UNIVERSITY OF SOUTHAMPTON

# Addressing The Computational Issues of the Shapley Value With Applications in The Smart Grid

by

Sasan Maleki

A thesis submitted in partial fulfillment for the degree of Doctor of Philosophy

in the
Faculty of Physical Sciences and Engineering
Electronics and Computer Science

May 2015

UNIVERSITY OF SOUTHAMPTON

# ABSTRACT

FACULTY OF PHYSICAL SCIENCES AND ENGINEERING
ELECTRONICS AND COMPUTER SCIENCE

Doctor of Philosophy

by Sasan Maleki

We consider the computational issues that arise in using the Shapley value in practical applications. Calculating the Shapley value involves computing the value of an exponential number of coalitions, which poses a significant computational challenge in two cases: (i) when the number of agents (players) is large (e.g., more than 20), and (ii) when the time complexity of the characteristic function is high. However, to date, researchers have aimed to address only the first case, although with limited success.

To address the first issue, we focus on approximating the Shapley value. In more detail, building upon the existing sampling-based approaches, we propose an improved error bound for approximating the Shapley value using simple random sampling (SRS), which can be used in any superadditive game. Moreover, we put forward the use of *stratified sampling*, which can lead to smaller standard errors. We propose two methods for minimising the standard error in supermodular games and a class of games that have a property that we call *order-reflecting*. We show that among others, *newsvendor* games, which have applications in the smart grid, exhibit this property. Furthermore, to evaluate our approach, we apply our stratified sampling methods to an instance of newsvendor games consisting of 100 agents using real data. We find that the standard error of stratified sampling in our experiments is on average 48% lower than that of SRS.

To address the second issue, we propose the characteristic function of the game be approximated. This way, calculating the Shapley value becomes straightforward. However, in order to maintain fairness, we argue that, in distributing the value of the grand coalition, agents' contribution to the complexity of the characteristic function must be taken into account. As such, we propose the *bounded rational Shapley value*, which, using the additivity axiom of the Shapley value, ensures that the share of each agent reflects its contribution to the difficulty of computing the coalition values. We demonstrate the usefulness of this approach in a demand response scenario where a number of apartments want to fairly divide the discount they receive for coordinating their cooling loads.

# Contents

# List of Tables

# List of Figures

# Abbreviations

| | |
|---|---|
| TU | Transferable Utility |
| MC-nets | Marginal Contribution Networks |
| SRS | Simple Random Sampling |
| CLT | Central Limit Theorem |
| i.i.d. | Independent and Identically Distributed |
| DER | Distributed Energy Resource |
| CVPP | Cooperative Virtual Power Plant |
| BETTA | British Electricity Trading Transmission Arrangements |
| PDF | Probability Density Function |
| CDF | Cumulative Distribution Function |
| TCL | Thermostatically Controlled Load |
| DR | Demand Response |
| DP | Dynamic Programming |
| SE | Standard Error of The Mean |

# Nomenclature

**Chapter 2**

| | |
|---|---|
| $A$ | The set of agents in a game. |
| $n$ | The number of agents in a game. |
| $C$ | A coalition, which is a subset of $A$. |
| $v$ | The characteristic function of a game. |
| $MC(a, C)$ | The marginal contribution of agent $a$ to coalition $C$. |
| $x$ | A payoff distribution. |
| $e(x, C)$ | The excess of the payoff distribution $x$ out of the total of $v(C)$. |
| $\chi$ | A map, which is a function that assigns a weight to each coalition. |
| $\mathcal{C}(A, v)$ | The core of a game. |
| $\pi(A)$ | The set of all permutations of agents. |
| $\mathcal{O}$ | A permutation of agents, which is a member of $\pi(A)$. |
| $Pre_a^O$ | The set of agents that precede $a$ in the permutation $O$. |
| $SV(a)$ | The Shapley value of agent $a$. |
| $\epsilon$ | The bound on the approximation error. |
| $\delta$ | The probability that the error bound fails, with 1-$\delta$ being the confidence on the bound. |
| $\overline{SV}(a)$ | The approximated Shapley value of agent $a$. |
| $r$ | The range of marginal contributions of agent $a$. |
| $X$ | A random variable that takes the values of agent $a$'s marginal contributions in any joining order with probability $1/n!$. |

**Chapter 3**

| | |
|---|---|
| $p$ | Unit price of electricity. |
| $\alpha$ | The expected system sell price. |
| $\beta$ | The expected system buy price. |
| $w$ | The actual generation of a generator. |
| $\nu$ | The contract volume of a generator. |
| $\nu^*$ | The optimal contract volume of a generator. |
| $\zeta(w, \nu)$ | The balancing cost of a generator, given the actual and contract volumes of $w$ and $\nu$, respectively. |
| $f_W$ | The probability distribution function of $W$. |

| | |
|---|---|
| $F_W$ | The cumulative distribution function of $W$. |
| $F_W^{-1}$ | The inverse cumulative distribution function of $W$, also called the quantile of $W$. |
| $M$ | The maximum generation capacity of a DER. |
| $P(\nu, \zeta)$ | A function representing the profit of a DER given a contract volume of $\nu$ and an imbalance of $\zeta$. |

**Chapter 4**

| | |
|---|---|
| $m$ | The sample size. |
| $\Phi$ | A random variable, $\Phi$, taking the $n!$ marginal contribution values of an agent, each with a probability of $1/n!$ |
| $\bar{\Phi}$ | The sample mean of $\Phi$, i.e., the estimated Shapley value. |
| $\bar{\Phi}_{SRS}$ | The sample mean of $\Phi$ under SRS. |
| $\bar{\Phi}_{STR}$ | The sample mean of $\Phi$ under stratified sampling. |
| $\phi^{max}$ | The maximum marginal contribution of an agent. |
| $\phi^{min}$ | The minimum marginal contribution of an agent. |
| $S$ | The number of strata. |
| $N_s$ | The size of stratum $s$. |
| $m_s^*$ | The optimal sample size for stratum $s$. |
| $\mathbb{B}$ | The set of branching agents. |
| $b_i$ | A branching agent. |
| $\mathcal{F}(C)$ | A function that reflects the order of marginal contributions of a given agent to $C$ among all coalitions of the same size. |
| $g$ | A production function in an output sharing game. |
| $q_a$ | The runway cost associated with aircraft type $a$ in an airport game. |
| $\ell_a$ | The input of agent $a$ in an output sharing game. |

**Chapter 5**

| | |
|---|---|
| $\omega(\|C\|)$ | The weight of the marginal contribution corresponding to coalition $C$ in the Shapley value formula (equation 2.2). |
| $v^{BR}$ | The characteristic function which gives the bounded rational value of a coalition. |
| $v^{RD}$ | The characteristic function which gives the rational discrepancy of a coalition. |
| $SV(a, v)$ | The Shapley value of agent $a$ in game $(A, v)$. |
| $p$ | The normal price of consuming 1 kWh of electricity. |
| $d$ | The discounted price of consuming 1 kWh of electricity. |
| $K$ | The set of equal-length time slots which represent an entire day. |
| $k$ | The last time slot in $K$ which also represents the size of $K$. |
| $\Delta t$ | The duration of one time slot in seconds. |
| $\psi$ | The target threshold of the neighbourhood. |

| | |
|---|---|
| $\eta_a^t$ | A cooling action, which is a binary variable that indicates whether the air conditioner at time $t$ is on or off. |
| $P_a$ | The electric power of the air conditioner in apartment $a$. |
| $l_a^t$ | The cooling load of apartment $a$ at time $t$. Refer to (5.4). |
| $l_A^t$ | The aggregate cooling load of all apartments in $A$. |
| $l_A^{*t}$ | The optimal cooling load of coalition $C$ at time $t$. |
| $l_a^{*t}$ | The optimal cooling load of apartment $a$ at time $t$. |
| $c$ | The function that gives the optimal consumption of coalitions. Refer to (5.5). |
| $T_{int}^t$ | The internal temperature of the apartment at time $t$ (measured in °C). Refer to (5.7) |
| $T_{env}^t$ | The temperature of the envelope of the apartment at time $t$ (measured in °C). Refer to (5.7) |
| $T_{ext}^t$ | The external temperature at time $t$ (measured in °C). |
| $\tau$ | The rate of leakage from the envelope to the inside (measured in 1/hr). |
| $\rho$ | The rate of leakage from the inside to the envelope (measured in 1/hr). |
| $\gamma$ | The rate of leakage from the outside to the envelope. |
| $T_{set}$ | The set-point temperature (in °C). |
| $T_{set}[a]$ | The set-point temperature of apartment $a$ as a member of a coalition (in °C). |
| $\theta$ | The tolerance of the apartment on deviations of the internal temperature from the set-point during the comfort period (in °C). |
| $\theta_a$ | The tolerance of apartment $a$ on deviations of the internal temperature from the set-point during the comfort period (in °C). |
| $\Omega$ | A parameter to limit the deviation of the average internal temperature from the set-point during the comfort period (in °C). |
| $CST$ | Comfort start time. |
| $CST_a$ | Comfort start time of apartment $a$ as a member of a coalition. |
| $CET$ | Comfort end time. |
| $CET_a$ | Comfort end time of apartment $a$ as a member of a coalition. |
| $H$ | The set of time slots that represent the comfort period specified by $CST$ and $CET$. $H = \{t \in K \mid CST \leq t \leq CET\}$ |
| $H_a$ | The set of time slots that represent the comfort period of apartment $a$ as a member of a coalition. $H_a = \{t \in K \mid CST_a \leq t \leq CET_a\}$ |
| $\Delta D_a$ | The discomfort of apartment $a$. Refer to equation (5.10) |
| $\Lambda$ | A set of congested time slots. |
| $\Lambda_\emptyset$ | The set of congested time slots that is found by optimising all apartments independently. |
| $\Lambda_C$ | The set of congested time slots that is given by re-optimising the most flexible apartment in $C$. |

| | |
|---|---|
| $\eta_m^{\Lambda_{C\setminus m}}$ | The vector of re-optimised cooling actions of apartment $m$ that is found based on $\Lambda_{C\setminus m}$. $m$ is the most flexible apartment in $C$. |
| $l_m^{\Lambda_{C\setminus m}}$ | The vector of best-found cooling load of apartment $m$ that is found using equation (5.4) and $\eta_m^{\Lambda_{C\setminus m}}$. |
| $l'_C$ | The vector of best-found aggregate cooling load of a non-empty coalition $C$. Refer to equation (5.12). |
| $l''_a$ | The vector of best-found cooling actions of apartment $a$ when it optimises its load independently. |

# Chapter 1

# Introduction

Von Neumann and Morgenstern (1944), in their seminal book *Theory of Games and Economic Behaviour*, sought to find a way for studying complex decisions that have uncertain outcomes in strategic environments. They referred to such an environment as a *game*, and suggested that the choices that a "player" (decision maker) is faced with be represented by a single number. They called this number the expected utility, which would be determined by a real-valued function called the utility function. This way, the behaviour of a rational player that seeks to make the best decision could be formulated as a problem of selecting the choice that maximises the expected utility, reducing the complexity of decision making to a numerical problem.

This foundational model, initially described for two-person games, was also considered by Von Neumann and Morgenstern in a class of $n$-person games, called *characteristic function form*, in which players could form coalitions by enforcing agreements among them. This game was defined by a set of players, $A$, and a characteristic function, $v$, assigning any coalition (a subset of $A$) a real number, called its worth or value. The characteristic function can be thought of as a utility function that expresses how much utility would be available to a coalition as a whole. This model was founded upon three basic assumptions, on account of which, the games in characteristic function form are now also called transferable utility (TU) games. These assumptions are as follows:

1. Utilities are exchangeable and fully trasnferable among the players.
2. The worth of a coalition is independent of the rest of the players in $A$.
3. A coalition can, without a loss, divide its worth among its members in a way that is agreed to by all of them.

Von Neumann and Morgenstern argued that for a TU game to be economical, the worths should be *superadditive*. That is, for any two disjoint coalitions, the worth of their union must be at least equal to the sum of the worth of the individual coalitions. The implication of superadditivity is that the *grand coalition* $A$ would be the only coalition that forms. However, while in two-person games making a simple evaluation of the game using utilities would be easy for each player, in $n$-person games, players could not evaluate the game unless they knew how much utility would be allocated to them. Von Neumann and Morgenstern proposed a solution for this problem, now known as the *stable set*, which is a set of allocations with certain properties that ensure players would not have an incentive to leave the coalition. However, the stable set is difficult to find, may not exist (Lucas, 1969), and furthermore, if it does exist, it is usually not unique (Lucas, 1992). Despite Von Neumann and Morgenstern's effort to reduce the complexity of decision making, the complexity and multiplicity of the solutions proved difficult for players to evaluate a game.

Lloyd Shapley, a co-winner of the Nobel Memorial Prize in Economic Sciences in 2012, proposed a way to enable individual players in an $n$-person game to evaluate "the prospect of having to play a game" (Shapley, 1953). Building on the success of Von Neumann and Morgenstern in simplifying the complexity of decision making by representing choices as a single number, he defined a function, which later became known as the *Shapley value*, that assigns a single number to each player of a game in characteristic function form. Shapley considered three axioms for this function. The first axiom, *symmetry*, states that if two players in the game have identical effects on the worths of coalitions, then their values are equal. The second axiom, *efficiency*, states that the full yield of the game, i.e., $v(A)$, is distributed among the players. Recognising the importance of evaluating games that might be interdependent, Shapley defined the third axiom, *the law of aggregation (additivity)*, to be a requirement that when two independent games are combined, their values must be added player by player. Remarkably, these three axioms characterise a unique value for each player, simplifying the evaluation of a game.

Shapley showed that the only function that satisfies the axioms is one that allocates to a player, $a$, an average of terms of the form $v(C \cup \{a\}) - v(C)$ over all coalitions that could possibly form without that player, i.e., $C \subseteq A \backslash \{a\}$. The aforementioned term is the *marginal contribution* that the player $a$ makes to the worth of the coalition $C \cup \{a\}$. Consequently, the value of each player would depend on how much it contributes to other players. An alternative interpretation of an $n$-person game is that a coalition of $n$ players can form in $n!$ possible ways, considering the different orders in which players can join the coalition. In each joining order, a given player makes a marginal contribution to the players joined before him, which can be seen as $n!$ possible events that occur with equal probabilities. The Shapley value of a player would be the expected value of a uniformly distributed random variable whose values are the marginal contributions of the player, each with $1/n!$ probability. As such, the Shapley value has also been interpreted as the *expected marginal contribution* of players (Young, 1988). Moreover, Shapley noted that the average marginal contribution function ensures that if a player, $a$, does not contribute to the worth of any coalition, it will be allocated only $v(\{a\})$. These properties along with Shapley's main axioms have been widely regarded as capturing the notion of fairness in distributing the worth of a grand coalition (Young, 1988; Moulin, 1992).

To date, a large literature has grown out of Shapley's original paper. The Shapley value has been applied to many interesting problems. Some of these applications include cost sharing problems such as airport runways (Littlechild and Owen, 1973) and railways infrastructure (Fragnelli et al., 2000), factorizing the risk of diseases (Land and Gefeller, 2000), multicast transmission (Feigenbaum et al., 2001), analysing customer satisfaction (Conklin et al., 2004), resolving political conflicts (Engelbrecht and Vos, 2011), and also identifying key members in terrorists networks (Lindelauf et al., 2013; Michalak et al., 2013). Undoubtedly, such a wide range of applications points to the importance of the Shapley value. However, the computational cost resulting from its combinatorial nature hinders its applications in practice.

Generally, the computational issues of the Shapley value is either or both of the following:

1. The average marginal contribution function requires calculation of marginal contributions of a player to an exponential number of coalitions– exactly, $2^{n-1}$ coalitions per player.
2. Computing the worth of coalitions using the characteristic function, $v$, requires a great computational effort.

In terms of the first issue, with a typical personal computer, calculating the Shapley value in a typical game with a few dozens of players could take several days. Since the time complexity is exponential in the number of players, for each additional player it takes twice as long. By contrast, in the second issue, the number of players is not problematic, but computing the worth of a coalition is computationally intensive (e.g., it involves a hard optimisation problem). Therefore, even when the number of players is relatively small, calculating the average of marginal contributions of a player can take very long.

Currently, three lines of research exist on addressing the computational issues of the Shapley value. Nevertheless, they all deal only with the first issue. A line of research has focused on computing the Shapley value using alternative representation formalisms (for example Ieong and Shoham (2006); Aadithya et al. (2011)), as opposed to the standard representation where the marginal contribution of a player $a$ to a coalition $C$ is simply obtained by calculating $v(C \cup \{a\}) - v(C)$. Moreover, another line has focused on certain, restricted classes of games for which some efficient exact algorithms have been developed (for example Ando (2012); Deng and Papadimitriou (1994a)). Other researchers have focused on some restricted classes of games, for which they have proposed bounded approximate solutions (for example, Owen (1972); Bachrach et al. (2008); Fatima et al. (2008)).

These approaches, however, suffer from major limitations (see Section 2.2.1 for more details). While alternative representation formalisms can, in certain circumstances, result in computational savings, using them in practice may require an effort that could well outweigh the potential benefits (see Section 2.2.1 in Chapter 2 for more details). Furthermore, the exact methods that are specific to restricted class of games are either impossible or difficult to be used in other classes. Among the existing approximate solutions there are three works, due to Castro et al. (2009), Bachrach et al. (2008) and Liben-Nowell et al. (2012), that directly, or with some extensions, can be applied to all superadditive games. These approximations exploit the fact that the Shapley value of a player is the expected value of the population of the player's marginal contributions, and use *simple random sampling*, a somewhat inefficient sampling technique, to approximate the Shapley value of players. However, they use different methods to bound the approximation error. Castro et al. (2009) use the Central Limit Theorem (CLT), which is an asymptotic bound. That is, the bound holds only when the sample size increases to infinity. Bachrach et al. (2008) use a powerful inequality from probability theory, known as Hoeffding's inequality, to bound the error in *simple* games, where the worth of a coalition is either zero or one. As Chapter 4 will show, this inequality can also be used to bound the error in other superadditive games. Liben-Nowell et al. (2012) use a similar inequality, known as Chebyshev's inequality, for bounding the error in an important and large subclass of superadditive games, called *supermodular* games. However, as Chapter 2 highlights, this bound is mathematically erroneous, and its time complexity is polynomial in the number of players in the game, which is not desirable.

Given the importance of the Shapley value, the shortcomings of the existing works in addressing the first computational issue, and lack of research on the second one, this thesis aims to fill the need for addressing these issues, with the objectives that are stated in the next section.

## 1.1 Research Objectives

It is important to note that, unless some assumption is made about the characteristic function of a given game, an exponential time complexity in calculating the exact Shapley value cannot be avoided. Therefore, any solution that would alleviate this problem has to approximate the Shapley value in some way. With this point in mind, our first objective is to develop a method for approximating the Shapley value. Where possible, this method should exploit the structure of the games, so as to improve upon the inefficiency of the existing approximation methods. Crucially, this method must also establish a worst-case bound on the potential error. That is, it should provide some form of guarantee that, in the worst case, the difference between the exact and the approximate Shapley value is not greater than a certain amount. Moreover, the approximation method should allow for aribitrarily small error bounds. Lastly, since in calculating the Shapley value the amount of computation grows with the number of players, the approximation method and the quality of the bound should have the least dependency on the number of players in a game.

The second objective is to investigate how the Shapley value, as a fair allocation, can be used in games whose characteristic function has a high computational complexity. Clearly, such a characteristic function would add an additional layer of complexity to the already difficult problem of calculating the Shapley value. Therefore, in order to address the issue that is due to the characteristic function, one should isolate the main complexity of calculating the Shapley value. This can be done by assuming that the number of players is such that calculating the Shapley value using a simpler characteristic function would be tractable.

Furthermore, without making assumptions about the structure of the characteristic function, in order to overcome the computational challenge, one can only approximate the worth of coalitions. By doing so, calculating the Shapley value would become straightforward. However, since the players would be deprived of their exact Shapley value as a consequence of the complexity of the characteristic function, to maintain fairness, one should also examine whether players could be responsible for this complexity. The following cases can be imagined regarding the players in any given coalition:

1. The complexity of the characteristic function is independent of the players.
2. All players in the coalition equally contribute to the complexity of the characteristic function.
3. Each player in the coalition contributes differently to the complexity of the characteristic function.

Fairness commands that if a player contributes to the complexity, his allocation should reflect this. However, the Shapley value is not directly concerned with this issue. Therefore, our objective is to find a way to incorporate the players' contribution to the complexity of the characteristic function in their Shapley values without violating Shapley's axioms. Note that, of the three cases mentioned above, the most challenging is when players do contribute to the complexity, but not equally. In fact, it can be argued that this is a general case of the other two, and thus, we aim to address this case which is sufficient to cover all possibilities.

Our third objective is to evaluate the results achieved from the first two objectives using games that model real world problems. As mentioned earlier, the Shapley value has many interesting applications. However, to demonstrate the usefulness of our results, we need settings where there is a large number of players, or the complexity of the characteristic function is high. For this purpose, the smart grid is a perfect domain, in which both of the aforementioned problems are frequently seen. The next section explains this in more detail.

## 1.2 The Smart Grid Application

According to Boyle (2004), over the course of the 20th century, the world's total consumption of all forms of primary energy (natural energy resources such as oil) has risen more than tenfold. Most of the burden of meeting this growing demand is on the power production industry, which is lagging behind other industries in terms of benefiting from computer technologies. For instance, utility companies still rely on sending workers out to read meters and collect data. Such inefficiencies that exist in many parts of the current grids, drive the need for a "smart grid", which is referred to a highly automated grid that takes great advantage of sensors, modern computers and communication technologies. The United States Department of Energy (2003) has outlined the purpose of the smart grid to be using information and communications technology to improve the efficiency, reliability, economics, and sustainability of producing power.

In parallel, there has been a global concern about the sustainability of relying on fossil fuels which constitute the main source of power production in most countries. In 2003, it was estimated that the world's oil reserves would last for approximately 40 years, natural gas for 60 years, and coal for around 200 years (Boyle, 2004). These resources will likely become increasingly expensive and more difficult to extract. Alarmed by the dwindling fossil fuel resources, and the environmental concerns associated with them, many governments have set targets to multiply their support for generating power using renewable resources such as wind. Therefore, renewable resources will have a large role in the smart grid. With the improved communication between various components of the grid, the smart grid will also pave the way for integration of *distributed energy resources (DERs)*, which are small to medium capacity renewable power sources that are scattered across a region. However, the power generated from renewable sources, and particularly DERs, is inherently variable, intermittent, and difficult to predict.

The two major proposals for addressing the issues with integration of DERs are the formation of *virtual power plants (VPPs)*, and *demand response (DR)*, which are concerned with managing supply and demand, respectively. VPPs aggregate the output of DERs to lower the overall uncertainty of generation, and improve reliability. Demand response allows utility companies to incentives consumers to shift their loads to match demand to the available supply, and also flatten peaks. An effective demand response minimises the need for fossil fuel power plants to meet spikes in demand.

A natural model for analysing the economics of DERs selling their aggregated power is an $n$-person game called the *newsvendor* game (Muller et al., 2002). This game models a number of retailers that want to stock a perishable commodity whose sale faces a random demand. If the demand is underestimated, the profit will not be maximised, and if it is overestimated, there will be a loss since the commodity perishes. By forming a coalition and aggregating their demand the retailers can reduce the uncertainty of their demand and increase their expected profit. The DERs in a VPP face a similar problem. Since electricity is perishable, and due to the special structure of electricity markets that require selling power ahead of generation, DERs can make binding agreements to aggregate their loads to reduce their overall uncertainty, and consequently increase their expected profit. Individual DERs can evaluate the prospects of playing this game by calculating their Shapley value, which would also reveal their fair share of the total profit. Nevertheless, the potentially large number of DERs in such a game is a major obstacle to applying the Shapley value. This issue can be mitigated by approximating the Shapley value as described in the first objective.

Similarly, demand response scenarios can be modelled by appropriate $n$-person games. While demand response typically involves interaction between utility companies and end users, some demand response programs can particularly depend on interaction among consumers, for instance, to coordinate their individual loads so as to flatten their total load. An optimal coordination among $n$ consumers can be highly complex, and due to strict preferences, lack of flexibility, etc. some consumers can have a larger role in the complexity. Thus, the characteristic function needed to represent the worth of coordinations can require intensive computation. Again, applying the Shapley value to find the consumers' individual contributions can be considerably challenging. The second objective of this thesis is to address such problems.

## 1.3 Research Contributions

This thesis advances the state of the art of the research on the Shapley value and cooperative games in the following ways:

1. We put forward a method for determining the minimum sample size required to approximate the Shapley value using sampling with an arbitrary confidence and error. The resulting bound, which is based on Hoeffding's and Chebyshev's inequalities, is superior to the existing ones in several respects.

    (a) Unlike the CLT-based bound provided by Castro et al. (2009) which is asymptotic, our bound holds with a finite sample size, and thus, truely bounds the error.

(b) It is a generalisation of the bound suggested by Bachrach et al. (2008) to superadditive games, and improves upon it by requiring smaller sample sizes for a certain range of confidence values.

(c) It corrects the erroneous bound put forward by Liben-Nowell et al. (2012), and unlike this bound which has a polynomial time complexity in the number of players, it does not depend on the number of players at all.

2. Instead of simple random sampling, we propose the use of a more efficient sampling technique, called *stratified sampling*, that stratifies the population of a player's marginal contributions into smaller populations (*strata*) so as to sample each stratum proportional to its impact on the overall average. Doing so can potentially result in a significantly improved accuracy as compared to simple random sampling which samples the population blindly.

3. We show that the standard error of approximating the Shapley value using stratified sampling is always lower than or equal to that of simple random sampling. Here, the standard error is the standard deviation of the approximated Shapley value which could vary with each new sample. This variation shows how accurate the approximation is.

4. We propose two stratification methods, namely *branching stratification* and *size-based stratification*.

   (a) Branching stratification divides the population using a given set of players $B \subseteq A$. For each player, $b$, in $B$, the population is divided into two sub-populations depending on whether the corresponding coalition of a marginal contribution includes $b$ or exclude it. This would result in $2^{|B|}$ strata.

   (b) Size-based stratification divides the population such that each stratum contains marginal contributions of the player to coalitions of equal size. This always results in $n$ strata.

5. Given the minimum sample size required to establish a desired error bound, we find the optimal number of marginal contributions that must be evaluated from each stratum such that the standard error is minimised. This requires that the range (or more preferably the variance) of each stratum is given. The total sample size would then be distributed among the strata proportional to their ranges. Therefore, a stratum whose values have higher ranges would be sampled more, leading to a smaller standard error and a more accurate approximation. Since finding the range of the strata requires knowledge about the population, we need to focus on more specific classes of superadditive games to exploit their properties. To this end, we show that supermodular games have a certain property that allows us to find the range of each stratum. Furthermore, we introduce a class of games that have a property that we call *order-reflecting*, using which we

can find the range of the strata in size-based stratification. Most importantly, we show that some games that have real world applications exhibit this property. These games include newsvendor games, *output-sharing* games (Moulin, 1992), and *airport* games (Littlechild and Owen, 1973). The order-reflecting property and its application in approximating the Shapley value of newsvendor games have appeared in our publication (Ramchurn et al., 2013).

6. We experimentally evaluate our stratified sampling methods on instances of newsvendor games (using real data), output-sharing, and airport games. Each of these instances consists of at least 100 agents. In these experiments, the average standard error of approximating the Shapley value of all players using stratified sampling is 48% lower than that of SRS.

7. We present an efficient implementation of the Shapley value, which, unlike the standard formula of the Shapley value, requires that the worth of each coalition be computed only once. Therefore, when the time complexity of the characteristic function of the game is high, a significant amount of redundant computation can be avoided.

8. To address the problem of players having different contributions to the complexity of a characteristic function, we model players as being *computationally bounded rational*, which means that their ability to make the best decisions is limited by their computational resources. The idea of bounded rationality in $n$-person games was originally put forward by Sandholm and Lesser (1997), who investigated the stability of coalitions of bounded rational players. In our model, the approximated worth of a coalition is regarded as its bounded rational worth, while the actual worth, the computation of which is intractable, would be the value that is given by the characteristic function. We show that by calculating the Shapley value using the approximated worth of coalitions, due to the additivity axiom of the Shapley value, the allocations are fair in the following sense: all players are rewarded for their contribution to the bounded rational worth of the grand coalition, and simultaneously penalised for their contribution to the discrepancy between the rational and bounded rational worth of coalitions.

9. We develop a demand response program where a block of apartments receive a discounted price of electricity if they coordinate their cooling loads such that at any point in time throughout the day their total load is below a certain threshold. The problem that the apartments face is dividing the total cost in a fair way. We model this problem as an $n$-person game, and use the Shapley value to obtain a fair allocation. However, since finding an optimal coordination of the air conditioning in all apartments would involve an intractable optimisation, we propose two algorithms for finding suboptimal coordinations, using which we can calculate the Shapley values readily. We also propose a dynamic programming algorithm that

exploits the recursive nature of the said algorithms to further speed up calculating the Shapley value.

## 1.4 Thesis Structure

In the remainder of this thesis, we give an overview of related concepts in cooperative game theory and their applications in the smart grid, and then present our approximation methods, as well as our approach regarding the Shapley value of bounded rational players. In more detail, the remaining chapters are structured as follows:

In Chapter 2, we begin by reviewing definitions and theorems from cooperative game theory. We then focus on the Shapley value and state its properties. This will be followed by an overview of the works in the literature that deal with efficient computation of the Shapley value. Next, we examine the existing work on bounded rationality, and in particular a model of bounded rationality that has been proposed for coalitions of bounded rational agents.

In Chapter 3, we give an overview of the related work on applications of cooperative game theory in the smart grid. This is followed by a description of electricity markets, and particularly the wholesale market of the Great Britain as one of the globally predominent market models. Inspired by this market, we present a formal model of a two settlement market, based on which we define the problem of selling the output of a DER and derive its expected profit. Next, we define a game where DERs form a coalition to sell their output in the market. Lastly, we examine the recent applications of cooperative game theory in demand response.

In Chapter 3, we propose an error bound for approximating the Shapley value using simple random sampling. Next, we describe stratified sampling in the context of approximating the Shapley value. This is followed by the description of branching and size-based stratification methods. Finally, we present the experimental evaluation of our approach.

In Chapter 5, we present a bounded rationality model for coalitions, and discuss why using the Shapley value based on this model results in a fair allocation. This discussion is followed by presenting the efficient implementation of the Shapley value. Next, we formally define the demand response program, and describe our algorithms for coordinating the cooling loads of apartments. We then show our experimental results of applying our bounded rationality proposition regarding the Shapley value to the demand response.

Finally, Chapter 6 concludes this thesis, summarises the contributions and limitations of our methods, and also outlines the direction for future work.

