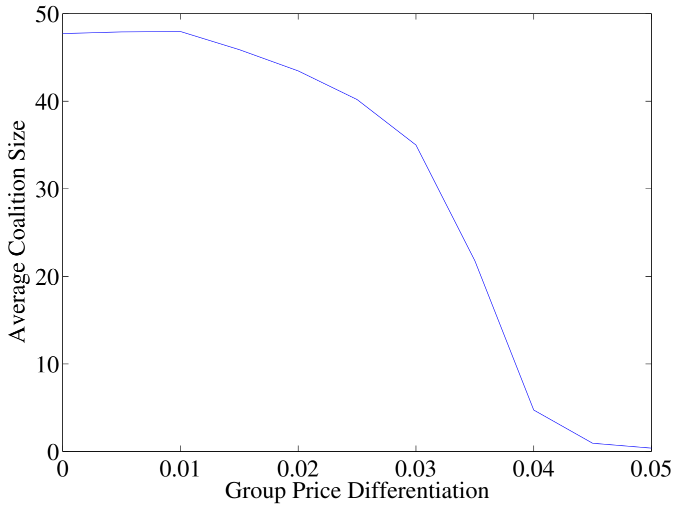

Figure 6: Avg. reducing coalition size vs. p_{group} increase.

6.2.3 *Cooperative Membership*

In a final set of experiments, we considered settings with considerably fewer consumers participating in the cooperative. With 30% of the agents participating, it is still possible to shift 98.52% of peak load, while 10% of the population manages to shift 93.82% of peak load. With 7% of agents participating, 75% of the total peak load is shifted; while 4% and 3% of all agents shift 51.86% and 12.66% of peak load respectively. Finally, 2.5% of all agents shift only 0.8% of the peak load. Thus, membership clearly has to reach a "critical mass" for the cooperative to be effective.

6.2.4 *Further Observations and Insights*

Here we provide further insights into benefits gained from cooperative scheme participation, given our simulations. These refer to the setting where all agents participate in the cooperative. We measured that an average number of agents participating into each reducing coalition at some t_h is 47.7 individuals. The 36.6% of these are *low shifting cost* agents, whereas 57.8% are *medium* and 5.6% *high* cost. When the reward transfer scheme is enforced, the actual amount transferred on average during a gain balancing operation at a given t_h is negligible, in the order of 10^{-5}€, and is granted by either *low cost* or *medium cost* consumers.

It is worth noting that the Grid each day grants back to consumers an average of *only* 895.56€, from its average daily income of 354064.3€. Note that we cannot account for the Grid profits emerging due to reduced generation costs from the evasion of peak intervals, because such an analysis would require information that is typically not disseminated by the Grid operators. However, since we observe that in an average simulation run an 98.616% of the peak load is "safely" shifted, we can infer that the Grid stands to gain from the shifting operations. Another positive side-effect is, of course, that power outages (and resulting costs) become more distant possibilities as the demand curve flattens out.

6.3 MONITORING AGENTS WITH STOCHASTIC FILTERING TECHNIQUES

In this section we employ the stochastic filtering techniques of the Chapter 5 to test their performance. First, we monitor single agents of specific behavior. Then we test each method at various settings in order to measure the relative error that accrues for each one. Finally we apply them on the electricity consumption shifting problem. Note that, for simplicity, agent forecasts over their expected relative error are given in the form of a $\mathcal{N}(0, \hat{\sigma})$, so the stated uncertainty forecasts from now on will be simply denoted by $\hat{\sigma}$.

6.3.1 *Monitoring the Performance of a Single Agent*

GP-UKF APPROACH For the GP generation we used the MATLAB OCTAVE GPML (Gaussian Processes for Machine Learning) toolbox as it is presented in [41]. The set of training points is $\mathcal{D} = \{(\hat{\sigma}_j, \alpha_j), j = 1, ..., 1000\}$, i.e. pairs of forecasted uncertainty and actual actions. We used expectation propagation as an inference method. The GP is trained over the past observations; and sampled to get the forecasted uncertainty values for the next day. Specifically, the GP output is the mean and variance of the expected α_j given the forecasts $\hat{\sigma}_j$.

For the UKF we used the EKF/UKF Toolbox for Matlab V1.3 developed by the department of Biomedical Engineering and Computer Science of Aalto University, tuned to the dynamic problem described in Section 5.5 by Eq. (37).

Non-linear function with noise

Initially, we use a setting with a non-linear function of a known form in order to test the performance of our approach. Here, $\hat{\sigma}$ is tuned to follow a $\mathcal{N}(0.5, 0.15)$ distribution. The nonlinear function that processes the samples and outputs α_i values is $f = 0.5 \sin^2(2\pi\hat{\sigma}) + 0.3$, with added noise following $\mathcal{N}(0, 0.1)$. Results are shown in Figure 7(a). Diamonds represent certain "test points" generated by f plus noise, to help illustrate the spread of the nonlinearly transformed input. As shown, trained means (blue crosses) and variances (gray shaded area) of the GP, successfully adapt to the nonlinear function f. UKF predictions are marked as squares, and are in fact very close to the GP trained means.

Samples from Beta distributions

Next, we let $\hat{\sigma}$ and α_i follow $\mathcal{B}(1, 5)$ and $\mathcal{B}(4, 2)$ respectively, as in the experimental setting explained in Section 6.1. Results are shown in Figure 7(b). In this case, there is no function of a specific known

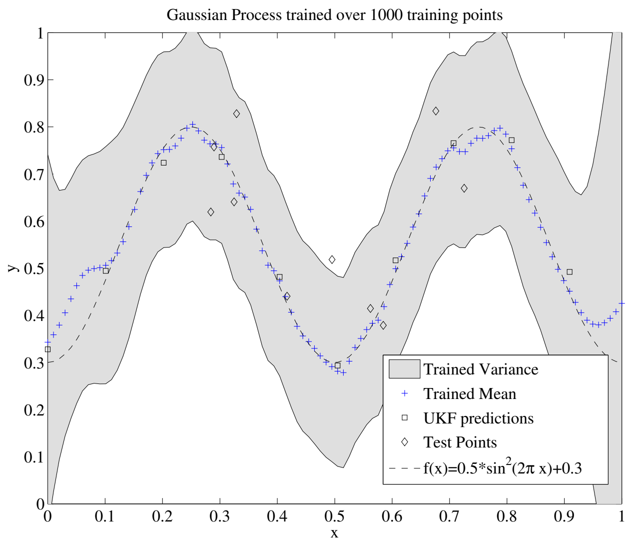

(a)

