Universal Privacy Gurantees for Smart Meters

The University Of Sheffield.

Miguel Arrieta

Department of Automatic Control and Systems Engineering
University of Sheffield

This dissertation is submitted for the degree of
Doctor of Philosophy

February 2020

Acknowledgements

I would like to express my gratitude to my advisor Dr. Iñaki Esnaola for the interesting discussions and honest feedback. I am indebted to my colleagues in room 305, for their invaluable support in the long hours writing this thesis. I hope they have received from me the same amount of support and good humour I have received from them. It is my pleasure to thank my Sheffield friends, and very specially the great people I have met in the cathSoc. From them I have learnt the most important things in this four years, and from them I take home the best memories and the deepest friendships. This thesis would not have been possible without the support from all those friends that, even if far away, encouraged me and gave me the motivation to complete this thesis. But above all, I owe my deepest gratitude to my family: Alvaro, Sandra, Valentin y Edurne, for being always there, for double and triple checking that I always had everything I needed, for looking after me in the good and in the bad moments, for standing my absence, THANK YOU!

Miguel

Abstract

Smart meters (SMs) provide advanced monitoring of consumer energy usage, thereby enabling optimized management and control of electricity distribution systems. Unfortunately, the data collected by SMs can reveal information about consumer activity, such as the times at which they run individual appliances. Two approaches have been proposed to tackle the privacy threat posed by such information leakage. One strategy involves manipulating user data before sending it to the utility provider (UP); this approach improves privacy at the cost of reducing the operational insight provided by the SM data to the UP. The alternative strategy employs rechargeable batteries or local energy sources at each consumer site to try decouple energy usage from energy requests. This thesis investigates the latter approach.

Understanding the privacy implications of any strategy requires an appropriate privacy metric. A variety of metrics are used to study privacy in energy distribution systems. These include statistical distance metrics, differential privacy, distortion metrics, maximal leakage, maximal α-leakage and information measures like mutual information. We here use mutual information to measure privacy both because its well understood fundamental properties and because it provides a useful bridge to adjacent fields such as hypothesis testing, estimation, and statistical or machine learning.

Privacy leakage under mutual information measures has been studied under a variety of assumptions on the energy consumption of the user with a strong focus on i.i.d. and some exploration of markov processes. Since user energy consumption may be non-stationary, here we seek privacy guarantees that apply for general random process models of energy consumption. Moreover, we impose finite capacity bounds on batteries and include the price of the energy requested from the grid, thus minimizing the information leakage subject to a bound on the resulting energy bill. To that aim we model the energy management unit (EMU) as a deterministic finite-state channel, and adapt the Ahlswede-Kaspi coding strategy proposed for permuting channels to the SM privacy setting.

Within this setting, we derive battery policies providing privacy guarantees that hold for any bounded process modelling the energy consumption of the user,

including non-ergodic and non-stationary processes. These guarantees are also presented for bounded processes with a known expected average consumption. The optimality of the battery policy is characterized by presenting the probability law of a random process that is tight with respect to the upper bound. Moreover, we derive single letter bounds characterizing the privacy-cost trade off in the presence of variable market price. Finally it is shown that the provided results hold for mutual information, maximal leakage, maximal-alpha leakage and the Arimoto and Sibson channel capacity.

Contents

List of Figures

Notation

UP	Utility Provider
SM	Smart Meter
EMU	Energy Management Unit
EHD	Energy Harvesting Device
AES	Alternative Energy Source
i.i.d	Independent and identically distributed
$[\![a, b]\!]$	Interval on the integers, e.g. $[\![a, b]\!] = \{a, a+1, \ldots, b-1, b\}$.
$[\![a, b]\!]^n$	The n-fold cartersian product of the interval, i.e. $[\![a, b]\!]^n = [\![a, b]\!] \times \ldots \times [\![a, b]\!]$
(a)	Positive part operator, e.g. $(a)^+ = \max(0, a)$
$\lceil a \rceil$	Ceil function, e.g. $\lceil a \rceil = \min\{n \in \mathbb{Z} : n \geq a\}$
$\lfloor a \rfloor$	Floor function, e.g. $\lfloor a \rfloor = \max\{m \in \mathbb{Z} : m \leq a\}$
$\mathbf{x}$	Bold notation is used for vectors
$\sigma(\mathbf{x})$	Sum over all the elements of the vector $\mathbf{x}$, e.g. $\sigma(\mathbf{x}) = \sum_i x_i$
X	Upper case denotes random variables
$\text{supp}(P_X)$	Support of the probability distribution P_X

Chapter 1

Introduction

1.1 Smart Grid

The electricity grid was designed around the prevailing production and consumption paradigms of the 20th century. At that time, the energy landscape was dominated by few large power stations, e.g. nuclear, coal, gas, whose production capacity could be increased or reduced on demand. This allowed for power generation to be adjusted to match the power consumption, with generation costs increasing with the peak power, as the more expensive stations were turned on. The small number of power sources allowed precise monitoring of the power grid by relatively few sensors placed primarily on the high and medium voltage grids [15]. Failure of any of the large power stations, or critical grid connections, could lead to a large scale failure of the grid [16]. However, due to the few players involved, grid failures and disruptions were relatively infrequent and could be solved in an ad-hoc manner [17]. This linear, one-way, behaviour is depicted in Figure 1.1.

Figure 1.1. Traditional model of the electricity grid showing a one-way system formed by (left to right) large power generation station, high voltage transmission grid, medium/low voltage distribution grid, consumers [1]

.

However, the evolving energy scenario and the emerging challenges of the 21st century urge for a modernization of the electricity grid. The growing climate crisis requires the reduction of greenhouse gas emissions, with specific targets set in the Paris agreement [18]. At the same time, the world energy consumption is predicted

to increase by 48% from 2012 to 2040 [19], driven by population growth, rise of the GDP per capita and increased penetration of electric vehicles among others [20]. In the UE for example, the 2030 Energy Strategy [21] includes the following targets:

- a 40% cut in greenhouse gas emissions compared to 1990 levels,
- at least a 32% share of renewable energy consumption, with an upward revisions clause for 2023,
- indicative target for an improvement in energy efficiency at EU level of at least 32.5%, following on from the existing 20% target for 2020,
- support the completion of the internal energy market by achieving the existing electricity interconnection target of 10% by 2020, with a view to reaching 15% by 2030.

Thus, large-scale introduction of renewable energy sources, and more efficient energy generation, distribution, and consumption are required [22]. However, the distributed and stochastic nature of renewable energy sources presents new challenges that call for a change of paradigm in the electricity grid. In the following we discuss those challenges.

The amount of energy generated by traditional power plants can be increased or reduced when required. However, the amount of energy generated by renewable energy sources is governed by stochastic factors outside engineering control such as solar radiation or wind speed. Moreover, there exists a strong geographical correlation of the atmospheric factors determining the power harvesting capabilities. This implies that the energy production of large geographical areas fluctuates randomly. For these reasons, the energy generation in this setting can no longer be adjusted to match the energy consumption. There are three main approaches to tackle this problem. The first approach relies on storing energy when production exceeds demand, and consuming it otherwise. A second approach relies on increasing the interconnection between distant areas. Thus, the weaker correlation between distant areas and the diversity of sources reduces the variability of the energy production. A third approach is to match the demand to the generation, incentivizing consumers to shift their energy demand to off-peak times. None of these solutions are perfect, and a combination of the three is required in the grid of the future.

Renewable energy sources, in particular wind and solar power that are predicted to account for 13.7% and 11.9% of the total energy generation by 2040 [20], are distributed in nature. Thus, the large scale adoption of renewable energy sources brings forth a new paradigm, where the grid is no longer dominated by a few large centralized power plants but by a large number of energy sources distributed across the

grid. Therefore, new schemes need to adapt to Distributed Energy Resources (DER), where multitude of agents generate electricity in different parts of the grid, as opposed to a few agents injecting power at very specific locations. At the same time, the advent of electric cars and batteries supporting available DER is predicted to increase the number and diversity of Distributed Energy Storage Systems (DESS). These small sources and storage systems are usually connected directly to the distribution grid (low voltage) whereas traditional electricity generators are connected to the transmission grid (high voltage). Large scale introduction of this sources requires fine management in the low voltage regime, increasing the need for precise energy consumption and production monitoring on the distribution grid.

Furthermore, the increasing dependency on the electricity supply and recent cyberattacks on the power grid call for a re-examination of the current failure and attack prevention mechanisms. This is exemplified by the December 2015 attack on Ukrainian electricity grid that switched off 30 substations leaving about 230 thousand people without electricity for multiple hours [23]. Examples of accidental grid failures are the 2003 US Northeast blackout affecting 55 million people for multiple days [16]; or more recently the 2012 Indian blackout affecting 620 million people, or around 9% of the world's population [24]. With an increasing number of players distributed across the grid and governed by diverse stochastic sources, the grid must be redesigned to include self healing capabilities and rapid recovery from failure, disconnection or sudden power injections by one or multiple correlated players.

The so called *smart grid* aims to update the existing grid to the new challenges of the 21st century described above. The smart grid is defined by the EU Commission Task Force for Smart Grids as follows.

Definition 1.1. *[25] A Smart Grid is an electricity network that can cost efficiently integrate the behaviour and actions of all users connected to it – generators, consumers and those that do both – in order to ensure economically efficient, sustainable power system with low losses and high levels of quality and security of supply and safety.*

Though elements of smartness also exist in many parts of existing grids, the difference between a today's grid and a smart grid of the future is mainly the grid's capability to handle more complexity than today in an efficient and effective way. A smart grid employs innovative products and services together with intelligent monitoring, control, communication, and self-healing technologies in order to:

- *Better facilitate the connection and operation of generators of all sizes and technologies.*
- *Allow consumers to play a part in optimising the operation of the system.*

- *Provide consumers with greater information and options for how they use their supply.*
- *Significantly reduce the environmental impact of the whole electricity supply system.*
- *Maintain or even improve the existing high levels of system reliability, quality and security of supply.*
- *Maintain and improve the existing services efficiently.*
- *Foster market integration towards European integrated market*

Figure 1.2 depicts a simple smart grid model with distributed and highly interconnected energy storage and stochastic energy sources.

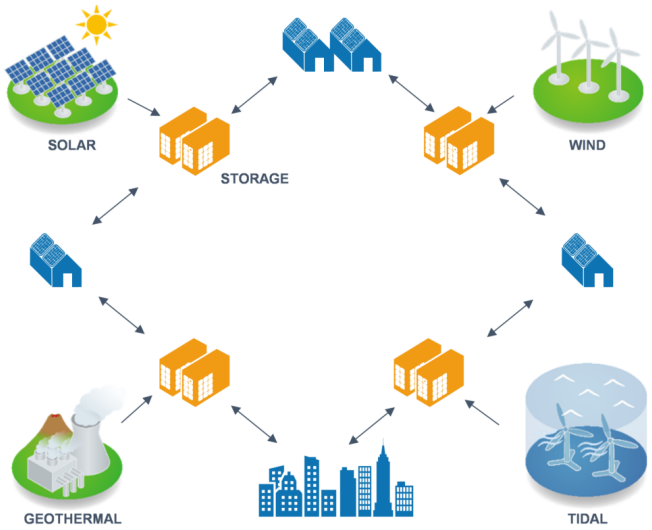

Figure 1.2. Smart grid model depicting a highly interconnected system with distributed energy storage capabilities and stochastic energy sources connected at different points of the grid [1].

We now describe the main actors, components and the energy market aiming to provide the natural self-balancing behaviour of the grid.

A wide variate of actors participate in the smart grid. The following describes some of the most relevant ones. The first actor is the consumer, sometimes called prosumer if they also generate electricity. Consumers demand energy, typically paying a fixed hourly rate for it. Another important actor is the Utility Provider (UP) who serves as the middleman between generators and consumers. UPs must balance their portfolio, matching the amount of energy they buy in the wholesale market to the the energy consumed by their users. Two key players managing the grid at a local level are the Meter Responsible Party (MRP) and the Distribution System

Operator (DSO). MRPs and DSOs are responsible for the low and medium voltage distribution network of a specific area, as well as for keeping quality parameters within the limits set by the regulator. The Transmission System Operator (TSO) is responsible for the high voltage transmission network, and for balancing the grid i.e. matching production and demand, at all times. Large generators and consumers are typically connected to the transmission network. Finally, the regulator audits the operation of the grid and designs technical regulations such as the European standard EN 50160 [26].

Energy markets are a key component of the electricity grid, enabling the self-balancing of the grid, matching generation and demand. There are three energy markets: the wholesale, the balancing and retail market. The retail market trades electricity between users and suppliers, market prices typically follows fixed tariff based on long term contract. The wholesale market trades large amounts of energy between suppliers and generators. The price is determined by competitive negotiation. Here electricity is traded for time for every settlement period, typically around half-hourly, with a submission deadline to negotiate before. The balancing market, managed by the TSO aims to match consumption and supply in real time. Therein, the TSO is able to buy energy in the balancing market or activate strategic reserves.

Smart meters are advanced metering devices that measure electricity consumption and generation data, as well as operational grid data to the MRP multiple times per hour. The MRPs then transfer the required data to the DSO and UP allowing them to automate grid management and bill the user for the consumed energy. SMs allow for managing the grid at a local level, balancing stochastic power injections with a responsive demand that reduces the requirements to increase the grid capacity and its associated cost. SMs also contain an off-switch, allowing the provider to disconnect the household from the grid in case of grid failure, in order to avoid a full-scale blackout. Moreover, SMs also send measurements about quality such as root mean square voltage variations (voltage drop-out, sags and swells, total harmonic distortion...) helping the DSO fulfil its duties. Furthermore, this infrastructure also enables dynamic energy pricing, shifting user demand to match energy generation [27, 28]. SMs also communicate their readings to the Home Area Network (HAN) gateway. This enables the introduction of energy consumption indicators for the user, raising awareness of the energy cost. Energy consumption indicators have been reported to reduce the energy consumption up to 15% [29].