# Chapter 2

# Background

This chapter provides the theoretical background and related work to this thesis. It begins by reviewing definitions and theorems from cooperative game theory that will be referred to throughout the thesis. It then moves on to the Shapley value and state its properties in Section 2.2. This will be followed by an overview of the works in the literature aiming to address the computational issues of the Shapley value. In particular, Section 2.2 gives a detailed description of the methods that use random sampling to approximate the Shapley value. The limitations of these methods provide the basis for the approach of this thesis in approximating the Shapley value, which will be discussed in Chapter 4. Section 2.3 examines the existing work on bounded rationality, and in particular a model of bounded rationality that has been proposed for coalitions of bounded rational agents. Chapter 5 builds upon this model to provide a method for fairly dividing the value of a coalition using the Shapley value. Finally, Section 2.4 summarises the chapter.

## 2.1 Cooperative Game Theory Definitions

A game is specified by the set of decision makers called players or *agents*[1], the set of all possible actions that the agents can take, and a *utility function* that associates any action to a payoff. The objective of a rational agent in a game is to take actions, or make decisions, that would maximise its own payoff. A game is called *cooperative* or *coalitional* when agents can coalesce (form coalitions) so as to achieve higher utilities than what they would achieve otherwise (Neumann and Morgenstern, 1944; Kahan and Rapoport, 1984; Osborne and Rubinstein, 1994; Peleg and Sudhölter, 2007). In what follows, some related concepts from cooperative game theory are reviewed.

[1] As is the convention among artificial intelligence researchers, we henceforth use the term agent.

Formally, given a set, $A$, of $n$ agents, a *coalition*, $C$, is defined as a subset of $A$. When $C$ consists of all agents in $A$, it is called the *grand coalition*. The worth or value of a coalition is expressed by a *characteristic function* $v$ (also known as *valuation function*), which maps each subset of $A$ to a real number, i.e., $v : 2^A \rightarrow \mathbb{R}$. In this thesis, the focus is on transferable utility (TU) games in which utilities can be losslessly transferred from one agent to another (Myerson, 1991). We also refer to these games as *cooperative games*, and specify them using the pair $(A, v)$.

**Definition 2.1.** A **coalition structure** is a partition of $A$, in which each agent belongs to exactly one coalition $C_i$, with some coalitions possibly being singletons (Larson and Sandholm, 2000). A coalition structure, $CS = \{C_1, \ldots, C_m\}$, as a set, satisfies two conditions: $\bigcup_{i=1}^{m} C_i = A$ and $i \neq j \Rightarrow C_i \cap C_j = \emptyset$. An optimal coalition is one that maximises $\sum_{i=1}^{m} v(C_i)$.

**Definition 2.2.** The **marginal contribution** of an agent $a$ to a coalition, $C \subseteq A \setminus a$, is the difference in $C$'s value that is achieved by adding $a$ to it, i.e., $v(C \cup \{a\}) - v(C)$. The marginal contribution of $a$ to $C$ will be denoted as $MC(a, C)$.

**Definition 2.3.** A game is **additive** when the value of a coalition achieved by merging two disjoint coalitions is equal to the sum of the values of those coalitions individually, i.e., $\forall C_1, C_2 \subseteq A \; C_1 \cap C_2 = \emptyset \quad v(C_1 \cup C_2) = v(C_1) + v(C_2)$

**Definition 2.4.** A game is **superadditive** when the value of a coalition achieved by merging two separate coalitions is at least equal to the sum of the values of those coalitions individually, i.e., $\forall C_1, C_2 \subseteq A \; s.t. \; C_1 \cap C_2 = \emptyset \; v(C_1 \cup C_2) \geq v(C_1) + v(C_2)$.

*Remark* 2.5. Unlike additive games, in superadditive games, the larger a coalition is, the higher its value is. Therefore, when a characteristic function is superadditive, the coalition structure that contains the grand coalition is always optimal.

**Definition 2.6.** A characteristic function is **supermodular** if it meets the following condition for any $C_1 \subseteq C_2 \subseteq A$:

$$\forall a \in A, \; s.t. \; a \notin C_1, C_2 \; : \quad v(C_1 \cup \{a\}) - v(C_1) \leq v(C_2 \cup \{a\}) - v(C_2).$$

The above supermodularity condition is equivalent to:

$$\forall C_1, C_1 \subseteq A \;\; v(C_1) + v(C_2) \leq v(C_1 \cup C_2) + v(C_1 \cap C_2).$$

*Remark* 2.7. Supermodularity is a special case of superadditivity, and means that given a coalition, any agent that is not its member makes a greater or equal marginal contribution to it than to any of its subsets.

**Definition 2.8.** A game with a supermodular characteristic function is called supermodular or convex.

**Definition 2.9.** A **simple game** is one in which the values of coalitions are either 0 or 1, i.e., $\forall C \subseteq A; v(C) \in \{0, 1\}$.

**Definition 2.10.** A **map** is a function $\chi : 2^A \rightarrow [0, 1]$ that assigns a weight to each coalition.

**Definition 2.11.** A **balanced map** is a map such that for any $a \in A$ it holds that $\sum_{C \in A} \chi(C)\mathbf{1}\{a \in C\} = 1$, where $\mathbf{1}\{.\}$ is either 0 or 1, indicating whether or not $a$ is a member of $C$.

**Definition 2.12.** A **balanced game** is one that, if for any balanced map $\chi$ the following inequality holds: $\sum_{C \in A} \chi(C)\, v(C) \leq v(A)$.

### 2.1.1 Solution Concepts

We now introduce a number of concepts that are related to how much of the value of a coalition is allocated to each agent.

**Definition 2.13.** Given a coalition $C$ of size $k$, a *payoff distribution* or a *solution concept* is a vector $x = [x_1 \ldots x_k]$, where $x_i$ represents how much of the value of $C$ should be allocated to the $i$th member.

**Definition 2.14.** The **excess** $e$ of a payoff distribution, $x$, is the remaining value of a coalition after allocating the shares of its members, i.e. $e(x, C) = v(C) - \sum_{i=1}^{k} x_i$.

*Remark* 2.15. The notion of excess can be used as a measure of *dissatisfaction* of a coalition, since if there is a positive excess it shows that some of the value of the coalition has not been allocated to the members. Based on this, the dissatisfaction about a payoff distribution can be evaluated using a vector of excesses over all coalitions.

**Definition 2.16.** A payoff distribution is **efficient** if and only if it has a zero excess. In other words, an efficient payoff distribution allocates the whole value of the coalition to the agents, i.e., $v(C) = \sum_{i=1}^{k} x_i$.

**Definition 2.17.** A payoff distribution is said to be **individually rational** whereby if all agents in $A$ can obtain a payoff that is at least equal to their payoff when they are alone in a coalition, i.e., $\forall a \in A \;\; x_a \geq v(\{a\})$.

**Definition 2.18.** A payoff distribution is said to be **group rational** if for every coalition $C \subseteq A$ the sum of its members' payoffs is at least equal to the value of the coalition, i.e., $\forall C \subseteq A \;\; \sum_{i=1}^{k} x_i \geq v(C)$.

**Definition 2.19.** A payoff distribution that is both efficient and individually rational is called an **imputation**.

A key concern in game theory is to understand what outcomes playing a game would have. To this end, different solution concepts have been proposed. Some solution concepts enable agents to determine whether or not a certain payoff distribution can ensure a stable coalition of agents. Below, a number of such concepts are reviewed.

**Definition 2.20.** The **core** is the set of imputations that are group rational. A payoff distribution $x$ is said to be in the core of a game $(A, v)$, if and only if it has the said properties. More formally,

$$\mathcal{C}(A, v) = \{x \mid \sum_{i=1}^{n} x_i = v(A) \ \wedge \ \forall C \subseteq A, \sum_{i=1}^{k} x_i \geq v(C)\}$$

*Remark* 2.21. A payoff distribution is in the core of a game when no set of agents, including the singletons, can achieve more by forming a different coalition. As such, no agent should have an interest in rejecting a payoff distribution that is in the core.

**Theorem 2.22.** ***Bondareva–Shapley theorem*** *states that a game's core is non-empty if and only if the game is balanced (Bondareva, 1963; Shapley, 1967).*

Note that determining the non-emptiness of the core using the original definition is in general an NP-hard problem (Deng and Papadimitriou, 1994b). The Bondareva–Shapley theorem offers an alternative approach to establish that by checking whether a game is balanced.

**Definition 2.23.** It is said that a vector $x \in \mathbb{R}^m$ is greater or equal to $y \in \mathbb{R}^m$ in the **lexicographic ordering**, denoted by $x \geq_{lex} y$, if and only if $x$ and $y$ are equal or

$$\exists t \ \ s.t. \ \ 1 \leq t \leq m \ \ s.t. \ \ \forall i \ \ s.t. \ \ 1 \leq i < t \ \ x_i = y_i \ \ \wedge \ \ x_t > y_t$$

**Definition 2.24.** A payoff distribution is said to be in the **nucleolus** (Schmeidler, 1969) if its corresponding vector of excess is minimal, meaning that this vector is preferable over all other excess vectors, due to being the smallest in the lexicographical ordering.

*Remark* 2.25. In general, the nucleolus has at most one element. If the set of imputations in a game is non-empty, then the nucleolus is also non-empty (Airiau, 2013).

**Theorem 2.26.** *If a game has a non-empty core, the nucleolus is always in the core.*

**Corollary 2.27.** *From the above theorem it follows that the nucleolus is a special subset of the core. This is because it minimises the vector of excesses over all coalitions, and thus has the lowest dissatisfaction, in lexicographic order, among all payoff distributions in the core.*

## 2.2 The Shapley Value

In the previous section, solution concepts that are concerned with the effect of payoff distributions on the stability of coalitions were examined. Here, the Shapley value (Shapley, 1953) is examined, which is a solution concept that offers a "value" for playing a game that also has some fairness properties. The fairness is due to the following axioms:

- **Symmetry**: Any two agents with equal marginal contributions to all coalitions, receive the same payoff. Formally, for any $a_1, a_2 \in A$, and for all $C \subseteq A$: $v(C \cup \{a_1\}) = v(C \cup \{a_2\})$, then $x_1 = x_2$, where $x_1$ and $x_2$ are the payoffs of $a_1$ and $a_2$, respectively.
- **Efficiency**: The value of the grand coalition is fully divided among the agents, i.e., $v(A) = \sum_{i=1}^{n} x_i$.
- **Additivity**: For any two TU games $(A, v)$ and $(A, w)$ with payoff distributions $x$ and $y$, respectively given by the Shapley value; the payoff distribution of the game $(A, v + w)$ is $x + y$.

In addition to the above axioms, the Shapley value exhibits the following property which is commonly known as the *dummy player* property: if the marginal contribution of an agent, $a$, to all coalitions is zero, then it only receives $v(\{a\})$ as payoff. This property along with Shapley's axioms provide a way of dividing the value of the grand coalition that is fair.

**Theorem 2.28.** *The Shapley value is the only value that satisfies the symmetry, efficiency and additivity axioms (Shapley, 1953).*

An advantage of using the Shapley value compared to other solution concepts is its uniqueness (Theorem 2.28). For instance, the core can contain multiple payoff distributions, which raises the question of which one should be picked. By contrast, the Shapley value offers a single payoff distribution, leaving only one option.

**Theorem 2.29.** *The Shapley value in superadditive games is always individually rational (Shapley, 1953).*

**Theorem 2.30.** *The Shapley value always exists, but is not necessarily in the core.*

**Theorem 2.31.** *In supermodular games, the Shapley value is the barycentre of the core (Shapley, 1971).*

In addition to the axiomatic approach, Shapley offered an alternative perspective to the value that he proposed. He noted that a coalition of $n$ agents can form in $n!$ ways, considering all orders that the agents can join the coalition. In each order, as a player steps in the coalition, it makes a marginal contribution to the agents joined before it. The Shapley value of each agent is the average of all marginal contributions that it makes to other coalitions. More formally, denote by $\pi(A)$ the set of all permutations of agents in $A$, each of which representing a distinct joining order. Furthermore, denote by $Pre_a^{\mathcal{O}}$ the set of agents that precede $a$ in the permutation $\mathcal{O}$. The Shapley value of $a$, denoted as $SV(a)$, is given by:

$$SV(a) = \frac{1}{n!} \sum_{\mathcal{O} \in \pi(A)} \left[ v(Pre_a^{\mathcal{O}} \cup \{a\}) - v(Pre_a^{\mathcal{O}}) \right] \tag{2.1}$$

Calculating the Shapley value using the above expression is of order $O(n!)$ and is also inefficient, as it leads to a significant amount of redundant computation. This occurs in computing the marginal contributions where the coalition values are computed based on the joining orders. Since the value of a coalition is independent of its members joining it, any marginal contribution that corresponds to the same coalition with different orders would have exactly the same value. To avoid calculating a marginal contribution more than once, Shapley introduced a factor to account for the multiplicities of each unique marginal contribution value. The improved formula is equivalent to a weighted average of marginal contributions which is given as:

$$SV(a) = \sum_{C \subseteq A \setminus \{i\}} \frac{|C|!\,(n - |C| - 1)!}{n!} \left( v(C \cup \{i\}) - v(C) \right) \tag{2.2}$$

The next section examines the existing approaches using which the Shapley value is computed.

### 2.2.1 Computation of the Shapley Value

Although equation (2.2), as compared to equation (2.1), reduces the time complexity of computing the Shapley value from $O(n!)$ to $O(2^n)$, it still requires intensive computation, and is generally intractable in games with even a few dozen agents. Several authors have put forward methods to mitigate this issue, which in general fall into the following three categories: (i) using alternative representation formalisms to compute the exact Shapley value efficiently, (ii) developing exact efficient methods for some specific classes of game, (iii) bounded approximate solutions. Below, an overview of each category is given, and approximate solutions are expaned on which are most related to this thesis.

#### 2.2.1.1 Using Alternative Representation Formalisms

In the characteristic function form of games, each coalition is assigned a value. Since there are exponentially many coalitions (i.e., subsets of $A$), it possibly requires an exponentially large memory space to even describe the coalition values, and in turn, the entire game. Some alternative representation formalisms have been proposed to address this issue. Accordingly, a line of research has focused on computing the Shapley value using these representations.

Conitzer and Sandholm (2004) developed a concise representation of characteristic functions for games where the agents' objective is to address a number of independent issues. Here, coalitions are modeled as groups of agents that together can address a number of distinct independent issues. Moreover, the characteristic function of the game, $v$, is represented as a vector of characteristic functions $(v_1, v_2, \ldots, v_T)$ with each $v_i$ being a decomposition $2^A \to \mathbb{R}$ over $T$ issues, such that for any $C \subseteq A, v(C) = \sum_{i=1}^{T} v_i(C)$. Using this representation Conitzer and Sandholm show that to calculate the Shapley value, one would need to average over the marginal contributions in permutations of *only a subset* of the agents.

Ieong and Shoham (2006) extended and generalised Conitzer and Sandholm's representation by developing *marginal contribution networks* (MC-nets) which are exponentially more concise. Under this representation, the characteristic function is decomposed into a set of rules that assign marginal contributions to groups of agents in the form of $Pattern \to value$ (e.g., $\{a_1 \wedge a_2\} \to 5, \{a_2\} \to 2$). If a group of agents *meet* the requirements of the *Pattern* in a rule, then it is said that the rule *applies* to each of those agents. The value of any given coalition is then the sum over the values of all rules that apply to the members of the coalition. The MC-nets is a fully expressive representation, in that it can represent an arbitrary cooperative game in characteristic function form. Furthermore, considering each rule as a separate game, the Shapley value of an agent is the sum over the Shapley value of the agent in each rule. The total computational time complexity of this computation is linear in the size of the input, since the time required to compute the Shapley value of an agent in any given rule is linear in the pattern of the rule. A class of MC-nets, called *read-once MC-nets*, were identified by Elkind et al. (2008) that are more compact than the original form of MC-nets. Under this representation, it is possible to compute the Shapley value in polynomial time.

Aadithya et al. (2010) developed another concise and fully expressive representation, called *Algebraic Decision Diagrams*, that is "highly optimised" for ordered decision trees on boolean decision variables. The conciseness is achieved by merging identical copies of duplicated subtrees, resulting in a potentially much smaller equivalent graph. The authors claimed that, under this representation, the Shapley value can be computed in polynomial in the size of the algebraic decision diagram.

Although using these representations can be useful in some circumstances, they have limitations of their own. For instance, the computational advantage of computing the Shapley value using MC-nets is limited to a special case, where the patterns in all rules are required to be conjunctions of literals (Elkind et al., 2008). More importantly, it is not clear that, in general, how efficiently an existing game with the standard characteristic function form can be be transformed to these representations. It is possible that, in the worst case, the computational cost required to do so would outweigh the potential benefit.

#### 2.2.1.2 Exact Efficient Methods For Specific Games

A number of researchers have proposed methods for computing the Shapley value that are specifically designed for some special classes of games. Some of these methods are as follows. Granot et al. (2002) developed a linear time method to compute the Shapley value in a cost allocation problem, called the extended tree game. Ando (2012) showed that it is #P-complete to compute the Shapley value of minimum cost spanning tree games, and identified a subset of these games in which the Shapley value can be computed in polynomial time. Moreover, Deng and Papadimitriou (1994a) proposed a method to compute the Shapley value in induced subgraph games with a time complexity of $O(n^2)$.

Since such methods have been developed for a restricted class of games, their applicability to other classes is either not possible or requires significant extension.

#### 2.2.1.3 Approximate Solutions

In another line of research, approximation methods have been proposed, again, for certain classes of games. In particular, voting games and their more general form, the $k$-majority voting games, have been the focus of most of the research on approximating the Shapley value. Mann and Shapley (1962) proposed a Monte Carlo method that approximates the Shapley value in voting games in linear time. However, the prominent approximation algorithm for voting games is the multi-linear extension method by Owen (1972), which was also able to calculate the exact Shapley value in a more efficient way than the direct enumeration (i.e., using brute-force). More recently, Fatima et al. (2008) put forward a different approximation approach for $k$-majority games, which improves upon the approximation error of Owen's multi-linear method. Furthermore, Bachrach et al. (2008) proposed a sampling-based algorithm for approximating the Shapley-Shubik and Banzhaf power indices, which can also be extended to the Shapley value. Castro et al. (2009) proposed a sampling-based algorithm for the case when the variance or the range of marginal contributions of an agent is known. Moreover, Liben-Nowell et al. (2012) proposed a sampling-based algorithm for supermodular games which establishes an error bound in time polynomial in the number of agents.

The vast majority of researches on approximating the Shapley value apply random sampling, which is a natural choice given the fact that the Shapley value of an agent is the mean of a special random variable, i.e., a uniformly distributed variable whose values are the marginal contributions of the agent. More specifically, forming a coalition of $n$ agents in different orders can be seen as $n!$ possible events that occur with equal probabilities. The Shapley value of an agent would then be the expected value of a random variable whose values are the marginal contributions of the agent, each with $1/n!$ probability. This view provides the intuition for sampling-based approaches.

The next subsection examines, in detail, three sampling-based approximation methods that are related to this thesis, namely, Castro et al. (2009), Liben-Nowell et al. (2012) and Bachrach et al. (2008), and state their limitations.

### 2.2.2 Sampling-Based Approximation of the Shapley Value

In statistics, estimating the mean of a set of numbers, called a *population*, is a very common task. This is typically done using a technique called random sampling, which is randomly choosing a subset of the population as a *sample*. When all elements in a population have the same probability of being chosen for a random sample, the process is called *simple random sampling (SRS)*. If the elements of the population are allowed to be chosen more than once, then sampling is called *with replacement*, otherwise it is *without replacement*. Often when the population is relatively small, sampling is done without replacement so as to obtain a more representative sample. Conversely, in large populations, sampling with replacement is typically used, and is approximately the same as without replacement under SRS. This is because the probability that an element in a large population is chosen more than once is low.

Once a sample is chosen, several parameters of the population can be estimated. In particular, the average of the elements in the sample, which is known as the *sample mean*, can be a good estimate of the population (Levy and Lemeshow, 2008). Similarly, the *sample variance*, which is the variance of the elements in the sample, can be used as an estimate of the population variance. Note that since different samples drawn from the same population would result in different sample mean values, the sample mean itself has a distribution, called the *sampling distribution of the mean*, which has its own mean and variance. Therefore, samples can be seen as random variables, and when they are taken with replacement, they are *independent and identically distributed (i.i.d.)*. As will be seen later in this section, this fact provides the basis of using some theorems from the probability theory that are used in approximating the Shapley value. Henceforth, unless otherwise stated, by sampling it shall be meant sampling with replacement.

Now, given that the Shapley value is the mean of marginal contributions of an agent to all coalitions, the existing sampling-based methods approximate it using SRS. However, SRS is not necessarily an efficient sampling technique, since it samples the population blindly (see Chapter 4 for more details and proof).

Of crucial importance as in any approximation is the quality of the solution, which is measured by the *approximation error* (or simply error). That is, the difference between the approximated value and the exact value. However, since the exact error cannot be measured in practice, often a theoretical bound on the error is provided instead, which guarantees that the error is bounded by a certain value, denoted by $\epsilon \in \mathbb{R}_+$. In all sampling-based approximations of the Shapley value that exist in the literature, the bound is probabilistic, meaning that the error is guaranteed to be within the bound, however, with a certain probability. This probability is known as the *confidence*, and is typically expressed as $1-\delta$, where $\delta \in [0,1]$ represent the probability that the bound fails (i.e., the actual error is greater than $\epsilon$). Formally, let $1-\delta$ and $\epsilon$ be an arbitrary confidence, and error bound, respectively, the probabilistic error bound is expressed in either of the following equivalent forms:

$$\Pr(|SV(a) - \overline{SV}(a)| \leq \epsilon) \geq 1-\delta, \quad \text{or} \tag{2.3a}$$

$$\Pr(|SV(a) - \overline{SV}(a)| \geq \epsilon) \leq \delta \tag{2.3b}$$

where $\overline{SV}(a)$ is the sample mean, i.e., the approximate Shapley value of agent $a$. It is important to note that, for the above bounds to hold, the sample size needs to be sufficiently large. To determine how large the sample size should be, three different ways have been proposed in the literature, which are explained below.

#### 2.2.2.1 Bounding the Error Using the Central Limit Theorem

Castro et al. (2009) used the *Central Limit Theorem (CLT)* (Stein, 1972) which states that when the sample size tends to infinity, the sampling distribution of the mean will be a normal distribution with a mean equal to the sample mean, and a variance equal to the popoulation variance, $\sigma^2$, divided by the sample size, denoted as $m$, i.e., $\overline{SV}(a) \sim \mathcal{N}(\mu, \sigma^2/m)$. To achieve a bound in the form of inequality (2.3), one can construct a confidence interval by transforming the normal sampling distribution of the mean into a standard normal distribution (i.e., $z \sim \mathcal{N}(0,1)$), and using the *z-score*. The $z$-score represents the central area under the curve of the standard normal distribution and equals how many standard deviations the values of a population are from its mean. Thus, the following inequality holds:

$$-Z_{\delta/2} \times \frac{\sigma}{\sqrt{m}} + \overline{SV}(a) \leq SV(a) \leq \overline{SV}(a) + Z_{\delta/2} \times \frac{\sigma}{\sqrt{m}} \tag{2.4}$$

where $Z_{\delta/2}$ is the $z$-score such that $\Pr(z \geq Z_{\delta/2}) = \delta/2$. Inequality (2.4) can be written as a bound on the error in the following way:

$$\epsilon = |SV(a) - \overline{SV}(a)| \leq Z_{\delta/2} \times \frac{\sigma}{\sqrt{m}}.$$

Given the variance of the marginal contributions of an agent, $\sigma^2$, one can determine the minimum sample size required in order to estimate the Shapley value of the agent using inequality (2.2.2.1). This is given as:

$$m \geq \frac{Z_{\delta/2}^2 \times \sigma^2}{\epsilon^2} \tag{2.5}$$

Calculating the above inequality requires *a priori* knowledge of the population variance. However, since it is often not known explicitly, it can be upper-bounded using *Popoviciu's inequality* (Popoviciu, 1935). Given the range of agent $a$'s marginal contributions (i.e., the distance between the minimum and maximum marginal contributions), denoted as $r$, this inequality is as follows:

$$\sigma^2 \leq \frac{r^2}{4} \tag{2.6}$$

Although the CLT provides a simple and easy way to bound the approximation error, it has a major drawback. This is due to the asymptotic nature of the CLT which establishes that the sample mean is normally distributed only when the sample size tends to infinity. Since an infinite sample size is impossible in practice, when the approximation is based on a finite sample, there will be an additional error due to the difference between the true distribution of the sample mean and the asymptotic normal distribution converge (Berry, 1941; Stein, 1972). Therefore, since this has not been factored in by Castro et al., their bound is not accurate.

#### 2.2.2.2 Bounding the Error Using Chebyshev's Inequality

Liben-Nowell et al. (2012) focused on approximating the Shapley value in supermodular games using SRS. Specifically, they used Chebyshev's inequality (Tchébychef, 1867) from probability theory to provide an error bound that requires a sample size polynomial in the number of agents. For any real number $k > 0$, Chebyshev's inequality guarantees that the probability that the values of a given random variable, $X$, are more than $k$ standard deviations away from the mean is at most $1/k^2$. More formally:

$$\Pr(|X - \mathbb{E}[X]| \geq k\sigma) \leq \frac{1}{k^2} \tag{2.7}$$

Recalling that the sample mean is a random variable, one can apply inequality (2.7) to bound the distance between the values of the sample mean and the expected value of the sample mean which is the actual Shapley value. However, to do this, one would need to know the variance of the marginal contributions of the agent. To this end, Liben-Nowell et al. exploit the supermodularity property to provide an upper bound on the variance. Let $X$ be a random variable representing the marginal contributions of agent $a$. They upper bound $\mathrm{Var}(X)$ using the variance of another random variable, $Y$, that supposedly represents the most extreme case of the values of $X$ in terms of variance. More specifically, they assume that $Y$ takes the maximum value of $X$, i.e., the marginal contribution to the grand coalition minus the agent (i.e., $A\backslash\{a\}$) with probability $1/n$, and 0 otherwise. However, it can be argued that $\mathrm{Var}(Y)$ is not ncessarily an upper bound on $\mathrm{Var}(X)$, since the maximum possible variance of a random variable occurs when it takes its minimum and maximum values with the same probability $1/2$ (Castro et al., 2009). This is also evident in Popoviciu's inequality (i.e., inequality (2.6)). Furthermore, in constructing the variable $Y$, Liben-Nowell et al. derive the maximum marginal contribution of agent $a$ as a factor of its Shapley value, i.e., $n \times SV(a)$. However, introducing $n$ into the variance of $Y$, and in turn in Chebyshev's inequality, would make the sample size dependent on the number of agents. Consequently, with higher number of agents the bound would require larger samples, which is undesirable. Chapter 4 will correct the issue of maximum variance, and use Chebyshev's inequality to provide a bound on the error that is independent of the number of agents.

#### 2.2.2.3 Bounding the Error Using Hoeffding's Inequality

Bachrach et al. (2008) focused on approximating the Shapley-Shubik and Banzhaf power indices (Shapley and Shubik, 1954; Banzhaf, 1965), which measure the powers of players in a voting game by averaging over the marginal contributions of an agent to all coalitions. The former is the common name for the Shapley value in simple games, and the latter differs from the Shapley value in the weight of coalitions. Bachrach et al. also use SRS to approximate the power indices, and use Hoeffding's inequality (Hoeffding, 1963) to bound the error. Hoeffding's Theorem 2 (Hoeffding, 1963), known as Hoeffding's inequality, states that if $S$ is the sum of $m$ independent random variables, $X_1, \ldots, X_m$, each of which almost surely bounded by two values, $\alpha_i$ and $\beta_i$, then the following inequality holds about $S$:

$$\Pr(|S - \mathbb{E}[S]| \geq t) \leq 2\exp\left(-\frac{2t^2}{\sum_{i=1}^{m}(\beta_i - \alpha_i)^2}\right) \tag{2.8}$$

This inequality provides a powerful means to bound the approximation error, since it implies that the probability of a large deviation from the mean is exponentially small. Although Bachrach et al. (2008) used Hoeffding's inequality to bound power indices, it can be readily applied to any game where the range of an agent's marginal contributions is available. Chapter 4 will show how this can be done.

## 2.3 Coalitions of Bounded Rational Agents

As stated in the previous section, each agent in a game, by taking actions or making decisions, receives a payoff. In most scenarios, it is assumed that agents are able to maximise their payoff by making optimal decisions. In other words, they are assumed to be fully rational. However, it has been suggested that this assumption may not always be realistic, since in reality, agents do not necessarily try to find an optimal decision (Rubinstein, 1997). For instance, Simon (1957) argues that humans beings as decision makers are only partly rational, and as such, he proposed *bounded rationality* as an alternative basis for modeling decision making. Rubinstein (1997) investigated bounded rationality further and argued that an agent's task often involves picking from a finite set of decisions that satisfies all the constraints, which is typically easier than "finding the optimal set of decisions", especially when multiple solutions exist. He also proposed to model bounded rationality by explicitly specifying decision making procedures.

Several other researchers have taken a computational point of view of bounded rationality. Kalai and Stanford (1988) proposed a method for measuring the possible difficulty that an agent may face in checking the rationality of a strategy combination. Neyman (1985) argues that in finitely repeated prisoners' dilemma problem, there are bounds to the complexity of the strategies that the players may use. Futia (1977) put forward an index number to measure various resource cost associated to the use of decision rules. More recently, Tsang (2008) proposed the Computational Intelligence Determines Effective Rationality (CIDER) theory based on Rubinstein's approach. Tsang's theory says that rationality involves computation, and as such, agents' rationality is bounded by the computational resources that are available to them. Therefore, a *computationally bounded rational agent* would choose good readily obtainable solutions rather thans intractable optimal solutions.

Against this background, Sandholm and Lesser (1997) extended the idea of bounded rationality to coalitional games. In more detail, they considered a setting where the value of a coalition is a cost that it incurs in a certain problem, which is given by a hard optimisation problem. However, since using computational resources is costly, each agent also has to pay for the computational resources that it uses for deliberation. As such, the value of a coalition is defined as follows: each coalition minimises the domain-specific cost that it incurs (which decreases as more computation is allocated) and computation cost (which increases as more computation is allocated). Therefore, as the unit cost of computation increases, agents need to pay more for the computation, or they have to use less computation and obtain worse solutions.

Based on this model, Sandholm and Lesser also investigated the optimal coalition structure that would maximise the total payoff of all agents, and used the concept of core to analyse the stability of this structure. However, they stopped short of providing a formalism for using the Shapley value using this model. Chapter 5 aims to fill this gap by presenting a similar model for dividing the value of a coalition of computationally bounded rational agents using the Shapley value.