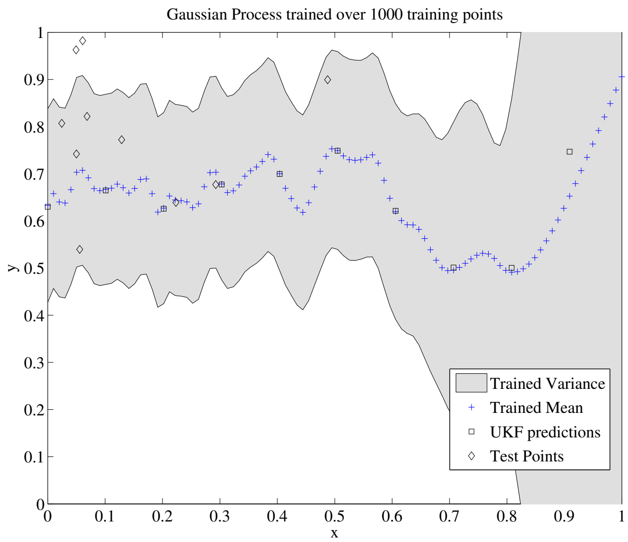

(b)

Figure 7: The Gaussian Process and UKF prediction points for a single agent.

form for the Gaussian Process to approximate, as the points are both random variables following different distributions. Despite that, the 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, the GP-UKF method is apparently able to produce meaningful predictions, even when the relationships between variables are governed by some highly complex function. 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.

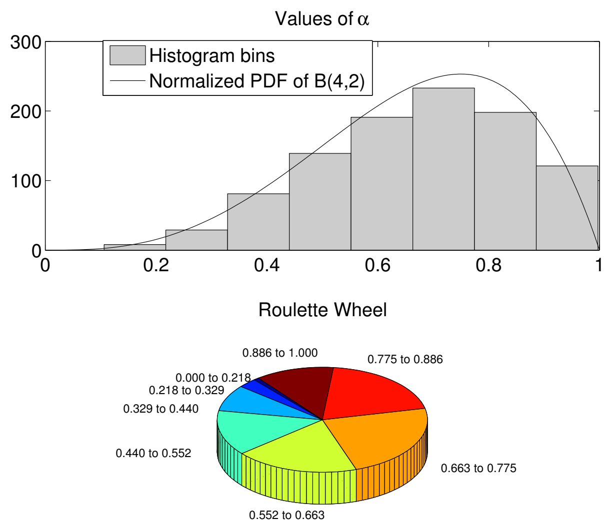

Figure 8: Histogram and corresponding roulette wheel (8 bins, 1000 observations).

HF APPROACH In contrast to UKF, the HF approach does not require $\hat{\sigma}_i$ forecasts to work. We test this technique in a setting with α_i following $\mathcal{B}(4, 2)$. For the histogram generation, MATLAB built-in function *hist* was used. Once again, we use 1000 training points. These are now simply α_i values, sampled from the aforementioned *beta* distribution. Next, a roulette wheel is constructed in order to perform stochastic universal sampling. According to the height of each bin of the histogram, a proportionate space on the roulette wheel is colored. The roulette wheel is spinned thereafter —i.e., a random sample is generated that maps to a specific bin, or *range*. The final prediction of the HF comes up by sampling a uniform distribution, normalized to the range that occurred from the roulette wheel spin.

The resulting histogram and the corresponding roulette wheel are shown in Fig. 8. One can clearly see the relevance between the probability density function (black curve) and the generated histogram; the histogram successfully converges on the true generating distribution. The accuracy of the approximation, however, depends on the number of training values and the number of bins.

6.3.2 *Monitoring Agent Classes*

In another set of experiments we tune our parameters to follow different distribution pairs. Three new agent classes are defined, the *accu-*

Class	Conservative		HF		GP-UKF	
	μ	σ	μ	σ	μ	σ
AP	0.023	0.017	0.022	0.017	0.015	0.012
UP	0.430	0.998	0.390	1.190	0.299	0.867
IP	0.157	0.117	0.021	0.016	0.014	0.012
B-N	1.626	3.836	0.554	0.328	0.272	0.392
B-B	0.904	1.921	0.547	0.277	0.281	0.394

Table 4: Means and variances of the absolute relative error for each class when using each method.

rate predictor, the *uncertain predictor* and the *inaccurate predictor*, apart from two additional test cases, as we explain below.

More specifically, accurate predictor's (AP) uncertainty forecasts and final actions both follow a normal distribution with low variance, $\mathcal{N}(0.5, 0.01)$, simulating consumers that act almost as they predicted to act. The inaccurate predictors (IP) on the other hand, seek to act 50% off of what they predicted, with their predictions following $\mathcal{N}(0.5, 0.1)$. The uncertain predictor (UP), 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. $\mathcal{N}(0.5, 0.15)$. Two additional scenarios are also tested, one with α_i following $\mathcal{B}(4, 2)$ and $\hat{\sigma}_i \sim \mathcal{N}(0.5, 0.15)$ (B-N) and the other, that described in Section 6.1 (B-B), using two *beta* distributions $\alpha_i \sim \mathcal{B}(4, 2)$ and $\hat{\sigma}_i \sim \mathcal{B}(1, 5)$. The first one makes more pessimistic estimates than the other; and both end up delivering the $r_i^{t_h}$ reduction value that they promised, but sometimes fail to do so. Our goal is to find the best method to obtain estimates of $\tilde{\alpha}_i$.

The first method, the *Conservative*, estimates final shifting capacity based on Eq. (13). The second, *HF*, utilizes a Histogram Filter constructed by past α_i observations[6] and does not take into account $\hat{\sigma}$ statements. The shifting capacity estimate is calculated via Eq. (26), where $\tilde{\alpha}_i$ is the prediction given by the HF. Finally, *GP-UKF* also keeps track of past observations, but takes into account $\hat{\sigma}$ statements. Shifting capacity estimates are calculated also by Eq. (26), but with $\tilde{\alpha}_i$ now corresponding to UKF predictions.

In order to test the aforementioned methods on the scenarios presented above, we first sample the distributions generating the $\hat{\sigma}_i$s and the α_is 1000 times in order to generate training points for the HF and GP-UKF methods. Next we sample the distributions again for 1000 times and measure the mean absolute relative error between the occurred α_i and predicted $\tilde{\alpha}_i$ from the proposed methods. Results are

6 In fact, α_i is calculated by $\frac{r_i^{t_h}}{\hat{r}_i^{t_h}}$, where $\hat{r}_i^{t_h}$ is the stated reduction capacity at t_h, and $r_i^{t_h}$ is the observed actual reduction by i there.