For these reasons the roll out of SM is currently taking place world-wide. The EU directive 2009/72/EC [30] states *"Where roll-out of smart meters is assessed positively, at least 80% of consumers shall be equipped with intelligent metering systems by 2020."*. Expenditure required on the UK electricity grid between 2012 and

2050 is estimated to drop from £46bn if only conventional technologies are employed, to £27bn if a smart grid is deployed [31].

1.2 Privacy

The University of Cambridge dictionary defines privacy as: *the right that someone has to keep their personal life or personal information secret or known only to a small group of people [32].* A selection of mathematical definitions of privacy are presented in Section 2.2. However, it is clear by the wide range of metrics and the lack of consensus inside the engineering community that finding an adequate mathematical definition of privacy is still an open problem.

Despite disagreements over how to measure privacy, there is a growing body of literature addressing privacy. This interest is partly fostered by some recent privacy breaches, such as the Cambridge Analytica scandals revealed in early 2018, that collected personal data from millions of Facebook users without their consent [33]. As another example, in 2019 it was revealed that some of the audio recorded by the voice assistant from Amazon, Google and Apple are listened and reviewed by humans, who have reported to overhear private conversations [34]. Another scandal is the one revealed in 2005, where it was discovered that, in an attempt to combat piracy, Sony BMG CDs were purposely infected with malware that sent the private listening habits and the IP of the user to the companies headquarters [35].

However, privacy breaches are not always intentional. In 2006 Netflix published 10 million film rankings by 500,000 as a research challenge. The data was anonymized by replacing users names and personal information with random numbers. However, correlations analysis with IMDB allowed for de-anonymization of some users, revealing the content they watched on Netflix [36]. On another example, using public anonymous data from the 1990 census, [37, 38] shows that 87% of the 248 million population in the United States is uniquely identified by their five-digit ZIP code, gender and date of birth, reducing the geographical resolution to city/towns or even counties uniquely identifies 50% and 18% of the population respectively. This is specially worrying if one takes into account that, at the time of that study, many states in the US published "anonymized" healthcare data for research purposes where ZIP code, gender and date of birth were linked with medical history [37].

Within the electricity grid, precise fine-grained information about energy consumption can be used to infer sensitive information about the users [39, 40]. Every electrical device has a characteristic energy consumption profile [41]. This consumption signature can be compared against the fine-grained data provided by the smart

meters. Techniques such as *Non-Intrusive Load Monitoring (NILM)* [42] use the information collected by smart meters to infer sensitive information about users [43–45]. In the case of households, human presence [43], patterns on domestic appliances [46], amount of sleep [47, 48], breakfast habits [46, 47], home presence during sick leave [47] or tuned TV channel [49, 50] are among the large list of events disclosed by smart meters. In the case of industry, it can leak acquisition of new machines, number of items manufactured in a day, broken machinery or out of hours working [47] among others. This information is of great personal, social and economical importance for companies or individuals.

Not surprisingly, leakages of energy consumption data have also been reported [43, 49]. In 2012 researchers in the USA revealed that the data collected by SMs in their area was sent to the UP as plain text [43]. This data contained real time consumption every 30 seconds as well as the ID of the user. It was also noted that the ID of SMs was printed on the front face of the meters, with the majority of SMs being placed outside the households, making de-anonymization simple. Researchers were able to monitor SMs in a range of up to 300 meters with a single antenna, identifying unoccupied residences and living routines. In 2009 two laws aimed to enforce the usage of SMs were blocked by the Senate of the Netherlands motivated by the privacy concerns that emerge as a result of the increased penetration of SMs [51]. This hinders the implementation of the smart grid. It is then paramount to understand and characterize the fundamental tradeoff between operational performance and user privacy.

1.3 Smart meter privacy

There is a different degree of overlapping between the data required by the service provider and the sensitive data that the user may not want to disclose. The local GP for instance requires access to the medical record of a patient, this information is the same that the user might want to keep private. However, in the smart grid, the information required by the grid operators, i.e. power generation/consumption, is different to the sensitive data that is required to keep private, i.e. what that power is used for. The operator does not need to know what electrical devices are used for, but a function of that data, e.g. the aggregated power consumption. This brings in a fundamental question, is it possible, and to what extend, to increase the privacy of the user while preserving the utility of the data for the provider.

The introduction of SMs has brought forth a growing body of literature addressing the conflict between efficient energy monitoring and privacy. In [52, 53] obfuscation of the knowledge that the utility provider (UP) has about the energy consumption

of the user is studied. Indeed, in the case in which the SM readings are the only source of information available to the UP obfuscation yields some degree of privacy. Obfuscation is achieved by different mechanisms, such as aggregating the consumption of multiple users [52], compression of the energy consumption sequences [53] or homomorphic encryption [54] among others. These techniques suffer from different shortcomings. Firstly, data aggregation and noise addition reduce the utility of the data provided to the DSO and UPs, preventing DSOs from accurately monitoring the grid, and limiting the benefits of SMs. Another shortcoming is the ability of the DSO or any eavesdropper to monitor the energy consumption of a user by other means [55]. Finally, many of these solutions still rely on a trusted party to aggregate or obfuscate the data.

A different approach to the problem arises in settings where the user has access to alternative energy sources [9, 56] or energy storage devices [57–59]. In this case, the UP has perfect knowledge of the energy provided to the user, but the user employs the alternative energy source and the energy storing capability of the system to dissociate the energy consumed by the appliances from the energy provided by the UP. This thesis focuses on the latter approach.

1.4 Overview and contributions

The aim of this thesis is to characterize the fundamental limits governing smart meter privacy in the presence of energy storage devices.

- In Chapter 2 we review the different system models proposed in the literature, presenting an integrated perspective of the different models. Subsequently, we introduce the main privacy metrics studied in the literature. Finally, we review the state of the art in smart meter privacy for scenarios in which local energy sources or local energy storage are available. This review shows the importance of characterizing privacy guarantees that hold for a wide class of energy consumption processes modelling the energy consumption of the user.

- In Chapter 3 we present universal privacy guarantees for EMUs with access to a finite capacity battery. Therein we provide mathematical guarantees that hold for any bounded energy consumption process. We further extend the analysis to characterize bounded processes with a given expected energy consumption. The bounds are shown to be tight by proposing an energy consumption process that achieve the upper bound for any feasible battery policy implemented by the EMU. The contents of this chapter were published on the 2017 *IEEE Proceedings on International Conference on Smart Grid and Communications*,

Dresden, Germany under the title "Smart meter privacy via the trapdoor channel" [60].

- In Chapter 4 we address the impact of variable market prices on the privacy guarantees. This allows us to characterize the impact of privacy optimization on the energy bill and on the self-balancing capacity of the grid. Therein we provide upper and lower bounds on the minimum achievable information leakage under a constraint on the energy bill. The contents of this chapter were published on 2019 *IEEE International Symposium on Information Theory*, Paris, France under the title "Universal Privacy Guarantees for Smart Meters" [61].

- In Chapter 5 we provide single letter bounds to the information leakage and extend the obtained results to maximal leakage and maximal α-leakage. To this aim, we provide upper and lower bounds on the cardinalities of the minimal covering and packing sets. This allows for a tight characterization of the information leakage when no privacy budget is available. The extension to other metrics is done by proving a more general result showing that under certain conditions, maximal leakage, maximal α-leakage and mutual information coincide.

Chapter 2

Literature review

In this chapter, we describe the state of the art in smart meter privacy when the energy management unit has access to energy storage devices or local energy sources. In Section 2.1, we present an integrated view of the different system models proposed across the literature. In Section 2.2, we focus on the most common privacy metrics studied in the literature. Subsequently, Section 2.3 studies the different solutions and techniques developed for smart meter privacy. Finally, in Section 2.4, we review the previous sections, identifying a research gap in the literature and presenting our main problem formulation for this thesis.

2.1 System model

In this section, we present an integrated perspective of the different system models employed across the literature. To this aim, we follow the main body of literature in employing a discrete time model. This is grounded on the fact that digital systems operate in discrete time and motivated by the better tractability presented by discrete time models. Within this setting, at time step i, the user consumes $X_i \in \mathcal{X} \subseteq \mathbb{Z}$ units of energy. To satisfy this consumption, the energy management unit (EMU) has access to four different sources of energy: the utility provider (UP), energy harvesting devices (EHD), energy storage devices (SD) and alternative energy sources (AES). This scheme is depicted in Figure 2.1. Moreover, the EMU is typically constrained to avoid any power outage or energy waste. Thus, the EMU must create a request sequence that meets the energy demands of the user and does not request energy it cannot use or store, i.e.

$$Y_i = X_i - V_i - E_i + \Delta S_i, \tag{2.1}$$

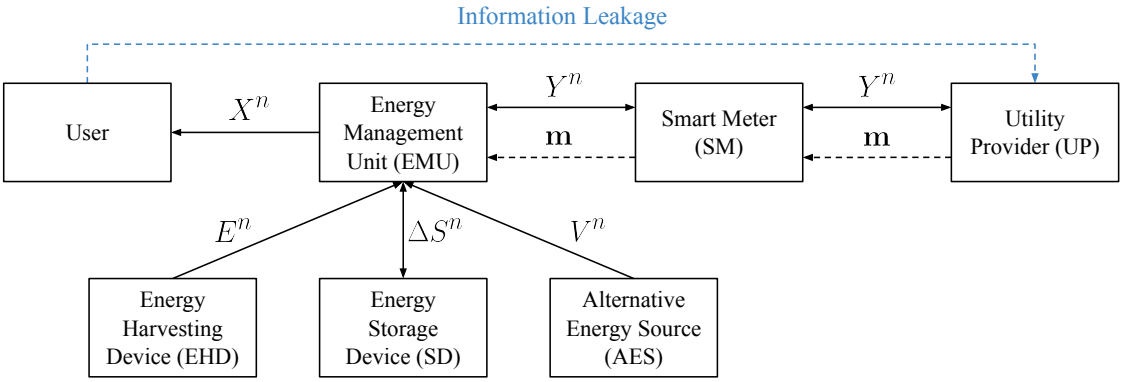

Figure 2.1. Integrated system model.

where Y_i denotes the energy requested from the UP, V_i and E_i denote the energy generated by the AES and EHD respectively, and $\Delta S_i = (S_{i+1} - S_i)$ denotes the energy stored or consumed from the battery.

Other alternatives explored in the literature include demand shifting and energy waste. Demand shifting implies leaving the demand of the user unsatisfied, forcing the consumption to be cancelled or rescheduled, this approach is studied in [62, 63] under different non information-theoretic metrics and joint energy bill optimization. Wasting energy, e.g. by dissipating it as heat, is also proposed in [13] as a way to increase the privacy of the user.Multiuser scenarios are proposed in [10]. The scenario of one user with access to two energy storage devices connected together is proposed in [64] for independent and identically distributed (i.i.d.) inputs.

Our interest on this model arises with the introduction of the smart meter (SM) and the variable market price $\mathbf{m} \in \mathbb{R}^n$. Variable market prices introduce an extra constraint on the EMU, as a budget constraint is typically imposed by the user, forcing energy requests Y_i to be shifted towards more economical time slots. Finally, the privacy risk posed by the fine grained energy consumption data that the SMs sends to the UP is centre to our attention in this thesis. The different elements of this system models are described in greater details in the following subsections.

2.1.1 Smart meters

Smart Meters periodically send electricity consumption and generation data to the MRP, which further distributes the necessary data to the DSO and the UP. MRP are also able to send commands to the SM, e.g. to request log files [15].

The European Union recommends smart meters should send integrated energy measurements every 15 minutes in order to enable the proper functioning of the

Smart Meter Model	Time Resolution
Itron Centron	1 min
REX2	5 min
Kamstrup Omnipower	5 min
Enel Open Meter	15 min

Table 2.1. Smallest time resolution of currently used SMs [65].

smart grid. Current specifications in the UK mandate resolutions of 30 minutes. Table 2.1 shows the smallest rime resolution of currently employed SMs. However, the real sampling capabilities of smart meters are above these values, with real time consumption data being displayed to users every 10 seconds in some devices [66]. Moreover, this resolution is expected to increase with the introduction of additional renewable energy sources and the diversification of sources and producers in the smart grid.

In our discrete time model, each time step represent a measurement from the SM. Thus, the time resolution of our model matches that of the SM. This comes with no loss of generality and enables the information metrics to capture the real information captured by the smart meter. More generally, this model generalizes to any metering device, such as the ones commercially available to monitor the consumption behaviour of the user [55].

2.1.2 Energy consumption

The energy consumption of the user, is modelled as a discrete time stochastic process X^n taking values in a real or discrete alphabet, i.e. $\mathcal{X} \subseteq \mathbb{R}$ or $\mathcal{X} \subseteq \mathbb{Z}$ respectively. The discrete nature of the consumption alphabet is usually assumed for the sake of simplicity and justified by the limited precision of energy devices in real power systems.

As seen in Figure 2.2, the energy consumption of typical users tends to exhibit non-stationary dynamics. A review of public datasets containing energy consumption profiles of real users is presented in [67]. This non-stationary behaviour, together with the high dependency of the consumption across different users calls for the adoption of a unified, simplified energy consumption model. Three main approaches are explored in the literature in order to model the statistics of energy consumptions. For the simpler and better understood model, the energy consumption of the user is assumed to be i.i.d.. Although the energy consumption of typical users exhibit strong time correlations [2], the i.i.d. assumption provides a first foundational step, enabling the single letter characterization of privacy measures [14]. A more refined

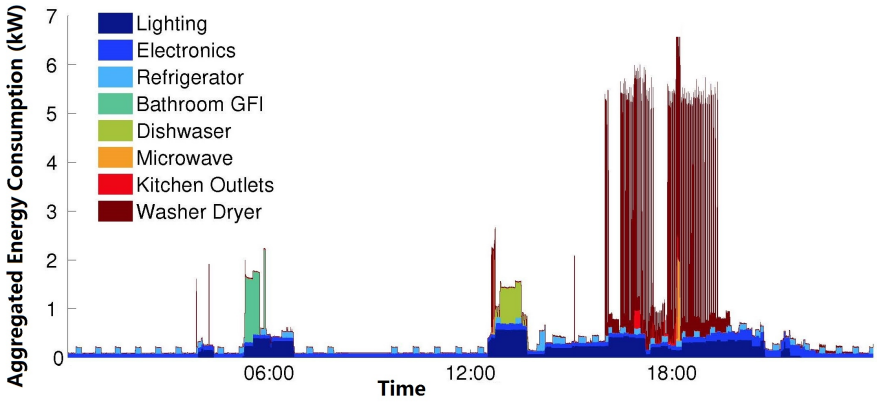

Figure 2.2. An example of energy consumption over the course of a day for one of the houses in REDD [2].

approach models the energy consumption as a Markov process. This model better captures the real dynamics of the system, opening the door to more complex models, such as hidden Markov models, which could arguably model the energy consumption with higher accuracy [2, 68]. The third family of approaches, relies on avoiding the need for probabilistic models. This is done by considering offline policies [7], adopting distribution independent metrics such as maximal leakage, or by learning an approximation of the system dynamics online [69].