## 2.4 Summary

This chapter began with an introduction to concepts in cooperative game theory. In particular, it examined different payoff distribution solution concepts. Specifically, the Shapley value and its special properties were stated. It was discussed that the Shapley value has an exponential time complexity, and to overcome this issue three approaches exist in the literature:

- Using alternative representation formalisms, which involve the difficult task of transforming a given game to conform to these representations.
- Developing exact efficient methods for specific class of games, which are unique to the spcial problems that they have been developed for.
- Approximate solutions, the majority of which are based on random sampling.

Furthermore, sampling-based approximation methods were examined in detail. These methods use SRS to approximate the Shapley value as the mean of a randomly chosen set of marginal contributions of an agent. However, since SRS samples the population blindly, it does not yield representative samples, and as such, it is not considered an efficient method. Furthermore, three different methods have been proposed to bound the approximation error of the Shapley value using SRS, which are as follows:

- Central Limit Theorem, which is an asymptotic bound. That is, the error bound holds as the sample size increases to infinity, which is impractical.
- Chebyshev's inequality, which provides a probabilistic bound on the deviations of the sample mean from the actual mean of a population. However, the sole use of this inequality in the literature for approximating the Shapley value is mathematically incorrect, and requires a sample size that depends on the number of agents in the game.
- Hoeffding's inequality, which also provides a probabilistic bound on the deviations of the sample mean from the actual mean of a population.

Lastly, the idea of bounded rationality of agents was examined. Specifically, it was discussed that agents, due to limited computational resources, may not be able to act with full rationality, and as such, they could be computationally rational. This concept was extended to coperative games by Sandholm and Lesser (1997), who investigated the concept of core to analyse the stability of coalitions of bounded rational agents. However, they did not investigate using the Shapley value to divide the social welfare in this setting, and this problem has remained unexplored to date.

# Chapter 3

# Cooperative Games In The Smart Grid

The phenomenal growth in energy demand, one the one hand, and dwindling oil and fossil fuel reserves on the other hand, have led governments to facilitate a greater role for renewable energy resources in current and future electricity grids. With the advent of a smart electricity grid (The United States Department of Energy, 2003), matching supply and demand will emerge as a crucial challenge, considering that the smart grid envisages large scale integration of distributed energy resources (DERs), which include, small to medium capacity, geographically dispersed renewable power generators. Particularly, the intermittency of these resources is highly problematic, since their generation mostly depends on weather conditions, and might not be available at all times. For example, photovoltaics rely on solar radiation for generating electricity, and their output has a direct relationship with how much sunlight is available. Similarly, the output of wind turbines depends solely on how much wind blows at any given time. Therefore, these generators are technically considered non-dispatchable, i.e. cannot be called in to generate power on demand. This makes DERs' power unsuitable for addressing surges in demand.

Furthermore, the amount of electricity that can be generated from renewable resources can be extremely difficult to predict, since they mostly depend on environmental and weather conditions which constantly vary (e.g., wind speeds, cloud cover, etc.). Given the current market structures that require selling power ahead of generation, the variability of generation is highly undesirable for the profitability of renewable generators, and also balancing supply and demand. From the generators' point of view, since their output cannot be accurately predicted, they will inevitably face penalties for not generating exactly the amount they commit to in their contracts. From the grid's point of view, it is impossible to know, ahead of time, how much renewable generation will be available so that supply and load can be balanced.

On balance, overcoming the intermittency and variability of DERs' outputs is key to a successful integration of DERs in a smart grid. Currently, the fluctuations resulting from these are offset using fossil-fueled generators (e.g., gas-fired stations), whose output can quickly ramp up to meet an incerase in demand. However, as these generators diminish in numbers (due to dwindling resources and phaseout-programs), addressing the issues of intermittency and variability of DERs requires new approaches in managing both supply and demand.

In this context, cooperative game theory can offer valuable insights, though its use has been very limited in the literature. For instance, it can provide models for groups of generators or consumers to mitigate the intermittency and variability issues in a way that cannot be achieved individually. Notably, creating *virtual power plants (VPPs)*, through grouping DERs and aggregating their outputs, is one such model that has high potentials for the supply side (Dimeas and Hatziargyriou, 2007; Pudjianto et al., 2007; Kok et al., 2010). On the demand side, coalitions of consumers can be formed to encourage more predictable demand profiles (Ramchurn et al., 2013), or to offer demand reduction services (Kota et al., 2012). To materialise these ideas, a fair division of profits and monetary incentives among the participants is essential. The Shapley value, due to its unique fairness properties, can fulfill this requirement. Nevertheless, applying the Shapley value can face a major computational challenge due to the potentially large sizes of groups, the time complexity of the characteristic functions of the games, and the time complexity of the Shapley value itself which is exponential in the number of agents.

This chapter gives an overview of the related work on supply-side and demand-side management, and provides the background for addressing the technical computational challenges of applying the Shapley value in two scenarios:

1. Dividing the expected profit of a VPP among its DER members. Here, the large number of DERs (e.g., hundreds) makes calculating the Shapley value difficult. This issue will be addressed in Chapter 4 by approximating the Shapley value.

2. Dividing a group discount among the participants of a demand-side management program that is designed to cap domestic loads from cooling (or heating) at a certain level. Here, the computational challenge is due to the time complexity of the characteristic function that involves an intractable optimisation of the loads. This issue will be addressed in Chapter 5 by adopting a bounded rationality model for coalitions.

The next section gives a description of electricity markets, and particularly the whole-sale market of the Great Britain as one of the globally predominent market models. Section 3.2 discusses the related works regarding VPPs. Section 3.3 presents a formal model of a two settlement market, based on which the problem of selling the output of a DER and its expected profit is defined. Section 3.4 defines a cooperative game where DERs form a coalition to sell their output in the market. Section 3.5 examines the recent applications of cooperative game theory, and the Shapley value, in demand-side management. Finally, Section 3.6 summarises the chapter.

## 3.1 Electricity Markets

The power industry has undergone fundamental changes in the past few decades. In particular, deregulation of electricity markets has created the potential for competition in retail and generation of electricity, marking an important milestone. In contrast to traditional structures of electricity grids, in which one organisation owns the generation, transportation and distribution systems and operates under a monopoly, in new structures those parts are separated and operate independently.

The properties of electricity as a commodity make an electricity market distinct from any other market. In economic terms, electricity is a perishable commodity, meaning it cannot be stored. As such, it has to be generated, delivered and used continuously in real-time. Furthermore, electricity is linked with a physical system that functions faster than any market (Kirschen and Strbac, 2004). Whereas in most markets there is some window for matching supply and demand, in electricity systems, the balance between supply and demand (i.e., generation and load) has to be maintained second by second. If this balance is not maintained, the system will collapse with potentially catastrophic consequences. Therefore, electricity markets are designed to ensure robustness of delivery. The design of electricity markets is an active area of research (see for example Marks, 2006; Anderson and Philpott, 2002; Stoft, 2002).

Generally, there are two types of electricity markets: *retail markets* and *wholesale markets*. The parties involved in the former are the end users of electricity (consumers), as well as the companies that sell the electricity to them, which are commonly referred to as *suppliers*. The aim of retail markets is to enable consumers to choose their suppliers from a competitive environment. In wholesale markets, the main parties involved can be divided into buyers, sellers and system operators. Sellers and buyers are mostly generators and suppliers. Generators can also be buyers, since they would need to buy electricity from other sellers when they cannot generate enough to meet the their contract volume.

### 3.1.1 The Great Britain Wholesale Electricity Market

The wholesale market under the current trading system in Great Britain, called the British Electricity Trading Transmission Arrangements (BETTA), allows bilateral trading between generators and suppliers in the market, which was largely designed to promote efficiency and competition between predictable fossil-fuelled and nuclear generation

(Energy and Climate Change Committee UK Parliament, 2011). The system splits the day into half-hour intervals called *settlement periods*. For each settlement period, contracts must be struck in a *forward market*, ahead of time, and deals must be finalised before the *gate closure*, which is one hour ahead of the settlement period. This requires the generators and suppliers to forecast their generation and demand accurately and strike deals accordingly.

Both generators and suppliers must fulfill their contracts for each settlement period by generating and consuming their agreed amounts. This is judged by the actual generation and consumption volumes that are metered over the course of each settlement period. At the end of the settlement period, a generator may have generated more or less than their contract volume. Similarly, a supplier may have consumed more or less than they contracted for. The difference between the contract volume and the actual volume is called *imbalance*, and is settled through a mechanism called *balancing mechanism*, which is as follows. Where a generator has generated less than they have committed to provide, they must purchase additional electricity in a spot market. The price at which this transaction is cleared is called the *system sell price*. Moreover, where a generator has generated more than their contract volume, they can sell the surplus in the spot market at a price called the *system buy price*. The suppliers must settle their imbalance in a similar fashion.

In order to minimise the overall imbalance of the system, generators and suppliers have incentive to strike contracts based on most accurate estimates of their generation and consumption. This is because the system sell price is always higher, and the system buy price is always lower than the forward market price. From a generator's point of view, this price difference means less profit in case of underestimating the generation, and an extra cost in case of overestimating it. As such, these prices can be seen as penalties for deviations from contracts, and thus, will be referred to as *balancing penalties* herein.

## 3.2 Virtual Power Plants

A VPP can be seen as a coalition of DERs, modeled as profit-maximising (i.e., rational) agents, that pool their generation so that they can be viewed as one large generator in the grid. Not only would this idea improve the reliability of individual generators, but also they would potentially benefit from cost-effective integration into the market (North American Electric Reliability Corporation, 2009; US National Renewable Energy Laboratory, 2010), and avoiding low-profit deals with third-party market participants (Pudjianto et al., 2007). Kok et al. (2005) proposed a novel agent-based system, called *PowerMatcher*, for automatic balancing of demand and supply in VPPs. Based on this system, a VPP would be able to optimally coordinate its DER members such that it can deliver almost real-time balancing services. Kok et al. (2010) suggested a structure for a PowerMatcher so that the VPP's actions can be guided by organising agents as a tree and assigning them roles. For instance, DERs would be represented as "local device agents", a specific agent type would concentrate the pricing bids of other agents into one single bid, and another agent would have the role of guiding the VPP's actions by implementing a business model.

Kok et al. (2010) envisaged that adopting a business model that can ensure the economical viability of VPPs is crucial. In particular, the problem of dividing the profit that the DERs would jointly make is the primary question that needs to be addressed, so that forming a VPP is justifiable to the DERs as well. This problem has remained somewhat underexplored thus far and the few solutions that have been proposed in the literature have some limitations. In what follows, some of these solutions will be examined.

Chalkiadakis et al. (2011) proposed formation of cooperatives of virtual power plants (CVPPs). A CVPP is essentially a coalition of DER agents that sell their electricity to the grid, ahead of time, based on a special pricing scheme agreed by the grid and the CVPP. This scheme consists of a Grid-to-CVPP and a CVPP-to-DER payment. The former determines the total profit that the coalition receives from the grid with respect to the volume of production as well as the accuracy of the production estimate. The latter determines the share of each agent from the total profit. The Grid-to-CVPP payments encourage accurate estimates as well as formation of larger VPPs by promising more profit, leading to a superadditive game. Moreover, the CVPP-to-DER payments are such that individual members would have the incentive to report their estimate of future generation to the highest possible degree of precision, improving the reliability of the CVPP as a whole. Assuming that all DERs in the VPP are able to estimate their generation equally well, the payoff distribution given by the CVPP-to-DER payments would lie in the core of the game. However, this scheme may only be useful in some circumstances where the grid operators would be willing to introduce a new trading model for DERs, and it is not clear that, under this scheme, how the DERs would be able to enjoy the kind of competition that wholesale market participants do.

Bitar et al. (2012) aimed to address the wholesale market participation of DERs by investigating how a wind power generator could optimally sell its variable power in a competitive electricity market. More specifically, they proposed a stochastic model for wind power generation, and a model for a two-settlement market consisting of a forward market, and a spot market. In the former, electricity is sold in advance of generation (e.g., one day ahead) based on generation estimates, while in the latter any difference between the actual and estimated generation is settled. Given this market model, they analysed optimal contract volumes and the corresponding optimal expected profit. Note that while Bitar et al. focused on wind power production, their model could also be applid to any DER with a power output that can be modeled as a stochastic variable. Based on the same model of power generation and market, Baeyens et al. (2011) investigated the merits of forming coalitions of independent wind power producers. To this end, they defined a cooperative game, where a number of wind power producers form coalitions, whose values are the maximum expected profit achievable through joint bidding of the aggregate wind power in a two-settlement market. This game is known as the *newsvendor game* (Muller et al., 2002), which is a cooperative form of the newsvendor problem (see Section 3.4 for more detail). Since the newsvendor game is superadditive, Baeyens et al. argued that a grand coalition of wind power producers would be an optimal coalition structure. Moreover, they showed that the game is balanced, and thus, according to Bondareva–Shapley theorem (Theorem 2.22) it has a non-empty core. The authors focused on the nucleolus as a payoff distribution, arguing that it provides fairness since it minimises the dissatisfaction of the coalition (Remark 2.15). However, calculating the nucleolus is computationally intensive, and the authors did not provide a method for addressing this issue, considering that coalitions are potentially large.

## 3.3 Selling Power In A Two-settlement Wholesale Market

Suppose that selling power takes place through a two-settlement market system, which consists of a forward market, and a spot market that deals with the balancing mechanism (as described in Subsection 3.1.1). These markets are assumed to be *perfectly competitive*, which means that no generator has a market power to influence the prices, and thus, they are all "price takers". Furthermore, to simplify the analysis, the costs of generating power is disregarded here, and thus, all revenue of a generator is assumed to be profit. Contracts are offered by the generators in the forward market at a unit price of $p \in \mathbb{R}_+$, which is known to all participants. The system sell and buy prices, however, are not known at the time of striking the contract, and thus their expected values are used instead. The difference between the actual generation, $w$, and the contracted volume of a generator, $\nu$, is referred to as the *imbalance* of the generator. For negative and positive imbalances, the expected system prices are denoted by $\alpha \in \mathbb{R}_+$ and $\beta \in \mathbb{R}_+$ respectively. Furthermore, due to the nature of the balancing penalties (as explained above), it is assumed that $\beta < p < \alpha$. Note that, in reality, $p$ is closer to $\beta$ than $\alpha$, since generating power at short notice is more expensive (and more carbon-intensive), and also negative imbalance is undesirable for the grid. Therefore, it is assumed that $\alpha - p > p - \beta$. Furthermore, note that although a positive imbalance would result in a generator to receive a payment for their surplus, this would be less than the profit they could make, had they sold the surplus in the forward market. Therefore, any imbalance could be seen as a cost to generators, and must be minimised. Herein, this cost is referred to as the *balancing cost*, and is formally defined as:

$$\zeta(w, \nu) \triangleq (w - \nu) \times \begin{cases} \alpha & w > \nu \\ 0 & w = \nu \\ \beta & w < \nu \end{cases} \tag{3.1}$$

Note that $\zeta$ is a convex function centred on zero, meaning that if there is no imbalance the balancing cost will be zero. The sign of the cost shows whether the generator has to pay or be paid. Moreover, due to asymmetric balancing penalties, the function is not symmetric.

### 3.3.1 Expected Cost and Profit of A DER

Like conventional generators, a DER must minimise its balancing cost in order to maximise its profit. However, unlike conventional generators, DERs cannot know their future generation in advance, and thus, instead of minimising the exact balancing cost defined above, they have to minimise their *expected balancing cost*. This cost would have a direct relationship with the uncertainty in their power generation. Therefore, in order to capture this uncertainty, the output of a DER are represented as a random variable, $W$, and the probability density function (PDF) of $W$, and its cumulative distribution function (CDF) are represented by $f_W$ and $F_W(w) = \int_0^w f_W(x)\,\mathrm{d}x$, respectively.

Given a power output, $W$, and a contract volume $\nu$, the corresponding expected balancing cost is the cost of a generator's expected negative imbalance, minus the profit it would make for its expected positive imbalance. More formally:

$$\begin{aligned}\zeta(\nu, W) &\triangleq \alpha\,\mathbb{E}[\nu - W] - \beta\,\mathbb{E}[W - \nu] \\ &= \alpha \int_0^{\nu} f_W(x)(\nu - x)\,\mathrm{d}x - \beta \int_{\nu}^{M} f_W(x)(x - \nu)\,\mathrm{d}x\end{aligned} \tag{3.2}$$

where $M$ is the maximum generation capacity of the generator.

The *expected profit* of a generator is the amount of profit that it expects to make by selling $\nu$ units in the forward market, minus its expected balancing costs. However, note that it is theoretically possible that the expected balancing cost is so high that the expected profit becomes negative, in which case, obviously, the DER will not participate in the market. More formally, let $P(\nu, \zeta)$ denote a function representing the profit given a contract volume of $\nu$, and a balancing cost of $\zeta$. The expected profit of a generator is defined as:

$$P(\nu, W) \triangleq \begin{cases} 0 & \zeta(\nu, W) \geq p\,\nu \\ p\,\nu\ - \zeta(\nu, W) & \zeta(\nu, W) < p\,\nu \end{cases} \tag{3.3}$$

Using equation (3.3), in the next subsection, the contract volume that maximises the expected profit is found.

### 3.3.2 Optimal Contract Volume: A Newsvendor Problem

As is clear from equation (3.3), the maximum expected profit occurs when $\zeta(\nu, W)$ is smallest. Therefore, to find the optimal contract volume $\nu^*$ the following optimisation must be performed:

$$\nu^* = \arg\min_{\nu} \zeta(\nu, W) \tag{3.4}$$

Note that due to the asymmetry in the balancing penalties (i.e., $p - \beta \neq \alpha - p$), simply contracting for the expected generation volume (i.e., $\mathbb{E}[W]$) would not result in the minimum expected balancing cost. Intuitively, a generator needs to contract for an amount less than its expected output so as to avoid the higher penalty of negative imbalance.

In operational research, this optimisation is known as the *newsvendor problem* (Whitin, 1955). In this problem, a person selling newspapers needs to decide how many newspapers he needs to stock, given that the sale of newspapers is subject to a stochastic demand. This decision is complicated by the fact that, on the one hand, an unsold newspaper will lose its original value, and on the other hand, stocking less than the necessary quantity results in losing sale opportunities. Thus, the newsvendor needs to optimise its stock quantity so as to minimise all these costs. Selling a DER's power in a two-settlement market is similar to the problem that the newsvendor faces, and thus, can be formulated as a newsvendor problem. The rest of this section is devoted to explaining how the optimal contract volume can be found.

If the expected profit function, $P(\nu, W)$, has a maximum value, a maximiser $\nu^*$ can be found by solving the derivative of the expected profit for $\nu$. The first derivative of equation (3.3) is as follows:

$$\frac{\mathrm{d}\, P(\nu, W)}{\mathrm{d}\nu} = p - \alpha \int_0^{\nu} f_W(x)\, \mathrm{d}x - \beta \int_{\nu}^{M} f_W(x)\, \mathrm{d}x = 0 \tag{3.5}$$

The second derivative reveals that the expected profit is a concave function, and thus, must have a maximum. This is because:

$$\begin{aligned} \frac{\mathrm{d}^2\, P(\nu, W)}{\mathrm{d}\nu^2} &= -\alpha f_W(\nu) + \beta f_W(\nu) \\ &= -\underbrace{(\alpha - \beta)}_{>0} \underbrace{f_W(\nu)}_{\geq 0} \\ &\leq 0 \end{aligned} \tag{3.6}$$

Considering the first derivative (equation (3.5)) again, the first integral is by definition equal to $F_X(\nu)$. Moreover, since by definition $F_X(M) = 1$, and the probability of negative generation is zero, the second integral simplifies to $1 - F_X(\nu)$. Consequently:

$$p - \alpha\, F_W(\nu) - \beta\, (1 - F_W(\nu)) = 0 \quad \Rightarrow \quad F_X(\nu) = \frac{p - \beta}{\alpha - \beta} \tag{3.7}$$

Therefore, the best volume to contract for in the forward market is given by the quantile, $F_W^{-1}$, of $(p-\beta)/(\alpha-\beta)$, which is the inverse of $F_W$. This is given by:

$$F_W(\nu) = \frac{p-\beta}{\alpha-\beta}\bigg|_{\nu=\nu^*} \quad \Rightarrow \quad \nu^* = F_W^{-1}\left(\frac{p-\beta}{\alpha-\beta}\right) \tag{3.8}$$

As an example, suppose that $W$ follows a Gaussian distribution $\mathcal{N}(\mu, \sigma^2)$, where $\mu = \mathbb{E}[W]$ and $\sigma^2 = \mathrm{Var}(W)$. The optimal contract volume in this case would be:

$$\nu^* = \mathbb{E}[W] - \sqrt{2\,\mathrm{Var}(W)}\,\mathrm{erfc}^{-1}\left[\frac{2(p-\beta)}{\alpha-\beta}\right],$$

where $erfc^{-1}$ is the inverse complementary error function. Plugging $\nu^*$ into equation (3.3) yields the maximum expected profit when $W$ follows a Gaussian distribution, as follows:

$$P(\nu^*, W) = p\,\mathbb{E}[W] - \frac{e^{-\mathrm{erfc}^{-1}\left[\frac{2(p-\beta)}{\alpha-\beta}\right]^2}(\alpha-\beta)}{\sqrt{2\pi}}\sqrt{\mathrm{Var}(W)} \tag{3.9}$$

In the second term of the above equation, the coefficient of $\mathrm{Var}(W)$ is a constant that, based on the balancing penalties, translates the uncertainty of generation, $\mathrm{Var}(W)$, into a cost. Montrucchio and Scarsini (2007) have shown that for an arbitrary distribution of $W$, the maximum expected profit can be concisely expressed as a function of $\mathbb{E}[W]$ and $\mathrm{Var}(W)$, similar to equation (3.9). More specifically, given the unit price $p$, and the balancing penalties $\alpha$ and $\beta$, there exists a $0 < \mathcal{K} \leq max\{\alpha, p+\beta\}$ such that the maximum expected profit can be written as:

$$P(\nu^*, W) = p\,\mathbb{E}[W] - \mathcal{K}\sqrt{\mathrm{Var}(W)} \tag{3.10}$$

This profit function will be the basis of forming coalitions of DERs, which is described in the next section.

## 3.4 Coalitions of DERs: A Cooperative Newsvendor Game

A VPP is modeled as a coalition, $A = \{a_1, a_2, \ldots, a_n\}$, consisting of $n$ DER agents, that sell their aggregate power in the market formalised in Section 3.3. Let $W_C$ denote a random variable representing the aggregate power of a coalition $C \subseteq A$, which is given by the sum of power output of individual members $W_{a1}, W_{a2}, \ldots$. More formally, the aggregate power of a coalition can be described using its expected generation volume and the corresponding variance as follows:

$$\mathbb{E}[W_C] = \sum_{a \in C} \mathbb{E}[W_a] \tag{3.11}$$

$$\mathrm{Var}(W_C) = \sum_{a \in C} \mathrm{Var}(W_a) + \sum_{a_i, a_j \in C,\, a_i \neq a_j} \mathrm{Cov}(W_{a_i}, W_{a_j}) \tag{3.12}$$

In this thesis, it is assumed that the geographical distances between the DERs of a VPP are so large that there is no correlation between their power outputs. As such, the covariance between the outputs of any pair of DERs in a coalition is zero. The case of correlated power outputs will be left for future work.

The value of a coalition is defined as its maximum expected profit, which can be determined using equations (3.11), (3.12) and (3.10). Given a coalition $C$, this value is given by the following characteristic function:

$$v(C) \triangleq \begin{cases} 0 & C = \emptyset \\ P(\nu^*, W_C) = p\,\mathbb{E}[W_C] - \mathcal{K}\sqrt{\mathrm{Var}(W_C)} & C \neq \emptyset \end{cases} \tag{3.13}$$

The cooperative game $(A, v)$ is called a *newsvendor game* (Muller et al., 2002). Montrucchio and Scarsini (2007) showed that this game is balanced (Definition 2.12), and for newsvendors with uncorrelated demand, it is supermodular. Here, two additional proofs of supermodularity for the game defined above are provided. First, observe the following lemmas.

**Lemma 3.1.** *For any variable $u > 0$ and constant $a > 0$, the function $\sqrt{u} - \sqrt{u + a}$ is monotonically increasing.*

*Proof.*

$$\frac{\mathrm{d}(\sqrt{u} - \sqrt{u + a})}{\mathrm{d}u} = \frac{1}{2\sqrt{u}} - \frac{1}{2\sqrt{u + a}} > 0$$

□

**Corollary 3.2.** *For an arbitrary agent $a_i \in A$, and any pair of coalitions $C_1, C_2 \subseteq A \backslash \{a_i\}$, with corresponding power outputs $W_{C_1}$ and $W_{C_2}$, since $\mathrm{Var}(W_{C_1}) \leq \mathrm{Var}(W_{C_2})$, from Lemma 3.1 it follows that:*

$$\sqrt{\mathrm{Var}(W_{C_1})} - \sqrt{\mathrm{Var}(W_{C_1}) + \mathrm{Var}(W_{a_i})} \leq \sqrt{\mathrm{Var}(W_{C_2})} - \sqrt{\mathrm{Var}(W_{C_2}) + \mathrm{Var}(W_{a_i})},$$

*which is equivalent to:*

$$\sqrt{\mathrm{Var}(W_{C_1})} - \sqrt{\mathrm{Var}(W_{C_1 \cup \{a_i\}})} \leq \sqrt{\mathrm{Var}(W_{C_2})} - \sqrt{\mathrm{Var}(W_{C_2 \cup \{a_i\}})}$$

**Lemma 3.3.** *Given a newsvendor game $(A, v)$, as defined above, for any two coalitions $C_1, C_2 \subseteq A$, it holds that $v(C_1 \cap C_2) \geq 0$.*

*Proof.* If $C_1 \cap C_2 = \emptyset$, by equation (3.13), $v(C_1 \cap C_2)$ is zero. Moreover, if $C_1 \cap C_2 \neq \emptyset$, by equation (3.3), $v(C_1 \cap C_2)$ is greater than or equal to zero. □

**Theorem 3.4.** *Given a set, $A$, of DERs with uncorrelated power outputs, and the characteristic function (3.13), the game $(A, v)$ is supermodular.*

*Proof* 1. Suppose that $C_1$ and $C_2$ are two coalitions such that $C_1 \subseteq C_2 \subseteq A \backslash \{a_i\}$. From Corollary 3.2, it follows that:

$$
\begin{aligned}
& -\mathcal{K}\sqrt{\mathrm{Var}(W_{C_1 \cup \{a_i\}})} + \mathcal{K}\sqrt{\mathrm{Var}(W_{C_1})} \leq -\mathcal{K}\sqrt{\mathrm{Var}(W_{C_2 \cup \{a_i\}})} + \mathcal{K}\sqrt{\mathrm{Var}(W_{C_2})} \\
\Leftrightarrow\ & p\mathbb{E}[W_{C_1}] - \mathcal{K}\sqrt{\mathrm{Var}(W_{C_1 \cup \{a_i\}})} - p\mathbb{E}[W_{C_1}] + \mathcal{K}\sqrt{\mathrm{Var}(W_{C_1})} \\
& \leq p\mathbb{E}[W_{C_2}] - \mathcal{K}\sqrt{\mathrm{Var}(W_{C_2 \cup \{a_i\}})} - p\mathbb{E}[W_{C_2}] + \mathcal{K}\sqrt{\mathrm{Var}(W_{C_2})} \\
\Leftrightarrow\ & \underbrace{p\mathbb{E}[W_{C_1}] + p\mathbb{E}[W_{a_i}] - \mathcal{K}\sqrt{\mathrm{Var}(W_{C_1 \cup \{a_i\}})}}_{v(C_1 \cup \{a_i\})} \underbrace{- p\mathbb{E}[W_{C_1}] + \mathcal{K}\sqrt{\mathrm{Var}(W_{C_1})}}_{-v(C_1)} \\
& \leq \underbrace{p\mathbb{E}[W_{C_2}] + p\mathbb{E}[W_{a_i}] - \mathcal{K}\sqrt{\mathrm{Var}(W_{C_2 \cup \{a_i\}})}}_{v(C_2 \cup \{a_i\})} \underbrace{- p\mathbb{E}[W_{C_2}] + \mathcal{K}\sqrt{\mathrm{Var}(W_{C_2})}}_{-v(C_2)} \\
\Leftrightarrow\ & v(C_1 \cup \{a_i\}) - v(C_1) \leq v(C_2 \cup \{a_i\}) - v(C_2)
\end{aligned}
$$

□

*Proof* 2. Suppose $C_1$ and $C_2$ are two coalitions such that $C_1, C_2 \subseteq A$.

Since $\sqrt{\mathrm{Var}(W_{C_1})} + \sqrt{\mathrm{Var}(W_{C_2})} \geq \sqrt{\mathrm{Var}(W_{C_1}) + \mathrm{Var}(W_{C_2})}$, the following holds:

$$
\begin{aligned}
& -\mathcal{K}\sqrt{\mathrm{Var}(W_{C_1})} + \sqrt{\mathrm{Var}(W_{C_2})} \leq -\mathcal{K}\sqrt{\mathrm{Var}(W_{C_1}) + \mathrm{Var}(W_{C_2})} \\
\Leftrightarrow\ & \underbrace{p\mathbb{E}[W_{C_1}] - \mathcal{K}\sqrt{\mathrm{Var}(W_{C_1})}}_{v(C_1)} + \underbrace{p\mathbb{E}[W_{C_2}] - \mathcal{K}\sqrt{\mathrm{Var}(W_{C_2})}}_{v(C_2)} \\
& \leq \underbrace{p\mathbb{E}[W_{C_1}] + p\mathbb{E}[W_{C_2}] - \mathcal{K}\sqrt{\mathrm{Var}(W_{C_1}) + \mathrm{Var}(W_{C_2})}}_{v(C_1 \cup C_2)}
\end{aligned}
$$

Therefore, $v(C_1) + v(C_2) \leq v(C_1 \cup C_2)$, which shows that the game is superadditive. Adding this inequality to $0 \leq v(C_1 \cap C_2)$ (Lemma 3.3), proves the supermodularity property:

$$v(C_1) + v(C_2) \leq v(C_1 \cup C_2) + v(C_1 \cap C_2)$$

□

Since the game is superadditive, the DERs overall, in expectation, make more profit by forming a grand coalition, as compared to selling their power in the market independently. Therefore, the DERs have the incentive to form a VPP and participate in the market as a single generator.

### 3.4.1 The Shapley Value For Dividing The Total Expected Profit

Once a coalition of DERs is formed, an important question will be what share of the total expected profit can each DER expect to receive. In this problem, the Shapley value gives a particularly useful payoff distribution, since the game is supermodular. Recall that in supermodular games, the Shapley value is in the core of the game (Theorem 2.31). That is, not only can the DERs expect to receive a fair share of the profit, but also this division guarantees a stable VPP, since no DER or groups of DERs would have an incentive to break away from the VPP. Therefore, the Shapley value is the only solution to this problem that can ensure fairness and stability simultaneously. However, given the potentially high numbers of DERs in a VPP, and the time complexity of computing the Shapley value (as discussed in Chapter 2), applying the Shapley value presents a significant computational challenge. Chapter 4 investigates algorithms that can address this issue.