Method	Setting	Expected Gain	Actual Gain	Expected Reduction	Final Reduction
Conserv.	B-B	2185.60	982.66	34375.497	25593.202
	UP	2358.64	1417.14	37480.805	30889.185
HF	B-B	2172.33	1458.31	34436.174	29327.052
	UP	2378.51	1436.78	37711.798	31171.310
GP-UKF	B-B	2184.96	1932.94	34556.859	33862.32
	UP	2391.65	2041.50	37809.045	36610.243

Table 5: Average (out of 10 days) results of electricity consumption shifting cooperative.

shown in Table 4. Judging from the mean absolute relative errors for each agent class, GP-UKF clearly outperforms both HF and the conservative method in both mean and variance. As results show, when dealing with *accurate predictors*, all three methods have similar performance. When the predictors are *uncertain*, however, absolute relative error rises, but GP-UKF maintains its better performance. When monitoring *inaccurate predictors*, the HF and GP-UKF methods capture the systematic inaccuracies of agents and perform very well, while the conservative does not. In the remaining two cases, one can observe that the conservative method performs far worse than the two proposed methods. This happens because the conservative method makes highly pessimistic estimates for agent final actions in these settings.

Therefore, our methods can successfully exploit historical data to learn and keep track of the actual agent performance, a feature that the conservative method does not possess. In addition, the GP-UKF can identify possible relations between agent forecasts and final actions, managing to predict the latter more precisely than the HF, by incorporating agent uncertainty forecasts $\hat{\sigma}_i$ in the prediction process.

6.3.3 *Cooperative Shifting Simulation*

In the final experimental setting, we apply our proposed monitoring methods to the real world scenario, considering *cooperative* consumption shifting efforts over a 10-day period.

User forecasts and final actions are set to follow two of the cases discussed in the previous subsection: B-B, which is the case actually presented in Section 6.1; and UP, which models a realistic scenario.[7]

7 AP and IP are less likely in practice, and so is B-N: if an agent originally unsure of its performance (that is, one sampling $\hat{\sigma}$ from $\mathcal{N}(0.5, 0.15)$) observes that its actual $(1-\alpha)$ error is small (as α originates from $\mathcal{B}(4,2)$), then it would most probably have corrected its uncertainty predictions already, to avoid penalties.

Thus, we run a 10-day simulation for each one of these two cases. For each day in each simulation, we apply the scheme three times, each using a different method (Conservative, HF or GP-UKF) for obtaining agent final action estimates —with the *same* initialization data per day, in order to be able to make a reasonable comparison. At the beginning of each 10 day simulation HF and GP-UKF are fed with 100 "past observed" pairs of forecasts and final actions, in order to infer an elementary model that will be used for their predictions. Numerical results are presented in Table 5. Expected and actual gain represent euros per day, while expected and final reduction quantities come in kWh per day.

Expected gain and expected reduction take close values at each setting for all methods. This is because, in all occasions, the underlying coalition formation mechanism employed by the cooperative keeps adding agents to reducing teams until the maximum reduction quantity requested at a given t_h is reached. However, each *prediction* method has its own estimates regarding each agents actual performance. Therefore, the performance of the three prediction methods varies. Clearly, GP-UKF achieves reduction that is closer to the expected (when compared against the other methods). One can observe that for the HF approach, the deviation between expected and final load reduced is at the same levels regardless of the distributions that generate final actions. This is due to approximation errors, induced by the number of bins of the histogram and uniform distribution sampling. Also, agent forecasts are not taken into account, so potentially important information is ignored. In terms of accuracy, we can thus safely conclude that the GP-UKF has the best performance. Intuitively, GP-UKF 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. Moreover, we can see that the GP-UKF performs significantly better *wrt*. economic benefits, generating as it does more actual gain in euros than any other method. Overall, GP-UKF appears to be a strong prediction tool for consumption management cooperatives.

6.4 DISCUSSION

To conclude this Chapter, our proposed mechanism appears to be highly effective at the proactive flattening of the electricity demand curve. Forming reducing coalitions by using Method 1 leads to significant monetary gains for the cooperative, incorporating even agents with negative expected gain, and is cheaper for the Grid to use. The CRPS score incentivizes agents to meet their stated commitments in their actual final shifting actions. Also, granting an appropriate p_{group} price helps agents overcome their shifting costs by participating in large reducing coalitions. Moreover, with the use of stochastic filtering methods, the performance and prediction accuracy are significantly improved and thus, increased financial gain for the cooperative are generated. Our filtering 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). This is the first time these methods are applied in this domain; and the potential value of this work to any real-world enterprise operating in the Smart Grid can be very high. Finally, as we will see in Appendix A, we have recently obtained more datasets, on which we plan to further experiment and test our proposed scheme.

Part IV

CONCLUSIONS AND FUTURE WORK

In this part, we review the contributions of our work, and discuss the positive side effects of the adoption of our proposed mechanism in real world Smart Grids. Finally, we provide the reader with insights regarding our future research efforts.

7 CONCLUSIONS AND FUTURE WORK

In this work, we proposed an entirely novel and directly applicable scheme for electricity consumption shifting, which promotes agent efficiency in the face of uncertainty. The scheme uses a strictly proper scoring rule, CRPS, to incentivize participants to accurately state their reduction capabilities and actually deliver their promised contribution; and promotes the formation of consumer cooperatives that can collectively significantly contribute to the trimming of the demand curve. This is achieved via the granting of a better group electricity price for large shifting operations. To incentivize agent participation, we put forward an incentive compatible and (weak) budget-balanced gain transfer scheme that awards variable group price rates to members, thus allowing agents with initially prohibitive shifting costs to participate in shifting operations. All our methods were tested and evaluated via appropriate simulations. We also applied two different methods for monitoring and predicting agent actions in the power consumption shifting scheme, a Histogram Filter and a Unscented Kalman Filter that uses Gaussian Processes to recognize possible underlying relationships between agent forecasts and final actions. Concluding, we believe that we have demonstrated that there are indeed good reasons to use the proposed demand management scheme, as its use can lead to the effective flattening of the electricity demand curve.