2.1.3 Energy storage device

Electric batteries, and more generally energy storage devices, are devices capable of storing energy and later converting it into electrical power, usually in the form of direct current. With the advent of local energy sources and the rapid adoption of electric cars, the availability of energy storage devices, and more specifically electric batteries, is predicted to increase [20]. Table 2.2 presents some of the batteries currently available in the market. Most providers guarantee 10.000 charge/discharge cycles of battery life or up to 10 years of life. Some exceptions to this rule such as the powervault ECO, which employ second hand car batteries, and thus provide shorter guarantees of 3 to 5 years. Most commercially available batteries are based on Lithium-ion technology. This batteries provide a lower life-time cost and provide higher efficiency levels than previous batteries based on lead acid and cheaper in terms of upfront cost.

Definition 2.1. *A battery is said to have capacity β, maximum power discharge rate $P_{\min}$ and maximum charging rate $P_{\max}$ when the amount of energy stored in the battery at time step $i \in [\![0, n-1]\!]$ satisfies*

$$0 \leq S_i \leq \beta, \tag{2.2}$$

Supplier	**Capacity (kWh)**	**Peak Power (kW)**	**Approximate Cost (£)**
BYD B-BOX	1.28 - 441.6	12.8	1,700
Enphase	1.2	0.27	1,700
LG Chem Resu	2.9 - 9.3	3 - 5	2,628
Moixa Smart Battery	2 - 3	2.4	2,950
Powervault 3	4.1 - 20.5	0.8 - 5.5	2,000
Puredrive ESS	4.8 - 9.6	3	3,492
Samsung SDI	3.24	2	3,500
Simpliphi	0.58 - 3.2	1.5	-
Solax Battery	3.3 - 6.5	6	1,920
Sonnen Batterie Eco	2 - 16	3 - 8	4,500
Tesla Powerwall 2.0	13.5	5	7,750
Varta Pulse	3.3 - 6.5	1.8 - 2.5	3,579
xStorage by Nissan	3.6 - 6	2.52	3000

Table 2.2. Battery models [65, 70].

and the amount of energy stored or consumed from the battery between two consecutive time steps satisfies

$$-P_{\min} \leq S_{i+1} - S_i \leq P_{\max}. \tag{2.3}$$

It is important to note that batteries are damaged by repeated charging and discharging cycles. Hence the usage of batteries for privacy optimization might shorten the life span of batteries. Therein privacy optimization might push batteries above the typical guarantees of $10,000$ cycles during 10 years, i.e. around 2.75 charge/discharge cycles per day. The wear and tear of the battery is considered in [71–74] where the damage produced to the battery by the charging and discharging process is considered. Therein, the damage is measured by the following cost function:

$$D(S^n) = \sum_{i=0}^{n-1} \mathbb{1}\{S_{i+1} \neq S_i\}, \tag{2.4}$$

where $\mathbb{1}\{S_{i+1} \neq S_i\}$ is one when the battery is being charged or discharged. The privacy optimization is then performed subject to a constraints on the damage inflicted to the battery. Finally, [7] considers more complex and realistic battery models, including thermal and chemical energy storage devices where multiple parameters such as capacity, minimum and maximum charging rates, initial battery state, and charging efficiency factor are considered in the model.

2.1.4 Local energy generation

In some scenarios, users have access to local energy sources. These sources can be broadly classified in two groups:

- Alternative Energy Sources (AES) or on-demand sources, where the user controls the dynamics of the energy generation process, e.g. a petrol generator.
- Energy Harvesting Devices (EHD) or stochastic sources, where the generation process is governed by external forces, e.g. a solar panel or wind turbine.

In the case of AES, the energy generated at time i, and denoted by V_i, is controlled by the user. This generation mechanism is subject to the physical constraints of the underlying source. These constraints are typically modelled as an average $\overline{P}$ and a peak power $\hat{P}$ generation constraint, i.e.

$$\Pi(\overline{P}, \hat{P}) = \left\{ P_{V^n} : \mathbb{E}\left[\sum_{i=0}^{n-1} V_i\right] = \overline{P} \text{ and } 0 \leq V_i \leq \hat{P} \text{ for all } i \right\}. \tag{2.5}$$

In the case of EHD the generated energy at time step i is modelled by E_i, with the stochastic process E^n taking values in alphabet $\mathcal{E}^n$ with distribution P_{E^n}. This stochastic process is usually governed by atmospheric factors, such as wind speed or sun radiation, which can be assumed to be similar across large areas. Thus, the UP has access to statistical information about the energy generated by the stochastic source, either by implementing their own sensors, checking publicly available data, or by inferring it from the consumption of nearby users. The type and installed capacity of alternative energy sources is not directly available to the UP. The level at which the UP knows the statistics of the energy generation process is modelled in a variety of ways across the literature.

2.1.5 Energy bill

Dynamic energy pricing, where the price for the energy varies along the day, plays a key role in the smart grid. In this setting, dynamic prices are expected to shift the demand of users adapting it to the generation peaks and valleys introduced by renewable energy sources [28, 29]. Energy storage devices are expected to play a key role helping user shift their demand towards cheaper, more environmentally friendly, time slots [75]. This sets a two objective optimization problem on the storage device, where joint privacy-cost optimization is required. Buying energy when it is cheaper reduces the power bill but might not be optimum from a privacy perspective (and

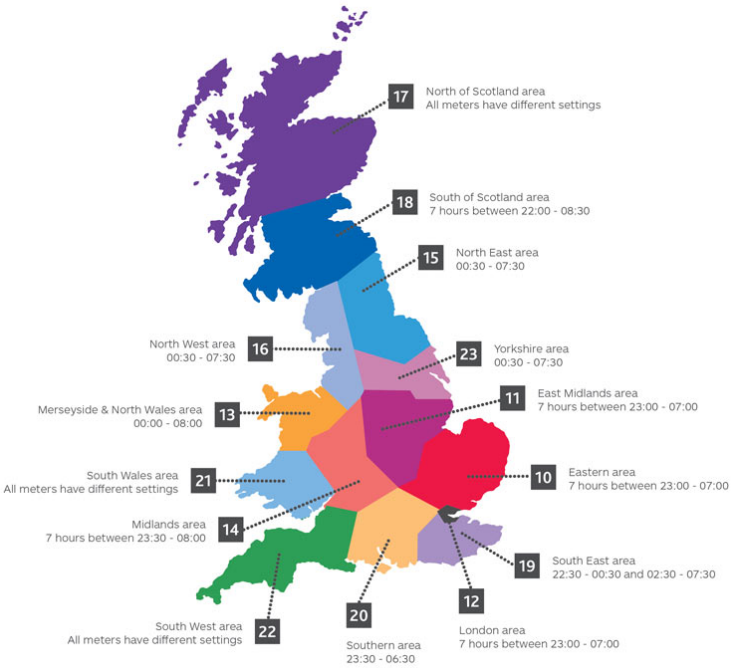

Figure 2.3. Example of market price evolution over time [3].

vice versa). Understanding the relationship between privacy and cost can help find interesting tradeoffs.

In the UK, the two most common variable tariffs, provided by the six major utility providers are UK Economy 7 and UK economy 10. UK economy 7 offers 7h/day of reduced off-peak prices over night, with the exact times depending on the region, as shown in Figure 2.3. Another popular tariff is UK economy 10, which offers 10 hours of reduced off-peak energy prices split in three blocks, usually between 1-4pm, 5-7pm and midnight to 5am. This variable price system, where the users pays a predefined price that varies in blocks of different lengths (always larger than the sampling period of the SMs) is the most common variable price tariff worldwide [15]. Interestingly, [76] offers real time prices for the consumers, which can register to receive notifications or program their smart appliances to be switched when market prices fall down.

Variable market prices are considered in [6, 77–80]. Therein, the energy bill is calculated by multiplying the energy request sequence Y^n, by the market price $\mathbf{m} \in \mathbb{R}^n$, i.e. the energy bill results in

$$B(Y^n) = \sum_{i=0}^{n-1} m_i Y_i = \mathbf{m}^T Y^n. \tag{2.6}$$

Mirroring the market price paid by users in real scenarios [81], it is often assumed, without loss of generality, that the market price is constant over each of the K blocks

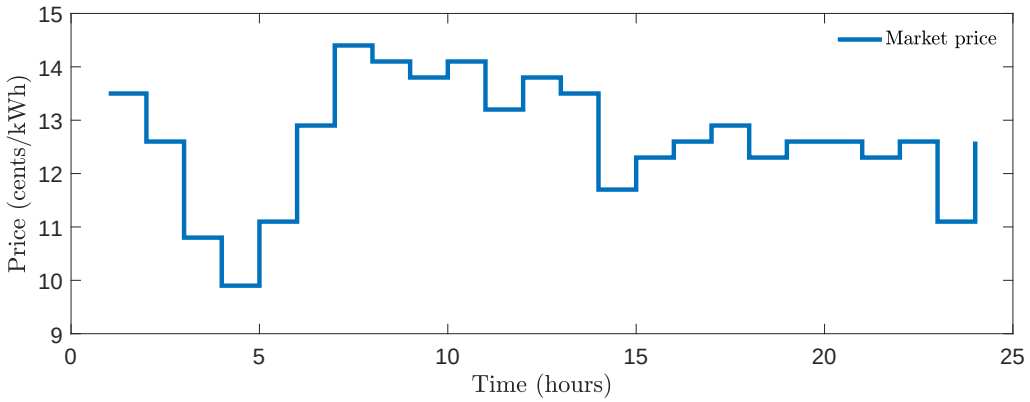

Figure 2.4. Example of market price evolution over time depicting prices on the Iberian Peninsula on Monday July 5, 2010 [3].

of time. The duration of the k-th block, with $k = 0, 1, \ldots, K-1$, is denoted by l_k, giving

$$\mathbf{m} = (\underbrace{m_0, \ldots, m_0}_{l_0}, \underbrace{m_1, \ldots, m_1}_{l_1}, \ldots, \underbrace{m_{K-1}, \ldots, m_{K-1}}_{l_{K-1}}), \tag{2.7}$$

where m_k denotes the market price during block k, and the k-th market change takes place at time step $t_k = t_{k-1} + l_{k-1}$ with $k = 1, 2, \ldots, K$ and $t_0 = 0$. An example of this model is pictured in Figure 2.4.

The privacy-cost tradeoff is analysed with a deviation from the average metric for scenarios in which the user energy consumption is known in advance [6] or is estimated by an online control algorithm [80]. In [79] bounds on the privacy-cost tradeoff are provided for scenarios where the user is allowed access to an energy storage device and privacy is measured with mutual information.

2.2 Privacy metrics

Before discussing the existing solutions aiming to maximize the privacy of the users, it is paramount to answer a more fundamental question: what do we mean by privacy? And in particular: How do we mathematically measure privacy? An overview of some of the most common privacy metrics proposed in the literature follows. A more comprehensive review of different privacy metrics describing over eighty privacy metrics is presented in [82].

2.2.1 Quadratic deviation

State of the art non-intrusive load monitoring (NILM) algorithms employ the differences in consecutive load measurements or *features*, i.e. $Y_{i+1} - Y_i$, to detect when appliances are switched on or off [83, 44]. Furthermore, it is clear that an EMU that maps every consumption sequence $\mathbf{x} \in \mathcal{X}^n$ to the same (typically constant) request $\mathbf{w} \in \mathcal{Y}^n$ leaks no information about the behaviour of the user. However, this mapping is generally infeasible, as it is not allowed by the battery constraints. A simple way of measuring how much the consumption of the user Y^n deviates from a non-informative output $\mathbf{w}$ is the quadratic deviation.

Definition 2.2. *Given a joint probability distribution P_{X^n,Y^n}, and a sequence $\mathbf{w} \in \mathbb{R}^n$, the* quadratic deviation *or* mean squared error *is given by*

$$\mathcal{V}(\mathbf{w}) \triangleq \frac{1}{n}\,\mathbb{E}\left[\|Y^n - \mathbf{w}\|_2^2\right], \tag{2.8}$$

where the expected value is taking with respect to P_{X^n,Y^n}.

In [71, 75, 84] the non-informative output $\mathbf{w}$ is considered constant over time, i.e. $w_i = \mathbb{E}[1/n \sum_{j=0}^{n-1} X_j]$ for all i, while [78] considers $\mathbf{w}$ to be a function of the market price, i.e. $w_i = f(m_i)$. Further joint variance-cost optimizations are considered in [6, 80].

The additive character of $\mathcal{V}(\mathbf{w})$ provided by the ℓ_2-norm simplifies the calculation of this privacy metric. However, it is unclear whether this metric captures the fundamental properties of privacy. As an intuitive example of this shortcoming, consider, for any finite $a \in \mathbb{R}$, the deterministic one to one mapping:

$$Y^n = X^n/a + \mathbf{w}. \tag{2.9}$$

Therein, the privacy leakage measured by the quadratic deviation

$$\mathcal{V}(\mathbf{w}) = \mathbb{E}[(X^n)^T X^n]/a, \tag{2.10}$$

can be made arbitrarily small by increasing the constant a. However, it is clear that the proposed mapping provides no privacy, as the UP can uniquely recover the consumption X^n from the request Y^n, i.e.

$$X^n = a(Y^n - \mathbf{w}). \tag{2.11}$$

While this metric provides some interesting insights, a stronger mathematical foundation is needed.