## 3.5 Cooperative Demand-side Management

Demand-side management is a way of influencing demand to follow supply through various means, most commonly through *demand response (DR)*. DR is defined by of Energy (2006); FERC (2009) as "changes in electric usage by end-use customers from their normal consumption patterns in response to changes in the price of electricity over time, or to incentive payments designed to induce lower electricity use at times of high wholesale market prices or when system reliability is jeopardized". It has been suggested that in the smart grid, DR can play an important role in improving the balance of supply and demand (Strbac, 2008; Su and Kirschen, 2009). For instance, it can help to reduce peak demand by disributing consumers' demand to off-peak times of the day. This lowers the capacity needed to deal with peak demand. It can also provide the flexibility needed to match demand to the available supply when unpredictable fluctuations occur. Furthermore, DR can keep the demand within the limitations of local network constraints. It is important that DR programs are designed in a way that are capable of managing large groups of consumers, since unless the consumptions are coordinated appropriately, peaks may occur at different times (Ramchurn et al., 2011). This need for coordination among consumers creates an opportunity for forming coalitions among them, and making use of the tools from cooperative game theory to study these coalitions.

Rose et al. (2012) designed a scheme for an aggregator (intermediaries between suppliers and end-users of electricity) to encourage accurate prediction of its customers' future demand. The aggregator is responsible for purchasing electricity for a set of homes in a two-settlement market similar to the model in Section 3.3. Due to the balancing penalties, the aggregator has to pay a cost for any difference between the amount it purchases in the market and the actual amount that its customers consume. The aggregator has an estimate of each house's consumption based on historical data, but also each house possibly has some information that can be used to make more precise estimates. The houses transmit this information to the aggregator in the form of a probability distribution over the their consumptions for a future period of time. In doing so, the houses ignore their own cost and minimise the sum of all houses' costs, effectively working as a coalition. By exploiting the information received from the houses, the aggregator's expected costs will be lower, and thus in expectation, it makes some savings. The aggregator then distributes these savings to the houses as a reward for their information. The rewards are such that, the more accurate each house's information is, the greater its share of the total reward, incentivising them to report truthfully and accurately.

In Ramchurn et al. (2013), we proposed a newsvendor game for collective purchasing of electricity for users of an agent-based system, called *AgentSwitch*, that receive a discounted price for having more predictable demand profiles. AgentSwitch predicts the annual power consumption of its users using a Bayesian quadrature (a numerical integration technique) that estimates the power consumption and gives a Gaussian probability distribution for any given period of time. These distributions are aggregated using equations (3.11) and (3.12), yielding an overall lower variance, which results in an improved predictability of the demand. In return for this, the supplier offers a discounted price, which is distributed among the users using the Shapley value. Due to the large number of users, the Shapley values are then approximated using a special property of newsvendor games. This property is explained in detail in Chapter 4.

Kota et al. (2012) designed a cooperative scheme for groups of consumers that can reduce their energy consumption when requested. The aim of this scheme is to aid participation of consumers in the market, and flattening their load profiles through shifting their consumption. The members of the coalition report their estimated baseline consumption to the operator of the scheme, and if a shift in their consumption is required, they will be notified one day ahead. Kota et al. proposed a mechanism for distributing the revenue that the coalition would generate as a result of reducing its demand. This mechanism is such that the members would have an incentive to participate in the scheme, and is incentive compatible, meaning the members would not be able to exaggerate their baseline consumption so as to show an artificial demand reduction and get rewarded consequently. However, the authors stopped short of providing a model for shifting consumption and implementing the demand reduction in practice. Akasiadis and Chalkiadakis (2013) aimed to fill this gap by proposing a method for shifting loads of residential, and industrial consumers. The consumers form a coalition to report their shiftable capacities so that their peak time consumptions can be shifted to to off-peak intervals where the price of electricity is lower. However, the authors were merely interested in mechanisms that would ensure the truthfulness of the reports, and like other works did not consider fairness and the Shapley value.

One of the areas where there is a high potential for managing demand is thermostatically controlled loads (TCLs), such as space and water heaters, which account for a significant percentage of domestic energy consumption (Administration, 2013; Department of Energy & Climate Change, 2014). Such appliances can act act as a buffer for increase in demand, in that the load from them can be shifted to off-peak times. Vasirani and Ossowski (2012) proposed a collaborative load management model where consumers with different types of shiftable loads (including TCLs) form coalitions and agree on a joint demand profile to be contracted with a supplier, and in return receive a discount. The authors propose a payment mechanism based on the core, and do not focus on fairness. The rest of the existing research on DR programs for TCLs mostly focuses on collective behaviour of users in a DR scenario (Zhang et al., 2012; Ihara and Schweppe, 1981), finding complementary load profiles among groups of consumers with TCL (Koch et al., 2011), and other technical requirements such as physical models for controlling TCLs (Kalsi et al., 2011).

In summary, most of the existing work that deal with payoff distribution in cooperative models for demand side management are concerned with efficiency, individual rationality, group rationality, and incentive compatibility. While the choice of solution concepts would depend on the scenarios, the Shapley value, which also has properties such as efficiency and individual rationality, has been widely neglected. Chapter 5 aims to bridge this gap by applying the Shapley value to a DR scenario where a group of apartments are offered a discount if they manage to reduce their load from air conditioning under a certain limit.

## 3.6 Summary

This chapter considered the applications of cooperative game theory in the smart grid which is an important research area in computer science and artificial intelligence. The idea of creating a smart grid faces the challenge of largely relying on distributed renewable energy resources, whose generation is intermittent and unpredictable. Overcoming this challenge requires a better management of both generators and consumers. An effective approach regarding the former is formation of VPPs, which can be viewed as coalitions of DERs. This chapter analysed the expected profit of a VPP that sells its aggregated output in a two settlement market consisting of a forward market and a spot market (for settling imbalances). It was shown that, in this market, the DER members can expect to make more profit compared to when they are not part of the VPP. Using the Shapley value to divide the expected profits is particularly advantageous, since it lies in the core of the game. That is, each member receives a fair share of the expected profit, and at the same time, no group of agents would have an incentive to leave the grand coalition and form a smaller VPP. However, given the exponential time complexity of the Shapley value, and the potentially large number of DERs in a VPP, computing the Shapley value poses a significant computational challenge.

Furthermore, applications of cooperative games in demand-side management were examined. Demand-side management is mostly conducted through demand response programs that incentivise consumers to alter their load profile so as to reduce peaks or match demand to the available supply. Running these programs requires coordinating groups of consumers to avoid generating further peaks at different times. This coordination among consumers can be modelled as cooperative games, in which the value of a coalition would be the reward it receives for altering its load. Various mechanisms proposed for dividing the rewards were examined in this chapter. These mechanisms are mostly concerned with efficiency, individual rationality, and group rationality. However, the Shapley value, which in addition to fairness boasts the efficiency and individual rationality properties, has not received much attention.

# Chapter 4

# Approximating the Shapley Value

In Chapter 2 the existing works regarding approximating the Shapley value were described. It was discussed that these works suffer from two important limitations. First, the bound proposed by Castro et al. (2009) is asymptotic (i.e., it holds only when the sample size increases to infinity), and Liben-Nowell et al. (2012)'s bound is mathematically incorrect, and also requires a sample size polynomial in the number of agents. Second, the existing approximation algorithms are all based on SRS which means that they are indifferent in picking between marginal contributions, and sample the population blindly. While SRS has the advantage of being easy to implement, it does not necessarily yield representative samples.

To address the shortcomings of the existing works regarding the estimation error of the Shapley value, this chapter proposes a bound based on Hoeffding's and Chebyshev's inequalities. Next, to overcome the computational challenge of applying the Shapley value in a newsvendor game (as discussed in Subsection 3.4.1), a sampling technique, called stratified sampling, is proposed which yields potentially more efficient samples as compared to SRS. Using this technique, the population of marginal contributions of an agent is stratified into smaller homogeneous sub-populations (called strata), and each stratum is sampled independently. This ensures that the samples would be spread over the entire population, yielding estimates with higher precisions.

The rest of this chapter is organised as follows. Section 4.1 discusses bounding the estimation error of SRS. Section 4.2 describes stratified sampling in the context of estimating the Shapley value, and presents two methods for stratifying populations. Section 4.3 is concerned with the experimental evaluation of our approach. Section 4.4 summarises the chapter.

## 4.1 Bounding the Estimation Error of SRS

Consider a random variable, $\Phi$, taking the $n!$ marginal contribution values of agent $a$ with a probability of $1/n!$. The Shapley value of $a$, i.e., $SV(a)$, can be estimated using the *sample mean*, denoted as $\bar{\Phi}$. Under SRS, the sample mean, $\bar{\Phi}_{SRS}$, is obtained by calculating the average of $m$ values of $\Phi$ that are randomly chosen with replacement, with all values having the same probability of being chosen. More formally, $\bar{\Phi}_{SRS} = 1/m \sum_{i=1}^{m} \phi_i$, where $\phi_i$ is the $i$-th element independently drawn from $\Phi$. The estimation error, i.e., $|\bar{\Phi}_{SRS} - SV(a)|$ (or $|\bar{\Phi}_{SRS} - \mathbb{E}[\Phi]|$), can be bounded using Chebyshev's or Hoeffding's inequalities. Given inequalities (2.7) and (2.8), one can find the smallest sample size required such that $\Pr(|\bar{\Phi}_{SRS} - \mathbb{E}[\Phi]| \geq \epsilon) \leq \delta$, where $\bar{\Phi}_{SRS}$ denotes the sample mean, $\epsilon \in \mathbb{R}_+$ is an arbitrary error bound, and $1-\delta$ is the confidence of the bound (i.e., the probability that the estimation error, $|\bar{\Phi}_{SRS} - \mathbb{E}[\Phi]|$, is actually less than $\epsilon$). Here, two lemmas are provided, based on Chebyshev's and Hoeffding's inequalities, that will help with bounding the estimation error of SRS, which is formalised in Theorem 4.3.

**Lemma 4.1.** *Given the variance of an agent's marginal contributions, $\sigma^2$, a bound, $\epsilon$, and a confidence $1-\delta$, the sample size required such that $\Pr(|\bar{\Phi}_{SRS} - \mathbb{E}[\Phi]| \geq \epsilon) \leq \delta$ is:*

$$m \geq \frac{\sigma^2}{\delta\, \epsilon^2} \tag{4.1}$$

*Proof.* Let $S = \sum_{i=1}^{m} \phi_i$. Since $\bar{\Phi}_{SRS} = 1/m\, S$, applying inequality (2.7) to $S$ yields:

$$\Pr\left(|S - \mathbb{E}[S]| \geq k\sqrt{\mathrm{Var}(S)}\right) = \Pr\left(|\frac{1}{m}S - \mathbb{E}[\Phi]| \geq \frac{k\sqrt{\mathrm{Var}(S)}}{m}\right)$$

Let $\epsilon$ be $(k/m)\sqrt{\mathrm{Var}(S)}$. Since each sample is actually a random variable, and also the variance of the sum of independent random variables is equal to the sum of variances of individual variables, we have:

$$\Pr(|\bar{\Phi}_{SRS} - \mathbb{E}[\Phi]| \geq \epsilon) \leq \frac{\mathrm{Var}(S)}{(m\,\epsilon)^2} = \frac{m\,\mathrm{Var}(\Phi)}{m^2\,\epsilon^2} \tag{4.2}$$

Since the right hand side of the above inequality has to be at most $\delta$, it follows that:

$$m \geq \frac{\sigma^2}{\delta\, \epsilon^2}$$

□

**Lemma 4.2.** *Given the range of an agent's marginal contributions, $r$, an error bound, $\epsilon$, and a confidence $1-\delta$, the sample size required such that $\Pr(|\bar{\Phi}_{SRS} - \mathbb{E}[\Phi]| \geq \epsilon) \leq \delta$ is:*

$$m \geq \frac{\ln(2/\delta)\, r^2}{2\,\epsilon^2} \tag{4.3}$$

*Proof.* Since in any game with finite number of agents, the population of marginal contributions of an agent is finite, there always exists a minimum marginal contribution, $\phi^{min}$, and a maximum marginal contribution, $\phi^{max}$. Therefore, $\Phi$ takes values in the range $r = \phi^{max} - \phi^{min}$. Now, let $S = \sum_{i=1}^{m} \phi_i$. From inequality (2.8), we have:

$$\Pr(|S - m\mathbb{E}[\Phi]| \geq t) = \Pr(|\bar{\Phi}_{SRS} - \Phi| \geq \frac{t}{m})$$
$$\Rightarrow \Pr(|\bar{\Phi}_{SRS} - \mathbb{E}[\Phi]| \geq \epsilon) \leq 2\exp\left(\frac{-2\,m^2\,\epsilon^2}{m\,r^2}\right)$$

Since we want the right hand side to be at most $\delta$, we have:

$$2\exp\left(\frac{-2\,m\,\epsilon^2}{r^2}\right) \leq \delta \;\Rightarrow\; m \geq \frac{\ln(2/\delta)\,r^2}{2\,\epsilon^2}$$

□

**Theorem 4.3.** *Given a superadditive game, an estimation error bound, $\epsilon$, and a confidence $1-\delta$, the minimum sample size required to estimate the Shapley value of agent a using SRS, such that $\Pr(|\bar{\Phi}_{SRS} - SV(a)| \geq \epsilon) \leq \delta$, is:*

$$m \geq \begin{cases} \dfrac{\ln(2/\delta)\,r^2}{2\,\epsilon^2} & \delta < 0.23 \\ \\ \dfrac{r^2}{4\,\delta\,\epsilon^2} & \delta \geq 0.23 \end{cases} \tag{4.4}$$

*where $r = MC(a,\, N\backslash\{a\}) - MC(a,\, \emptyset)$.*

*Proof.* The first part of inequality (4.4) was proved in Lemma 4.2. Here, the range of $\Phi$ is initially found. Then, the second part is derived using Lemma (4.1), and it is proved that for $\delta > 0.23$, Chebyshev's inequality requires a smaller sample size than Hoeffding's to achieve the same bound on the error. Figure 4.1 depicts the sample size required by each of these inequalities to establish the same error bound.

From superadditivity, it follows that $\Phi$ takes its maximum value when $a$ is the last agent in the permutation, which occurs in at least $(n-1)!/n!$ of permutations, and takes its minimum value when $a$ is the first agent in the permutation, which also occurs in at least $(n-1)!/n!$ of permutations. Therefore, the range of $\Phi$ (i.e., $r$) is $MC(a,\, N\backslash\{a\}) - MC(a,\, \emptyset)$. Moreover, $\Phi$ takes its extreme values with the same probability $1/n$. Given this, the maximum possible variance of $\Phi$ occurs when half of the rest of the values take the value of the maximum, $\phi^{max}$, and the other half take the value of the minimum, $\phi^{min}$. In other words, the maximum variance of $\Phi$ is when half of the marginal contributions are $\phi^{min}$, and half of them are $\phi^{max}$. Let $Y$ be a random variable that represents this extreme case. Therefore, $Y$ takes $\phi^{min}$ and $\phi^{max}$, each with probability $1/2$. By definition, the variance of $Y$ is greater than or equal to the variance of $\Phi$. Therefore, we

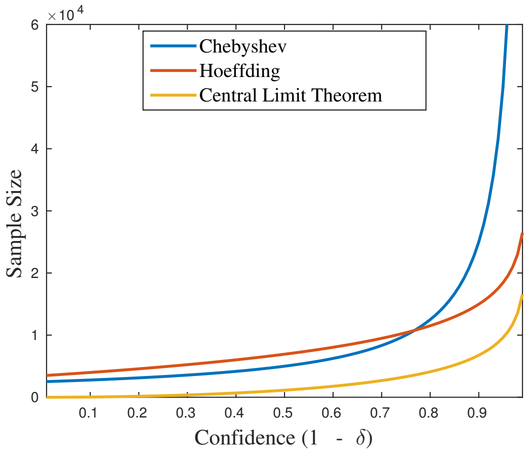


FIGURE 4.1: Comparison of sample sizes required by different methods for the same error bound and confidence.

have the following relation regarding $\sigma^2$:

$$\sigma^2 \leq \mathrm{Var}(Y) = \mathbb{E}[Y^2] - \mathbb{E}[Y]^2 = \frac{(\phi^{max} - \phi^{min})^2}{4}$$

Combining the above with inequality (4.2) yields the second part of inequality (4.4):

$$\frac{\sigma^2}{m\,\epsilon^2} \leq \frac{r^2}{4\,m\,\epsilon^2} \leq \delta \;\;\Rightarrow\;\; m \geq \frac{r^2}{4\,\delta\,\epsilon^2}. \tag{4.5}$$

Now, we would like to find $\delta$ such that $r^2/(4\,\delta\,\epsilon^2) \leq \ln(2/\delta)\,r^2/(2\,\epsilon^2)$. Simplifying and then multiplying both sides by $\delta/2$ yields $\delta/2 \times \ln(\delta/2) \leq -1/4$. With respect to the definition of the Lambert W function, since $e^{W(z)}\,.\,W(z) = z$, we have:

$$e^{\ln\frac{\delta}{2}}.\ln(\frac{\delta}{2}) \leq -1/4 \;\Rightarrow\; \ln\frac{\delta}{2} \geq W(\frac{-1}{4}) \text{ or } W_{-1}(\frac{-1}{4})$$

Since $\delta \leq 1$ only the lower branch, i.e., $W_{-1}(-1/4)$ is valid, and thus:

$$\delta \geq 2\,e^{W_{-1}(-1/4)} \approx 0.232$$

□

## 4.2 Stratified Sampling

As was metioned eariler, SRS does not always yield representative samples, since it dismisses any useful information about the population (e.g., how spread the values are), which could lead to highly different sample mean values with every new sample. The degree to which the sample mean varies with different samples is measured by the *standard error of the mean (SR)*, which is the standard deviation of the sampling distribution of the mean. Given the variability of the sample mean, under SRS, the standard error of the Shapley value is:

$$SE(\bar{\Phi}_{SRS}) = \sqrt{\frac{\mathrm{Var}(\Phi)}{m}} = \frac{\sigma}{\sqrt{m}} \tag{4.6}$$

In contrast to SRS, *stratified random sampling* potentially results in samples with an improved standard error over SRS. In this method, the population is stratified into $S$ mutually exclusive sub-populations (strata), each of which containing $N_s$ elements with values close to one another. Then, each stratum is sampled independently using SRS. This way, if the number of elements drawn from each stratum is proportionate to its size and variance, the sample mean would have a lower variability as compared to SRS. It has been shown that, given a total sample size, $m$, the optimal number of elements that should be drawn from each stratum, $m_s^*$, is given by:

$$m_s^* = m \frac{\frac{N_s\,\sigma_s^2}{\sqrt{c_s}}}{\sum_{i=1}^{S} \frac{N_i\,\sigma_i^2}{\sqrt{c_i}}}\ , \tag{4.7}$$

where $\sigma_s^2$ is the variance of stratum $s$, and $c_s$ is the computational cost of evaluating an element in stratum $s$ (Levy and Lemeshow, 2008). This allocation is particularly useful in some games where the characteristic function has different computational costs for different coalitions (see Aziz et al. (2014) and Alam et al. (2013) for example). For the purpose of this thesis, it is assumed that the costs are all equal. This special case is called *Neyman allocation* and is given as:

$$m_s^* = m \frac{N_s\,\sigma_s^2}{\sum_{i=1}^{S} N_i\,\sigma_i^2} \tag{4.8}$$

Under stratified sampling, the estimated Shapley value, and its standard error are respectively:

$$\bar{\Phi}_{STR} = \sum_{s=1}^{S} \left(\frac{N_s}{N}\right) \sum_{i=1}^{m_s} \phi_{i,s}\ , \tag{4.9}$$

$$SE(\bar{\Phi}_{STR}) = \sqrt{\mathrm{Var}(\bar{\Phi}_{STR})} = \sqrt{\sum_{s=1}^{S} \left(\frac{N_s}{N}\right)^2 \frac{\sigma_s^2}{m_s}}\ , \tag{4.10}$$

where $\phi_{i,s}$ is the $i$-th element in stratum $s$, and $m_s$ is the number of elements drawn from stratum $s$. Note that $m_s$ must be at least 1 for equation (4.10) to hold.

**Theorem 4.4.** *Given the same sample size, the estimated Shapley value under stratified sampling with Neyman allocation has at most the same standard error as under SRS,* *i.e.,* $SE(\bar{\Phi}_{STR}) \leq SE(\bar{\Phi}_{SRS})$.

*Proof.* Intuitively, this theorem holds because the total variance of the sample mean under stratified sampling is the variance within the strata, plus the variance between the strata. Unlike SRS, the variance between the strata under stratified sampling could be zero. Thus, when the strata are heterogeneous between, a higher precision than SRS is gained. We now provide a formal proof.

First, observe that when the number of strata is exactly one (i.e., $S = 1$), $SE(\bar{\Phi}_{STR})$ (equation 4.10) becomes equal to $SE(\bar{\Phi}_{SRS})$ (equation 4.6). For $S > 1$, the most extreme case occurs when the variance of all strata are equal to the variance of the population (they cannot be higher). Therefore, we assume that for any given $s$, $\sigma_s^2 = \sigma^2$. In what follows, we show that in that case, the inequality $SE(\bar{\Phi}_{STR}) \leq SE(\bar{\Phi}_{SRS})$ would still hold. We begin by observing the fact that $\sum_{s=1}^{S} N_s \leq N$, from which it follows:

$$\sum_{s=1}^{S} \frac{N_s \sigma_s^2}{N} \leq \sigma^2 \quad \Leftrightarrow \quad \sum_{s=1}^{S} \frac{N_s \sigma_s^2 \sum_{i=1}^{S} N_i}{N^2} \leq \sigma^2 \quad \Leftrightarrow \quad \sum_{s=1}^{S} \frac{N_s \sum_{i=1}^{S} N_i \, \sigma_s^2}{N^2 \, m} \leq \frac{\sigma^2}{m}$$

$$\Leftrightarrow \quad \sum_{s=1}^{S} \frac{N_s}{N^2 \, m} \frac{N_s \sigma_s^2}{\frac{N_s \sigma_s^2}{\sum_{i=1}^{S} N_i \, \sigma_i^2}} \leq \frac{\sigma^2}{m} \quad \Leftrightarrow \quad \sum_{s=1}^{S} \frac{N_s}{N^2} \frac{N_s \sigma_s^2}{m \frac{N_s \sigma_s^2}{\sum_{i=1}^{S} N_i \, \sigma_i^2}} \leq \frac{\sigma^2}{m}$$

$$\Leftrightarrow \quad \sum_{s=1}^{S} \frac{N_s^2 \, \sigma_s^2}{N^2 \, m_s^*} \leq \frac{\sigma^2}{m} \quad \Leftrightarrow \quad \sqrt{\sum_{s=1}^{S} \left(\frac{N_s}{N}\right)^2 \frac{\sigma_s^2}{m_s^*}} \leq \frac{\sigma}{\sqrt{m}} \quad \Leftrightarrow SE(\bar{\Phi}_{STR}) \leq SE(\bar{\Phi}_{SRS})$$

□

Note that equation (4.8) relies on the variance of each stratum. However, rarely would we have enough information from a game to determine the variance of the population, not to mention the variance of each stratum. Fortunately, the range of marginal contributions can typically be found more readily. Using the range of a random variable, one can find its maximum variance, since the variance is at most equal to the range squared divided by 4, i.e., $\sigma^2 \leq r^2/4$. This is known as Popoviciu's inequality (Popoviciu, 1935). As such, one can replace the variance of a stratum in equation (4.8) with $r_s^2/4$, where $r_s$ is the range of stratum $s$. This way, the sample size of the strata will still be proportional to their variance, as in the optimal allocation. Therefore, when only the range of the strata can be identified, the sample sizes can be calculated as:

$$m_s = m \frac{N_s \, r_s^2}{\sum_{i=1}^{S} N_i \, r_i^2} \tag{4.11}$$

Likewise, in equations (4.6) and (4.10), one can replace $\sigma_s^2$ with $r_s^2/4$ to obtain the maximum possible standard errors.

It should be noted that while the variance is typically estimated using the sample variance, doing so requires further sampling of the population, the cost of which would outweigh the gain. Furthermore, the role of variance in the above allocation is simply to divide the sample size proportionally. Therefore, the range can serve the purpose well (Tippett, 1925; Noether, 1955).

Having described stratified sampling for estimating the Shapley value, we now explain how the population of marginal contributions of an agent can be stratified. It can be shown that, finding an optimal stratification is at least $O(N)$. However, with that much computation, one could compute the exact Shapley value. Therefore, an optimal stratification is not possible, and instead, heuristics must be used to achieve a computationally efficient solution. The objective is to divide the population of agent $a$'s marginal contribution values into a number of mutually exclusive strata, such that the range of each stratum can be found with little computation. Next, we propose two methods for achieving this, namely branching stratification, and size-based stratification. Specifically, we explain these methods in the context of supermodular games and a new class of games which we define in Subsection 4.2.2.

### 4.2.1 Branching Stratification

Branching stratification works similarly to a branch and bound algorithm. Consider a set of agents $\mathbb{B} = \{b_1, \dots, b_k\}$, where $\mathbb{B} \subseteq A \backslash \{a\}$. By iterating through the members of $\mathbb{B}$ (which we refer to as a *branching agent*), the population can be divided into two groups in each iteration: those marginal contributions whose corresponding coalition *includes* $b_i$, and those whose corresponding coalition *excludes* $b_i$. Based on this, we would achieve $2^k$ mutually exclusive strata, each of which corresponding to one subset of $\mathbb{B}$ which is denoted as $B_s$. For instance, if $\mathbb{B} = \{b_1, b_2, b_3\}$, the stratum that corresponds to $B_6 = \{b_2, b_3\}$ contains those marginal contributions whose corresponding coalitions contain $b_2$ and $b_3$ but not $b_1$. Now, we need to find the range of each stratum. Let us observe the following lemma, which exposes a property of supermodular games that can be exploited to find the ranges.

**Lemma 4.5.** *In a supermodular game, given an arbitrary permutation, $O$, the marginal contribution of agent $a$ in $O$ is a lower bound on the marginal contribution of $a$ to any coalition containing $Pre_a^O$, and an upper bound on the marginal contribution of $a$ to any subset of $Pre_a^O$.*

*Proof.* Let $C_1$ and $C_2$ be two coalitions such that $C \subseteq Pre_a^O \subseteq D \subseteq A$. By supermodularity, we have: $MC(a, C_1) \leq MC(a, Pre_a^O) \leq MC(a, C_2)$ □

From Lemma 4.5 it follows that, of all the values in stratum $s$, the minimum marginal contribution (i.e., $\phi_s^{min}$) is to the coalition that only includes all members of $B_s$. Likewise, the maximum marginal contribution in stratum $s$ (i.e., $\phi_s^{max}$) is to the largest subset of $A\backslash\{a\}$ that includes all agents in $B_s$, but excludes any branching agent that is not in $B_s$. Given these values, we obtain the range of stratum $s$ as $r_s = \phi_s^{max} - \phi_s^{min}$.

Note that depending on the information available from the game, it might be possible to find an optimal set of branching agents such that it results in the most efficient stratification. Alternatively, one can choose the agents experimentally or randomly.

### 4.2.2 Size-based Stratification

Size-based stratification divides the population of the marginal contributions based on the size of their corresponding coalitions, and will always result in $n$ strata. We now explain how the range of each stratum can be found. In particular, we do so by focusing on a class of games that have a certain property: all subsets of $A$ can be ordered such that, if a coalition, $C_1$, is before another coalition of the same size, $C_2$, then the marginal contribution of any given agent $a \in A$ to $C_1$ is greater than, or equal to, its marginal contribution to $C_2$. Formally, we define a function, $\mathcal{F} : 2^A \rightarrow \mathbb{R}$, that assigns to each coalition a value reflecting its position among other coalitions of the same size according to the aforementioned ordering. We call this function *order-reflecting*, and say a game for which an order-reflecting function exists has the order-reflecting property. Formally, we define it as:

$$\mathcal{F}(C_1) \geq \mathcal{F}(C_2) \Leftrightarrow MC(a, C_1) \geq MC(a, C_2) \tag{4.12}$$

Now, of all the coalitions in a given stratum, the coalition to which an agent has the greatest marginal contribution would be the one that maximises $\mathcal{F}$. Similarly, the minimal marginal contribution would be given by the coalition that minimises $\mathcal{F}$. More formally, for a given stratum, $s$, these coalitions can be found by solving the following equations:

$$C_s^{min} = \underset{C:|C|=s-1}{\arg\min}\ \mathcal{F}(C), \quad C_s^{max} = \underset{C:|C|=s-1}{\arg\max}\ \mathcal{F}(C) \tag{4.13}$$

Then, $r_s = MC(a, C_s^{max}) - MC(a, C_s^{min})$. Note that Property (4.12) is particularly useful when equations (4.13) can be solved efficiently (e.g, in linear or constant time). For instance, if they can be solved in constant time, we can find the ranges of the strata in linear time. Now, let us examine a few games, for which an order-reflecting function exists, and equations (4.13) can be solved in constant time for each stratum.

### 4.2.3 Games with the order-reflecting property

We now examine some supermodular games, for which an order-reflecting function exists, and equations (4.13) can be solved in constant time for each stratum. Figuers 4.2, 4.3

and 4.4 depict the outcome of applying size-based stratification to an instance of each of these game (details of the experiments are provided in Section 4.3).

#### 4.2.3.1 Newsvendor games

Newsvendor games were introduced in Section 3.4, and it was shown that when the outputs of the DERs were uncorrelated, the game was supermodular. We now show that under the same assumption, newsvendor games exhibit the order-reflecting property.

**Theorem 4.6.** *In newsvendor games with uncorrelated outputs, an order-reflecting function is* $\sum_{a\in C} \text{Var}(W_a)$.