Future work includes running extensive simulations on larger scale environments. Ideally, we would like to test our model in an environment that includes hundreds of thousands of actual consumers (as opposed to using simulated data generated from only a few dozens of real consumers). This would also enable us to study the behavior of various types of agents, classified according to particular industry type; and allow us to devise coalitional policies specifying the best possible mixture of members to include in shifting operations. In particular, a sizable real world consumption dataset from a municipality of Crete has already been preprocessed and is presented in Appendix A, along with a plan for future experiments concerning this dataset. We also intend to devise methods for distributed, combined performance monitoring, so that predictions are effectively cross-validated. Moreover, we will try to model the problem of *trust* in this domain. The sequential selection of participants for reducing coalitions in particular, could require solving a belief-state MDP, much in the spirit of what is done in [48].

Last but not least, the study of prosumer cooperatives (i.e., cooperatives whose members can simultaneously consume and produce electricity) is definitely an interesting and challenging research direction. Allowing the members to have dual roles (i.e., consumers and producers) adds a significant level of complication that will most probably have to be dealt with separately. Prosumer agents have the choice to strategize about short term decisions (such as choosing not to reduce consumption, via, for instance, shifting their energy production activities now) in anticipation of achieving greater benefits, e.g. by producing and storing energy to sell to the Grid with an expected higher profit on a later date.

Part V

APPENDIX

This part shows the results from a preprocessing of annual real power consumption data coming from the municipality of Kissamos, in western Crete.

A

APPENDIX: ANALYSIS OF A REAL DATASET PROVIDED BY THE PUBLIC POWER CORPORATION OF GREECE

The spreadsheet containing consumption data of the municipality of Kissamos during the year 2012, came to our possession on March 31, 2013. We are grateful to the Department of Informatics of the greek P.P.C. utility company. One can find there the total consumption for specific time periods, of 8.487 consumers of various types. More specifically, data for the fields shown in Table 6 are provided:

Voltage level
ID of connection
Street name
Street number
City
Municipality
Type of usage category
Invoice code
Invoice type
Date of measurement
No. of days measurements correspond to
KWH consumed
Reactive power demand
Apparent power demand
True power demand
Impedance phase angle

Table 6: Fields covered in the Kissamos dataset.

By processing the above, we should end up with consumption patterns of various types of consumers depicting as much as possible the real situation of Kissamos.

A.1 REDUNDANT FIELDS

Some of the fields included in the spreadsheet can be considered redundant for our initial approach. These are *Street name, Street number, City, Municipality* and *Date of measurement*.

ID of connection
Type of usage category
Invoice code
Invoice type
Number of days
KWH consumed

Table 7: The reduced set of fields of interest.

Table 8: Number of consumers of each type.

Type	Count	Count without zero users
Residential	6263	5889
Commercial	1497	1381
Agricultural	298	271
Municipal	302	295
Public	75	68
Industrial	39	38
Public Law Entity	13	12
Total	8487	7954

While *Reactive* and *Apparent* power demand might give us some information about the types of the consumers, that would be beyond the scope of the current project, as we cope most with actual power consumption, thus the only field needed is *KWH consumed*. The aforementioned exclusions reduce the fields to the set shown in Table 7.

A.2 TYPES OF CONSUMERS

The next issue is categorizing consumers into specific and discrete types. To achieve this, we must take a look at the 40 different value combinations of Usage category, Invoice code and Invoice type, as shown in Fig 9. However, we choose to categorize them according to the first column of Fig 9, the *usage category*. The numbers of each type are shown in Table 8.

A.3 MEAN CONSUMPTION CALCULATION

From each consumer, multiple measurements are provided during the year. Based on this, we can derive the distributions that describe daily consumption patterns and proceed further with the analysis.

ΑΓΡΟΤΙΚΗ	33	Τ33ΑΡ
ΑΓΡΟΤΙΚΗ	20	ΑΕΡΓ
ΑΓΡΟΤΙΚΗ	21	Γ21
ΑΓΡΟΤΙΚΗ	22	Γ22
ΑΓΡΟΤΙΚΗ	94	Τ33ΜΤ
ΑΓΡΟΤΙΚΗ	Ν3	ΜΤ-Ν
ΒΙΟΜΗΧΑΝΙΚΗ	40	Γ21Β
ΒΙΟΜΗΧΑΝΙΚΗ	41	Γ22Β
ΒΙΟΜΗΧΑΝΙΚΗ	5	Φ/Β
ΒΙΟΜΗΧΑΝΙΚΗ	42	Γ22Β
ΒΙΟΜΗΧΑΝΙΚΗ	20	ΑΕΡΓ
ΔΗΜΟΣΙΑ / ΔΗΜΟΤΙΚΑ / ΦΟΠ	25	Γ22
ΔΗΜΟΣΙΑ / ΔΗΜΟΤΙΚΑ / ΦΟΠ	21	Γ21
ΔΗΜΟΣΙΑ / ΔΗΜΟΤΙΚΑ / ΦΟΠ	22	Γ22
ΔΗΜΟΣΙΑ / ΔΗΜΟΤΙΚΑ / ΦΟΠ	90	Β1
ΔΗΜΟΣΙΑ / ΔΗΜΟΤΙΚΑ / ΦΟΠ	20	ΑΕΡΓ
ΔΗΜΟΣΙΑ / ΔΗΜΟΤΙΚΑ / ΦΟΠ	Ν5	ΜΤ-Ν
ΔΗΜΟΤΙΚΑ / ΦΟΠ	53	Γ4/ΛΛ
ΔΗΜΟΤΙΚΑ / ΦΟΠ	54	Γ4/ΛΛ
ΔΗΜΟΤΙΚΑ / ΦΟΠ	22	Γ22
ΕΜΠΟΡΙΚΗ	22	Γ22
ΕΜΠΟΡΙΚΗ	20	ΑΕΡΓ
ΕΜΠΟΡΙΚΗ	21	Γ21
ΕΜΠΟΡΙΚΗ	25	Γ22
ΕΜΠΟΡΙΚΗ	24	Γ23Ν
ΕΜΠΟΡΙΚΗ	23	Γ23Η
ΕΜΠΟΡΙΚΗ	5	Φ/Β
ΕΜΠΟΡΙΚΗ	91	Β2
ΕΜΠΟΡΙΚΗ	Ν2	ΜΤ-Ν
ΝΠΔΔ/ΔΗΜ.ΕΠΙΧΕΙΡ.	21	Γ21
ΟΙΚΙΑΚΗ	10	Γ1
ΟΙΚΙΑΚΗ	11	Γ1
ΟΙΚΙΑΚΗ	36	ΓΤ
ΟΙΚΙΑΚΗ	74	ΓΠ
ΟΙΚΙΑΚΗ	6	ΓΤ
ΟΙΚΙΑΚΗ	5	Φ/Β
ΟΙΚΙΑΚΗ	12	Γ1Ν
ΟΙΚΙΑΚΗ	75	ΓΠ
ΟΙΚΙΑΚΗ	7	ΓΤ
ΟΙΚΙΑΚΗ	37	ΓΤ

Figure 9: Categories of Kissamos consumers.