2.2.2 Fisher information

Fisher information is a measure of information typically used in mathematical statistics. Given a random variable Y, with probability distribution parametrized by parameter θ, Fisher information measures the information that Y contains about the parameter θ. Intuitively, for any given realization y and parameter θ the *score*, i.e.

$$\ell'_{\theta,\mathbf{y}} = \frac{\partial}{\partial\theta} \log P_{Y|\Theta}(y|\theta), \tag{2.12}$$

measures how sensitive the probability distribution is to changes in the parameter θ. Low score values imply that it is hard to accurately estimate θ upon observation of y. Following this intuition, the Fisher information is defined as the variance of the score, i.e.

$$\mathcal{I}(\theta) \triangleq \mathbb{E}\left[\left(\ell'_{\theta,\mathbf{y}} - \mathbb{E}[\ell'_{\theta,\mathbf{y}}]\right)^2\right], \tag{2.13}$$

where the expected value is taken over the conditional probability distribution $P_{Y|\Theta=\theta}$. It can be shown that the expected value of the score satisfies $\mathbb{E}[\ell'_{\theta,\mathbf{y}}] = 0$, and therefore, the Fisher information is given by

$$\mathcal{I}(\theta) = \mathbb{E}\left[\left(\ell'_{\theta,\mathbf{y}}\right)^2\right] = \int \left(\frac{\partial}{\partial\theta} \log P_{Y|\Theta}(\mathbf{y}|\theta)\right)^2 \log P_{Y|\Theta}(\mathbf{y}|\theta) d\mathbf{y}. \tag{2.14}$$

In the smart meter privacy context the interest is not in estimating a single parameter θ, but a sequence $\mathbf{x} \in \mathcal{X}^n$. Thus, the matrix form extension of the Fisher information is used in [74] as a measure of privacy.

Definition 2.3. *The Fisher information matrix is given by*

$$\mathcal{I}(\mathbf{x}) \triangleq \int P_{Y|X}(\mathbf{y}|\mathbf{x}) \left[\frac{\partial \log P_{Y|X}(\mathbf{y}|\mathbf{x})}{\partial \mathbf{x}}\right] \times \left[\frac{\partial \log P_{Y|X}(\mathbf{y}|\mathbf{x})}{\partial \mathbf{x}}\right]^T. \tag{2.15}$$

Interestingly, as a result of the Cramér-Rao bound [85], the Fisher information matrix provides a lower bound on the variance of any unbiased estimator, specifically:

$$\mathrm{Tr}\left(\mathcal{I}(\mathbf{x})^{-1}\right) \leq \mathbb{E}\left\|\hat{\mathbf{x}}(Y) - \mathbf{x}\right\|_2^2, \tag{2.16}$$

where $\mathrm{Tr}(\cdot)$ denotes the trace operator. In the context of SM privacy, this provides an operational meaning to the Fisher information, as described in [74].

2.2.3 Hypothesis testing

Intuitively, one can say that a system achieves a high level of privacy if it is hard for an attacker to determine which of a set of hypothesis about the user $\mathcal{H} = \{h_0, h_1, ..., h_K\}$ is true. Examples of possible hypothesis include determining whether the user is at home or not $\mathcal{H} = \{h_0 : \text{user is at home},\ h_1 : \text{user is not at home}\}$ or whether the user is watching TV or not $\mathcal{H} = \{h_0 : \text{TV is on},\ h_1 : \text{TV is off}\}$. In the most common case of binary hypothesis testing, the attacker is subject to two types of errors, i.e. rejecting the null hypothesis h_0 when it is true, and failing to reject it when it is false. These errors are described in the following table:

	h_0 is true	h_1 is true
Fail to reject h_0	Right decision	Type II error
Reject h_0	Type I error	Right decision

The probabilities of Type I error and Type II error are denoted by p_I and p_{II} respectively. The Neyman-Pearson lemma [86] shows that the optimal decision region is given by the likelihood ratio test:

$$C_{h_0} = \left\{ \mathbf{x} \in \mathcal{X} : \frac{P_{X|H}(\mathbf{x}|h_0)}{P_{X|H}(\mathbf{x}|h_1)} \geq \lambda \right\}. \tag{2.17}$$

Therein, hypothesis h_0 is rejected when $\mathbf{x} \notin C_{h_0}$ and vice versa. For a fixed Type I probability of error, the minimum achievable Type II probability, and its asymptotic exponential decay rate, are used in [8, 87–89] as privacy metrics, yielding the following definition.

Definition 2.4. *Given a joint probability distribution $P_{X,Y}$, a set of two mutually exclusive hypothesis $\mathcal{H} = \{h_0, h_1\}$, and a constraint $p_I \leq \alpha$ on the Type I probability of error, the privacy level is given by the probability of Type II error, i.e.*

$$p_{II}^* \triangleq \mathbb{P}[\ \textit{Fail to reject } h_0 \mid h_0 \textit{ is false}\ ], \tag{2.18}$$

the asymptotic exponential decay rate of p_{II}^* is given by

$$r_{II}^* = \lim_{n \to \infty} -\frac{\log p_{II}^*}{n}. \tag{2.19}$$

Interestingly, in [90], the minimum Type II probability of error p_{II}^*, is used as a utility metric instead of a privacy metric. Therein, the disclosure mechanism must allow the UP to infer information about one specific hypothesis, while minimizing the global concept of privacy captured by mutual information.

2.2.4 Differential privacy

Differential privacy, as originally proposed in [91] arises in the context of database privacy. The underlying intuition behind differential privacy is that, for a disclosure mechanism to be considered private, the outcome of any query A over two datasets differing only in one element should be indistinguishable. The following definition captures this notion.

Definition 2.5. *Given an n-dimensional dataset X^n, a randomized algorithm to answer query A is (δ, ϵ)-differentially private if for all $\mathbf{x}, \mathbf{x}' \in X^n$ that differ only in one element, and all $\mathcal{S} \subseteq range(A)$*

$$\mathbb{P}[A(\mathbf{x}) \in \mathcal{S}] \leq e^{\epsilon}\mathbb{P}[A(\mathbf{x}') \in \mathcal{S}] + \delta. \tag{2.20}$$

Thus, for two datasets differing only on one element, the probability distribution of the response to any query should be similar. Note that setting $(\delta, \epsilon) = (0, 0)$ implies perfect privacy, while letting $\delta = 1$ or ϵ be sufficiently large implies no privacy guarantee. Differential privacy in the smart meter context is proposed as a privacy measure in [92].

Interestingly, in [93] an ordering between privacy metrics is presented. Therein, a randomized mechanism $P_{Y|X^n}$, randomly mapping the input X^n to an output variable Y, is defined as ϵ-mutual information differentially private (ϵ-MI-DP) if:

$$\sup_{i, P_{X^n}} I(X_i; Y|X^{-i}) \leq \epsilon \log e. \tag{2.21}$$

Furthermore, a randomized mechanism $P_{Y|X^n}$ is defined as ϵ-Kullback-Leibler differentially private (ϵ-KL-DP) if for all neighbouring data sets $\mathbf{x}, \mathbf{x}'$:

$$\mathcal{D}(P_{Y|X^n=\mathbf{x}} \| P_{Y|X^n=\mathbf{x}'}) \leq \epsilon \log e. \tag{2.22}$$

Finally, the following ordering is established:

$$\epsilon\text{-DP} \succeq \text{KL-DP} \succeq \text{MI-DP} \succeq (\delta)\text{-DP} \succeq (\epsilon, \delta)\text{-DP}, \tag{2.23}$$

where α-DP $\succeq$ β-DP implies that for all $\beta' > 0$ there exist an $\alpha' > 0$ such that α'-DP implies β'-DP. That is, α-DP is a stronger privacy metric in the sense that for any set of parameters defining the metric β-DP, there exist a set of parameters defining α-DP such that α-DP implies β-DP. Furthermore, it is noted that

$$\text{MI-DP} = (\epsilon, \delta)\text{-DP}, \tag{2.24}$$

when the cardinality of the database entries or the query response are bounded. Finally, links between differential privacy and Hypothesis testing are established in [94]. Therein, the degradation of differential privacy when multiple queries are responded by the disclosure mechanism, i.e. the composition theorem, is characterized.

2.2.5 Maximal leakage and maximal α-leakage

Maximal leakage [95] considers an adversary that upon observation of a random variable Y tries to guess a function of a related random variable X. This captures an attacker that is not interested in estimating the consumption of the user X, but some other related property U satisfying the Markov chain $U - X - Y$, e.g. whether the user is at home. Specifically, the maximal leakage measures the logarithm of the ratio between the probability of a correct guess of U when Y is observed, i.e.

$$\mathbb{P}[U = \hat{U}|Y] = \max_{P_{\hat{U}|Y}} \mathbb{E}\left[P_{\hat{U}|Y}(\hat{U} = U|U, Y)\right] = \sum_{\mathbf{y}} \max_{\mathbf{u}} P_{U,Y}(\mathbf{u}, \mathbf{y}), \tag{2.25}$$

where the equality follows by using maximum a posteriori probability (MAP) detection, and the probability of a correct guess without observing Y, i.e.

$$\mathbb{P}[U = \tilde{U}] = \max_{P_{\tilde{U}}} \mathbb{E}\left[P_{\tilde{U}}(\tilde{U} = U|U)\right] = \max_{\mathbf{u}} P_U(\mathbf{u}). \tag{2.26}$$

This ratio aims to capture the multiplicative gain in estimation power provided by observation of Y and is computed for the worst functional U, in order to capture any possible function of interest. The following definition captures this idea.

Definition 2.6. *Given a joint distribution P_{XY} on finite alphabets $\mathcal{X}$ and $\mathcal{Y}$ the maximal leakage from X to Y is defined as*

$$\mathcal{L}(X \to Y) \triangleq \sup_{U-X-Y} \log \frac{\mathbb{P}[U = \hat{U}|Y]}{\mathbb{P}[U = \tilde{U}]}, \tag{2.27}$$

where U, $\tilde{U}$ and $\hat{U}$ take values on an arbitrary finite alphabet $\mathcal{U}$.

Interestingly, [95, Theorem 1] shows that the maximal leakage equals:

$$\mathcal{L}(X \to Y) = \log \left(\sum_{\mathbf{y} \in \mathcal{Y}} \max_{\mathbf{x} \in \mathcal{X}: P_X(\mathbf{x}) > 0} P_{Y|X}(\mathbf{y}|\mathbf{x}) \right). \tag{2.28}$$

This shows that the maximal leakage $\mathcal{L}(X \to Y)$ depends only on the support of X and not on its distribution P_X. Maximal α-leakage, a generalization of maximal leakage with a stronger dependency on the distribution P_X, is presented in [96] and given by the following definition.

Definition 2.7. *Given a joint distribution P_{XY} on finite alphabets $\mathcal{X}$ and $\mathcal{Y}$ the maximal α-leakage from X to Y is defined as*

$$\mathcal{L}_{\alpha}^{\max}(X \to Y) \triangleq \sup_{U-X-Y} \frac{\alpha}{\alpha-1} \log \frac{\max_{P_{\hat{U}|Y}} \mathbb{E}\left[P_{\hat{U}|Y}(\hat{U}=U|U,Y)^{\frac{\alpha}{\alpha-1}}\right]}{\max_{P_{\tilde{U}}} \mathbb{E}\left[P_{\tilde{U}}(\tilde{U}=U|U)^{\frac{\alpha}{\alpha-1}}\right]}, \quad (2.29)$$

with U and $\hat{U}$ taking values on the same finite, but arbitrary alphabet.

The definition of Maximal α-leakage is equivalent to that of maximal leakage when α goes to ∞, while it is equivalent to mutual information when α goes to 1 as discussed in [96].

While the Shannon capacity of a channel $P_{Y^n|X^n}$ captures the number of messages that can be reliably reconstructed at the receiver, maximal leakage impose no reliability constraint. This idea is captured, for i.i.d. inputs, by the following equality [97]:

$$C = \lim_{n\to\infty} \max_{P_{X^n}} \frac{1}{n} I(X^n;Y^n) = \lim_{\epsilon\to 0} \lim_{n\to\infty} \frac{1}{n} \mathcal{L}_{\epsilon}^{C}(X^n \to Y^n) \quad (2.30)$$

with the *recoverable leakage defined by*

$$\mathcal{L}_{\epsilon}^{C}(X \to Y) \triangleq \sup_{\substack{U-X-Y \\ \mathbb{P}[U=\hat{U}|Y]\geq 1-\epsilon}} \log \frac{\mathbb{P}[U=\hat{U}|Y]}{\mathbb{P}[U=\tilde{U}]}. \quad (2.31)$$

Interestingly, as noted in [95], the maximal leakage is equivalent to another notion of leakage previously introduced in [98] under the same name and defined as

$$\mathcal{ML}(X \to Y) \triangleq \sup_{X} \log \frac{\mathbb{P}[X=\hat{X}|Y]}{\mathbb{P}[X=\tilde{X}]} = \sup_{U-X-Y} \log \frac{\mathbb{P}[U=\hat{U}|Y]}{\mathbb{P}[U=\tilde{U}]}. \quad (2.32)$$

2.2.6 Mutual information

Another approach to measure the information leakage, or the information obtain by the UP, about the consumption X, upon observation of the request Y, is the mutual information. In [99] Shannon proposes three fundamental properties that any measure of information must satisfy. Specifically, any information metric H of a random variable X on alphabet $\mathcal{X}$ must satisfy

- Breaking independent choices into successive choices does not increase nor decrease the uncertainty of the outcome, i.e. for any independent X_1 and X_2 random variables $H(X_1 \times X_2) = H(X_1) + H(X_2)$

- For equally likely events, i.e. $P(x) = \frac{1}{|\mathcal{X}|}$ for all $x \in \mathcal{X}$, the uncertainty increases with the number of choices, i.e. H is a monotonically increasing function of $|\mathcal{X}|$.
- H should be a continuous function in $P(X)$.