*Proof.* Consider two coalitions $C_1, C_2 \subseteq A\backslash\{a\}$, such that $|C_1| = |C_2|$.

$$\begin{aligned} MC(a, C_1) &= p\left(\sum_{a'\in C_1} \mathbb{E}[W_{a'}] + \mathbb{E}[W_a]\right) - \mathcal{K}\sqrt{\sum_{a'\in C_1} \text{Var}(W'_a) + \text{Var}(W_a)} \\ &\quad - p\sum_{a'\in C_1} \mathbb{E}[W_{a'}] + \mathcal{K}\sqrt{\sum_{a'\in C_1} \text{Var}(W_{a'})} \\ &= p\,\mathbb{E}[W_a] - \mathcal{K}\left(\sqrt{\text{Var}(W_{C_1}) + \text{Var}(W_a)} - \sqrt{\text{Var}(W_{C_1})}\right) \end{aligned}$$

Similarly, it is trivial to show that:

$$MC(a, C_2) = p\,\mathbb{E}[W_a] - \mathcal{K}\left(\sqrt{\text{Var}(W_{C_2}) + \text{Var}(W_a)} - \sqrt{\text{Var}(W_{C_2})}\right)$$

Now, in order to show that $MC(a, C_1) \geq MC(a, C_2)$, it is sufficient to prove that:

$$\sqrt{\text{Var}(W_{C_1})} - \sqrt{\text{Var}(W_{C_1}) + \text{Var}(W_a)} \geq \sqrt{\text{Var}(W_{C_2})} - \sqrt{\text{Var}(W_{C_2}) + \text{Var}(W_a)}$$

The only condition under which the inequality above would hold is when $\sqrt{\text{Var}(W_C)} - \sqrt{\text{Var}(W_C) + \text{Var}(W_a)}$ is monotonic with respect to $\text{Var}(W_C) = \sum_{a'\in C} \text{Var}(W_{a'})$. Indeed, this was proved in Lemma 3.1. Therefore, in newsvendor games, the function $\sum_{a'\in C} \text{Var}(W_{a'})$ is oreder-reflecing, i.e.:

$$\sum_{a'\in C} \text{Var}(W_{a'}) \geq \sum_{a''\in C_2} \text{Var}(W_{C''}) \Leftrightarrow MC(a, C_1) \geq MC(a, C_2)$$

□

**Theorem 4.7.** *Assuming that the agents in $A$ are sorted in the ascending order of the variance of their outputs, the range of the strata in the newsvendor game $(A, v)$ with uncorrelated outputs can be found in linear time using the order-reflecting function* $\sum_{a\in C} \text{Var}(W_a)$.

*Proof.* Let us assume that the agents in $A$ are sorted such that agent $a_1$ becomes the one with the smallest variance, agent $a_2$ with the second smallest variance, and so on, i.e., $\text{Var}(W_{a1}) \leq \text{Var}(W_{a2}) \leq \cdots \leq \text{Var}(W_n)$. Equation (4.13) can be solved in constant time for each stratum. This is because, out of all the coalitions of size 1, the two coalitions that minimise and maximise $\mathcal{F}$ are $\{a_1\}$ and $\{a_n\}$, respectively. Similarly, out of all the coalitions of size 2, the coalitions that minimise and maximise $\mathcal{F}$ are $\{a_1, a_2\}$ and $\{a_{n-1}, a_n\}$, respectively, and so on. Thus, given $C_s^{min}$ and $C_s^{max}$, in order to find the minimum and maximum values of stratum $s+1$, we would only need to add one new agent to each of those two coalitions, yielding a time complexity of $O(2)$ per stratum. Since under size-based stratification there are exactly $n$ strata, the overall complexity of finding the ranges will be $O(2n)$. □

#### 4.2.3.2 Output sharing games

Consider a group of companies who collectively own a technology that is used to produce a good, and each company, $a$, contributes an input, $\ell_a \in \mathbb{R}_+$, to the production which is modeled as a function, $g$, of the inputs. This is known as an *output sharing game* (Moulin, 1995; Corchon and Puy, 2000). Given a coalition of companies, $C$, when $g$ is a non-decreasing superadditive function of $\sum_{a \in C} \ell_a$, the output sharing game $(A, g)$ exhibits the order-reflecting property. This is because since $g$ is non-decreasing, given any two coalitions $C_1$ and $C_2$, such that $|C_1| = |C_2|$, the inequality $MC(a, C_1) \geq MC(a, C_2)$ holds by assumption. Therefore, the function $\mathcal{F}(C) = \sum_{a \in C} \ell_a$ satisfies property (4.12).

Again, to achieve a constant time complexity per stratum, the set of agents, $A$, must be sorted in the ascending order of the agents' inputs. Then, the minimum marginal contribution in stratum $s$ corresponds to the coalition that consists of the first $s$ agents in $A \backslash \{a\}$. Likewise, the maximum marginal contribution in $s$ corresponds to the coalition that consists of the last $s$ agents in $A \backslash \{a\}$.

#### 4.2.3.3 Airport games

An *airport game* is concerned with dividing the cost of an airport runway among several aircraft types, each of which requiring different lengths of the runway (Littlechild and Owen, 1973). Each aircraft type has a cost, $q_a \in \mathbb{R}_+$, and the value of a coalition is given as $v(C) = \max_{a \in C} q_a$. It can be easily shown that an order-reflecting function in airport games is $v(C)$. Furthermore, if the agents in $A$ are sorted in the *descending* order of the costs, the range of the strata can be found in constant time, in the same way as newsvendor and output sharing games.

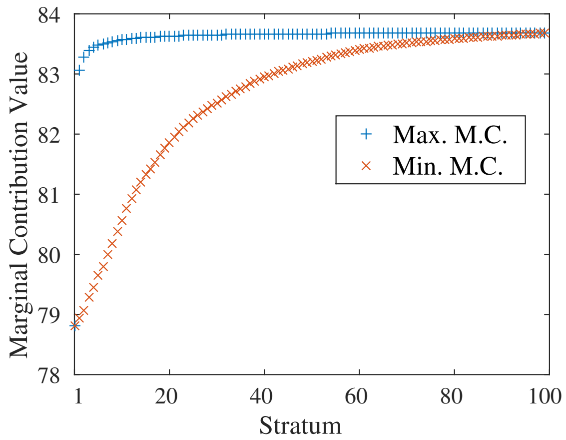


FIGURE 4.2: The range of the strata given by size-based stratification for agent 50 in the newsvendor game

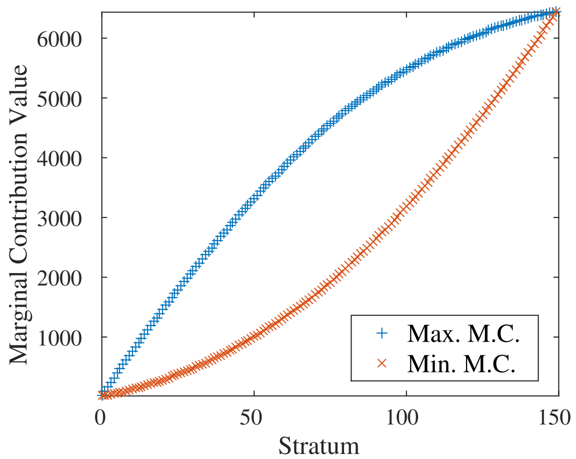


FIGURE 4.3: The range of the strata given by size-based stratification for agent 50 in the output sharing game

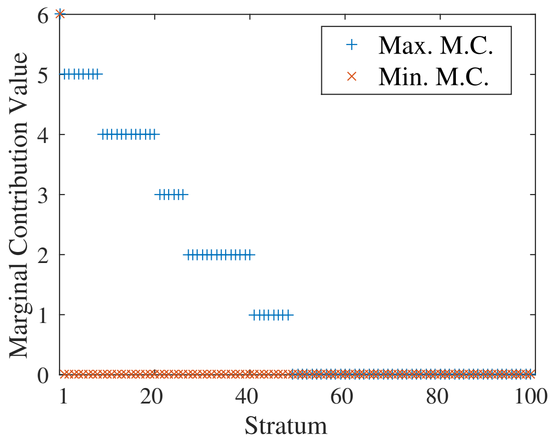


FIGURE 4.4: The range of the strata given by size-based stratification for agent 50 in the airport game

## 4.3 Experimental Results

We now evaluate stratified sampling our approach by applying it to instances of newsvendor games, output sharing games, and airport games.

First, we consider a newsvendor game consisting of 100 wind power producing DERs spread across the United States of America, that as a coalition, sell their next day's worth of generation ahead of time. We use the Shapley value to fairly divide the total expected profit of the coalition. To find coalition values, we calculate representative values for the amount of electricity each producer expects to generate in a day, by using real local wind speed forecasts at the location of each wind farm, and calculating the mean and variance of power generation over a day. Based on average spot prices provided by www.ferc.gov, we assume that the electricity unit price is \$42 per MWh, and surplus electricity is bought from the generators at \$21/MWh, and shortfall is penalized at \$84/MWh, which result in a $\mathcal{K}$ value of 14.6357.

Second, we consider an output sharing game consisting of 150 agents, with the characteristic function $v(C) = (\sum_{a \in C} \ell_a)^2$, where the agent inputs are randomly drawn from a uniform distribution on $[1, 10]$.

Finally, we consider an airport game with 100 aircraft types, identical to the one used by Castro et al. (2009) in evaluating their approximation algorithm that uses the CLT to bound the approximation error. In this game, the costs associated with the aircraft types are defined as follows:

$$q = \{q_{a1}, q_{a2}, \ldots, q_{a100}\} = \{\underbrace{1,\ldots,1}_{\text{8 times}}, \underbrace{2,\ldots,2}_{\text{12 times}}, \underbrace{3,\ldots,3}_{\text{6 times}}, \underbrace{4,\ldots,4}_{\text{14 times}}, \underbrace{5,\ldots,5}_{\text{8 times}}, \underbrace{6,\ldots,6}_{\text{9 times}}, \underbrace{7,\ldots,7}_{\text{13 times}}, \underbrace{8,\ldots,8}_{\text{10 times}}, \underbrace{9,\ldots,9}_{\text{10 times}}, \underbrace{10,\ldots,10}_{\text{10 times}}\}$$

For each game, we estimate the Shapley value of all agents with a 99% confidence, and an $\epsilon$ value equal to 1% of the range of each agent's population range. With these parameters, the sample size given by Theorem 4.3 is always 26492. Using this sample size, we compare the average of maximum standard errors (as discussed in Section 4.2) across all agents. Specifically, we evaluate three methods: (i) SRS, (ii) branching stratified sampling with three different sets of branching agents, consisting of 5 consecutive agents from the beginning, middle, and end of $A \backslash \{a\}$, respectively, and (iii) size-based stratified sampling. Furthermore, in each stratified sampling, of the 26492 observations, we dedicate 2 to each stratum to account for the marginal contributions required to find the ranges, and then distribute the rest according to equality (4.11).

In Figures 4.5, 4.6 and 4.7 we show the average standard error with sample sizes ranging from 2 per stratum to 26492. As expected from Theorem 4.4, SRS consistently has a

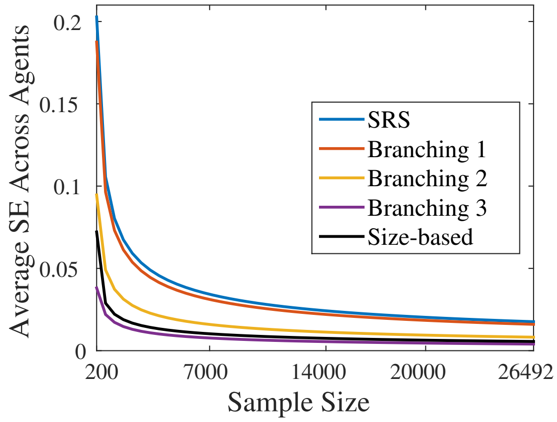


FIGURE 4.5: Comparison of average standard error across all agents in the newsvendor game

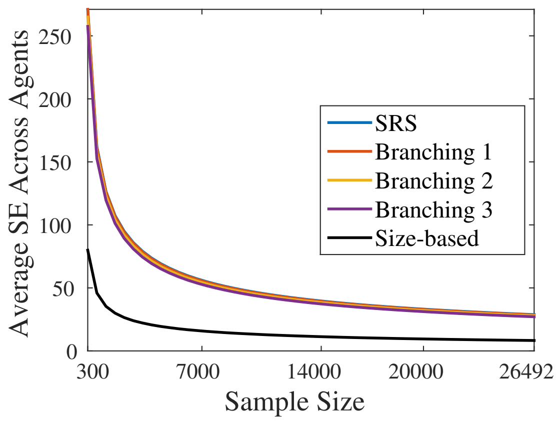


FIGURE 4.6: Comparison of average standard error across all agents in the output sharing game

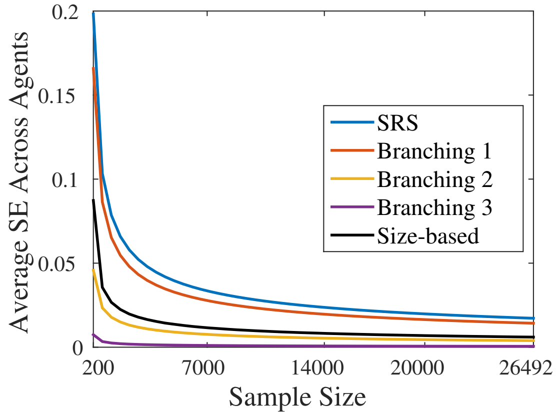


FIGURE 4.7: Comparison of average standard error across all agents in the airport game

| Method | Newsvendor | Output Sharing | Airport |
|---|---|---|---|
| Branching 1 | 9% | 1% | 16% |
| Branching 2 | 53% | 3% | 97% |
| Branching 3 | 77% | 6% | 97% |
| Size-based | 68% | 71% | 74% |

TABLE 4.1: Improvement of stratified sampling upon SRS

higher standard error in all games. Also, as can be seen, even with smaller sample sizes one could obtain a significantly lower standard error using stratified sampling.

To further evaluate our approach, for each stratification method, we calculate the improvement percentage as $(SE(\Phi_{SRS}) - SE(\Phi_{STR}))/SE(\Phi_{SRS})$, and report the results in Table 4.1. Each number represents how much lower the standard error of its corresponding method is as compared to SRS. On average, the stratified sampling methods have a standard error 48% lower than that of SRS.

## 4.4 Summary

In this chapter, we considered sampling-based approximation of the Shapley value. Specifically, we addressed the issues with the existing works in the that were highlighted in Section 2.2.2, namely erroneous and asymptotic error bounds, as well as the inefficiency of SRS. To address the former, we proposed an error bound based on Chebyshev's and Hoeffding's inequalities. To address the latter we proposed the use of stratified sampling, and theoretically proved that it can result in approximations with potentially smaller standard error due to more efficiently chosen samples. Furthermore, we proposed two stratification methods, namely branching stratification and size-based stratification. The former can be readily applied in supermodular games, a large an important class of games, while the latter can be applied to a class of games that have the order-reflecting property. We showed that some interesting class of games including newsvendor games, whose application was extensively discussed in Chapter 3, exhibit this property. Finally, we experimentally evaluated our stratified sampling approach and benchmark it against the existing methods in 12 instances of newsvendor, output sharing, and airport games. The results showed that, on average, the stratified sampling approaches have a standard error 48% lower than that of SRS.

# Chapter 5

# The Shapley Value In Games With Bounded Rational Agents

The idea of bounded rationality in cooperative games was put forward by Sandholm and Lesser (1997). They investigated a case where agents with finite computational resources form coalitions whose values are determined by a hard optimisation problem. Specifically, Sandholm and Lesser showed how the value of a grand coalition of agents can be distributed among its members such that the coalition remains stable. However, they did not consider the issue of dividing the this value in a fair way, and this problem has remained uninvestigated to date.

In order to investigate this issue, this chapter builds upon the bounded rationality model proposed by Sandholm and Lesser (1997), and extends it to the Shapley value in order to divide the value of a game with a characteristic function that has a high computational cost. Recall that the Shapley value is calculated using a formula, the input of which is the value of all possible coalitions. Here, the focus is on settings where the computational challenge arises from the hardness of computing the coalition values themselves; calculating the formula itself is relatively easy in comparison. This is quite the opposite of the typical computational issues considered in the literature, where the underlying assumption is that a coalition's value can be obtained in constant time, and the main challenge is in calculating Shapley's formula. In the setting considered here, even a single coalition value can be hard to compute. As such, approximation methods such as those in Chapter 4 cannot overcome the computational challenge, and a fundamentally different approach is needed. Furthermore, approximating the Shapley value would not necessarily result in a fair division, simply because the approximation would not be equal to the Shapley value—the unique value satisfying Shapley's axioms (Theorem 2.28).

The rest of this chapter is organised as follows. Section 5.1 presents the bounded rationality model and how the Shapley value based on this model can be calculated. In Section 5.2 an efficient implementation of the Shapley value is presented, where the value of each coalition is computed only once. Section 5.3 describes a demand response program as a platform for studying the Shapley value of bounded rational agents in a real world problem. Section 5.4 summarises the chapter.

## 5.1 The Shapley Value of Bounded Rational Agents

A common assumption in the literature is that the characteristic function has a negligible computational cost, e.g., it can be done in constant time. However, in many real world problems, such as the one considered in Section 5.3, and those considered by Alam et al. (2013); Aziz et al. (2014), computing the value of a coalition $C$ involves solving a hard optimisation problem. Henceforth, this will be referred to as the *optimisation problem of coalition* $C$. Since $v(C)$ is hard to compute, it will be referred to $v(C)$ as the *rational value* (rather than simply the *value*) of $C$. The corresponding *bounded rational game* of $(A, v)$ is a new game, denoted by $(A, v^{BR})$, where $v^{BR}(C)$ is the best solution to the optimisation problem of $C$ that is obtained given the available computational resources. Moreover, $v^{BR}(C)$ will be referred to as the *bounded rational value* of coalition $C$.

Under the full rationality assumption, the rational value of every coalition is known, and thus, the value of the game can be fairly divided using the Shapley value. However, with bounded rational agents, the rational values of the coalitions are unknown, and it is not immediately clear how a fair division of the value of the game can be obtained. One way to deal with this issue is proposed next.

Let us assume that it is possible to find a suboptimal solution to the optimisation problem of every coalition in reasonable time (i.e., the bounded rational value), and that the value of the coalitions are computed using the same hardware. Furthermore, in order to treat all agents indiscriminately, assume that the algorithm whereby the coalition values are found is the same for all coalitions.

**Proposition 5.1.** *Given a game, $(A, v)$, by allocating to each agent its Shapley value of the corresponding bounded rational game, $(A, v^{BR})$, one would obtain a payoff division that is fair in the following sense: (i) it fairly rewards each agent for its contribution to $v(A)$, and (ii) fairly penalises each agent for its contribution to reducing $v(A)$ to $v^{BR}(A)$.*

To better understand the intuition behind Proposition 5.1, some additional notation are needed. Denote by $SV(a, v)$, the Shapley value of agent $a$ that is calculated using a characteristic function $v$. Furthermore, for every $C \subseteq A$, let $v^{RD}(C)$ denote the *rationality discrepancy* of coalition $C$, defined as the difference between the rational and bounded rational values of $C$. More formally,

$$v^{RD}(C) = v(C) - v^{BR}(C). \tag{5.1}$$

The rationality discrepancy of $C$ can be viewed as the payoff that $C$ loses due to its members' lack of full rationality. Clearly, the more rational a coalition's members are (i.e., more computational resources are used to compute its value), the smaller its rationality discrepancy is. As such, if the members of $C$ were fully rational, they would be able to diminish their loses completely, i.e., we would have: $v^{RD}(C) = 0$.

Observe that the game $(A, v)$ is the sum of the games $(A, v^{BR})$ and $(A, v^{RD})$ (see equation (5.1)). Therefore, based on the additivity axiom of the Shapley value (defined in Section 2.2), the following holds:

$$SV(a, v^{BR}) = SV(a, v) - SV(a, v^{RD}). \tag{5.2}$$

Note that, by assumption, neither $SV(a, v)$ nor $SV(a, v^{RD})$ can be computed. However, $SV(a, v^{BR})$, which indeed can be computed, is agent $a$'s fair share of the value of the game (i.e, $v(A)$) combined with its fair share of the loss in value due to the agents' bounded rationality (i.e., $v^{RD}(A)$).

## 5.2 An Efficient Implementation of the Shapley Value

Calculating the Shapley value of all agents using the standard formula, i.e., equation (2.2), requires computing the value of each coalition multiple times. This is because for each $C \subseteq A$, the value of $C$, i.e., $v(C)$, is used in calculating the marginal contribution of all agents that are not members of $C$. This is not an issue in games where the characteristic function does not have considerable computational complexity. However, given the bounded rationality assumption, computing $v(C)$ more than once is highly costly. One trivial way to overcome this issue is to store all coalition values in memory. Clearly, however, doing so requires exponential memory space. Alternatively, one can compute the values of all coalitions one at a time and, for each value, update *all* marginal contributions of *all* agents that require that value. This ensures that the value of any given coalition is computed exactly once. What follows explains this process in more detail.

First, observe that calculating the marginal contribution of agent $a$ to a coalition, $D$, requires the following two terms: $v(D \cup \{a\})$ and $v(D)$, which represent the value of the coalition *with* and *without* the agent, respectively. More formally:

$$MC(a, D) = v(D \cup \{a\}) - v(D). \tag{5.3}$$

Now, for each $C \subseteq A$, one can use $v(C)$ to calculate the marginal contributions of two groups of agents:

- For all $a \in C$, $v(C)$ can be taken as the value of a coalition with the agent, i.e., the $v(D \cup \{a\})$ in equation (5.3);
- For all $a \in A \backslash C$, $v(C)$ can be taken as the value of a coalition without the agent, i.e., the $v(D)$ term in equation (5.3).

Furthermore, recall that the Shapley value in equation (2.2) is a weighted sum of marginal contribution. Let the weight of the marginal contribution of an agent to a coalition $D$ be denoted by $\omega(|D|)$. Formally, this weight is given as:

$$\omega(|D|) = \frac{|D|!\,(n - |D| - 1)!}{n!}$$

Multiplying each term of the marginal contribution terms by its corresponding weight yields $\omega(|D|)\,v(D \cup \{a\})$ and $-\omega(|D|)\,v(D)$. Therefore, to calculate the Shapley value, we can sum all coalition values multiplied by their corresponding weight just as in equation (2.2):

- For every non-empty coalition $C \subset A$, $v(C)$ is multiplied by $\omega(|C| - 1)$ for every $a \in A \backslash C$, and multiplied by $-\omega(|C|)$ for every $a \in C$, $v(C)$.
- For $C = \emptyset$, $v(C)$, which only represents the value of a coalition without the agent, is multiplied by $-\omega(0)$.
- For $C = A$, $v(C)$, which only represents the value of a coalition with the agent, is multiplied by $\omega(n - 1)$.

The pseudocode of this process is presented in Algorithm 1.

**Algorithm 1** Implementation of the Shapley Value which computes every coalition value exactly once

```
function ShapleyValue(A, v)
SV ← [];
∀a ∈ C, SV(a) ← 0;
for all C ⊆ A do
  coefficientWhenAgentIsIn ← ω(max(|C| − 1, 0));
  coefficientWhenAgentIsOut ← −ω(min(|C|, |A| − 1));
  for all a ∈ A do
    if a ∈ C then
      SV(a) ← SV(a) + (coefficientWhenAgentIsIn × v(C));
    else
      SV(a) ← SV(a) + (coefficientWhenAgentIsOut × v(C));
    end if
  end for
end for
return SV
```

## 5.3 Cooling Load Demand Response Program

Having described the bounded rational Shapley value proposition, this section illustrates how this can be applied in a real world problem. The context that is chosen for this purpose is the problem of dividing a discount that a group of apartments in a block obtain by ensuring that their aggregate load from cooling always remains below a certain threshold. In more detail, the apartments coordinate the periods during which their air conditioners (AC) are tuned on. The goal of this coordination is to satisfy the temperature preferences of each apartment, while ensuring that the aggregate load does not exceed a certain predetermined limit. Using a thermal model of an apartment to model the evolution of internal temperature over time, this problem is formulated as a binary integer program. Due to the substantial time required to compute the optimal load of even a single coalition, it is not possible to use the Shapley value to fairly divide the total discounted cost that the group is charged. Instead, one should use the bounded rational Shapley value as a fair division of the cost.

The rest of this section formalises the problem, describes the thermal model of an apartment, and presents the algorithms required to compute the bounded rational values of coalitions. Finally, the experimental results show how the bounded rational Shapley value can divide the discounted costs.

### 5.3.1 The Discount Scheme

Consider a set of $n$ apartments in a block, $A$. Denote by $l_a^t$ the cooling load of apartment $a$ at time $t$ (measured in kW), and denote by $p$ the price at which every kWh is charged. In order to encourage consumers to use less energy for air conditioning, which constitutes a significant amount of the domestic load in warm-climate countries (McNeil and Letschert, 2008; Hsu and Su, 1991), the electricity supplier offers a discount to the block. Specifically, each apartment is offered a binary option of signing up to the scheme or not. Let $K = \{1, \ldots, k\}$ represent an entire day divided into $k$ equal-length time slots, and $l_a = [l_a^1 l_a^2 \ldots l_a^k]$ represent the vector of cooling loads of apartment $a$ in all time slots in $K$. If at any point in time throughout the day, the cooling load of the block is not more than $\psi$ kW, i.e., $\forall t \in K; \sum_{a \in A} l_a^t \leq \psi$, then those apartments that have signed up are charged at $d < p$ per kWh of usage, and the rest are charged at $p$ per kWh. The cooling load of apartment $a$ at time $t$ is:

$$l_a^t = P_a \times \eta_a^t, \tag{5.4}$$

where $P_a$ is the electric power of the AC (in kW), and $\eta_a^t \in \{0, 1\}$ represents a *cooling action*, which is a binary variable that indicates whether or not air conditioning has been used at time $t$.

Since the discount is offered only when the whole block's load is below $\psi$, the price at which a coalition $C \subseteq A$ is charged is influenced by the behaviour of the apartments that are not members of the coalition, i.e., $A \backslash C$. In other words, the value of a coalition is influenced by the load of other apartments outside that coalition. If these apartments could form other coalitions, then it would be a game with *externalities* (also known as a partition function game), which is a game where the value of a coalition depends on how other agents are structured (Thrall and Lucas, 1963). However, since in the discount scheme, the agents that do not sign up cannot form any other coalitions, this is a special case of externalities where the agents outside the coalition can only be structured as singletons, and thus, the game is reduced to a characteristic function game.

Naturally, each apartment, whether signed up to the scheme or not, would want to optimise its use of the AC such that its internal temperature preferences are satisfied with minimal electricity consumption. In order to secure the discount, those apartments that sign up need to coordinate their loads so that the aggregate load of the block will be kept below the threshold and their internal temperatures remain as they individually deem comfortable. Clearly, if an apartment decides not to sign up to the scheme, and desires to optimise its load, it can only do so independently, without any coordination with other apartments.

Assuming that some of the $n$ apartments form a coalition, $C$, and sign up to the scheme, the aggregate cooling load of all $n$ apartments (in kW) at time $t$ is given by:

$$l_A^t = l_C^{*t} + \sum_{a \in A \backslash C} l_a^{*t},$$

where $l_C^{*t}$ represents the aggregate optimal cooling load at time $t$ of the apartments that have signed up, and $l_a^{*t}$ represents the optimal cooling of apartment $a$ at time $t$ which has not signed up to the scheme.

Note that a coalition can meet the threshold only by running its members' ACs for longer periods (mostly during the off-peak times). This is because when an apartment has to avoid running its AC during its comfort period, it has to run the AC earlier to cool down the apartment enough so that at the start of the comfort period (when the AC is off) the internal temperature equilibrates the set-point. Therefore, meeting the threshold requires extra consumption of electricity, which increases the cost. As such, an outcome of collective optimisation of loads is desirable unless the extra consumption is so high that the discounted cost becomes higher than the normal cost. Consequently, if the discounted cost turns out to be higher, or if a feasible solution to the collective optimisation cannot be found, then the apartments optimise their loads independently. Based on this, the optimal consumption of a coalition (measured in kWh) can be written as:

$$c(C) = \begin{cases} \sum_{t\in K} l_C^{*t} \times \Delta t & \sum_{t\in K} l_C^{*t}\, \Delta t\, f \leq \sum_{a\in C}\sum_{t\in K} l_a^{*t}\, p \\ \\ \sum_{a\in C}\sum_{t\in K} l_a^{*t} \times \Delta t & \sum_{t\in K} l_C^{*t}\, \Delta t\, f > \sum_{a\in C}\sum_{t\in K} l_a^{*t}\, p \end{cases}, \qquad (5.5)$$

where $\Delta t$ is the duration of a time slot (in seconds). Based on the above consumption function, we now define the characteristic function, $v$, of the cooperative game $(A, v)$ that represents the above discount scheme. In more detail, $v(C)$ is equal to the total cost of consumption of its members. More formally, $v(c)$ is given by:

$$v(C) = c(C) \times \begin{cases} d & \forall t \in K \; ; \; l_A^t \leq \psi \\ p & \forall t \in K \; ; \; l_A^t > \psi \end{cases} \qquad (5.6)$$

With respect to the above characteristic function and the binary choice that the apartments are offered in the scheme, it is clear that it would be in the interest of each apartment to sign up and benefit from the potential discount, because each apartment will not be worse off by joining the grand coalition.

The next subsection explains how the cooling load of each apartment can be optimised such that the temperature preferences of the apartments are satisfied while minimising consumption. First, a formal model of thermal dynamics of an apartment is presented, which governs the evolution of the apartments' internal temperature.

### 5.3.2 Thermal Dynamics of An Apartment

We use a standard thermal model in which heat is assumed to enter an apartment (by thermal conduction and ventilation) at a rate that is proportional to the temperature difference between the cold air inside and the hot air outside (Y. Guo and Zeman, 2008; Rogers et al., 2011; Andersen et al., 2000). This model also incorporates the thermal capacity of the building structure, since through experimentation on real data collected from apartments in Jeddah, Saudi Arabia it was found that this model best explains the observed data. This thermal model is represented as a set of coupled difference equations as per:

$$\begin{aligned} T_{int}^{t+1} &= T_{int}^{t} - r\eta^{t}\Delta t + \tau\Delta t \left(T_{env}^{t} - T_{int}^{t}\right) \\ T_{env}^{t+1} &= T_{env}^{t} + \rho\Delta t \left(T_{int}^{t} - T_{env}^{t}\right) + \gamma\Delta t \left(T_{ext}^{t} - T_{env}^{t}\right), \end{aligned} \qquad (5.7)$$

where $T_{int}^t \in \mathbb{R}^+$ denotes the internal temperature (measured in °C) of apartment $a$ at time $t$, $T_{env}^t \in \mathbb{R}^+$ denotes the temperature of the building structure, or envelope, (measured in °C), and $T_{ext}^t \in \mathbb{R}^+$ denotes external temperature (measured in °C). Assume that $T_{ext}^t$ is the same for all apartments in $A$. Moreover, $r$ (measured in °C/hr) represents the rate at which the AC reduces the internal temperature, and $\tau, \rho$ and $\gamma$ (measured in 1/hr) are the rates of leakage from the envelope to the inside, from the inside to the envelope, and from the outside to the envelope, respectively. Henceforth, notations are indexed by the name of the agents.

Equation (5.7) is the discrete equivalent to a set of coupled differential equations which has been used previously to model data collected from real buildings (Bacher and Madsen, 2011). In this model, an envelope is introduced to act as an additional thermal mass to minimise temperature deviations inside the apartment due to extremes of temperature outside. Given historical observations of internal, $T^t_{int}$, and external, $T^t_{ext}$, temperatures and the times during which the AC was on (which were collected from a number of apartments in Jeddah, Saudi Arabia) the evolution of the internal temperature, $\overline{T}^t_{int}$, was predicted. The error in this prediction is given by $\sum_{t \in K} \left( \overline{T}^t_{int} - T^t_{int} \right)^2$. Consequently, the best estimates of the parameters are those that minimise this error and can subsequently be learned through recursive least squares (LSQ) (Simon, 2006).