Table 9: Mean, minimum and maximum consumption for each type.

Type	Mean KWH (zero\no zero users)	Min (without zero users)	Max
Residential	6.8588\7.2944	0.00107	83.69257
Commercial	23.1370\25.0804	0.00271	1022.49635
Agricultural	101.3737\111.4737	0.01219	1985.01832
Municipal	13.4569\13.7762	0.03003	84.49058
Public	69.2345\76.3616	0.00269	1937.73004
Industrial	104.5073\107.2575	0.20202	1830.61383
Public Law Entity	13.7733\14.9211	1.66666	64.45923

Note that measurements regarding reactive power are not taken into consideration for simplicity.

For each couple of KWH and Days values, we come up with a mean of each month (or four months) and afterwards with the mean of the year, according to the procedure shown in Algorithm 3.

Algorithm 3 Computation of μ_i and σ_i^2

Input: $\vec{KWH}, No.of\vec{D}ays$ (Multiple values for each consumer)
Output: $\vec{\mu}, \vec{\sigma}^2$ (One value for each consumer)
for $\forall i$ (each consumer) **do**
  for $\forall j$ (each measurement of the consumer) **do**
    $\frac{KWH}{No.ofDays} \rightarrow temp_j$
  end for
  $\mu_i = \frac{\sum_{\forall j} temp_j}{|j|}$
  for $\forall j$ (each measurement of the consumer) **do**
    $(\frac{KWH}{No.ofDays})^2 \rightarrow temp2_j$
  end for
  $\sigma_i^2 = \frac{\sum_{\forall j} temp2_j}{|j|} - \mu_i^2$
end for

The results for every usage category are shown in Fig. 10. The mean consumption for each category is shown in Table 9, and the mean variance is shown in Table 10.

To provide a better visualization, we transform our data x to $\log(x)$, after *removing consumers with zero consumption from our dataset* since such consumers cannot contribute to the scheme, neither affect consumption patterns. Results are shown in Figure 10. Histograms of means and variances for each class is summarized in Figures, 11 and 13. Another non-parametric way of obtaining a probability density function of observed data is the *Kernel density estimate*. Kernel density estimates are closely related to histograms, and can be smooth and continuous, based on the kernel we use. The kernel density estimates of means and variances for each class using a normal kernel function, is shown in Figures 12 and 14.

Table 10: Mean, minimum and maximum variance σ^2 of consumption for each type.

Type	Mean σ^2 (zero\no zero users)	Min (without zero users)	Max
Residential	9.6054\10.3687	0	2200.371
Commercial	2560.2\2811.9	0	1769926.260
Agricultural	36689\40493	0	1759549.917
Municipal	19.9674\ 21.3080	0	1573.370
Public	9058.1\9990.6	0.000014452538	324108.367
Industrial	8227.2\8443.7	0.000332684800802729	138245.986
Public Law Entity	80.6619\87.3837	0.00444097754475337	865.332

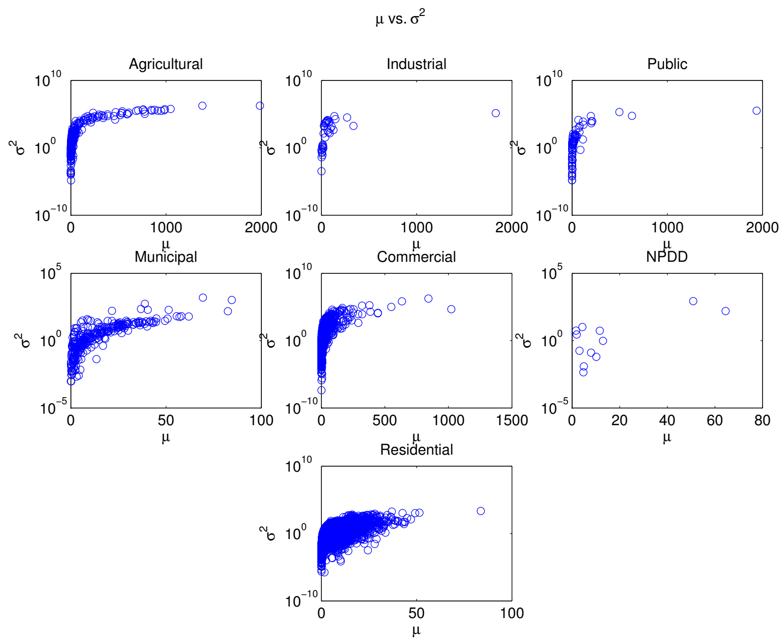

Figure 10: Means and variances of consumers based on their usage category.

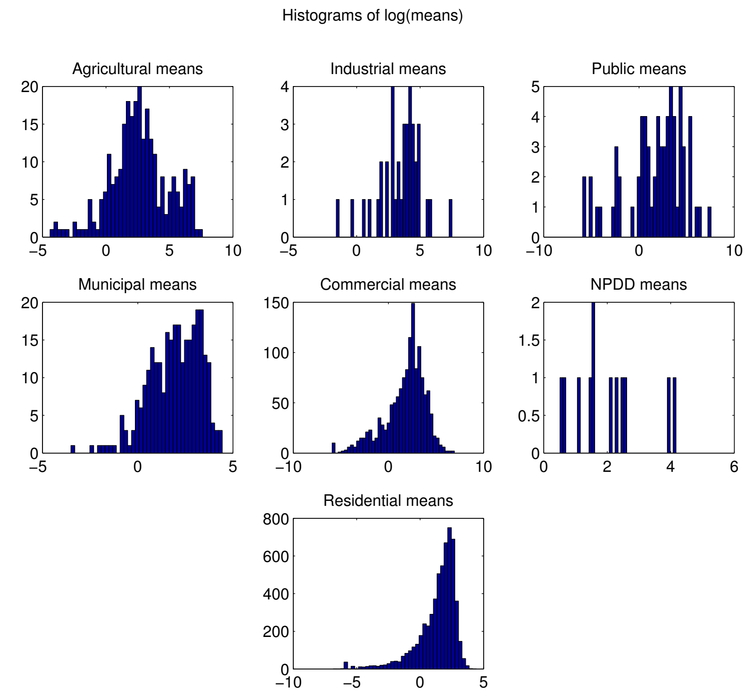

Figure 11: Histograms of means for each usage category.