In [99, Theorem 2] Shannon shows that the only function satisfying the three axioms is

$$H(X) = -K \sum_{x \in \mathcal{X}} P_X(x) \log P_X(x), \tag{2.33}$$

where K is an arbitrary constant typically set to 1, and log is usually the logarithm in base 2 or e. Intuitively, the information metric or entropy function (2.33), measures the average rate needed to describe X. The notion of entropy allows the following generalization to joint entropy of X and Y:

$$H(X,Y) = -\sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} P_{X,Y}(x,y) \log P_{X,Y}(x,y). \tag{2.34}$$

Moreover, the conditional entropy, or entropy of X given that Y is known, is given by

$$H(X|Y) = -\sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} P_{X,Y}(x,y) \log P_{X|Y}(x|y). \tag{2.35}$$

A natural extension capturing the amount of information about X obtained upon measuring Y follows.

Definition 2.8. *Given a joint distribution $P_{Y,X}$, the mutual information is given by*

$$I(X;Y) \triangleq H(X) - H(X|Y), \tag{2.36}$$

where $H(X)$ and $H(X|Y)$ capture the average rate needed to describe X when Y is not known and when it is known respectively

Within this definition, the mutual information captures how much of the entropy of X is explained by Y. The Venn diagram in Figure 2.5 helps contextualize the different measures of entropy and mutual information explained above.

We now describe some fundamental properties shown by the mutual information. More fundamental properties as well as detailed proofs are presented in [100]. Interestingly, by the chain rule, the entropy of a collection of random variables is the sum of the conditional entropies. Specifically, given n random variables $X_1, X_2, \cdots, X_n$

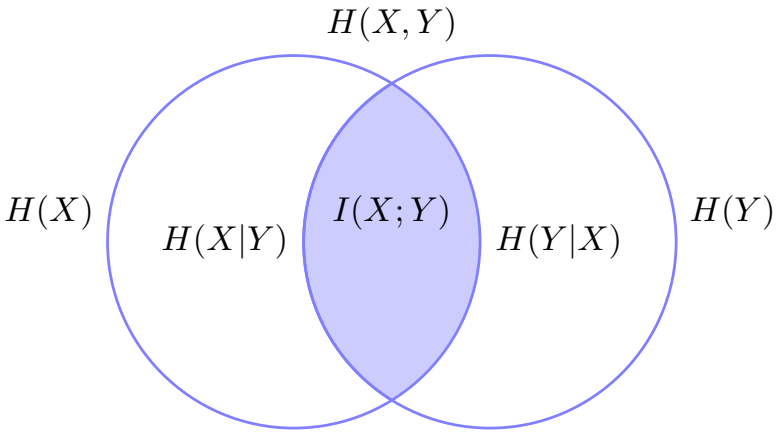

Figure 2.5. Venn diagram of mutual information

with mass probability $p(x_1, x_2, \cdots, x_n)$, the entropy satisfies

$$H(X^n) = \sum_{i=1}^{n} H(X_i|X^{i-1}). \tag{2.37}$$

This property is preserved by the mutual information

$$I(X^n; Y) = \sum_{i=1}^{n} I(X_i; Y|X^{i-1}) \tag{2.38}$$

which follow by definition of entropy and mutual information.

Moreover, the entropy of a random variable is always positive, i.e.

$$H(X) = -\sum_{x\in\mathcal{X}} P_X(x) \log P_X(x) \geq 0. \tag{2.39}$$

This property follows by noting that $P_X(x) \in [0, 1]$ for all $x \in \mathcal{X}$, and therefore $\log P_X(x)$ is always negative, yielding the negative sum greater or equal to zero. On the other extreme, the entropy is always bounded by the logarithm of the cardinality of the output alphabet, i.e.

$$H(X) \leq \log |\mathcal{X}|. \tag{2.40}$$

This follows from the concavity of the logarithm and the Jensen's Inequality:

$$H(X) = \sum_{x\in\mathcal{X}} P_X(x) \log \frac{1}{P_X(x)} \tag{2.41}$$

$$\leq \log \sum_{x\in\mathcal{X}} P_X(x) \frac{1}{P_X(x)} = \log |\mathcal{X}|. \tag{2.42}$$

Interestingly, the mutual information can be expressed in terms of the Kullback-Leibler divergence [101] between two probability distributions

$$I(X;Y) = \mathcal{D}(P_{X,Y}(x,y) \| P_X(x)P_Y(y)), \tag{2.43}$$

where the Kullback-Leibler divergence is given by

$$\mathcal{D}(p\|q) = -\sum_{x\in\mathcal{X}} p(x) \log \frac{p(x)}{q(x)}. \tag{2.44}$$

This relation shows a different perspective on the mutual information, where mutual information is a measure of divergence between the joint $P_{X,Y}(x,y)$ and the product distributions $P_X(x)P_Y(y)$. The more independent or unrelated X and Y are, the more the joint distribution resembles the product distribution, the smaller the divergence, and the smaller the information about X contained in Y. A simple corollary of this identity shows the non-negativity of the mutual information

$$I(X;Y) = \mathcal{D}(P_{X,Y}(x,y) \| P_X(x)P_Y(y)) \geq 0. \tag{2.45}$$

This implies that conditioning reduces entropy, i.e.

$$H(X) = I(X;Y) + H(X|Y) \geq H(X|Y). \tag{2.46}$$

Interestingly, by the data processing inequality, post-processing cannot increase the amount of information. Specifically, given three variables satisfying the Markov chain $\mathcal{X} \to \mathcal{Y} \to \mathcal{Z}$, the mutual information satisfies $I(X;Y) \geq I(X;Z)$, i.e. no transformation $\mathcal{Y} \to \mathcal{Z}$ can increase the information about X. This follows from noting that by the chain rule

$$I(X,Y,Z) = I(X,Z) + I(X,Y|Z) = I(X,Y) + I(X,Z|Y), \tag{2.47}$$

and using the fact that $I(X,Z|Y) = 0$ by assumption and $I(X,Y|Z) \geq 0$.

In some scenarios, as the one captured by maximal leakage, the provider is not interested on the energy consumption sequence X, but on a function of it (e.g. whether the user is at home). Thus, it is interesting to note that by the data processing inequality, for any Markov chain $U - X - Y$ such that $U \perp Y|X$ it holds that

$$I(U;Y) \leq I(X;Y). \tag{2.48}$$

Thus, the mutual information provides an upper bound on the information leaked about any possible function U.

Another concern arising with the usage of mutual information as a privacy measure is its heavy dependency on the input distribution P_X. In other words, it captures the information leaked by a specific distribution of the energy consumption. However, the energy consumption of real users greatly varies between users and over time, often presenting non-stationary behaviours.

2.3 Existing solutions and techniques

In the following we present some of the existing privacy preserving solutions and techniques available in the literature. Some surveys of the current state of the art are also available in [65, 47, 102, 103].

2.3.1 Heuristic policies

Heuristic policies are constructed based on human intuition of what good policies should resemble. These policies are not derived in a systematic manner, as the solution of a mathematical optimization process. Although in the absence of lower bounds, they provide no guarantee of optimality, the heuristic origin of these policies does not hinder the characterization of the resulting leakage according to different privacy measures.

In the best effort (BE) algorithm [58], the EMU attempts to preserve the energy requested form the grid constant over time. Therein, whenever allowed by the battery constraints, the energy request at time i matches the request at the previous time step, i.e. $Y_i = Y_{i-1}$. When such request is not possible, the battery is fully charged or discharged. The authors of [58] proposed three different measures of information leakage: Kullback-Leibler divergence, cluster classification, and regression analysis. Intuitively, the larger the difference between the energy consumption and the request the more private a system is. Thus, the authors suggest the usage of the empirical Kullback-Leibler divergence between the input P_{X^n} and output P_{Y^n} distributions. The second metric proposed is based on cluster classification. This metric is based on the idea that the power consumption $\{X_i\}$ fluctuates around a set of values depending on the underlying system, thus cluster classification of the output load $\{Y_i\}$ is also employed as a measure of privacy. The third metric proposed is regression analysis, and in particular the coefficient of determination R^2, i.e. which proportion of the residuals are explained by the model. These metrics are further discussed in

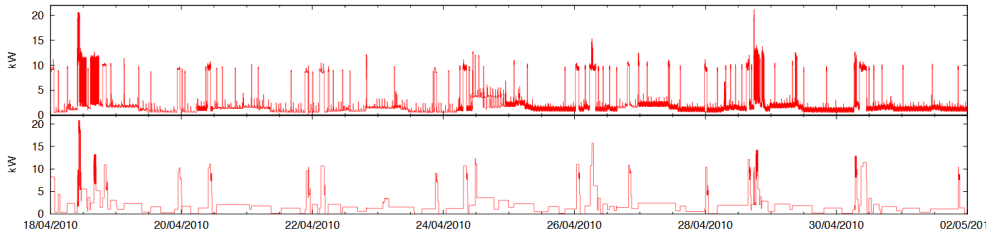

Figure 2.6. Energy consumption X^n (top) and energy request Y^n (bottom) when NILL proposed in [4] is used with a realistic 6kWh battery.

[104], where the load balancing of the grid is also considered in the optimization process.

Another approach, named *non-intrusive load levelling* (NILL), is proposed in [4]. Therein, when allowed by the battery constraints, the energy request is set to a steady state value K_{SS}. Otherwise, the request is set to a charging K_L or discharging K_H level. Thus the energy request Y_i lies in the finite alphabet $\{K_H, K_{SS}, K_L\}$, where the middle state K_{SS} is given preference over the others. This approach aims to reduce the number of so called features, or variations on the output load used by state of the art non-intrusive load monitoring algorithms to infer information about the user [83, 44]. Thus, the reduction of the number of features is used in [4] as a measure of privacy. Simulation results using complex realistic battery models and real energy consumption data are provided in Figure 2.6. The authors show that with their proposed algorithm the number of features decreases from thousands to $1-6$ per day. Comparison of the empirical entropies of the consumption and request, i.e. $H(X^n)$ and $H(Y^n)$ is also used in [4] as a measure of privacy. A similar strategy is compared in [79] with the solution of a greedy Markov Decision Process.

The two algorithms above, i.e. BE and NILL, are reviewed and generalised in [5]. Therein, the energy requests take values in a finite alphabet $\{0, d, 2d, \cdots\}$, where the value of d is selected so that maximum charging and discharging constraints of the battery are always satisfied, i.e. $d = \min\{P_{\min}, P_{\max}\}$. Thus, given a consumption $X(t)$, the EMU is allowed to request $Y(t) = \lceil X(t)/d \rceil$ or $Y(t) = \lfloor X(t)/d \rfloor$. Three so called stepping algorithms are proposed in [5] to decide when to round up or down. The lazy stepping (LS1 and LS2) algorithms, aim to keep the request constant, unless otherwise forced by the battery constraints. The lazy charging (LC) algorithm aims to reduce the number of charging/discharging cycles of the battery, by always rounding up/down, until the battery is fully discharged/charged, when the strategy is inverted and repeated. A third stepping algorithm is the random charging (RC), where the rounding direction is selected at random. Numerical approximations of the mutual information between particular functions of the consumption and request

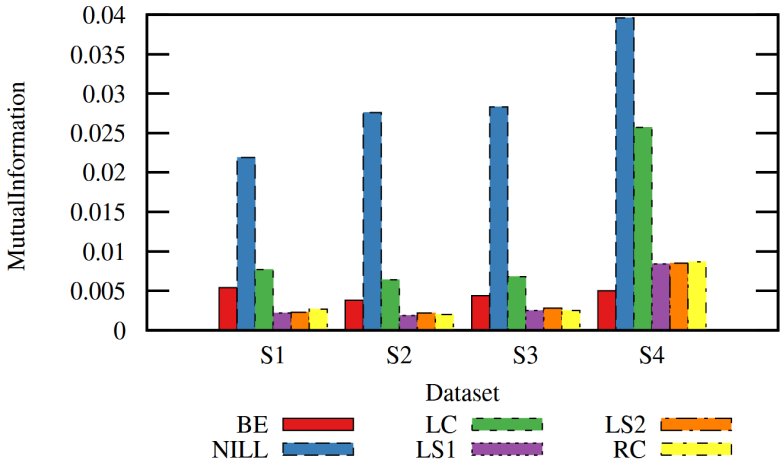

Figure 2.7. Comparison of the privacy achieved by different heuristic algorithms over multiple data sets [5].

sequences i.e. $I(f(X); f(Y))$ are studied as a measure of privacy and compared between different algorithms, Figure 2.7.