### 5.3.3 Comfort Model

What follows outlines a few assumptions that underpin the operation of the cooling system in an apartment. Each apartment is assumed to have a central air conditioning driven by a heat pump that transfers heat from a lower temperature heat source (the apartment) into a higher temperature heat sink (external ambient air). This system is connected to a thermostat within the apartment, where a user can set a desired temperature to be maintained, i.e., the set-point temperature, which is denoted as $T_{set}$ (°C).

The user in each apartment can specify the time interval during which he or she desires "comfort". That is, the time slots when the user wants the internal temperature to be maintained at, or close to, $T_{set}$. This interval is referred to as the *comfort period*, and define it as: $H = \{t \in K | \ CST \leq t \leq CET\}$, where $CST \in K$ and $CET \in K$ are the comfort start time and comfort end time, respectively. A tolerance level is specified by the user to limit deviations of the internal temperature from the set-point temperature, during the comfort period. Let this tolerance be denoted by $\theta \in \mathbb{R}^+$ (°C). Furthermore, the AC in each apartment ensures that the *average internal temperature* during the comfort period is at most $\Omega \in \mathbb{R}^+$ (°C) different from the set-point temperature. This will ensure that the internal temperature during the comfort period is on average close to $T_{set}$. Note that $\Omega$ is a parameter of the model, which is equal for all apartments, and is not set by users. However, $T_{set}$ and $\theta$ are determined individually by each user. Lower values of $T_{set}$ suggest that a user feels more comfortable at lower temperatures. Similarly, smaller values of $\theta$ indicate that a user is sensitive to large deviations in the internal temperature from the set-point.

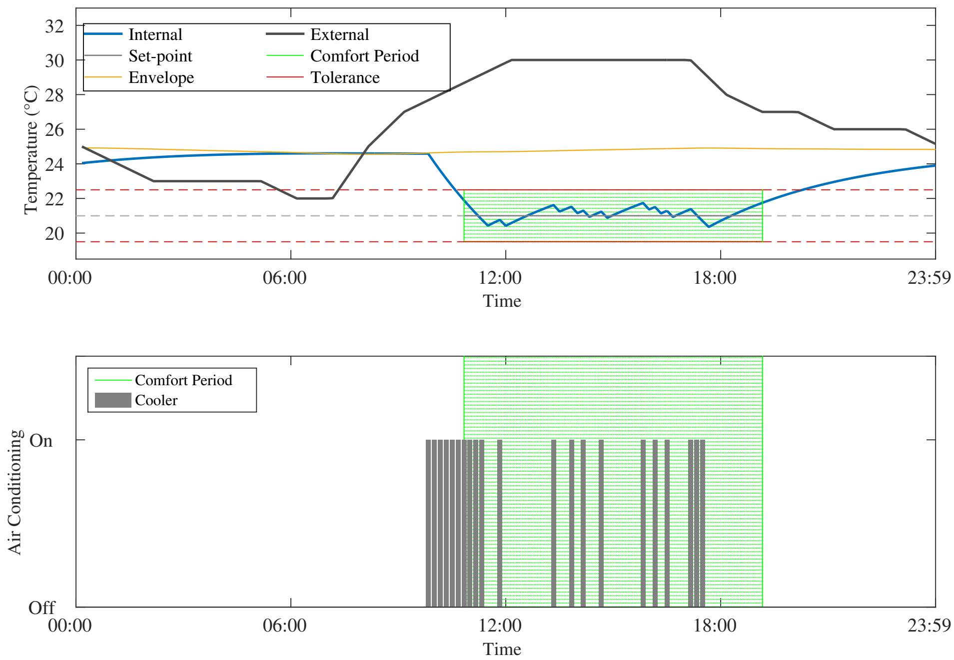


FIGURE 5.1: Understanding the impact of $T_{set}$=21°C.

Intrinsically, the above preferences have an impact on the cooling load (equation (5.4)). As $T_{set}$ is gradually lowered, the amount of cooling required increases proportionately to achieve lower temperatures. Figures 5.1 and 5.2 show the internal temperature profile in an apartment when $T_{set}$ equals 21°C and 23°C, respectively. The bottom sub-plots in these figures show the cooling actions over the course of a day. It is clear that more cooling is required in Figure 5.1 compared to Figure 5.2. The total time when the AC is on is 52% less when $T_{set}$ equals 23°C, as opposed to when $T_{set}$ equals 21°C.

Similarly, when $\theta$ is small, a user is more sensitive to deviations of the internal temperature from $T_{set}$. Consequently, the AC is turned on for longer to ensure that the deviation of the internal temperature from $T_{set}$ lies within the tolerance level, resulting in higher energy consumption. This is clearly shown in Figures 5.1 and 5.2, where the $\theta$ is set to 0.5°C and 1°C, respectively. The bottom sub-plots in these figures show the cooling actions over the course of a day. As can be seen in Figures 5.3 and 5.4, more cooling is required for a larger $\theta$. The total time when the AC is on is 12% less when $\theta$ is set to 1.5°C compared to when $\theta$ is set to 0.6°C.

As per equation (5.7), an apartment that is well-insulated will have a small value of $\gamma$. In contrast, a leaky apartment will have a high value of $\gamma$. This is of interest as more cooling is required to maintain a leaky apartment at a certain temperature. The effect of varying values of $\gamma$ are shown in Figures 5.5 and 5.6. The value of $\gamma$ used to generate Figures 5.5 and 5.6 are 0.36 1/hr and 0.48 1/hr, respectively. Consequently, the total

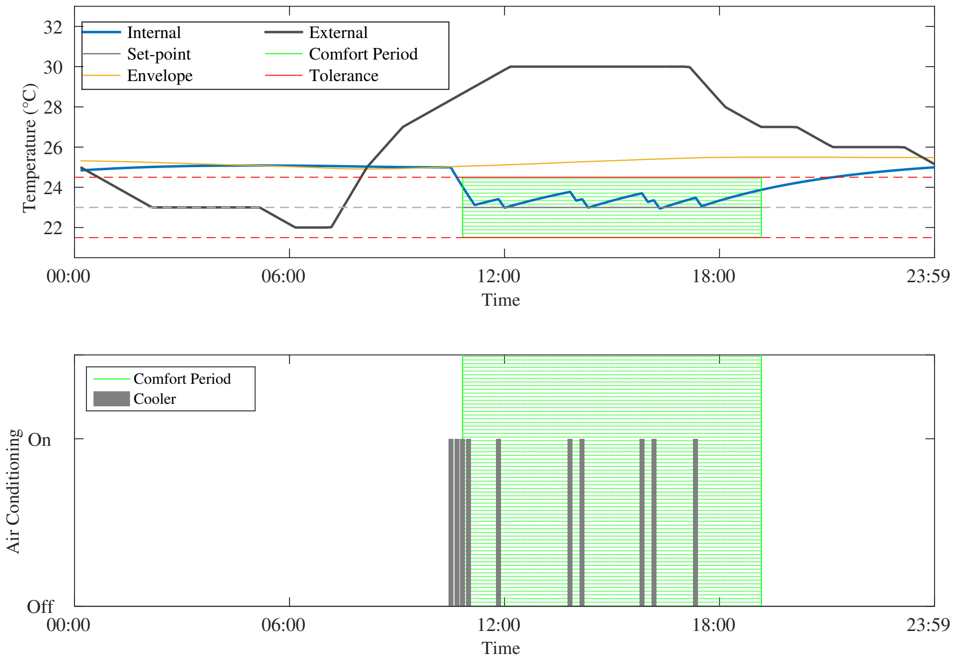


FIGURE 5.2: Understanding the impact of $T_{set}$=23°C.

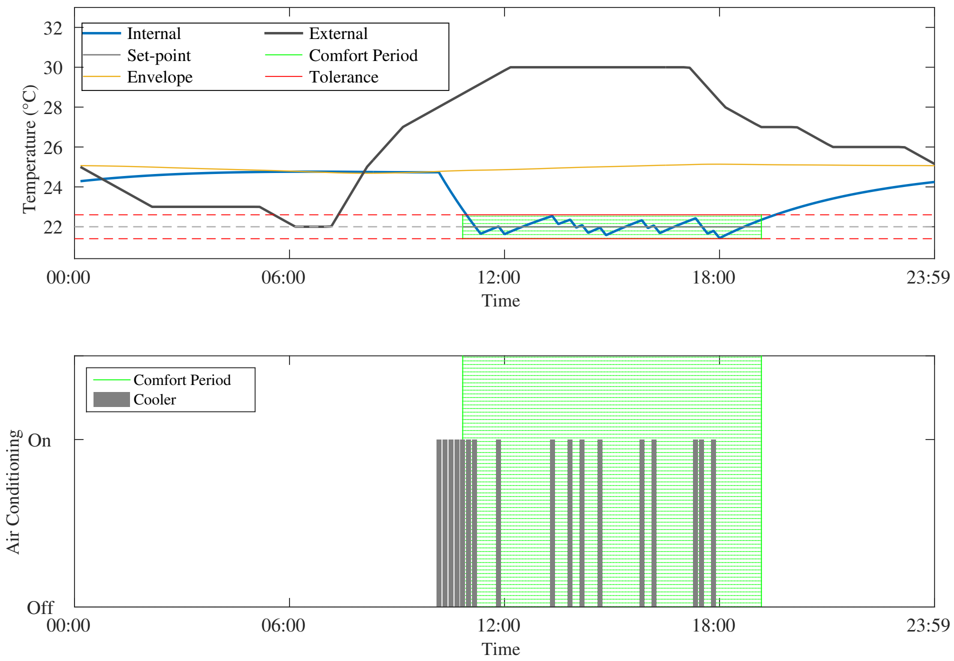


FIGURE 5.3: Plotting the effect of $\theta$=0.6°C, represented as a band around the set-point temperature in the top sub-plot, on the thermal dynamics.

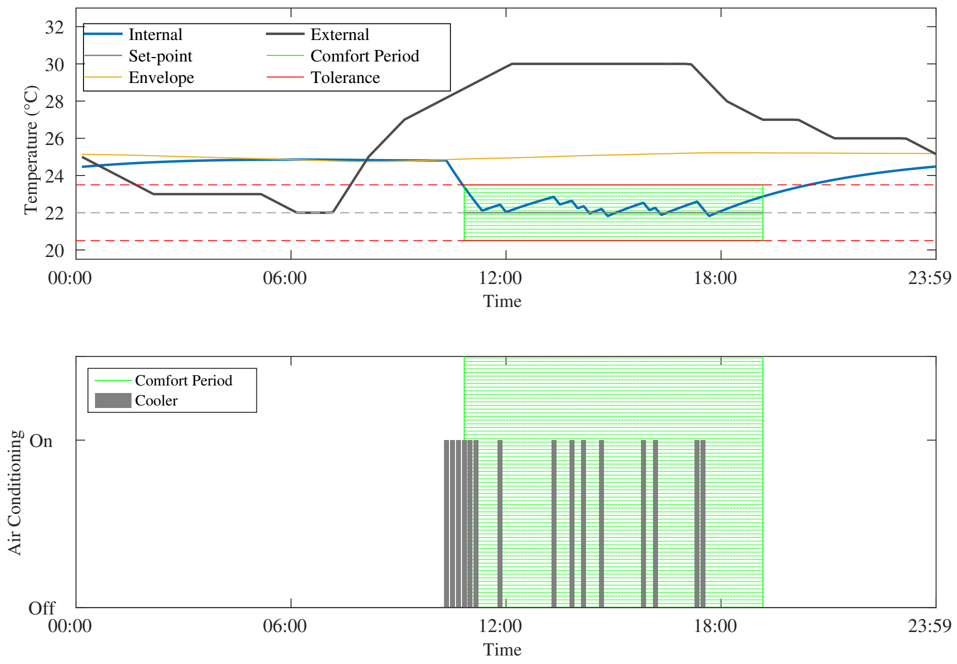


FIGURE 5.4: Plotting the effect of $\theta$=1.5°C, represented as a band around the set-point temperature in the top sub-plot, on the thermal dynamics.

amount of time when the AC is on, as shown in the bottom sub-plots, is 24% more in Figure 5.6, when compared with Figure 5.5.

### 5.3.4 Independent Optimisation of Loads

Having introduced the model for thermal dynamics of an apartment, this subsection describes how an apartment optimises its use of the AC so as to satisfy only its own comfort preferences. In summary, as described in Subsection 5.3.3, the preferences of an apartment are: (i) the desired set-point temperature, denoted by $T_{set}$, (ii) a tolerance level on the deviation of the internal temperature from the set-point, denoted by $\theta \in \mathbb{R}^+$, and (iii) the comfort start time ($CST$) and comfort end time ($CET$) which determine $H$—the set of time slots representing the comfort period. Based on these, an optimal cooling plan for apartment $a$ is defined to be a vector of cooling actions $[\eta^1\eta^2\ldots\eta^k]$, that result in meeting the above preferences as well as the following requirements: (i) the overall energy consumption is minimised, i.e., all constraints are satisfied with the AC running in as few time slots as possible, (ii) the deviation of the average internal temperature from the set-point temperature during the comfort period is limited by $\Omega \in \mathbb{R}^+$, to ensure that the internal temperature in this period stays close to the set-point temperature, not the highest acceptable temperature, (iii) the internal and envelope temperatures at the start and end of the day converge. Note that the last requirement is to ensure that the cooling plan is optimised over an infinite horizon, which prevents

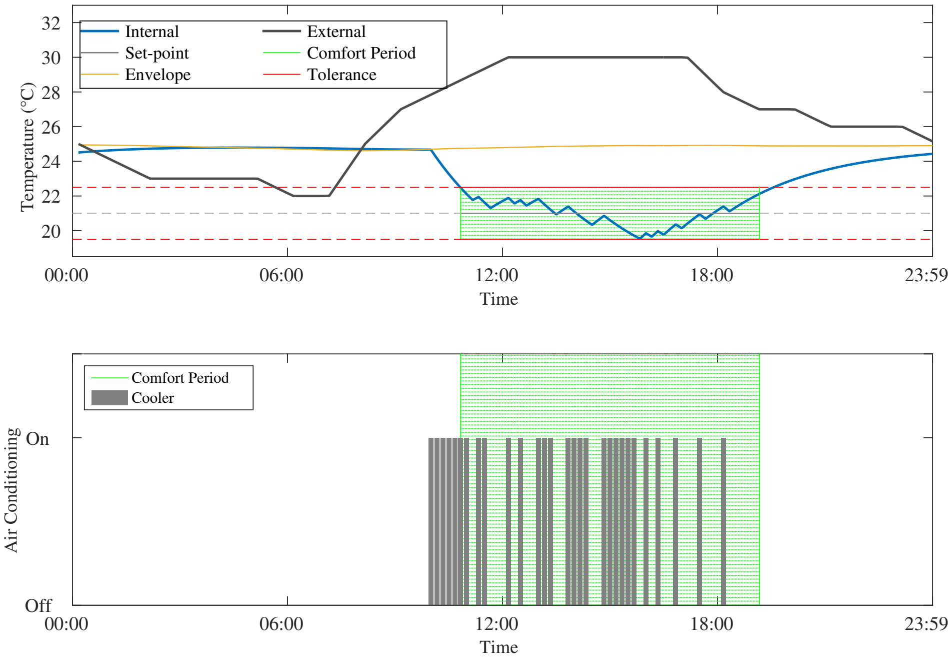


FIGURE 5.5: Plotting the effect of $\gamma = 0.36$ 1/hr on the thermal dynamics.

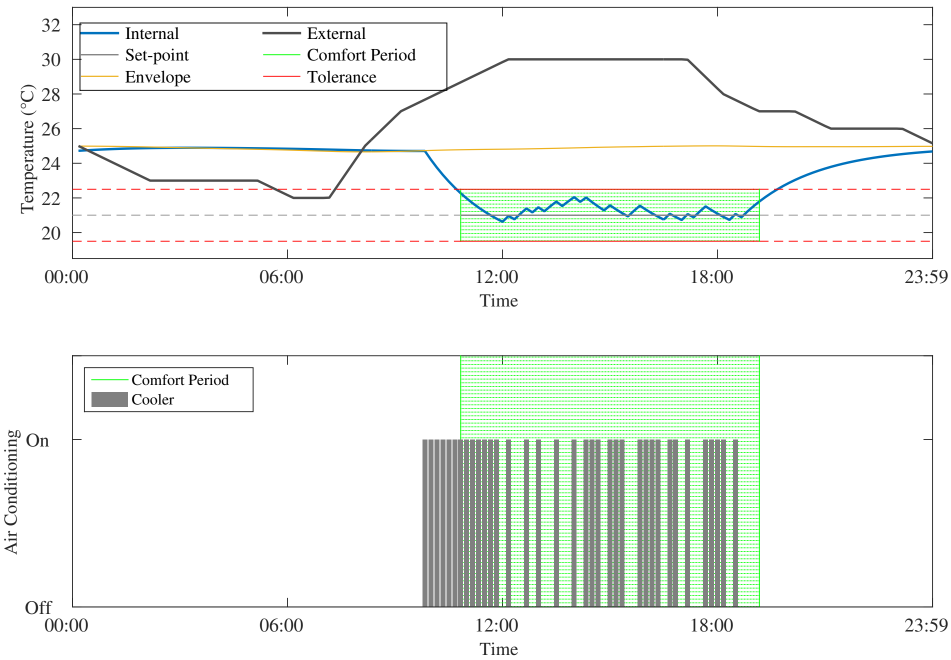


FIGURE 5.6: Plotting the effect of $\gamma$=0.48 1/hr on the thermal dynamics.

erroneous solutions that minimise AC use in the short term, but require additional cooling later, as would be the case if a finite planning horizon were used. Given the above preferences and requirements, the optimal cooling load of an apartment throughout the day can be computed as per equation (5.4) using $\eta^1$, $\eta^2$, ..., $\eta^k$ found by solving the following optimisation problem:

$$\begin{aligned}
&\text{minimise} \sum_{t \in K} \eta^t \\
&\text{subject to} \left| \frac{\sum_{h \in H} T^h_{int}}{|H|} - T_{set} \right| \leq \Omega, \\
&\forall h \in H, \ \ |T^h_{int} - T_{set}| \leq \theta, \\
&T^1_{int} = T^k_{int}, \\
&T^1_{env} = T^k_{env}.
\end{aligned} \tag{5.8}$$

Observe that the above formulation avoids the explicit trade-off between consumption and comfort within a single objective function, which is dependent on specifying appropriate weights for both objectives. This is because, in practice, there is no principled way to specify such weights (Aswani et al., 2012).

### 5.3.5 Collective Optimisation of Loads

This subsection describes how a coalition of apartments, $C$, collectively optimise their cooling loads. Similar to the single apartment case, the user in each apartment in $C$ specifies their individual cooling preferences. These include their desired set-point temperature, $T_{set}[a]$, their tolerance on the deviation of the internal temperature during the comfort period, $\theta_a$, and their comfort start and end times which determine $H_a$. Again, we ensure that in finding the optimal cooling plan of the coalition the following requirements are also met: (i) the overall energy consumption of each apartment should be minimised, (ii) the deviation of the average internal temperature of each apartment from its set-point temperature during the comfort period is not more than $\Omega$°C, (iii) the internal and envelope temperatures at the start and end of the day in each apartment converge. Additionally, we introduce a key constraint, which ensures that at all times, the total load of all apartments in the coalition, plus the total load of the apartments who optimise their loads independently, is less than or equal to $\psi$. More formally, the vector of optimal cooling actions, $[\eta^1_a \eta^2_a \dots \eta^k_a]$, for every apartment $a \in C$ is given by the following optimisation problem:

$$
\begin{aligned}
\forall a \in C \ \text{ minimise} \sum_{t \in K} \eta_a^t \text{ subject to:} \\
& \left| \frac{\sum_{h \in H_a} T_{int}^h(k)}{|H_a|} - T_{set}[a] \right| \leq \Omega, \\
& T_{int}^1[a] = T_{int}^k[a], \\
& T_{env}^1[a] = T_{env}^k[a], \\
& \forall h \in H \ \ |T_{int}^h[a] - T_{set}[a]| \leq \theta_a, \\
& \forall t \in K \ \ l_A^t = \sum_{a \in C} P_a \, \eta_a^t + \sum_{a' \in A \setminus C} l_{a'}^{*t} \leq \psi.
\end{aligned}
\tag{5.9}
$$

Note that if a feasible solution to the above optimisation does not exist, then the apartments optimise their loads individually as per equation (5.8). Furthermore, note that it is possible for an individual apartment within $C$ to have a significant impact on the feasibility of $C$ satisfying the constraint on the aggregate load. For instance, if $T_{set}[a]$ is set to a particularly low temperature, or $\theta_a$ is particularly small, the corresponding energy consumption in that apartment will be greater, which in turn increases the likelihood of the aggregate load exceeding the threshold. As the individual apartments become more flexible and less stringent with their preferences, the aggregate cooling load, $l_A^t$, is more likely to satisfy the constraint on the threshold.

### 5.3.6 Example

Having established the theoretical underpinnings of how cooling loads are independently and collectively optimised, this subsection illustrates how they work in practice through a simple example. Consider a 3-agent game ($A = \{1, 2, 3\}$), where apartments are located in the same block, and all three agree to participate in the discount scheme. Apartment 1 desires that the temperature be maintained at 21°C ($T_{set}[1] = 21$°C) for 6 hours from $CST = 10:00$ to $CET = 16:00$, and is satisfied with wide swings of temperature ($\theta_1 = 1.5$°C). Apartment 2 desires the temperature to be at 22°C ($T_{set}[1] = 22$°C) for 8 hours from $CST = 09:00$ to $CET = 17:00$, and has very strict preferences over temperature ($\theta_2 = 0.5$°C). Apartmentlat 3 too desires the temperature to be at 21°C ($T_{set}[3] = 21$°C) for 6 hours from $CST = 10:00$ to $CET = 16:00$, and is satisfied with wide swings of temperature ($\theta_1 = 1.5$°C).

The AC in each apartment operates on a 10-minute cycle, i.e., a cooling decision is made for each 10-minute interval in a day ($\Delta t = 600$s). As a result, $K = [1, \ldots, 144]$ and the decision variable is $\eta_a^t, \forall t \in K$. The ACs in all apartments are similar and consume at a rate of 3 kW when on, i.e. $P_a = 3$ kW. Consequently, when the AC is on for a an hour the energy consumed is 3 kWh. Hence, the total possible energy load

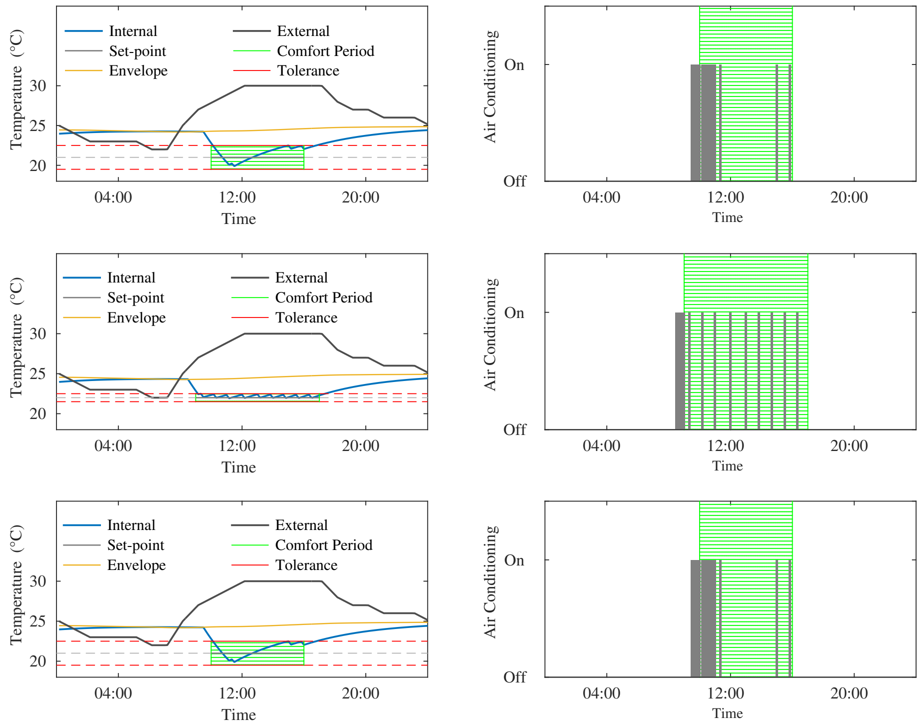


FIGURE 5.7: The temperature profiles (in the left column) and the cooling profiles (right column) for each apartment when they optimise their cooling loads individually.

of all three apartments if they optimise their apartments independently is 9 kW. Now, when the apartments participate in the scheme, a threshold at 3 kW (representing a 2/3 reduction), is set on their total load.

Let us first consider the case where all 3 apartments optimise their cooling loads independently. The plots in the second column of Figure 5.7 show the individual cooling profiles within each apartment for a single day. Each profile is obtained by solving the optimisation problem in equation (5.8) using CPLEX, to yield $\eta_a^t$ $(\forall t \in K)$. Also shown in the first column of Figure 5.8 are the corresponding internal temperature profiles, which are estimated by iterating equation (5.7) for each apartment, using the cooling actions, $\eta_a^t$ as inputs. It is evident from the plot that the internal temperature is, on average, maintained close to the desired set-point temperatures at times when a user desires cooling in each apartment. Also, the deviation from the set-point temperature is greater in in Apartment 1 and Apartment 2, as they are less sensitive to large swings in temperature. Finally, the optimisation ensures that the temperature at the start and end of each day is the same, as required.

Now consider the case where all 3 apartments form a coalition and collectively optimise their loads to ensure that their aggregate load does not exceed the threshold during the course of a day. The plots in the second column of Figure 5.8 show the cooling

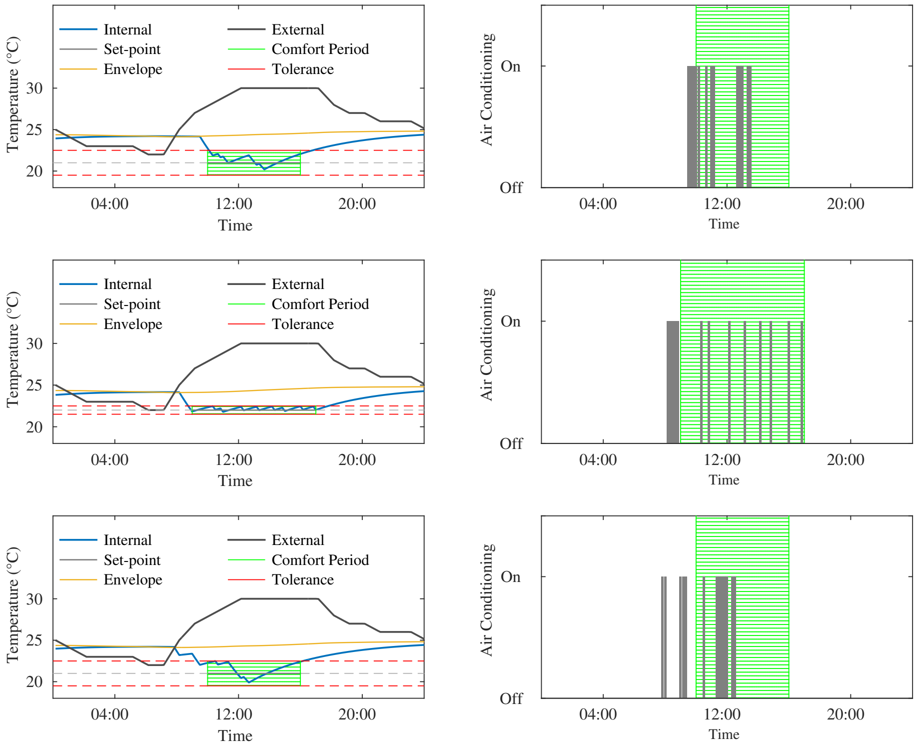


FIGURE 5.8: The temperature profiles (in the left column) and the cooling profiles (right column) for each apartment, when they optimise their cooling loads collectively.

profiles within each apartment for a single day. Each profile is obtained by solving the optimisation problem in equation (5.9) using CPLEX, to yield $\eta_a^t$ ($\forall t \in K$), which in turn generate a temperature profile based on the thermal model as per equation (5.7).

More importantly, as shown in Figure 5.9, collective optimisation of the loads results in the aggregate load never exceeding the threshold. In contrast, when the apartments independently optimise their cooling loads, as shown in Figure 5.9, the aggregate load does exceed the threshold. Thus, the apartments can receive the discount.

Since optimising the loads of the three apartments is tractable (optimising the grand coalition takes approximately 4 minutes), the total cost of the apartments can be divided based on the true Shapley value (as opposed to the bounded rational value). Assume that the electricity cost is £0.15 per kWh in the case when agents optimise their cooling loads independently. Now, as per the discount scheme, if they ensure that their aggregate load does not exceed 3 kW, then the electricity cost will be reduced to £0.06 per kWh. The payment of the apartments in this example, according to the Shapley value, is

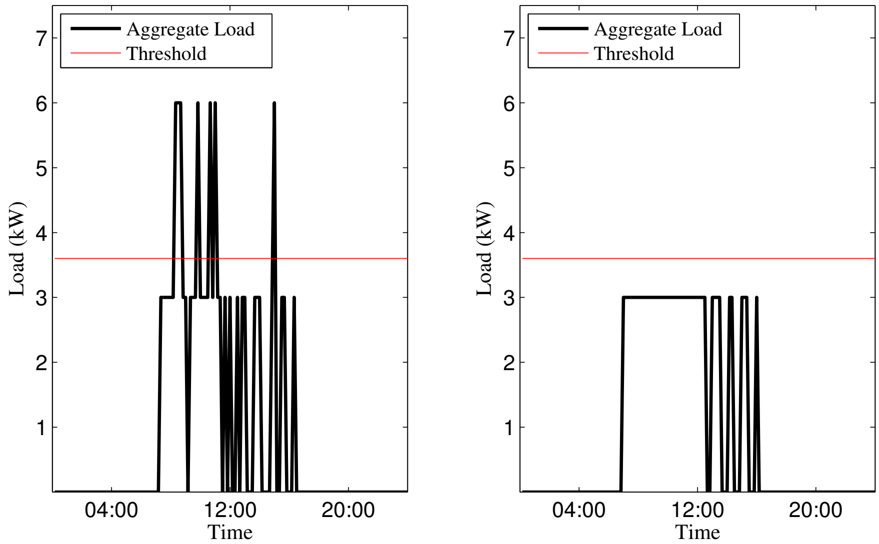


FIGURE 5.9: Total load when apartments optimise their cooling loads individually and collectively.

calculated as follows:

$$\phi[1,v] = \frac{1\times 5.85}{3} - \frac{1\times 1.17}{6} - \frac{1\times 1.35}{6} + \frac{1\times 2.34}{3} = £2.31$$
$$\phi[2,v] = \frac{1\times 5.85}{3} - \frac{1\times 1.17}{6} - \frac{1\times 1.17}{6} + \frac{1\times 2.52}{3} = £2.40$$
$$\phi[1,v] = \frac{1\times 5.85}{3} - \frac{1\times 1.17}{6} - \frac{1\times 1.35}{6} + \frac{1\times 2.34}{3} = £2.31$$

| Coalition ($C$) | Threshold | Cost | Independent Apartments ($A\backslash C$) | Cost |
|---|---|---|---|---|
| {} | Not Satisfied | £0 | {1},{2},{3} | £17.55 |
| {1} | Not Satisfied | £5.85 | {2}, {3} | £11.70 |
| {2} | Not Satisfied | £5.85 | {1}, {3} | £11.70 |
| {3} | Not Satisfied | £5.85 | {1}, {2} | £11.70 |
| {1, 2} | Satisfied | £4.68 | {3} | £5.85 |
| {1, 3} | Satisfied | £4.50 | {2} | £5.85 |
| {2, 3} | Satisfied | £4.68 | {1} | £5.85 |
| {1,2,3} | Satisfied | £7.02 | {} | £0 |

TABLE 5.1: Example showing the cost of coalitions when members optimise their loads collectively and independently.