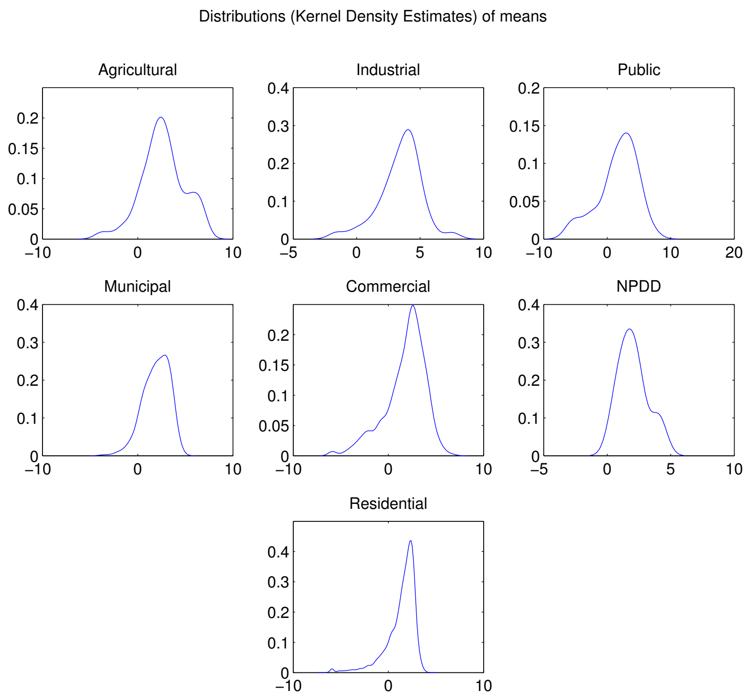

Figure 12: Kernel density estimates of logarithm of means for each usage category.

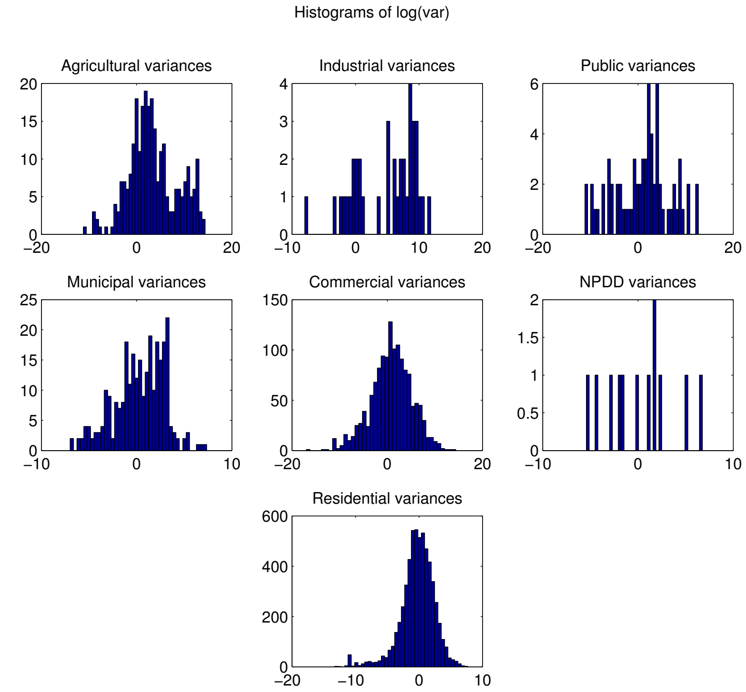

Figure 13: Histograms of variances for each usage category.

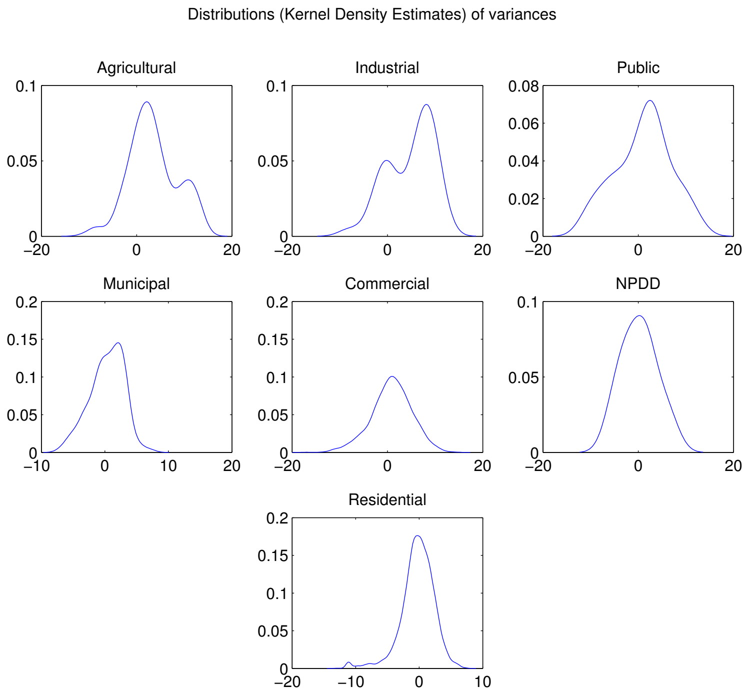

Figure 14: Kernel density estimates of logarithm of variances for each usage category.

A.4 PLAN FOR FUTURE EXPERIMENTS

Taking into consideration the aforementioned real world agent classes, we can assume scenarios of consumption shifting operations where:

- every category participates
- residential, commercial and industrial consumers participate, while public companies do not
- only residential
- every other category apart from residential

Also, the modeling of *prosumers* is quite easy, simply by assigning negative values at the field of KWH consumed. Unfortunately, such profiles are not provided in this dataset.

[1] Thomas Ackermann, Göran Andersson, and Lennart Söder. Distributed generation: a definition. *Electric Power Systems Research*, 57(3):195 – 204, 2001. ISSN 0378-7796. doi: http://dx.doi.org/10.1016/S0378-7796(01)00101-8. URL http://www.sciencedirect.com/science/article/pii/S0378779601001018.