2.3.2 Quadratic deviation: Lagrangian and Lyapunov

While arguably not the best available metric to capture the privacy leakage, the tractability of quadratic deviation makes this metric worth exploring. Moreover analysis of the solutions yield by the optimization of this metric can lead to interesting insights. In [71, 6, 84, 105] the quadratic deviation from average is proposed as a privacy metric, i.e.

$$\mathcal{V}(\mathbf{w}) = \frac{1}{n}\,\mathbb{E}\left[\|Y - \mathbf{w}\|_2^2\right], \tag{2.49}$$

with $w_i = \mathbb{E}[1/n \sum_{j=0}^{n-1} X_j]$ for all i. In [78] the above idea is relaxed, and the objective level w_i is allowed to fluctuate with the market price m_i, i.e. $w_i = f(m_i)$. Therein, the user is also interested in minimizing the price paid for the energy, $B = \frac{1}{n}\mathbf{m}^T Y^n$, for a given market price $\mathbf{m} \in \mathbb{R}^n$. In order to preserve privacy, the user has access to an energy storage device with capacity β. This joint optimization of $\mathcal{V}$ and B under the battery constraints is then posed as a weighed average optimization problem:

$$\min_{\Pi}\Big(\theta\mathcal{V}(\mathbf{w}) + (1-\theta)B\Big) = \min_{\Pi}\frac{1}{n}\left(\theta\,\mathbb{E}\left[\|Y - \mathbf{w}\|_2^2\right] + (1-\theta)\mathbf{m}^T Y^n\right), \tag{2.50}$$

where Π denotes the set of feasible battery polices such the battery constraints are satisfied, i.e. $0 \leq \sum_{i=0}^{n-1}(Y_i - X_i) \leq \beta$ for all i. Note that this is a convex optimization problem, that is solved by standard Lagrangian multipliers. The solution is a backward water-filling algorithm, where the water level is limited by the battery

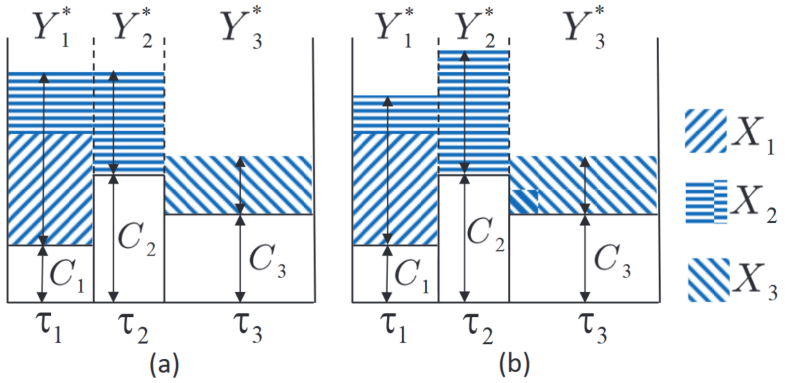

Figure 2.8. Backward water-filling algorithm resulting from joint cost-variance optimization in [6], with (a) infinite and (b) finite capacity batteries, X_i, C_i and τ_i respectively denoting the energy consumption, the price paid of the energy, and the length of the i-th market block. Figure is obtained for $\theta = 1/3$.

capacity. Figure 2.8 shows an instance of this backward water-filling algorithm, where the water level tends to be flat subject to the constraints imposed by the battery, i.e. the energy request must precede the consumption (water can not be moved forward) and only a limited amount of water can be moved backwards (in the case of finite capacity batteries). In [105] mutual information is numerically evaluated for the resulting policy obtained from the minimization of the load variance.

One of the main advantages of the above model is that it makes no assumptions over the statistics of X^n or M^n. Contrary to usual information-theoretic metrics no a priori knowledge on P_{X^n} or P_{M^n} is required. Moreover this technique holds for arbitrary input X^n and output Y^n alphabets. However, the energy consumption is required to be known in advance. In [80] a relaxation that does not require future knowledge of the energy consumption and asymptotically achieves optimality is proposed. The solution is achieved by deriving the Lyapunov function with a drift-plus-penalty [80].

2.3.3 Finite state machine

Another possible approach to capture the memory introduced by batteries are finite state machines. Figure 2.9 depicts a simple FSM, representing binary system, with a battery of capacity one, and binary input X^n and output Y^n alphabets, i.e. $\mathcal{S} = \mathcal{X} = \mathcal{Y} = \{0, 1\}$. Therein, at any time step i, the state of the FSM captures the state of the battery, while the different consumption profiles P_{X^n} and the battery policies $P_{Y^n|X^n}$ determine the transition probabilities between states. This is shown on Table 2.3 and Figure 2.9. Note that this model allows for larger battery sizes and input/output alphabets at the expense of increased number of states and complexity.

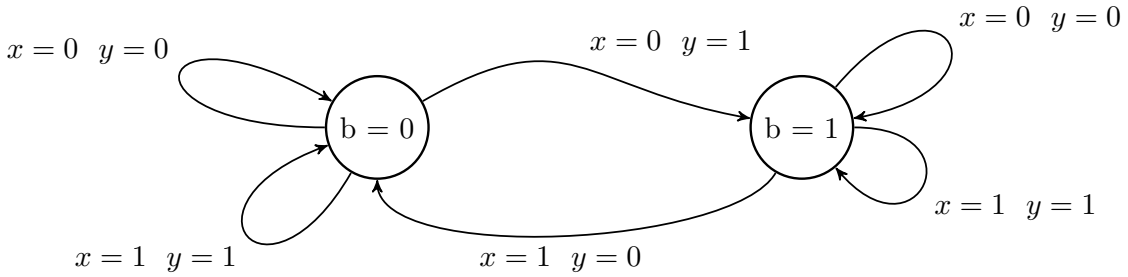

Figure 2.9. Finite State Machine modelling a binary battery

x_i	s_i	y_i	$P(y_i\|x_i, s_i)$	s_{i+1}
0	0	0	p_a	0
0	0	1	$1-p_a$	1
0	1	0	1	1
0	1	1	0	-
1	0	0	0	-
1	0	1	1	0
1	1	0	$1-p_b$	0
1	1	1	p_b	1

Table 2.3. FSM transition table for SMs with battery.

In [57] a simple FSM with a binary i.i.d. input and battery is studied. In [13, 56] this model is extended to include other elements such as Energy Harvesting Devices (EHDs) or schemes allowing energy waste. Extensions to non-binary batteries or any arbitrary size input, output and battery alphabets are studied in [13]. For the continuous alphabet case, the usage of simulation based solutions [106] is proposed in [57]. Simulation based solutions allow the computation of upper and lower bounds on the information leakage rate. Note that in theory the restriction to i.i.d. inputs can also be relaxed, extending the problem to include Markov sources, by including the previous input X_{i-1} in an augmented state space of the FSM.

In order to characterize the privacy leakage numerical approximations of the mutual information are employed in [56, 57, 13]. These approximations rely on the Asymptotic Equipartition Property (AEP) [100]. The AEP states that when $n \to \infty$ the logarithm of the joint probability $\frac{1}{n}\log P(r_1, \cdots, r_n)$ limits in probability to the entropy of the process $H(R^n)$. This leads to

$$\begin{aligned} nI(X^n;Y^n) &= H(X^n) + H(Y^n) + H(X^n, Y^n) && (2.51)\\ &\approx H(X^n) - \log P(y_1, y_2, \cdots, y_n) \\ &\quad + \log P(x_1, x_2, \cdots, x_n, y_1, y_2, \cdots, y_n), && (2.52) \end{aligned}$$

where the values of $P(y_1, y_2, \cdots, y_n)$ and $P(x_1, x_2, \cdots, x_n, y_1, y_2, \cdots, y_n)$ are recursively computed by running the forward pass of the Bahl-Cocke-Jelinek-Raviv (BCJR) algorithm [107]. Exhaustive search among a discrete version of all the possible policies is then employed in [13, 57, 56] to find the best battery policy.

Interestingly, in [108] a similar approach is used to compute the event-based information theoretic privacy:

$$\sup_i I(X_i; Y^n). \tag{2.53}$$

The aim of this metric is to capture adversaries that are not interested in the whole sequence X^n but only on certain event X_i. This formulation is one step closer to that of ϵ-mutual information differential privacy (2.21), and its links to differential privacy.

FSM models provide a simple approach, enabling the tools to calculate information-theoretic privacy figures in simple scenarios. However characterization of optimum points of tradeoff requires exhaustive search over all the possible strategies. At the same time, the computational complexity of FSM models grows with the size of the alphabets $\mathcal{S}$, $\mathcal{X}$ and $\mathcal{Y}$. This exponential growth poses serious difficulties in the extension of this model to continuous or arbitrary size alphabets.

A more fundamental drawback of the FSM model is that it fails to capture the complete set of actions the user can take in order to protect the privacy. The scheme proposed in Table 2.3 assigns the output probabilities as a function of the current input and battery state. However, in a system with memory, optimal privacy policies decide the output taking all previous inputs and outputs into account. This can be partially solved by including last outputs [57] and inputs into the state of the FSM. However, including all past inputs and outputs in the state space renders the problem computationally infeasible.

2.3.4 Markov decision process

Markov decision processes (MDP) provide a mathematical framework to model decision processes satisfying the Markov property. In particular, let S_i denote the state of the system at time i. For any time i and state S_i, the controller is allowed to take a possibly random decision A_i out of the set of possible decisions $\mathcal{A}(S_i)$. For a given action and state, the system then moves into a new state with probability $P(S_{i+1}|A_i, S_i)$. That is, the state transition probability depends on the current action and state but not on past actions or states, i.e. the system satisfies the Markov property. This transition generates a cost for the controller $c(S_i, A_i)$, which is a

function of the decision taken and the current and future state. The objective of the MDP is to minimize the total cost obtained by the decision maker over a finite interval (finite horizon MDP) or the infinite future (infinite horizon MDP) [109]. In the SMs context, when the energy consumption and the energy generated by the available energy sources satisfy the Markov property, it is possible to formulate the optimization of the energy request process as an MDP.

Quadratic deviation

In [7, 69] the privacy optimization problem is formulated as an MDP. Therein, the aim of the user is to minimize the privacy leakage measured in terms of quadratic deviation while simultaneously minimizing the energy bill. Unlike most models in the literature, the user buys and sells energy at different prices m_i^c and m_i^d, respectively. Thus, the cost function is given by

$$J(\pi) = \mathbb{E}\left[\sum_{i=0}^{n-1} c(s_i, a_i)\right] = \mathbb{E}\left[\sum_{i=0}^{n-1} \lambda\mathcal{V}(s_i, a_i) + (1-\lambda)B(s_i, a_i)\right], \tag{2.54}$$

where $\mathcal{V}(s_i, a_i)$ and $B(s_i, a_i)$ respectively denote the privacy leakage and the energy bill induced by (s_i, a_i), and π denotes a feasible battery policy. Interestingly, in [7, 69], two possible models for the energy storage device are considered, namely: thermal and chemical energy storage devices. Within those models maximum capacity, minimum and maximum charging rates, initial battery state, and charging efficiency factor are considered.

To determine the optimal policy, the transition probabilities of the energy consumption, i.e. $P_{X_{i+1}|X_i}$ are required. However, the authors note the non stationary nature and the large variability of the energy consumption. Thus a model-free learning method is employed to learn the cost associated with each state-action pair. In particular an $\eta-$greedy Q-Learning algorithm [110] that explores a random action with probability η and exploits a greedy action with probability $1-\eta$ is employed. After each action, the system learns the incurred cost $c(s_i, a_i)$, and updates the learned cost function $Q : \mathcal{S} \times \mathcal{A} \to \mathbb{R}$ as

$$Q(s_i,a_i) \leftarrow Q(s_i,a_i) + \alpha\left(c(s_i,a_i) + \min_{a\in\mathcal{A}(s_{i+1})}(Q(s_{i+1},a) - Q(s_i,a))\right), \tag{2.55}$$

where $\alpha \in [0, 1]$ is the learning rate. In order to characterize the performance of the learning algorithm, the authors of [7] compare it with the corresponding offline algorithm, where the future energy consumption is known in advance to the algorithm. Figure 2.10 compares this two algorithms, showing that both algorithms yield very

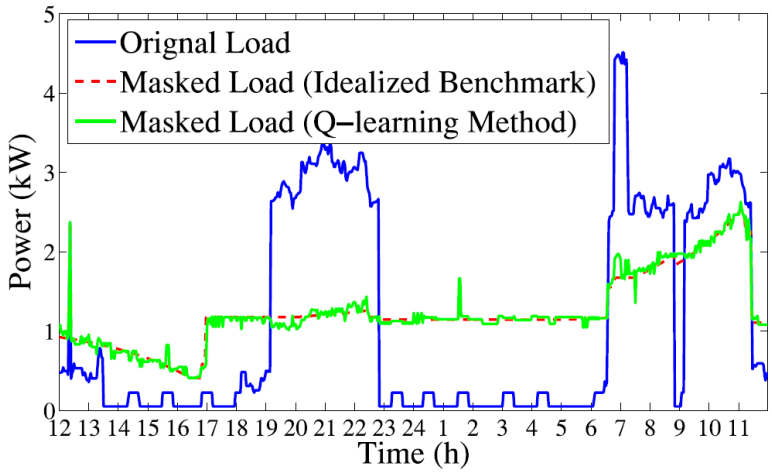

Figure 2.10. Energy consumption and requests under the online (Q-learning) and the offline (Benchmark) algorithms proposed in [7] for a battery of capacity 10kWh.

similar energy requests. The privacy obtained by this method is further analysed by computing the empirical mutual information.

Mutual information

In [14, 111] mutual information is considered as a privacy measure when the user has access to an energy storage device. Thus, the cost function is given by

$$J(\pi) = I(S_0, X^n; Y^n) = \sum_{i=0}^{n-1} I(S^i, X^n; Y_i | Y^{i-1}), \tag{2.56}$$

where the equality follows from the chain rule and noting that the battery states are determined by the past input and outputs. At any time step i, the feasible values of Y_i are determined by the current energy consumption X_i and battery state S_i. Noticing also the following inequality,

$$\sum_{i=0}^{n-1} I(S^i, X^n; Y_i | Y^{i-1}) \geq \sum_{i=0}^{n-1} I(S_i, X_i; Y_i | Y^{i-1}), \tag{2.57}$$

the authors of [14] show that there is no loss of optimality in restricting to battery policies that decide the request Y_i based only on the current consumption X_i, battery state S_i and past requests Y^{i-1}, i.e. $f : \mathcal{X} \times \mathcal{S} \times \mathcal{Y}^{i-1} \to \mathcal{Y}$.

This additive formulation of the cost function, allows to formulate the problem as a MDP by defining the system state as (X_i, S_i, Y^{i-1}). However, under this formulation, the number of states grows exponentially with time, rendering the computation

infeasible for large n. This is solved by showing that the belief state:

$$\pi_i(x,s) = \mathbb{P}[X_i = x, S_i = s | Y^{i-1} = \mathbf{y}^{i-1}], \tag{2.58}$$

is a sufficient statistic for (2.57). The discrete approximation of the belief state can be calculated recursively. Defining the system state as (X_i, S_i, π_i) allows the formulation of the problem as an MDP, for which numerical simulations yield the optimal policies.