Apartment 1 is charged £2.31, Apartment 2 is charged £2.40, and Apartment 3 is charged £2.31. When optimised collectively to keep the aggregate load below the threshold, Apartment 2's preferences are so strict that a somewhat small change is only possible. However, Apartment 1 and Apartment 3 are able to shift their cooling loads to satisfy their preferences as well as the threshold constraint, but to do so, they have to run their AC earlier and longer. Interestingly, as can be seen in Table 5.2, although all

| Apartment | Payment in the Scheme | Payment Outside the Scheme |
|---|---|---|
| 1 | £2.31 | £5.85 |
| 2 | £2.40 | £5.85 |
| 3 | £2.31 | £5.85 |
| Total | £7.02 | £17.55 |

Table 5.2: The comparison of the amount that the apartments would pay if they participated in the scheme or not.

apartments would incur the same cost if they did not participate in the scheme, Apartment 2's share of the discount is slightly less than that of the other two apartments. This is due to its stricter preferences that are harder to satisfy in the collective optimisation. Also note that the optimal loads of Apartment 1 and Apartment 3, who have identical preferences, are the same, hence their Shapley values are equal.

### 5.3.7 Computationally Efficient Optimisation of Apartments

In the previous subsection, optimising the apartments in the simple case of 3 agents did not require a considerable amount of computation, and calculating the Shapley value was easy. However, as more apartments are added to the game, satisfying the constraints takes more and more time. Considering the fact that, in calculating the Shapley value, an exponential number of coalitions need to be optimised, even using the efficient implementation of Algorithm 1, the time it takes to calculate the share of only a few agents can be very long. For instance, using IBM ILOG CPLEX, computing the value of even a single coalition of size 5, i.e., computing the optimal load of a coalition of 5 apartments as per equation (5.9), takes approximately 5 minutes on a typical desktop computer. In order to overcome this issue, this subsection presents two greedy algorithms, namely, iOPT and cOPT, for optimising the cooling load of the apartments independently and collectively. Due to the difficulty of finding the optimal solutions in reasonable time, these algorithms trade off optimality for computation speed. Thus, a feasible cooling plan given by these algorithms are considered to be the best solution that can be found given the limited computational resources, and as such, such a solution will be called *best-found*. Clearly, these solutions may or may not coincide with the actual optimal solution.

First, the workings of iOPT will be explained, the pseudocode of which is presented is in Algorithm 2.

Using a heuristic, iOPT searches for a set of cooling actions that satisfy the constraints of an individual apartment as per equation (5.8). This heuristic, called the *discomfort* of an apartment, represents the discrepancy between an apartment's preferences (as outlined in Subsection 5.3.3) and the temperature profile resulting from a cooling plan found by

**Algorithm 2** Greedy Algorithm For Optimising Apartments Independently

```
function iOPT(T_set, θ, Ω, T_ext, Λ, maxIterations)

  η ← [ ]
  for iteration = 1 to maxIterations do
    bestTime ← −1
    minDiscomfort ← ∞
    avgTemp ← 0
    maxDeviation ← 0
    for t = 0 to k do
      if t ∈ Λ then
        continue
      end if
      η_test ← η
      η^t_test ← 1
      ∀t ∈ K update T_int[test] and T_env[test] as per equation (5.7) based on η_test
      Calculate ΔD_test
      if ΔD_test < minDiscomfort then
        bestTimeToSwitchOn ← t
        minDiscomfort ← ΔD_test
      end if
    end for
    if bestTime > −1 then
      η^bestTime ← 1
      update T_env
      update T_int
    end if
  end for

  return (avgTemp − T_set) ≤ Ω ∧ (maxDeviation ≤ θ)
```

the algorithm. More formally, the discomfort of apartment $a$, denoted by $\Delta D_a$, is the largest deviation of the internal temperature from the set-point temperature during the comfort period, plus the deviation of the average internal temperature from the set-point during the comfort period. This is given as:

$$\Delta D_a = \left( \frac{\sum_{h \in H} T^h_{int}[a]}{|H_a|} - T_{set}[a] \right) + \max_{h \in H} \left( T^h_{int}[a] - T_{set}[a] \right) \tag{5.10}$$

The algorithm incrementally finds the time slots where switching the AC on results in the largest discomfort reduction. Initially, the AC is off in all time slots (i.e., $\forall t \in K, \eta^t_a = 0$), and is then switched on only if it results in a reduction of the discomfort. This way, in addition to searching for a feasible solution, the consumption is also minimised (as required in the optimisation problem in equation (5.8)). However, as soon as the constraints of the apartment are satisfied, the algorithm will not seek to minimise the consumption further. Furthermore, recall that one of the constraints in equation (5.8) is that the internal and envelope temperatures at the start and end of the day should

**Algorithm 3** Algorithm For Optimising Apartments Collectively

```
function cOPT(C, A, ψ, T_set, θ, Ω, T_ext, maxIterations)
for all a ∈ A do
   iOPT(T_set[a], θ_a, Ω, T_ext, ∅, maxIterations)
end for
if ∀t ∈ K l^t_A ≤ ψ then
   return true
end if
Sort C based on θ_a/|H_a|
for all a ∈ C do
   Λ ← FindCongestedTimeSlots()
   successfullyOptimised ← iOPT(T_set[a], θ_a, Ω, T_ext, Λ, maxIterations)
   if successfullyOptimised then
      if ∀t ∈ K l^t_A ≤ ψ then
         return true
      end if
   else
      iOPT(T_set[a], θ_a, Ω, T_ext, ∅, maxIterations)
   end if
end for
return false
```

converge (i.e., $T^1_{int} = T^k_{int}$ and $T^1_{env} = T^k_{env}$). In order to ensure this, one can run Algorithm 2 repeatedly, and in each iteration calculate $T^1_{int}$ based on $T^k_{int}$ from the previous iteration. It was found through experiments that, this way, no more than 4 iterations are typically needed for the internal and envelope temperatures at the end of the day to be within 0.1°C of the start of the day. Moreover, since iOPT may not always find a feasible solution, the algorithm terminates after $maxIterations$ iterations. Lastly, iOPT takes a set $\Lambda$ as input, which, as we explain later, is used in the collective optimisation to indicate the time slots in which a member of the coalition should avoid running its AC. However, when an apartment is optimised independently, this set is empty.

If iOPT finds a feasible solution for apartment $a$, it gives a vector of $\eta^t_a$ values, based on which we calculate the best-found cooling load of the apartment as an independent apartment. This load is then used to calculate the consumption of apartment $a$ as per equation (5.5), based on which we obtain the bounded rational value of the singleton $\{a\}$, i.e., $v^{BR}(\{a\})$. Next, we describe the workings of cOPT, which optimises the cooling load of the members of a coalition. The pseudocode for this algorithm is given in Algorithm 3.

Given a coalition, all members are first independently optimised using iOPT. If by doing so the constraints of all apartments, as well as the threshold constraint are already satisfied, then the best-found cooling plan of the members of the coalition, in this case, is the same as when the apartments optimise their loads independently. However, if the threshold is not satisfied, it means that at least in one time slots there is congestion, i.e., the aggregate load is higher than the threshold. Denote the set of *congested time slots* by $\Lambda$, which is formally defined as: $\{t \in K | l^t_A \geq \psi\}$. The objective of the algorithm is to *decongest* these time slots by *re-optimising*[1] the apartments such that they do not run their ACs in these time slots. Obviously, those members of the coalition that are stringent with their temperature preferences may not be able to avoid the congested time slots. As such, the algorithm performs decongestion with respect to the flexibility of the load of the apartments. The idea is that the more flexible an apartment is, the more likely it can satisfy its constraints without having to run its AC in the congested time slots. Observe that the longer the comfort period is, the more cooling an apartment needs. Furthermore, as was seen in Figures 5.3 and 5.4, the higher the tolerance on the set-point temperature of an apartment is, the less cooling it requires, and thus, it can be considered more flexible than an apartment that has a lower tolerance. Based on these observations, cOPT uses the following ratio as a heuristic to determine the severity of the preferences of the apartments relative to one another:

$$\frac{\theta_a}{P_a \times |H_a|} \tag{5.11}$$

Using the above heuristic, the algorithm sorts the apartments in the coalition in ascending order, so that the least flexible apartment is dealt with first. Next, given a set of congested time slots, the apartments are iteratively re-optimised using iOPT. In each iteration, the set of congested time slots, $\Lambda$, is computed anew. If $\Lambda$ is not empty (i.e, the threshold constraint has not been satisfied yet), iOPT will be called again to optimise the apartment in the current iteration such that it does not run its AC in the congested time slots. If the constraints of the apartment are successfully satisfied this way, the algorithm moves on to the next apartment, and repeats this procedure until the threshold is satisfied or all apartments have been re-optimised. In any iteration, if the constraints of the apartment is not successfully satisfied, its best-found cooling plan (as a member of the coalition) will be the same as its independently optimised one. Similarly, if the threshold constraint is not satisfied at the end of the process, the best-found cooling plan of each apartment will be its independently optimised plan.

The above process can potentially result in an incremental reduction of the congestions, until the threshold is eventually satisfied. The effect of this collective optimisation on the individual apartments is that those apartments that are more flexible turn out to lower their internal temperature far ahead of their comfort period, so that they will not have to run their AC in the congested time slots. Note that if the algorithm fails to find a feasible solution, joining the grand coalition will not make the apartments worse off, since their best-found cooling plan will be the same as when they are independent. Therefore, although a solution given by cOPT algorithm may not be optimal, all apartments can still sign up to the discount scheme. The best-found cooling plan given by cOPT can then be used to compute the bounded rational value of a coalition $C$, i.e., $v^{BR}(C)$.

[1] By re-optimising an apartment given a set of congested time slots it shall be meant that after all apartments are initially independently optimised, the apartment is again optimised using iOPT, such that its AC is not turned on in any of the congested time slots.

### 5.3.8 Efficient Calculation of the Shapley Value Using DP

Recall from Subsection 5.3.1 that the value of a coalition is given by the sum of the consumption of its members. Therefore, once the best-found cooling plan of a coalition is obtained, the bounded rational value of the coalition can be calculated using equation (5.6). Now observe that in calculating the Shapley value, when cOPT is sequentially applied to the subsets of the grand coalition, some steps of optimising one coalition are repeated in optimising subsequent coalitions. By taking advantage of this recurrence, the Shapley value can be calculated using a dynamic programming (DP) algorithm more efficiently.

In order to calculate the Shapley value using Algorithm 1, the cooling plans of all subsets of the grand coalition need to be optimised one by one. Based on the flexibility heuristic in equation (5.11), cOPT arranges all apartments in $A$ such that apartment 1 and apartment $n$ are the least and most flexible apartments, respectively. Then, all apartments are independently optimised using iOPT. Here, if the result of the individual optimisations were stored in the memory, one could avoid re-computing them $2^n$ times. Furthermore, note that for optimising any coalition, the congested time slots are always initially the same. This is because the congested time slots are always identified after all apartments are independently optimised, which always results in the same aggregate load profile. Moreover, when cOPT optimises any coalition that contains apartment 1, it always performs decongestion starting from apartment 1, since it is always the least flexible apartment in any coalition. Therefore, if apartment 1 were present in a coalition, it would always be the first apartment to be re-optimised. Likewise, since apartment 2 is the next least flexible apartment, it is always the next candidate to be re-optimised, and so on. More importantly, each re-optimisation results in a $\Lambda$ that will be repeated in optimising subsequent coalitions. The following example will illustrate this recurrence relation.

Suppose that in a game with 5 apartments, we would like to optimise $\{1\}$. Optimising all apartments independently yields the set of congested time slots $\Lambda_\emptyset$. If apartment 1 can be re-optimised such that its temperature preferences are satisfied, a new set of congested time slots, $\Lambda_{\{1\}}$, will be yielded. Let us assume that re-optimising apartment 1 indeed results in satisfying its preferences, but $\Lambda_{\{1\}}$ is not empty (i.e., there are some congested time slots). Here, although the best-found cooling plan of apartment 1 will be reverted to its independently optimised cooling plan (since the threshold is not satisfied), we can re-use the result of re-optimising apartment 1 as well as $\Lambda_{\{1\}}$ in optimising any other coalition that contains apartment 1. This is because re-optimising apartment 1 given the initial congested time slots always yields the same cooling plan and the same $\Lambda_{\{1\}}$. Now, suppose that we would like to optimise $\{1, 2\}$. After re-optimising apartment 1, since $\Lambda_{\{1\}}$ is not empty, we need to re-optimise apartment 2 which will yield $\Lambda_{\{1,2\}}$. Regardless of whether $\Lambda_{\{1,2\}}$ is empty or not, every time a coalition that contains apartment 1 and apartment 2 (e.g., $\{1, 2, 4, 5\}$) is optimised, re-optimising apartment 1 and apartment 2 will result in the same cooling plans for these two apartments, and the same $\Lambda_{\{1\}}$ and $\Lambda_{\{1,2\}}$. Therefore, if after each re-optimisation the cooling plans along with the resulting set of congested time slots were stored, it would not be necessary to compute them again in optimising the subsequent coalitions. This way, for each coalition we would need to re-optimise only one apartment, which is essentially the most flexible member. Note that when calculating the Shapley values using Algorithm 1, to compute the coalition values efficiently, it is important to visit the subsets of the grand coalition such that optimising one coalition would depend only on the previously optimised ones. To this end, one can use the natural order of coalitions in the binary representation of coalitions. In this representation, a non-empty coalition $C = \{c_1, c_2, \ldots, c_m\}$ is represented by the binary equivalent of $2^{c_1-1} + 2^{c_2-1} + \ldots + 2^{c_m-1}$, where each bit indicates whether or not the corresponding agent is a member of the coalition. For instance, $\{2, 3\}$ comes immediately before $\{1, 2, 3\}$ as their corresponding binary numbers are 110 and 111, respectively. Table 5.3 shows the recurrence relation of the collective optimisation using this representation. For example, for optimising $\{1, 2, 4\}$, the right column shows that one can re-use the result of optimising $\{1, 2\}$, which is in turn optimised using the result of optimising $\{1\}$.

| Coalition | Apt. 5 | Apt. 4 | Apt. 3 | Apt. 2 | Apt. 1 | Retrieved from memory |
|---|---|---|---|---|---|---|
| $\{\}$ | 0 | 0 | 0 | 0 | 0 | |
| $\{1\}$ | 0 | 0 | 0 | 0 | <u>**1**</u> | |
| $\{2\}$ | 0 | 0 | 0 | <u>**1**</u> | 0 | |
| $\{1,2\}$ | 0 | 0 | 0 | <u>**1**</u> | 1 | $\{1\}$ |
| $\{3\}$ | 0 | 0 | <u>**1**</u> | 0 | 0 | |
| $\{1,3\}$ | 0 | 0 | <u>**1**</u> | 0 | 1 | $\{1\}$ |
| $\{2,3\}$ | 0 | 0 | <u>**1**</u> | 1 | 0 | $\{2\}$ |
| $\{1,2,3\}$ | 0 | 0 | <u>**1**</u> | 1 | 1 | $\{1,2\} \rightarrow \{1\}$ |
| $\{4\}$ | 0 | <u>**1**</u> | 0 | 0 | 0 | |
| $\{1,4\}$ | 0 | <u>**1**</u> | 0 | 0 | 1 | $\{1\}$ |
| $\{2,4\}$ | 0 | <u>**1**</u> | 0 | 1 | 0 | $\{2\}$ |
| $\{1,2,4\}$ | 0 | <u>**1**</u> | 0 | 1 | 1 | $\{1,2\} \rightarrow \{1\}$ |
| $\{3,4\}$ | 0 | <u>**1**</u> | 1 | 0 | 0 | $\{3\}$ |
| $\{1,3,4\}$ | 0 | <u>**1**</u> | 1 | 0 | 1 | $\{1,3\} \rightarrow \{1\}$ |
| ... | ... | ... | ... | ... | ... | ... |
| $\{1,2,3,4,5\}$ | <u>**1**</u> | 1 | 1 | 1 | 1 | $\{1,2,3,4\} \rightarrow \{1,2,3\} \rightarrow \{1,2\} \rightarrow \{1\}$ |

TABLE 5.3: Binary representation of coalitions in a game consisting of 5 apartments. For each coalition, <u>**1**</u> indicates the only apartment that may be required to be re-optimised. The best-found cooling plans of all other apartments in the coalition are retrieved from memory.

Based on the recurrence relation described above, we can construct a DP algorithm to calculate the Shapley value of the apartments in a more efficient manner. This will enable us to (re-)optimise only one apartment per coalition—the most flexible apartment according to equation (5.11)—since the cooling plan of the rest of the members can be used from the previously optimised coalitions. This recurrence relation is formalised next.

Let a coalition of apartments sorted in the ascending order of flexibility (according to equation (5.11)) be $C = \{1, 2, \ldots, m\}$, such that apartments 1 and $m$ are the least and most flexible apartments, respectively. Furthermore, let $\Lambda_C$ denote the set of congested time slots obtained by re-optimising apartment $m$, given the set of congested time slots, $\Lambda_{C\backslash\{m\}}$, obtained by re-optimising the most flexible apartment in $C\backslash\{m\}$ (i.e., $m-1$). As such, $\Lambda_\emptyset$ is the set of congested time slots after optimising all apartments independently, and $\Lambda_A$ is the set of congested time slots obtained by re-optimising apartment $n$ given $\Lambda_A\backslash\{n\}$. Moreover, let a vector of re-optimised cooling actions of the most flexible apartment in coalition $C$ over an entire day be denoted by $\eta_m^{\Lambda_{C\backslash\{m\}}}$, which is obtained by re-optimising $m$ given $\Lambda_{C\backslash\{m\}}$. Note that if $m$ cannot be re-optimised based on $\Lambda_{C\backslash\{m\}}$ such that its temperature preferences can be satisfied, then $\eta_m^{\Lambda_{C\backslash\{m\}}}$ will simply be the independently optimised plan that is found by $iOPT$. Denote by $l_m^{\Lambda_{C\backslash\{m\}}}$ the vector of best-found cooling load of apartment $m$ that is given by equation (5.4) using $\eta_m^{\Lambda_{C\backslash\{m\}}}$. We can now compute the vector of best-found aggregate cooling load, $l'_C = [l'^1_C l'^2_C \ldots l'^k_C]$, of a non-empty coalition, $C$, using the following recursive formula:

$$l'_C = \begin{cases} \sum_{a \in C} l''_a & \text{if } |C| = 1 \text{ or } \Lambda_C \neq \emptyset \\ l'_{C\backslash\{m\}} + l_m^{\Lambda_{C\backslash\{m\}}} & \text{if } \Lambda_C = \emptyset \end{cases}, \tag{5.12}$$

where $l''_a$ is the vector of best-found cooling actions of apartment $a$ when it optimises its load independently. Using equations (5.12) and (5.5) the bounded rational values of $C$, i.e., $v^{BR}(C)$ can be found.

### 5.3.9 Evaluation of the Coalitional Cooling Discount Scheme

This subsection undertakes an evaluation of applying the bounded rational Shapley value to the discount scheme. For this purpose, a block of 15 apartments is considered. Similar to the example case with 3 agents discussed previously, the AC system in every apartment operates on a 10-minute cycle, i.e., a cooling decision is made for each 10-minute interval in a day ($\Delta t = 1/6$hr), and thus, $K = [1, \ldots, 144]$. The AC systems in all apartments are similar and consume at a rate of 3.0 kW when fully on (i.e., $P_a = 3.0$ kW). That is, when the AC system is on for an hour, the total energy consumed is 3.0 kWh. Consequently, the maximum possible load of all apartments if they do not optimise their apartments collectively is 45 kW. Now when the apartments participate in the scheme, a threshold ($\psi$) at 18 kW (representing a 60% reduction), is set on their total load. If the aggregate load of the block is always 18 kW, the price per kWh of energy consumed over the entire day is £0.05, otherwise it is £0.16.

| Apartment $a$ | $CST_a$ | $CET_a$ | $T_{set}[a](°C)$ | $\theta_a(°C)$ |
|---|---|---|---|---|
| 1 | 11:40 | 20:00 | 22.0 | 1.5 |
| 2 | 13:20 | 20:00 | 22.0 | 1.2 |
| 3 | 14:00 | 19:40 | 22.0 | 1.1 |
| 4 | 10:10 | 18:30 | 22.0 | 1.1 |
| 5 | 13:00 | 19:00 | 22.0 | 1.1 |
| 6 | 11:40 | 22:30 | 22.0 | 1.1 |
| 7 | 12:00 | 21:40 | 22.0 | 1.2 |
| 8 | 10:20 | 18:20 | 22.0 | 1.1 |
| 9 | 12:00 | 21:40 | 22.0 | 1.1 |
| 10 | 12:00 | 21:20 | 22.0 | 1. |
| 11 | 12:00 | 21:40 | 22.0 | 1.2 |
| 12 | 12:00 | 22:20 | 22.0 | 1.2 |
| 13 | 13:00 | 19:00 | 22.0 | 1.1 |
| 14 | 13:00 | 19:00 | 22.0 | 1.1 |
| 15 | 12:00 | 18:40 | 22.0 | 1.1 |

TABLE 5.4: The cooling preferences and leakage rates of each apartment.

The cooling preferences of the individual apartments, which include $CST_a$, $CET_a$, $T_{set}[a]$ and $\theta_a$, are shown in Table 5.4. The leak rates $\tau$, $\rho$, and $\gamma$ for all apartments are 0.36 (1/hr), 0.024 (1/hr), and 0.024 (1/hr), respectively. An $\Omega$ value of 0.8°C is considered to limit the deviation of the average internal temperature during the comfort period from the set-point.

The aggregate cooling load when the 15 apartments, with the above settings and parameters, optimise their load independently and collectively are shown in Figure 5.10 and 5.11, respectively. The total energy consumption in the latter case is 2% more than the former, which represents the extra energy that is consumed by the apartments in off-peak times so as to keep the aggregate load below the threshold. As will be seen in the following experiments, this extra cost is offset by the discount that the apartments would receive under the scheme.

The experiments presented here show how the payments incurred by an apartment vary, as its cooling preferences are changed, on a single day. The payments that an apartment can be charged for its consumption are compared in four different cases: (i) when the apartment does not sign up to the discount scheme and optimises its load independently, (ii) when the apartment optimises its load as a mapartmentr of the coalition, but the discount is not taken into account (obtained from the consumption of the apartment in the grand coalition at the normal price), (iii) when the apartment optimises its load as a member of the coalition and each apartment receives an equal share of the total saving from the discount (calculated in the same way as the previous case, but the difference between the payment of the grand coalition at the discounted and normal prices is equally divided and deducted from the payment of each apartment), and (iv) when the apartment optimises its load as a member of the coalition and receives its

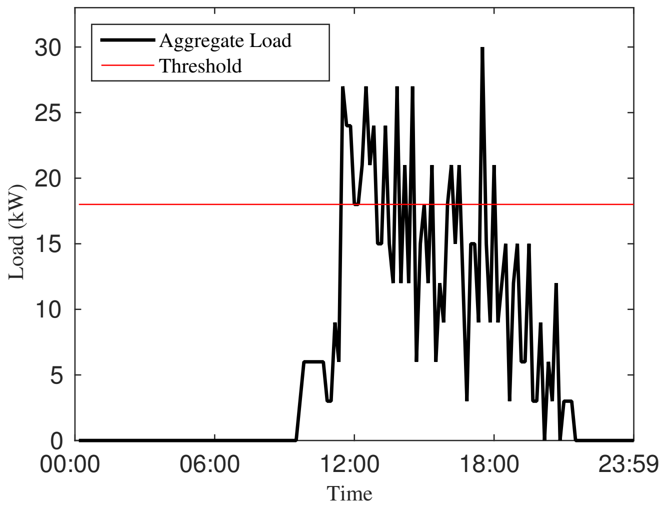


FIGURE 5.10: The aggregate cooling load when 15 apartments in a block optimise their loads independently using Algorithm 2.

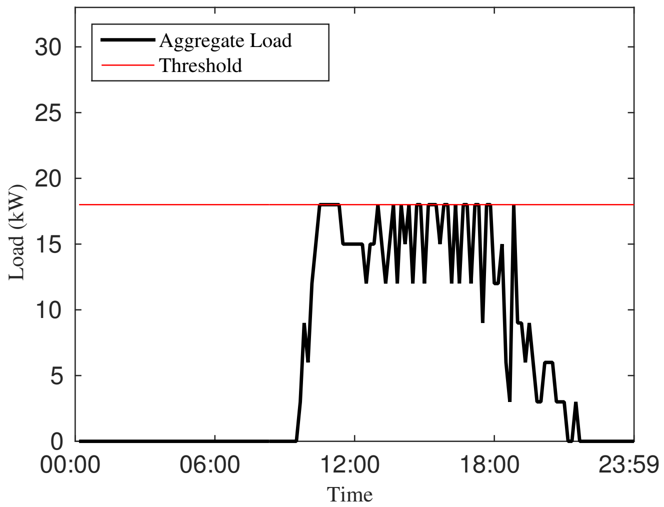


FIGURE 5.11: The aggregate cooling load when 15 apartments in a block optimise their loads collectively using Algorithm 3.

bounded rational Shapley value. In this experiment, the total amount that a coalition is charged is calculated as per equation (5.6), which represents its bounded rationality value, and the optimal loads of the apartments are given by Algorithms 2 and 3. Due to the intractability of calculating the true Shapley values of the 15 apartments, the actual amount that each apartment is charged is its bounded rational Shapley value (i.e., equation (5.2)), which is calculated using the DP method described in Subsection 5.3.8. Figure 5.12 shows a comparison of the time it takes to calculate the Shapley value of 8 through 15 apartments using the DP method for computing the coalition values, and without it. In both cases, the Shapley value is calculated using Algorithm 1, which is a more efficient implementation of equation (2.2). It is evident that a significant gain in computation time is achieved when Algorithm 1 is used with the DP method.

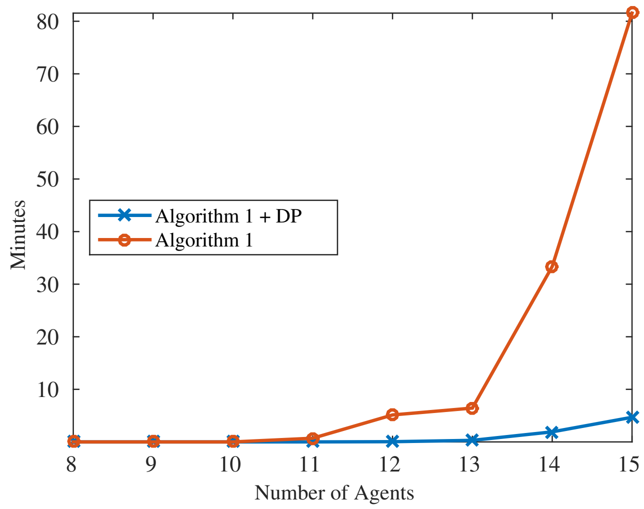


FIGURE 5.12: Comparison of computation times of calculating the Shapley value of different number of apartments.

The comparisons show that for any preference, the payment that is incurred by an apartment is more favourable when it signs up to the scheme than when it does not. They also demonstrate how sensitive the bounded rational Shapley value and the other naive payment mechanisms are to the variations in the set-point, $T_{set}$ (Subsection 5.3.10), the tolerance level, $\theta$ (Subsection 5.3.11), and the leakage rate, $\gamma$ (Subsection 5.3.12).

### 5.3.10 Set-point Temperature vs Payments

To explore the relationship between the set-point temperature settings and payments, $T_{set}[6]$ is varied from 19°C to 24°C, while leaving the other preferences unchanged. Then, the four payment cases described above are calculated for each set-point temperature setting, and the corresponding payment that apartment 6 incurs for a single day of cooling is depicted in Figure 5.13.

As can be seen the relationship between the set-point temperature and payments is almost linear. Intuitively, this is because when the set-point temperature is increased, less cooling is required, and thus, the payment reduces. It is also evident that an apartment stands to benefit from its participation in the scheme by receiving its bounded rational Shapley value, since the payment it incurs for a set-point setting is consistently lower than what it would incur if it chose to optimise its cooling load independently. Moreover, the fact that the collectively optimised payment for the set-point of 21°C is higher than that of the independently optimised, shows that apartment 6 has to use more energy in the coalition to help satisfy the threshold. However, since the bounded rational Shapley value for this set-point is lower, the cost of the extra consumption is offset by the discount. When the collectively and independently optimised curves match, it means that the energy consumption in and out of the coalition are equal. Furthermore, the

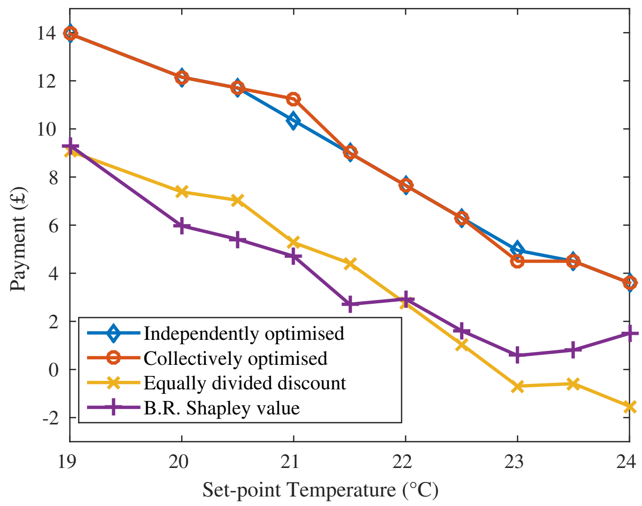


FIGURE 5.13: Relationship between set-point temperature ($T_{set}[6]$) settings and the payments for a block of 15 apartments

fact that the independently optimised payment for the set-point of 23°C is higher than that of the collectively optimised shows that the solution that is found for Apartment 6 when it optimises its load collectively is closer to the optimal than the solution found by the greedy algorithm (Algorithm 2) when it optimises its load independently. This is because, from the payment it is obvious that the consumption of the apartment in the coalition is lower, which is contrary to expectation.