[2] M.H. Albadi and E.F. El-Saadany. A summary of demand response in electricity markets. *Electric Power Systems Research*, 78 (11):1989 – 1996, 2008. ISSN 0378-7796. doi: 10.1016/j.epsr.2008.04.002.

[3] C. Anger, R. Schrader, and U. Klingauf. Unscented Kalman Filter with Gaussian Process Degradation Model for Bearing Fault Prognosis. In *European Conference of the Prognostics and Health Management Society*, 2012.

[4] P. Asmus. Microgrids, virtual power plants and our distributed energy future. *The Electricity Journal*, 23(10):72 – 82, 2010. ISSN 1040-6190. doi: 10.1016/j.tej.2010.11.001.

[5] M. Blum and M. Riedmiller. Electricity Demand Forecasting using Gaussian Processes. In *Workshops at the Twenty-Seventh AAAI Conference on Artificial Intelligence*, 2013.

[6] 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.

[7] 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.

[8] Georgios Chalkiadakis, Edith Elkind, and Michael Wooldridge. Computational aspects of cooperative game theory. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 5(6):1–168, 2011.

[9] Javier Contreras, Felix Wu, Matthias Klusch, and Onn Shehory. Coalition formation in a power transmission planning environment. In *in In Proc. of 2nd Intl. Conference on Practical Applications of Multi-Agent Systems, PAAM*, pages 21–23. Online]. Available: citeseer.nj.nec.com/12028.html.

[10] D. Cooray, R. Malek, S.and Roshandel, and D. Kilgore. RESISTing reliability degradation through proactive reconfiguration. In *Proc. of the IEEE/ACM International Conference on Automated Software Engineering*, ASE '10, 2010.

[11] B. Davito, H. Tai, and R. Uhlaner. The smart grid and the promise of demand-side management. *McKinsey*, 2010.

[12] K. Dielmann and A. van der Velden. Virtual power plants (vpp) - a new perspective for energy generation? In *Modern Techniques and Technologies, 2003. MTT 2003. Proceedings of the 9th International Scientific and Practical Conference of Students, Post-graduates and Young Scientists*, pages 18–20, 2003. doi: 10.1109/SPCMTT.2003.1438108.

[13] AL Dimeas and ND Hatziargyriou. Agent based control of virtual power plants. In *Intelligent Systems Applications to Power Systems, 2007. ISAP 2007. International Conference on*, pages 1–6. IEEE, 2007.

[14] A. Elkhodary, N. Esfahani, and S. Malek. FUSION: a Framework for Engineering Self-Tuning Self-Adaptive Software Systems. In *Proc. of the 18th ACM SIGSOFT International Symposium on Foundations of Software Engineering*, FSE '10, pages 7–16, New York, NY, USA, 2010. ACM. ISBN 978-1-60558-791-2. doi: 10.1145/1882291.1882296. URL http://doi.acm.org/10.1145/1882291.1882296.

[15] N. Esfahani and S. Malek. Uncertainty in Self-Adaptive Software Systems. In *Software Engineering for Self-Adaptive Systems II*, volume 7475 of *LNCS*, pages 214–238. 2013. ISBN 978-3-642-35812-8. doi: 10.1007/978-3-642-35813-5_9. URL http://dx.doi.org/10.1007/978-3-642-35813-5_9.

[16] A. Y. Alanis, L. J. Ricalde, C. Simetti and F. Odone. Neural Model with Particle Swarm Optimization Kalman Learning for Forecasting in Smart Grids. *Mathematical Problems in Engineering*, vol. 2013. doi: 10.1155/2013/197690.

[17] X. Fang, S. Misra, G. Xue, and D. Yang. Smart grid - the new and improved power grid: A survey. *Communications Surveys Tutorials, IEEE*, PP(99):1 –37, 2011. ISSN 1553-877X. doi: 10.1109/SURV.2011.101911.00087.

[18] T. Gneiting and A. E. Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007. doi: 10.1198/016214506000001437.

[19] 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 Math-

ematical Systems. Springer, 2012. ISBN 9781447143840. URL http://books.google.gr/books?id=BJAsrYbW9PIC.

[20] 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. doi: 10.1109/SMARTGRID.2010.5622091.

[21] 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. URL http://ideas.repec.org/p/gwc/wpaper/2012-003.html.

[22] S. J. Julier and J. K. Uhlmann. A new extension of the kalman filter to nonlinear systems. In *Proc. of AeroSense: The 11th Int. Symp. on Aerospace/Defence Sensing, Simulation and Controls.*, pages 182–193, 1997.

[23] Daniel Kirschen and Goran Strbac. *Fundamentals of Power System Economics*. Wiley, 2004.

[24] J. Ko, D. J. Klein, D. Fox, and D. Haehnel. GP-UKF: Unscented Kalman Filters with Gaussian Process Prediction and Observation Models. In *Intern. Conf. on Intelligent RObots and Systems - IROS*, 2007.

[25] J.K. Kok, M.J.J. Scheepers, and I.G. Kamphuis. Intelligence in electricity networks for embedding renewables and distributed generation. In Rudy R. Negenborn, Zofia Lukszo, and Hans Hellendoorn, editors, *Intelligent Infrastructures*, volume 42 of *Intelligent Systems, Control and Automation: Science and Engineering*, pages 179–209. Springer Netherlands, 2010. ISBN 978-90-481-3597-4. doi: 10.1007/978-90-481-3598-1_8. URL http://dx.doi.org/10.1007/978-90-481-3598-1_8.

[26] Ramachandra Kota, Georgios Chalkiadakis, Valentin Robu, Alex Rogers, and Nicholas R. Jennings. Cooperatives for demand side management. In *Proceedings of the 20th European Conference on Artificial Intelligence (ECAI-2012)*, pages 969–974, August 2012.

[27] Junghoon Lee, Gyung-Leen Park, Sang-Wook Kim, Hye-Jin Kim, and Chang Oan Sung. Power consumption scheduling for peak load reduction in smart grid homes. In *Proceedings of the 2011 ACM Symposium on Applied Computing*, SAC '11, pages 584–588, New York, NY, USA, 2011. ACM. ISBN 978-1-4503-0113-8. doi: 10.1145/1982185.1982310. URL http://doi.acm.org/10.1145/1982185.1982310.

[28] 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. ISBN 1-58113-788-5. doi: 10.1145/948005.948015.

[29] Radu Casian Mihailescu, Matteo Vasirani, and Sascha Ossowski. Towards agent-based virtual power stations via multi-level coalition formation. In *Proc. of the 1st Int. Workshop on Agent Technologies for Energy Systems*, pages 107–108, 2010.