For the i.i.d. case, [14] shows that there is no lost of optimality in further restricting the search to invariant policies i.e. policies such that

$$P(X_i, S_i, Y_i | Y^{i-1}) = P(X_1, S_1, Y_1). \tag{2.59}$$

Marginalizing the above distribution shows that under invariant policies, the request $\{Y_i\}$ is i.i.d.. The authors of [14] further restrict the search to so called structured policies, under which the marginal distribution of Y_i is P_X, yielding consumption and request statistically indistinguishable. After this structural simplifications and by further restricting the system state to $W_i = S_i - X_i$ a single letter characterization is obtained. The information leakage under the optimal policy is then given by

$$J(\pi^*) = \min_{P_S \in \mathcal{P}_S} I(S - X; X), \tag{2.60}$$

and is achieved by the memoryless structured policy

$$q(y|x,s,\pi_i) = \frac{P_X(y) P_S^*(y - x + s)}{\sum_{x'} P_X(x') P_S^*(x' - x + s)}, \tag{2.61}$$

where P_S^* denotes the optimal distribution achieving (2.60). The authors of [14] note that extending the MDP formulation to incorporate extra additive costs such as the price paid for the energy is rather immediate. However, the approach used to characterize the single letter expression for the case of the i.i.d. consumption may not generalize well. The above also appears in [111–114]. In [115] this idea is extended from i.i.d. to independent periodically time varying input distributions. It is shown that the straight forward application of the above scheme is suboptimal for periods of length two.

In [79] the price paid for the energy is considered by characterizing upper and lower bounds on the privacy-cost tradeoff. The corresponding MDP is formulated by means of belief state (2.58) for scenarios where the user is allowed access to an energy storage device. In [116, 117] the MDP formulation of systems with access to an EHD and a finite capacity battery is presented. Interestingly, a model predictive

control (MPD) and a dynamic programming approach are presented in [118] and [119] respectively.

Hypothesis testing

In [8], the behaviour of the user at time i is given by $H_i \in \mathcal{H}$ with probability law determined by the time invariant kernel $P_{H_{i+1}|H_i}$, i.e. H^n is a Markov process. Based on that behaviour, the user consumes $X_i \in \mathcal{X}$ units of energy with probability determined by the time invariant kernel $P_{X_{i+1}|H_{i+1},X_i}$. To preserve the privacy of the user, the EMU has access to a finite capacity energy storage device with internal state denoted by S_i.

Within this setting, at every time step i, the attacker performs an n-ary hypothesis test in order to infer the behaviour h_i. Let $c(\hat{h}_i, h_i) > 0$ denote the cost incurred by the adversary on making decision $\hat{h}_i$ when h_i is true. It is then shown in [8] that the optimal strategy $P^*_{\hat{H}|Y_i}$ from a greedy adversary is to perform a deterministic likelihood-ratio test (LRT). The minimal Bayesian risk r_i^* is defined by

$$r_i^* = \sum_{(\hat{h}_i,h_i)\in\mathcal{H}^2} c(\hat{h}_i, h_i) P_{\hat{H}_i,H_i}(\hat{h}_i, h_i). \tag{2.62}$$

The privacy measure is defined as the accumulated discounted minimal Bayesian risk, i.e.

$$V = \sum_{t=0}^{\infty} \beta^t r_i^*, \tag{2.63}$$

with $\beta \in [0, 1)$. This metric suits scenarios where privacy concerns degrades over time. The aim of the EMU is to implement the policy that maximizes privacy given by

$$P^*_{Y_i|X_i,S_i} = \underset{\Pi}{\operatorname{argmax}}\, V \tag{2.64}$$

for Π denoting the set of feasible battery policies. This problem is then formulated as an MDP problem with state (h_i, x_i, s_i) and belief state $b_i = P_{H_i,X_i,S_i}$. Noting the complexity of solving the above model, a simplified version is also proposed. Thus an instantaneously optimal control strategy maximizing the instantaneous Bayesian risk is proposed, and given by

$$P^*_{Y_i|X_i,S_i} = \underset{\Pi}{\operatorname{argmax}}\, r_i^*(b_i) \tag{2.65}$$

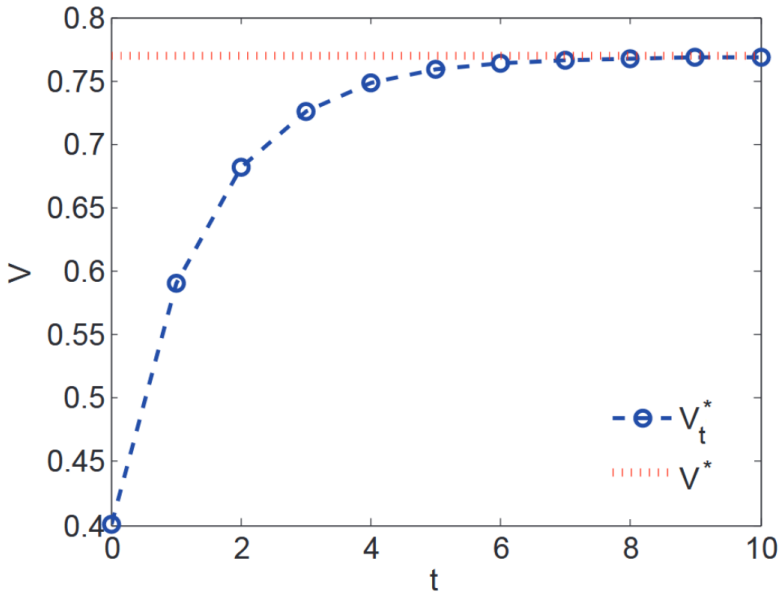

Figure 2.11. Accumulated discounted minimal Bayesian risk V against time step $t = i$, for optimal policies V^* and instantaneously optimal (greedy) policies V_t^* [8].

Numerical results comparing both strategies are then computed and shown in Figure 2.11.

2.3.5 Instantaneous constraints - rate distortion theory

The approach presented in [120, 9, 121, 11, 10, 122, 53] relies on noticing that when users have access to an alternative energy source with instantaneous power constraints, the information leakage, measured in terms of mutual information, resembles the well known information rate distortion function. This link allows access to concepts and information theoretic tools developed for the calculation of rate distortion functions. Within this setting, the user has access to an alternative energy source, with peak power constraint $\hat{P}$ and average power constraint $\overline{P}$. In [120], the information leakage is then given by

$$\mathcal{I}(\overline{P}, \hat{P}) = \inf_{\Pi(\overline{P}, \hat{P})} I(X^n, Y^n), \tag{2.66}$$

where $\Pi(\overline{P}, \hat{P}) = \{P_{Y^n|X^n} : y_i \in \mathcal{Y}, \mathbb{E}[(X_i - Y_i)] \leq \overline{P}, 0 \leq X_i - Y_i \leq \hat{P},\ \forall i\}$ denotes the set of feasible battery policies that satisfy the instantaneous power constraints. Similarly, the information rate distortion function is given by

$$R^I(D) = \inf_{\mathbb{E}[\sum_i d(X_i, \hat{X}_i)] \leq D} I(X^n; \hat{X}^n). \tag{2.67}$$

The analogy between mutual information, information leakage and rate distortion is shown by setting the distortion function to

$$d(x_i, y_i) = \begin{cases} x_i - y_i, & \text{when } 0 \leq x_i - y_i \leq \hat{P} \\ \infty, & \text{otherwise.} \end{cases} \tag{2.68}$$

Interestingly, by the rate distortion theorem [100], the information rate distortion function $R^I(D)$ is equivalent to the rate distortion function $R(D)$. This rate determines the size $2^{nR(D)}$ of the minimum codebook needed to compress a source X^n, such that the distortion between the original X^n and the reconstructed copy $\hat{X}^n$ satisfies $\mathbb{E}[\sum_i d(X_i, \hat{X}_i)] \leq D$. However, a main difference between both formulations, as noted by the authors of [120], is that the rate distortion function is considered for offline lossy compression, while the policies they derived for the smart meter context typically work in an online manner. Different privacy results are obtained in the literature by exploiting this link on i.i.d. scenarios, where the rate distortion function is better understood.

It is shown in [9, Theorem 1] that for i.i.d. inputs and additive fidelity criteria (2.68) the optimal battery policy is memoryless, i.e.

$$\inf_{\Pi(\overline{P},\hat{P})} I(X^n; Y^n) = \inf_{\Pi(\overline{P},\hat{P})} I(X; Y). \tag{2.69}$$

Furthermore, for discrete inputs, it is shown in [10, 120] that under i.i.d. assumptions and instantaneous power constraints there is no loss of optimality in restricting to discrete output alphabets of the form $\mathcal{Y} = \mathcal{X}$. The proof is done by constructing the mapping:

$$P_{\hat{Y}|X}(\hat{y}|x) = \begin{cases} 0, & \text{when } x_i \leq \hat{y} \leq x_{i+1}, \\ \int_{(x_i, x_{i+1}]} P_{Y|X}(y|x)dy, & \text{when } \hat{y} = x_{i+1}, \\ P_{Y|X}(y|x), & \text{otherwise,} \end{cases} \tag{2.70}$$

for each i and ordered set $\mathcal{X}$. The mapping is then shown to preserve the feasibility, and the data processing inequality guarantees that the mutual information does not increase. This reduces the optimization to discrete alphabets, moreover by the convexity of the privacy metric over $P_{Y^n|X^n}$, the privacy optimization is reduced to a convex problem with linear constraints. This allows direct application of efficient algorithms such as the Blahut-Arimoto algorithm, to numerically compute the information leakage [100].

In the case of continuous input loads, the well known rate distortion Shannon Lower Bound, i.e.

$$\inf_{\Pi(\overline{P},\hat{P})} I(X;Y) = \inf_{\Pi(\overline{P},\hat{P})} \Big(H(X) - H(X|Y) \Big) \tag{2.71}$$

$$\geq H(X) - \sup_{\Pi(\overline{P},\hat{P})} H(X-Y). \tag{2.72}$$

is proposed in [10]. The distribution of $V = X - Y$ that maximizes the entropy among those random variables V with mean $\overline{P}$ and satisfying $0 \leq V \leq \hat{P}$ is the truncated exponential distribution P_V [100, Chapter 11], defined as

$$P_V(v) = \begin{cases} \frac{1}{\lambda_0} e^{-\frac{v}{\lambda_1}} & 0 \leq v \leq \hat{P} \\ 0 & \text{otherwise,} \end{cases} \tag{2.73}$$

where the values of λ_0 and λ_1 are set so that $\mathbb{E}[V] = \overline{P}$ and $\int_{\infty}^{\infty} P_V(v)dv = 1$. Thus, the Shannon Lower bound, as provided in [9], is given by

$$\inf_{\Pi(\overline{P},\hat{P})} \frac{1}{n} I(X^n;Y^n) \geq h(X) - \ln(\lambda_0) - \frac{\overline{P}}{\lambda_1}, \tag{2.74}$$

where $h(\cdot)$ stands for differential entropy. For input distributions continuous on $\mathcal{R}_+$ the optimal battery policy is given by

$$P_{Y|X}(y|x) = P_V(x-y)\frac{P_Y(y)}{P_X(x)}, \tag{2.75}$$

where the output distribution $P_Y(y)$ is obtained by Laplace transformation [9]. This bound is known to be tight for i.i.d. exponential distribution on the input. The privacy-average power tradeoff obtained in [9] for different peak power constraints is depicted in Figure 2.12.

In [10] these results are generalized for multiuser scenario, where multiple users share access to a common AES. This is done by setting $X_i = (X_{0,i}, X_{1,i}, \cdots, X_{N-1,i})$ and $Y_i = (Y_{0,i}, Y_{1,i}, \cdots, Y_{N-1,i})$. For independent distributions over different users, and under no peak power constraints, the privacy function is shown to be

$$\mathcal{I}(\overline{P}, \infty) = \inf_{\sum_{j=0}^{N-1} \overline{P_j} \leq \overline{P}} \sum_{j=0}^{N-1} \mathcal{I}(\overline{P_j}, \infty). \tag{2.76}$$

This approach leads to a reverse waterfilling algorithm. This is in line with well known results obtained for problems with similar formulations as the case of rate distortion function for parallel Gaussian sources [100, Theorem 10.3.3]. Comparison of numerical results with those of non-collaborative scenarios show the benefit

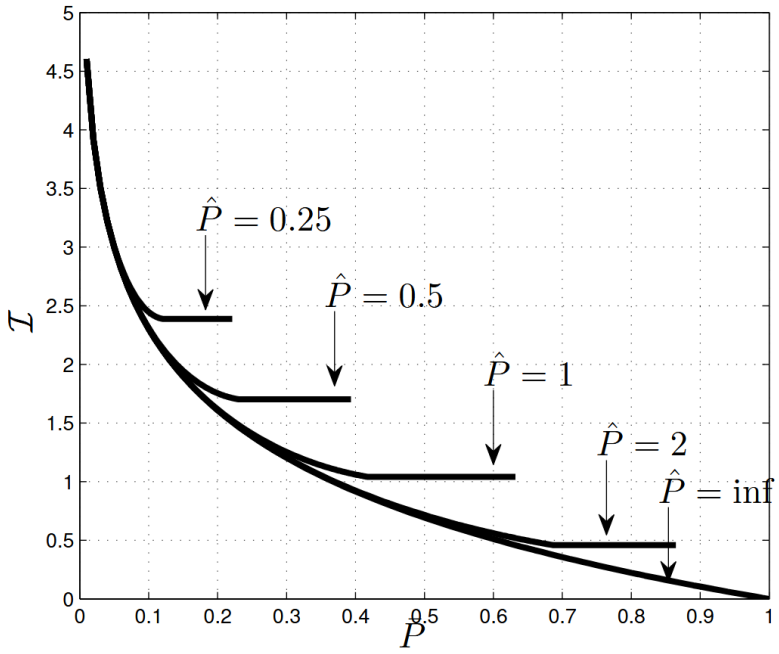

Figure 2.12. Privacy-Average power tradeoff for different peak power constraints and $X \sim \exp(1)$ as shown in [9].