### 5.3.11 Set-point Deviation Tolerance vs Payments

Figure 5.14 shows how the payment of Apartment 6 in each of the four cases changes as $\theta_6$ is varied from 0.5°C to 2.3°C.

As is evident, as $\theta_6$ is increased, the payments initially decrease but remain almost constant from 1.1°C onwards, which shows that the amount of energy required to satisfy the set-point constraint does not increase after this point. Since the internal temperature is also restricted by the average internal temperature constraint (i.e., $\Omega$), the relationship between the set-point deviation tolerance and the energy consumption may not be linear. This observation is true both when the apartment participates in the discount scheme and when it does not, as the bounded rational Shapley value and the independently optimised curves follow a similar trend.

### 5.3.12 Leakage Rate vs Payments

Previously, it was established how the leakage rate, $\gamma$ is related to the level of insulation of an apartment. It was also mentioned how an apartment with a high $\gamma$ value will

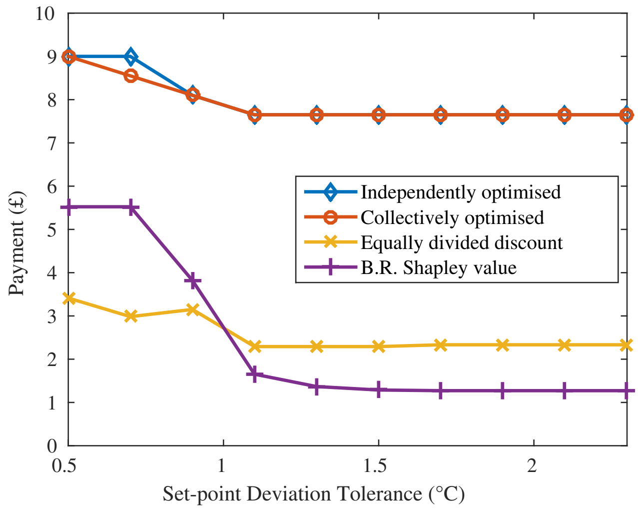


FIGURE 5.14: Relationship between the tolerance on the deviation of the internal temperature from the set-point ($\theta_6$) and the payments of Apartment 6 in a block of 15 apartments.

typically incur higher energy consumption. Following this trend, now the relationship between $\gamma$ and the payment an apartment incurs for a single day of cooling is explored. To do so, $\gamma_6$ is varied from 0.24°C/hr, representing a relatively high level of insulation to to 0.78°C/hr, representing a poorly insulated apartment, in apartment 6. The real $\gamma_6$ value is shown in Table 5.4. For each $\gamma_6$ value, the payment that apartment 6 incurs for a single day of cooling is then calculated based on the four different cases.

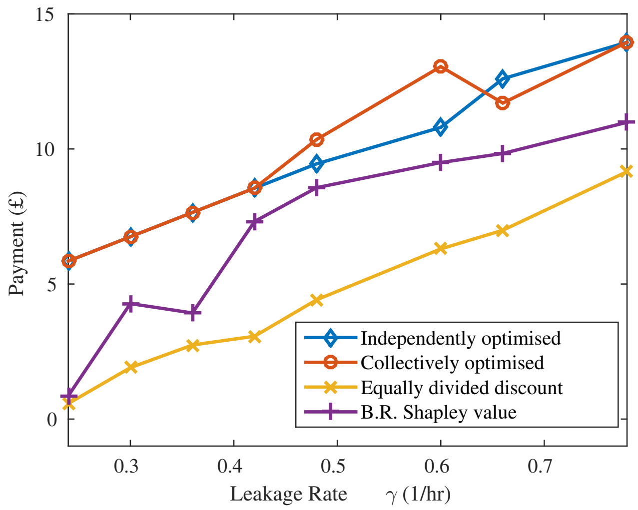


FIGURE 5.15: Relationship between the thermal leakage rate ($\gamma_6$) and the payments of Apartment 6 in a block of 15 apartments.

Figure 5.15 shows that the relationship between the leakage rates and the payments is almost linear when the apartment optimises its load independently. However, the bounded rational Shapley value curve exhibits less linearity, since after the leakage rate of 0.36 the gap between these two curves rapidly shrinks. This indicates the higher sensitivity of the payments to the thermal leakage when the apartment is in the coalition than when it is not. This reduction in the gap is significant compared to the set-point and set-point deviation tolerance plots (i.e., Figures 5.13 and 5.14). On balance, it can be concluded that a well-insulated apartment with a low leakage rate incurs a very low payment, and the payment grows as the apartment becomes leakier.

## 5.4 Summary

This chapter investigated the problem of fairly dividing the value of a coalition in settings where computing the value involves solving a hard optimisation problem. In such settings, calculating the Shapley value entails an extra computational challenge, namely solving an exponential number of hard optimisation problems. Since solving the optimisation problem of each coalition can take a considerable amount of time, and that the agents do not have infinite computational resources to find an optimal solution, they are considered to be computationally bounded rational. For these reasons, using the Shapley value as a fair division of the value of the grand coalition may not be possible in practice. Based on the additivity axiom of the Shapley value, it was proposed that it is still possible to obtain a fair division, without solving the optimisation problems optimally. If the optimisation problems can be solved using a faster but not necessarily optimal algorithm (e.g., greedy), the coalition values can be computed at a higher speed. The Shapley value given the suboptimal coalition values (called the bounded rational Shapley value), is not only easier to calculate, but also provides a division of the grand coalition that is fair in the following sense: all agents are rewarded for their contribution to the value of the grand coalition, and simultaneously penalised for their contribution to the discrepancy between the suboptimal and optimal solutions.

This approach was applied to a real world problem where a number of apartments in a block participate in a demand response program to receive a discount for coordinating their cooling loads so that the aggregate load does not exceed a certain threshold. In this problem, a coalition of apartments needs to optimise its members' use of air conditioning subject to the individual temperature preferences of each apartment and the given threshold. Due to the magnitude of the constraints involved, computing the optimal load of a coalition is computationally intensive. Instead of solving this problem optimally, a greedy algorithm was proposed which produced suboptimal solutions at a higher speed. Consequently, a suboptimal value for each coalition could be found, allowing for computation of the Shapley value in a reasonable time. It was demonstrated how, using the bounded rational Shapley value approach, the apartments could obtain a division of the discount, which is fair from the perspective described above.

# Chapter 6

# Conclusions and Future Work

The Shapley value is a widely used solution concept in cooperative game theory, which determines a fair allocation for agents in a game. Unlike other solution concepts such as the core and the stable set, the Shapley value is unique and always exists, which make it usable in any cooperative game. However, the use of Shapley value in games with more than a few dozens of agents faces a significant computational challenge, and thus, its use in practice is limited. In this thesis, we reviewed the existing proposals for mitigating this issue, which all have some merits but suffer from major drawbacks.

Building upon the existing approximation methods, we proposed an improved error bound for approximating the Shapley value using SRS. Furthermore, we proposed the use of stratified sampling, which unlike SRS that samples the population blindly, it exploits the structure of the population and results in a potentially more efficient and more accurate approximation. We showed how this method can be applied in some important class of games, including supermodular games, and some other games which have interesting real world applications such as newsvendor and output-sharing games. We also experimentally evaluated our method on instances of the the aforementioned games, each with at least 100 agents, using real and randomly generated data. Our experiments showed that the average standard error across all agents in all these games were %48 lower than that of SRS.

The limitation of our stratified sampling is in the way the minimum sample size is found. While our stratified sampling methods can be used to approximate the Shapley value in any game, our method of finding an optimal allocation of samples assumes that the game is supermodular or exhibits the order-reflecting property. Moreover, since sampling-based approximations are inherently randomised, the corresponding error bounds are in most cases probabilistic. This means that with a certain probability, the bound may not hold. Nevertheless, since the confidence can be arbitrarily large, one can set it very close to %100.

In addition to addressing the exponential time complexity of the Shapley value, we investigated how the Shapley value can be used when the complexity of the characteristic function is high, in which case calculating the Shapley value becomes even more difficult. We proposed to approximate the value of coalitions, and calculate the Shapley value using the approximated coalition values. Remarkably, due to the additivity axiom of the Shapley value, doing so not only would result in a fair division of the value of the grand coalition, but also each agent would be penalised for its contribution to the complexity of the characteristic function. We applied this approach to a demand response program where a number of apartments in a block coordinate their cooling loads to receive a discounted price. Since an optimal coordination of loads is computationally intensive, we used greedy algorithms to approximate the value of coalitions, and consequently calculate the Shapley value in a reasonable time. We also demonstrated the intuitive relationship between the bounded rational Shapley value and different sets of thermal model parameters.

In investigating the bounded rational Shapley value, we assumed that the number of agents in the game is relatively small. This was to separate the complexity of the Shapley value from the complexity of the characteristic function. In games with large number of agents, we can obviously approximate the bounded rational Shapley value in the same fashion as the standard Shapley value. However, an approximate Shapley value may not satisfy all of Shapley's axioms, because it may not be equal to the exact Shapley value that is unique. Consequently, it can be argued that approximating the bounded rational Shapley value may not result in a fair allocation. One way to alleviate this problem could be to relax the efficiency axiom, which is perhaps not a prerequisite for fairness. This would generalise the Shapley value to a set of all "values" that satisfy the symmetry and additivity axioms. Note that the symmetry axiom is fundamental for fairness, and the additivity axiom is the cornerstone of the bounded rational Shapley value, and thus, they cannot be dropped. However, by relaxing efficiency, we can still scale the allocations (e.g., through normalisation) such that they add up to the bounded rational value of the grand coalition. The more important task, however, would be to ensure that the approximated Shapley value satisfies symmetry and additivity, which is an interesting research problem in its own right.

In future work we will aim to address the limitations of our approach. In addition, we will concentrate on the following:

1. Recall that under CLT, the sampling distrubution of mean is assumed to be normal, which may be different than the actual distribution. Stein (1972) proposed a method for bounding the distance between two probability distributions with respect to a probability metric. We would like to investigate whether and how Stein's method can be used to incorporate the distance between the actual and the asymptotic sampling distribution in the CLT-based error bound proposed by Castro et al. (2009), so as to achieve an accurate bound.

2. Given that the stratified sampling can be more efficient than SRS, we would like investigate how the same error bound can be established with smaller sample sizes.

3. We would like to extend the order-reflecting property to newsvendor games with correlated output.

# References

Karthik. V. Aadithya, Tomasz P. Michalak, and Nicholas R. Jennings. Representation of coalitional games with algebraic decision diagrams. In *The 10th International Conference on Autonomous Agents and Multiagent Systems - Volume 3*, AAMAS '11, pages 1121–1122, Richland, SC, 2011. International Foundation for Autonomous Agents and Multiagent Systems. ISBN 0-9826571-7-X, 978-0-9826571-7-1.

Karthik V. Aadithya, Balaraman Ravindran, Tomasz P. Michalak, and Nicholas R. Jennings. Efficient computation of the shapley value for centrality in networks. In *Internet and Network Economics*, pages 1–13. Springer, 2010.

Energy Information Administration. Residential energy consumption survey (recs), 2013.

Stphane Airiau. Cooperative games and multiagent systems. *The Knowledge Engineering Review*, 28:381–424, 12 2013. ISSN 1469-8005.

Charilaos Akasiadis and Georgios Chalkiadakis. Agent cooperatives for effective power consumption shifting. In *Proceedings of the 27th Conference on Artificial Intelligence (AAAI-13)*, Bellevue, Washington, July 2013.

Muddasser Alam, Sarvapali Ramchurn, and Alex Rogers. Cooperative energy exchange for the efficient use of energy and resources in remote communities. In *Twelfth International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2013)*, May 2013.

K. K. Andersen, H. Madsen, and L. H. Hansen. Modelling the Heat Dynamics of a Building using Stochastic Differential Equations. *Energy and Buildings*, 31(1):13–24, 2000.

Edward J. Anderson and Andy B. Philpott. Using supply functions for offering generation into an electricity market. *Operations Research*, 50(3):pp. 477–489, 2002.

Kazutoshi Ando. Computation of the Shapley value of minimum cost spanning tree games: #p-hardness and polynomial cases. *Japan Journal of Industrial and Applied Mathematics*, 29(3):385–400, 2012. ISSN 0916-7005.

Anil Aswani, Neal Master, Jay Taneja, David E. Culler, and Claire Tomlin. Reducing Transient and Steady State Electricity Consumption in HVAC Using Learning-Based Model-Predictive Control. *Proceedings of the IEEE*, pages 240–253, 2012.

Haris Aziz, Casey Cahan, Charles Gretton, Phillip Kilby, Nicholas Mattei, and Toby Walsh. A study of proxies for shapley allocations of transport costs. In *The Fifth Workshop on Cooperative Games in Multiagent Systems (CoopMAS 2014)*, May 2014.

Peder Bacher and Henrik Madsen. Identifying suitable models for the heat dynamics of buildings. *Energy and Buildings*, 43(7):1511 – 1522, 2011. ISSN 0378-7788.

Yoram Bachrach, Evangelos Markakis, Ariel D. Procaccia, Jeffrey S. Rosenschein, and Amin Saberi. Approximating power indices. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems - Volume 2*, AAMAS '08, pages 943–950, Richland, SC, 2008. International Foundation for Autonomous Agents and Multiagent Systems. ISBN 978-0-9817381-1-6.

Enrique Baeyens, Eilyan Bitar, Pramod Khargonekar, and Kameshwar Poolla. Wind energy aggregation: A coalitional game approach. In *Proceedings of the 50th IEEE Conference on Decision and Control and European Control Conference, CDC-ECC 2011, Orlando, FL, USA, December 12-15, 2011*, pages 3000–3007, 2011.

John F. Banzhaf. Weighted Voting Doesn't Work: A Mathematical Analysis. *Rutgers Law Rev.*, 19:317–343, 1965.

Andrew C. Berry. The accuracy of the gaussian approximation to the sum of independent variates. *Transactions of the American Mathematical Society*, 49(1):pp. 122–136, 1941. ISSN 00029947.

Eilyan Bitar, Ram Rajagopal, Pramod Khargonekar, Kameshwar Poolla, and Pravin Varaiya. Bringing wind energy to market. *Power Systems, IEEE Transactions on*, 27 (3):1225–1235, Aug 2012. ISSN 0885-8950.

Olga N. Bondareva. Some applications of linear programming methods to the theory of cooperative games. *Problemy Kybernetiki*, 10:119139, 1963.

Godfrey Boyle, editor. *Renewable Energy: Power For A Sustainable Future*. Oxford University Press, Oxford, 2004.

Javier Castro, Daniel Gómez, and Juan Tejada. Polynomial calculation of the shapley value based on sampling. *Computers & Operations Research*, 36(5):1726–1730, May 2009. ISSN 0305-0548.

Georgios Chalkiadakis, Valentin Robu, Ramachandra Kota, Alex Rogers, and Nicholas R. Jennings. Cooperatives of distributed energy resources for efficient virtual power plants. In *The Tenth International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2011)*, pages 787–794, 2011.

Vincent Conitzer and Tuomas Sandholm. Computing Shapley values, manipulating value division schemes, and checking core membership in multi-issue domains. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 219–225, San Jose, CA, USA, 2004. AAAI Press.

Michael Conklin, Ken Powaga, and Stan Lipovetsky. Customer satisfaction analysis: Identification of key drivers. *European Journal of Operational Research*, 154(3):819–827, 2004.

Luis Corchon and M. Socorro Puy. Existence and nash implementation of efficient sharing rules for a commonly owned technology. Working papers. serie ad, Instituto Valenciano de Investigaciones Econmicas, S.A. (Ivie), 2000.

Xiaotie Deng and Christos H. Papadimitriou. On the complexity of cooperative solution concepts. *Mathematics of Operations Research*, 19(2):pp. 257–266, 1994a. ISSN 0364765X.

Xiaotie Deng and Christos H. Papadimitriou. On the complexity of cooperative solution concepts. *Mathematics of Operations Research*, 19(2):257–266, 1994b. ISSN 0364-765X.

Department of Energy & Climate Change. Energy consumption in the uk, 2014.

Aris L. Dimeas and Nikos D. Hatziargyriou. Agent based control of virtual power plants. *2007 International Conference on Intelligent Systems Applications to Power Systems*, pages 1–6, 2007.

Edith Elkind, Leslie Ann Goldberg, Paul W. Goldberg, and Michael Wooldridge. A tractable and expressive class of marginal contribution nets and its applications. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems - Volume 2*, AAMAS '08, pages 1007–1014, Richland, SC, 2008. International Foundation for Autonomous Agents and Multiagent Systems. ISBN 978-0-9817381-1-6.

Energy and Climate Change Committee UK Parliament. Energy and climate change - fourth report, April 2011.

G.N. Engelbrecht and A.P. Vos. On the use of the Shapley value in political conflict resolution. *Scientia Militaria - South African Journal of Military Studies*, 37(1), 2011. ISSN 2224-0020.

Shaheen Fatima, Michael Wooldridge, and Nicholas R. Jennings. A linear approximation method for the Shapley value. *Artificial Intelligence Journal*, 172(14):1673–1699, 2008.

Joan Feigenbaum, Christos H. Papadimitriou, and Scott Shenker. Sharing the cost of multicast transmissions. *Journal of Computer and System Sciences*, 63:21–41, 2001.

FERC. A national assessment of demand response potential staff report, 2009.

Vito Fragnelli, Ignacio García-Jurado, Henk Norde, Fioravante Patrone, and Stef Tijs. How to share railways infrastructure costs? In *Game Practice: Contributions from Applied Game Theory*, pages 91–101. Kluwer Academic, 2000.

Carl Futia. The complexity of economic decision rules. *Journal of Mathematical Economics*, 4(3):289 – 299, 1977. ISSN 0304-4068.

Daniel Granot, Jeroen Kuipers, and Sunil Chopra. Cost allocation for a tree network with heterogeneous customers. *Math. Oper. Res.*, 27(4):647–661, November 2002. ISSN 0364-765X.

Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):pp. 13–30, 1963. ISSN 01621459.

Yuan-Yih Hsu and Chung-Ching Su. Dispatch of direct load control using dynamic programming. *Power Systems, IEEE Transactions on*, 6(3):1056–1061, Aug 1991. ISSN 0885-8950.

Samuel Ieong and Yoav Shoham. Marginal contribution nets: A complact representation scheme for coalitional games. *ACM EC-06*, pages 170–179, 2006.

Satoru Ihara and Fred C. Schweppe. Physically based modeling of cold load pickup. *Power Apparatus and Systems, IEEE Transactions on*, PAS-100(9):4142–4150, Sept 1981. ISSN 0018-9510.

James P. Kahan and Amnon Rapoport. *Theories of Coalition Formation*. L. Erlbaum Associates, 1984.

Ehud Kalai and William Stanford. Finite Rationality and Interpersonal Complexity in Repeated Games. *Econometrica*, pages 397–410, 1988.

K. Kalsi, F. Chassin, and D. Chassin. Aggregated modeling of thermostatic loads in demand response: A systems and control perspective. In *Decision and Control and European Control Conference (CDC-ECC), 2011 50th IEEE Conference on*, pages 15–20, Dec 2011.

Daniel S. Kirschen and Goran Strbac. *Fundamentals of Power System Economics*. Wiley, 2004.

S. Koch, J. L. Mathieu, and D. S. Callaway. Modeling and control of aggregated heterogeneous thermostatically controlled loads for ancillary services. In *Proceedings of The 17th Power system Computation Conference*, Stockholm, Sweden, August 2011.

J K Kok, M J J Scheepers, and I G Kamphuis. Intelligence in Electricity Networks for Embedding Renewables and Distributed Generation. *Intelligence*, 42(2):179–209, 2010.

J. K. Kok, C. J. Warmer, and I. G. Kamphuis. Powermatcher: Multiagent control in the electricity infrastructure. In *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multiagent Systems*, AAMAS '05, pages 75–82, New York, NY, USA, 2005. ACM. ISBN 1-59593-093-0.

Ramachandra Kota, Georgios Chalkiadakis, Valentin Robu, Alex Rogers, and Nicholas R. Jennings. Cooperatives for demand side management. In *The Seventh Conference on Prestigious Applications of Intelligent Systems (PAIS @ ECAI)*, pages 969–974, August 2012.

Matthias Land and Olaf Gefeller. A multiplicative variant of the Shapley value for factorizing the risk of disease. In *Game Practice: Contributions from Applied Game Theory*, pages 143–158. Kluwer Academic, 2000.

Kate S. Larson and Tuomas W. Sandholm. Anytime coalition structure generation: An average case study. *Journal of Experimental and Theoretical AI*, 12:40–47, 2000.

Paul S. Levy and Stanley Lemeshow. *Stratification and Stratified Random Sampling*, pages 121–142. John Wiley & Sons, Inc., 2008. ISBN 9780470374597.

David Liben-Nowell, Alexa Sharp, Tom Wexler, and Kevin Woods. Computing shapley value in supermodular coalitional games. In Joachim Gudmundsson, Julin Mestre, and Taso Viglas, editors, *Computing and Combinatorics*, volume 7434 of *Lecture Notes in Computer Science*, pages 568–579. Springer Berlin Heidelberg, 2012. ISBN 978-3-642-32240-2.

R.H.A. Lindelauf, H.J.M. Hamers, and B.G.M. Husslage. Cooperative game theoretic centrality analysis of terrorist networks: The cases of jemaah islamiyah and al qaeda. *European Journal of Operational Research*, 229(1):230–238, 2013.

Stephen Littlechild and Guillermo Owen. A simple expression for the shapley value in a special case. *Management Science*, 20(3):370–372, 1973.

William F. Lucas. The proof that a game may not have a solution. *Transactions of the American Mathematical Society*, 137:pp. 219–229, 1969. ISSN 00029947.

William F. Lucas. Von Neumann-Morgenstern stable sets. In R.J. Aumann and S. Hart, editors, *Handbook of Game Theory with Economic Applications*, volume 1 of *Handbook of Game Theory with Economic Applications*, chapter 17, pages 543–590. Elsevier, 1992.

Irwin Mann and Lloyd S. Shapley. Values for large games iv: Evaluating the electoral college exactly. Technical Report RM-3158-PR, The RAND Corporation, Santa Monica, 1962.

Robert Marks. Market Design Using Agent-Based Models. In Leigh Tesfatsion and Kenneth L. Judd, editors, *Handbook of Computational Economics*, volume 2 of *Handbook of Computational Economics*, chapter 27, pages 1339–1380. Elsevier, 00 2006.

Michael A. McNeil and Virginie E. Letschert. Future air conditioning energy consumption in developing countries and what can be done about it: The potential of efficiency in the residential sector, 04 2008.

Tomasz Michalak, Talal Rahwan, Piotr L. Szczepanski, Oscar Skibski, Ramasuri Narayanam, Michael Wooldridge, and Nicholas R. Jennings. Computational analysis of connectivity games with applications to the investigation of terrorist networks. In *Proceedings of the 23rd International Joint Conference on AI (IJCAI)*, 2013.

Luigi Montrucchio and Marco Scarsini. Large newsvendor games. *Games and Economic Behavior*, 58(2):316 – 337, 2007.

Hervé Moulin. An application of the shapley value to fair division with money. *Econometrica*, 60(6):pp. 1331–1349, 1992. ISSN 00129682.

Hervé Moulin. *Cooperative microeconomics*. Prentice Hall, 1995. ISBN 0132276798.

Alfred Muller, Marco Scarsini, and Moshe Shaked. The newsvendor game has a nonempty core. *Games and Economic Behavior*, 38(1):118 – 126, 2002. ISSN 0899-8256.

Roger B. Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, 1991.

John Von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1944. ISBN 0691119937.

Abraham Neyman. Bounded complexity justifies cooperation in the finitely repeated prisoners' dilemma. *Economics Letters*, 19(3):227–229, 1985.

Gottfried E. Noether. Use of the range instead of the standard deviation. *Journal of the American Statistical Association*, 50(272):pp. 1040–1055, 1955. ISSN 01621459.

North American Electric Reliability Corporation. Accommodating High Levels of Variable Generation, 2009.

Department of Energy. Benefits of demand response in electricity markets and recommendations for achieving them. a report to the unit states congress pursuant to section 1252 of the energy policy act of 2005., 2006.

Martin Osborne and Ariel Rubinstein. *A Course in Game Theory*. MIT Press, 1994.

Guillermo Owen. Multilinear extensions of games. *Management Science*, 18(5):pp. P64–P79, 1972. ISSN 00251909.

Bezalel Peleg and Peter Sudhölter. *Introduction to the Theory of Cooperative Games (Theory and Decision Library C)*. Springer, 2nd edition, 2007.

Tiberiu Popoviciu. Sur les équations algébriques ayant toutes leurs racines réelles. *Mathematica (Cluj)*, 9:129–145, 1935.

Danny Pudjianto, Charlotte Ramsay, and Goran Strbac. Virtual power plant and system integration of distributed energy resources. *IET Renewable Power Generation*, 1(1): 10, 2007.

Sarvapali Ramchurn, Michael Osborne, Oliver Parson, Talal Rahwan, Sasan Maleki, Steve Reece, Trung Dong Huynh, Muddasser Alam, Joel Fischer, Tom Rodden, Luc Moreau, and Sephen Roberts. Agentswitch: towards smart electricity tariff selection. In *12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2013)*, May 2013.

Sarvapali D. Ramchurn, Perukrishnen Vytelingum, Alex Rogers, and Nicholas R. Jennings. Agent-based control for decentralised demand side management in the smart grid. In *The 10th International Conference on Autonomous Agents and Multiagent Systems - Volume 1*, AAMAS '11, pages 5–12, Richland, SC, 2011. International Foundation for Autonomous Agents and Multiagent Systems.

Alex Rogers, Sasan Maleki, Siddhartha Ghosh, and Nicholas R. Jennings. Adaptive Home Heating Control Through Gaussian Process Prediction and Mathematical Programming. In *Second Agent Technologies for Energy Systems workshop, Tenth International Autonomous Agents and Multi-agent Systems Conference, AAMAS*, 2011.

Harry Rose, Alex Rogers, and Enrico H. Gerding. A scoring rule-based mechanism for aggregate demand prediction in the smart grid. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 2*, AAMAS '12, pages 661–668, Richland, SC, 2012. International Foundation for Autonomous Agents and Multiagent Systems.

Ariel Rubinstein. *Modeling Bounded Rationality*, volume 1 of *MIT Press Books*. The MIT Press, January 1997.

Tuomas W. Sandholm and Victor R. Lesser. Coalitions among computationally bounded agents. *Artificial Intelligence*, 94(1-2):99–137, July 1997.

David Schmeidler. The nucleolus of a characteristic function game. *SIAM Journal on Applied Mathematics*, 17(6):1163–1170, 1969.

Lloyd S. Shapley. A value for n-person games. *Contributions to the theory of games*, 2: 307–317, 1953.

Lloyd S. Shapley. On balanced sets and cores. *Naval Research Logistics Quarterly*, 14 (4):453–460, 1967. ISSN 1931-9193.

Lloyd S. Shapley. Cores of Convex Games. *International Journal of Game Theory*, 1971.

Lloyd S. Shapley and Martin Shubik. A method for evaluating the distribution of power in a committee system. *The American Political Science Review*, 48(3):pp. 787–792, 1954. ISSN 00030554.

Dan Simon. *Optimal State Estimation: Kalman, H Infinity, and Nonlinear Approaches.* Wiley, 2006.

Herbert A Simon. *Models of Man.* John Wiley & Sons, 1957.

Charles Stein. A bound for the error in the normal approximation to the distribution of a sum of dependent random variables. *Proceedings of the Sixth Berkeley Symposium on Mathematical Statistics and Probability*, pages 583–602, 1972.

Steven Stoft. *Power System Economics.* Wiley-Interscience, New York, NY, 2002. ISBN 0471150401.

Goran Strbac. Demand side management: Benefits and challenges. *Energy Policy*, 36 (12):4419 – 4426, 2008. ISSN 0301-4215. Foresight Sustainable Energy Management and the Built Environment Project.

Chua-Liang Su and D. Kirschen. Quantifying the effect of demand response on electricity markets. *Power Systems, IEEE Transactions on*, 24(3):1199–1207, Aug 2009. ISSN 0885-8950.

Pafnuty Tchébychef. Des valeurs moyennes (traduction du russe, n. de khanikof. *Journal de Mathmatiques Pures et Appliques*, pages 177–184, 1867.

The United States Department of Energy. "Grid 2030" - A National Vision for Electricity's Second 100 Years. Technical report, United States Department of Energy Office of Electric Transmission and Distribution, 2003.

R. M. Thrall and W. F. Lucas. N-person games in partition function form. *Naval Research Logistics Quarterly*, 10(1):281–298, 1963. ISSN 1931-9193.

L. H. C. Tippett. On the extreme individuals and the range of samples taken from a normal population. *Biometrika*, 17(3/4):pp. 364–387, 1925. ISSN 00063444.

EdwardP.K. Tsang. Computational intelligence determines effective rationality. *International Journal of Automation and Computing*, 5(1):63–66, 2008. ISSN 1476-8186.

US National Renewable Energy Laboratory. *Eastern wind integration and transmission study.* National Renewable Energy Laboratory, Golden, CO, revised edition edition, 2010.

Matteo Vasirani and Sascha Ossowski. A collaborative model for participatory load management in the smart grid. In *Proceedings of the Workshop on AI Problems and Approaches for Intelligent Environments*, pages 21–26, Montpellier, France, 2012.

John Von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior.* Princeton University Press, 1944. ISBN 0691119937.

Thomson McLintock Whitin. Inventory control and price theory. *Management Science*, 2(1):61–68, 1955.

G. Poulton Y. Guo, R. Li and A. Zeman. A simulator for self-adaptive energy demand management. In *Proceedings of the 2nd IEEE International Conference on Self-Adaptive and Self-Organizing Systems*, pages 64–73, 2008.

Hobart Peyton Young. Individual contribution and just compensation. In Alvin E. Roth, editor, *The Shapley value*, pages 267–278. Cambridge University Press, 1988. ISBN 9780511528446. Cambridge Books Online.

Wei Zhang, K. Kalsi, J. Fuller, M. Elizondo, and D. Chassin. Aggregate model for heterogeneous thermostatically controlled loads with demand response. In *Power and Energy Society General Meeting, 2012 IEEE*, pages 1–8, July 2012.