[30] MIT authors. Engaging electricity demand. In *MIT Interdisciplinary Study on the future of the electric grid*. MIT, 2011.

[31] A. 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. *Smart Grid, IEEE Transactions on*, 1(3):320 –331, dec. 2010. ISSN 1949-3053. doi: 10.1109/TSG.2010.2089069.

[32] A.-H. Mohsenian-Rad and A. Leon-Garcia. Optimal residential load control with price prediction in real-time electricity pricing environments. *Smart Grid, IEEE Trans.*, 1(2):120 –133, sept. 2010. ISSN 1949-3053. doi: 10.1109/TSG.2010.2055903.

[33] Athanasia Mpourgkoudi. Comparison and evaluation of the application of renewable energy sources. Master's thesis, National Technical University of Athens, University of Piraeus, October 2009.

[34] Roger B Myerson. *Game theory: analysis of conflict*. Harvard university press, 2013.

[35] 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. URL http://eprints.soton.ac.uk/341638/.

[36] Guido Pepermans, Johan Driesen, and Dries 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. URL http://ideas.repec.org/p/ete/etewps/ete0308.html.

[37] D. Pudjianto, C. Ramsay, and G. Strbac. Virtual power plant and system integration of distributed energy resources. *Renewable Power Generation, IET*, 1(1):10–16, 2007. ISSN 1752-1416. doi: 10.1049/iet-rpg:20060023.

[38] J. Quiñonero Candela and C. E. Rasmussen. A Unifying View of Sparse Approximate Gaussian Process Regression. *J. Mach. Learn. Res.*, 6:1939–1959, December 2005. ISSN 1532-4435. URL http://dl.acm.org/citation.cfm?id=1046920.1194909.

[39] S. 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*, 2012.

[40] Sarvapali D. Ramchurn, Dong Huynh, and Nicholas R. Jennings. Trust in multi-agent systems. *Knowl. Eng. Rev.*, 19(1):1–25, March 2004. ISSN 0269-8889. doi: 10.1017/S0269888904000116. URL http://dx.doi.org/10.1017/S0269888904000116.

[41] C. E. Rasmussen and H. Nickisch. The gpml toolbox. http://www.gaussianprocess.org/, January 2013.

[42] C.E. Rasmussen and C.K.I. Williams. *Gaussian Processes for Machine Learning*. Adaptative computation and machine learning series. University Press Group Limited, 2006. ISBN 9780262182539. URL http://books.google.gr/books?id=vWtwQgAACAAJ.

[43] V. Robu, R. Kota, G. Chalkiadakis, A. Rogers, and N. R. Jennings. Cooperative Virtual Power Plant Formation using Scoring Rules. In *Proc. of the 26th Conference on Artificial Intelligence (AAAI-12)*, pages 370–376, August 2012.

[44] Valentin Robu, Enrico H. Gerding, S. Stein, David C. Parkes, Alex Rogers, and Nick R. Jennings. An online mechanism for multi-unit demand and its application to plug-in hybrid electric vehicle charging. *J. Artif. Intell. Res. (JAIR)*, 48:175–230, 2013.

[45] C. (R.) Sastry, V. Srivastava, R. Pratt, and S. Li. Use of residential smart appliances for peak-load shifting and spinning reserves - cost/benefit analysis. *U.S. Department of Energy*, December 2010.

[46] Yoav Shoham and Kevin Leyton-Brown. *Multiagent systems: Algorithmic, game-theoretic, and logical foundations*. Cambridge University Press, 2009.

[47] Goran Strbac. Demand side management: Benefits and challenges. *Energy Policy*, 36(12):4419–4426, December 2008. URL http://ideas.repec.org/a/eee/enepol/v36y2008i12p4419-4426.html.

[48] W. T. L. Teacy, G. Chalkiadakis, A. Rogers, and N. R. Jennings. Sequential decision making with untrustworthy service providers. In *Proceedings of the 7th International Conference on Autonomous Agents and Multiagent Systems*, pages 755–762, May 2008. URL http://eprints.soton.ac.uk/265168/. Event Dates: 12th-16th May, 2008.

[49] S. Thrun, W. Burgard, and D. Fox. *Probabilistic Robotics*. MIT Press, 2005. ISBN 0262201623.

[50] 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. doi: 10.1080/09613218.2010.494832.

[51] Perukrishnen Vytelingum, Sarvapali D. Ramchurn, Thomas D. Voice, Alex Rogers, and Nicholas R. Jennings. Trading agents for the smart electricity grid. In *The Ninth International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, pages 897–904, 2010. URL http://eprints.soton.ac.uk/268361/. Event Dates: May 10-14, 2010.

[52] Perukrishnen Vytelingum, Thomas D. Voice, Sarvapali D. Ramchurn, Alex Rogers, and Nicholas R. Jennings. Agent-based micro-storage management for the smart grid. In *The Ninth International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2010) - Won the Best Paper Award*, pages 39–46, 2010. URL http://eprints.soton.ac.uk/268360/. Winner of the Best Paper Award Event Dates: May 10-14, 2010.

[53] E.A. Wan and R. Van der Merwe. The unscented kalman filter for nonlinear estimation. In *Adaptive Systems for Signal Processing, Communications, and Control Symposium 2000. AS-SPCC. The IEEE 2000*, pages 153–158, 2000. doi: 10.1109/ASSPCC.2000.882463.

[54] Michael Wooldridge. *An introduction to multiagent systems*. Wiley.com, 2008.

[55] Jie Xiong. *An Introduction to Stochastic Filtering Theory*. Number 9780199219704 in OUP Catalogue. Oxford University Press, 2008. URL http://ideas.repec.org/b/oxp/obooks/9780199219704.html.

[56] 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. ISBN 1-58113-326-X. doi: 10.1145/375735.376452.

[57] J. Yen, Y.H. Yan, B. J. Wang, P. K H Sin, and F.F. Wu. Multi-agent coalition formation in power transmission planning. In *System Sciences, 1998., Proceedings of the Thirty-First Hawaii International Conference on*, volume 4, pages 433–443 vol.4, 1998. doi: 10.1109/HICSS.1998.655300.

[58] C.S.K. Yeung, A.S.Y. Poon, and F.F. Wu. Game theoretical multi-agent modelling of coalition formation for multilateral trades.

Power Systems, IEEE Transactions on, 14(3):929–934, 1999. ISSN 0885-8950. doi: 10.1109/59.780905.