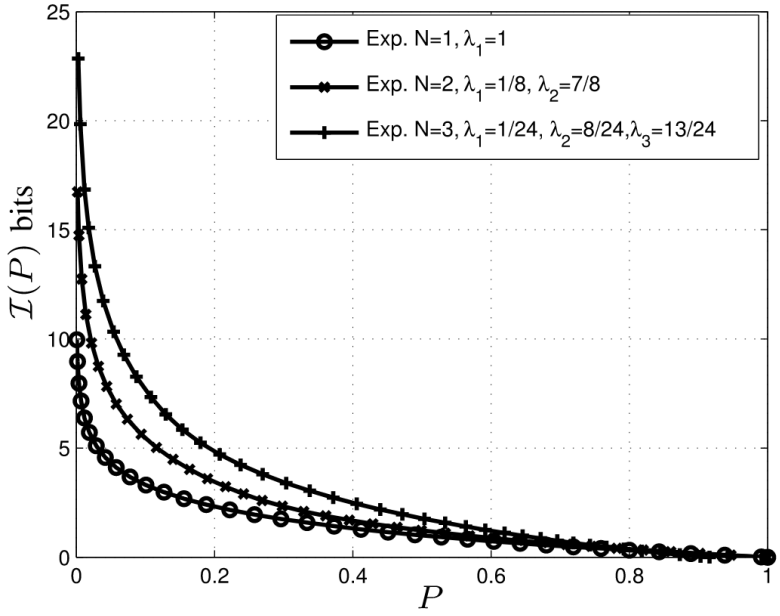

Figure 2.13. Privacy-Average power tradeoff for different number of users as shown in [10].

of multiuser collaboration in Figure 2.13. Extensions including energy waste are presented in [121].

Interestingly, in [11] systems with EHD and infinite capacity batteries are shown to be equivalent to systems with an average power constrained AES. Specifically, under infinite capacity battery, i.i.d. inputs, and EHD generating an average power

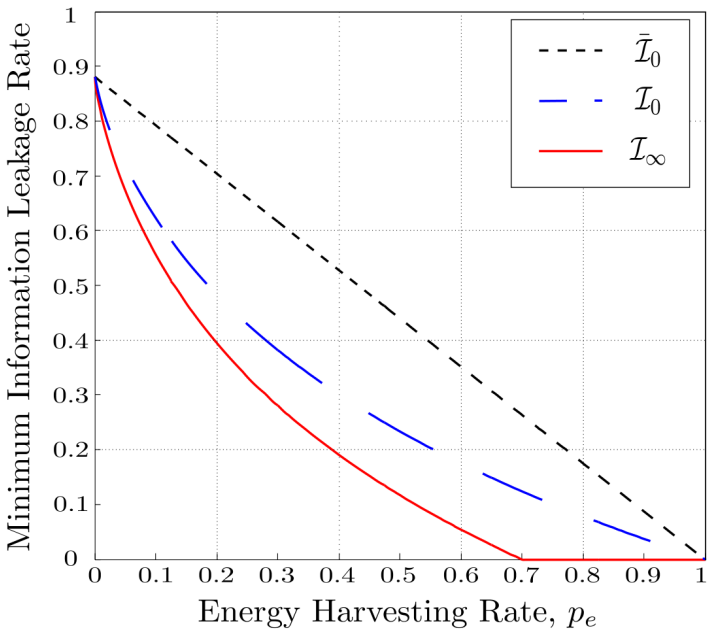

Figure 2.14. Leakage rate with an EHD and an battery with infinity capacity $\mathcal{I}_\infty$, and zero capacity for a UP knowing $\tilde{\mathcal{I}}_0$ and not knowing $\mathcal{I}_0$ the energy generated by the EHD [11].

of $\overline{P}_E$ the leakage is shown to be

$$\mathcal{I}_\infty = \mathcal{I}(\overline{P}_E, \infty). \tag{2.77}$$

For the case of no battery, i.i.d. inputs and i.i.d. EHD generating E units of energy, the leakage is shown to be characterized by

$$\mathcal{I}_0 = \inf_{P_{Y|X,E}: 0 \leq X - Y \leq E} I(X;Y), \tag{2.78}$$

where $I(X;Y)$ is conditioned on E, i.e. $I(X;Y|E)$ when the amount of energy harvested is known to the provider. The three cases above are compared in Figure 2.14, providing upper and lower bounds for cases with finite capacity batteries.

Another interesting approach comes from [53, 122], where instead of a battery, measurement are distorted before being sent to the provider. Therein, the utility of the data sent to the provider limits the amount of distortion that can be applied to the data. The distortion is measured in terms of quadratic distortion while the privacy is measured in terms of mutual information.

2.3.6 Hypothesis testing: Chernoff-Stein lemma

This family of strategies relies on the application of the Chernoff-Stein Lemma in order to characterize the asymptotic decay rate of the p_{II}^* probability of error when the privacy is measured in terms of Hypothesis testing. Specifically, the Chernoff-Stein lemma [100] shows that for i.i.d. distributions the asymptotic exponential decay rate of p_{II}^* is given by

$$r_{II}^* = \lim_{n\to\infty} -\frac{\log p_{II}^*}{n} = \mathcal{D}(P(Y|h_0)\|P(Y|h_1)), \tag{2.79}$$

where $\mathcal{D}(p\|q)$ denotes the Kullback-Leibler divergence between two probability distributions p and q, and is given by (2.44).

In [87] the user is allowed access to an EHD generating Z_i units of energy with probability P_Z. Thus, the EHD is i.i.d. and independent of the consumption process X_i. The user has also access to an energy storage device with "sufficiently large" capacity that is always able to satisfy the energy demands of the user, as long as the total amount of energy stored or consumed from the battery is asymptomatically zero, that is,

$$\lim_{n\to\infty} \sum_{i=0}^{n-1} (y_i + z_i - x_i) = 0. \tag{2.80}$$

The energy consumption of the user is assumed to be i.i.d. with distribution $P_{X|H}$. Thus, the optimization problem becomes

$$\min_{\mathcal{P}_{Y|H}} \mathcal{D}(P_{Y|H}(\cdot|h_0)\|P_{Y|H}(\cdot|h_1)), \tag{2.81}$$

with $\mathcal{P}_{Y|H}$ the set of memoryless policies such that

$$\mathbb{E}(Y|h_0) = \mathbb{E}(X|h_0) - \mathbb{E}(Z) = f_0, \tag{2.82}$$

$$\mathbb{E}(Y|h_1) = \mathbb{E}(X|h_1) - \mathbb{E}(Z) = f_1. \tag{2.83}$$

The above problem is a convex optimization problem with linear constraints. Thus the optimal solution has to satisfy the Karush-Kuhn-Tucker (KKT) conditions [123]. Careful study of the KKT conditions under different cases allowed the authors of [87] to show that the alphabet of the optimal policy satisfies the cardinality constraints $|\mathcal{Y}^*| \leq 2$. The trivial case $|\mathcal{Y}^*| = 1$, occurs only when both hypothesis have the same expected value $f_0 = f_1$, yielding $r_{II}^* = 0$. On the non-trivial scenario $|\mathcal{Y}^*| = 2$, the optimal alphabet is shown to satisfy $\mathcal{Y}^* = \{\min \mathcal{Y}, \max \mathcal{Y}\}$, yielding

$$r_{II}^* = \frac{f_0 - \min\mathcal{Y}}{\max\mathcal{Y} - \min\mathcal{Y}} \log \frac{f_0 - \min\mathcal{Y}}{f_1 - \min\mathcal{Y}} + \frac{\max\mathcal{Y} - f_0}{\max\mathcal{Y} - \min\mathcal{Y}} \log \frac{\max\mathcal{Y} - f_0}{\max\mathcal{Y} f_1}. \tag{2.84}$$

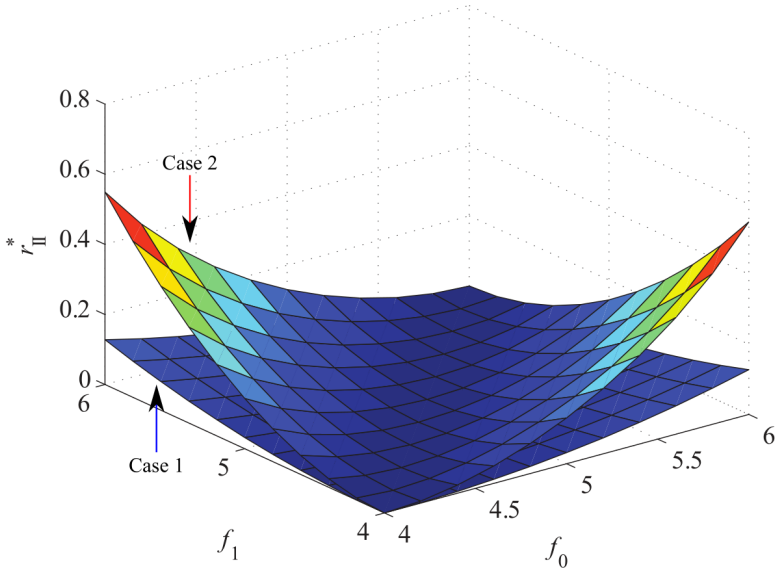

Figure 2.15. Minimum KL distance r_{II} against the constraints on the average energy supply from the energy provider (f_0, f_1). Case 1: $\min \mathcal{Y} = 1$ and $\max \mathcal{Y} = 9$. Case 2: $\min \mathcal{Y} = 3$ and $\max \mathcal{Y} = 7$).

The numerical example in Figure 2.15 shows two ways to enhance the privacy of smart meters: increasing the difference $\max \mathcal{Y} - \min \mathcal{Y}$ and decreasing the difference $|f_0 - f_1|$. It is further shown that the leakage does not increase when the attackers knows Z^n. In [88] memoryless hypothesis-aware and hypothesis-unaware policies with memory are considered, showing that memoryless hypothesis-aware policies cannot outperform hypothesis-unaware policies. Furthermore, it is shown that there is no loss of optimality in reducing the energy demand alphabet $\mathcal{Y}$ to the energy consumption alphabet $\mathcal{X}$.

2.3.7 Fisher information

In [74] Fisher information is used as a proxy to bound the variance of the estimation error of unbiased estimators of the energy consumption of the user. This is based on the Cramér-Rao bound [85], showing that the Fisher information matrix provides a lower bound on the variance of any unbiased estimator:

$$\mathbb{E}\left\|\hat{\mathbf{x}}(Y) - \mathbf{x}\right\|_2^2 \geq \mathrm{Tr}\left(\mathcal{I}(\mathbf{x})^{-1}\right), \tag{2.85}$$

where $\mathrm{Tr}(\cdot)$ denotes the trace and the Fisher information matrix is given in 2.3. In [74] the energy consumption $\mathbf{x}$ is assumed to be a deterministic vector on $\mathbb{R}^+$.

Interestingly, the battery constraints are defined as constraints on the energy drawn from the battery $b_i = s_i - s_{i+1}$, with the sequence of steps $\mathbf{b}$ in the set

$$\mathcal{B} = \left\{ \mathbf{b} \in \mathbb{R}^T : 0 \leq s_0 + \sum_{k=0}^{i} b_k \leq \beta, \text{ for all } i \right\}. \tag{2.86}$$

Feasible battery policies are then defined as a random transformation $\gamma(\mathbf{b}|\mathbf{x}) = \mathbb{P}_{\mathbf{x}}[B^n = \mathbf{b}]$, where we recall the energy consumption $\mathbf{x}$ is assumed to be deterministic. The first problem considered is the characterization of optimal battery policies when the consumption is not known to the EMU, i.e. $\gamma(\mathbf{b}|\mathbf{x}) = \gamma(\mathbf{b})$. The resulting problem is given by

$$\mathcal{J}(\gamma) = \operatorname{Tr}\left(\mathcal{I}(\mathbf{x})^{-1}\right) \geq n^2 \operatorname{Tr}\left(\mathcal{I}(\mathbf{x})\right). \tag{2.87}$$

Lagrangian multipliers show that the optimal battery policy γ^* is the greedy algorithm, yielding

$$\mathbb{E}\left\|\hat{\mathbf{x}}(Y) - \mathbf{x}\right\|_2^2 \geq n^2 \operatorname{Tr}\left(\mathcal{I}(\mathbf{x})\right) = n^2 \kappa \beta^2, \tag{2.88}$$

with $\kappa = \operatorname{Tr}\left(\mathcal{I}(\mathbf{x})^{-1}\right)$ for $\beta = 1$.

The second problem is the characterization of optimal battery policies when the consumption is known in advance to the EMU, i.e. $\gamma(\mathbf{b}|\mathbf{x})$. The objective function is then

$$\mathcal{J}(\gamma) = \int_{\mathbf{x}\in\mathcal{X}^n} \operatorname{Tr}\left(\mathcal{I}(\mathbf{x})^{-1}\right) f(\mathbf{x}) d\mathbf{x} \geq n^2 \left(\int_{\mathbf{x}\in\mathcal{X}^n} \operatorname{Tr}\left(\mathcal{I}(\mathbf{x})\right) f(\mathbf{x}) d\mathbf{x} \right)^{-1} \tag{2.89}$$

where $f(\mathbf{x})$ denotes the weight associated with $\mathbf{x}$, or the interest to make the estimation of $\mathbf{x}$ hard. Note $f(\mathbf{x})$ can be set without loss of generality so that $\int_{\mathbf{x}\in\mathcal{X}^n} f(\mathbf{x}) d\mathbf{x} = 1$. Lagrangian optimization shows than when the weight is uniform across $\mathbf{x}$, i.e. $f(\mathbf{x}) = f$ for all $\mathbf{x} \in \mathcal{X}^n$, then the solution to this problem is equivalent to that of optimizing when no knowledge of the consumption is available.

Constraint optimization considering market cost and battery's wear and tear is further studied in [74]. Lagrangian optimization of the problem show the optimal solution is given by the solution of a linear partial differential equation. Moreover, charging and discharging rates are studied, the greedy algorithm is showed to no longer be optimal, and the solution is then computed numerically. Finally, numerical characterization of the f-score obtained by state of the art NILM algorithms are presented for different battery sizes under the proposed battery policies.

2.4 Conclusions and main problem formulation

The smart meter privacy problem has been formulated in a variety of ways. Finite state machine models have been proposed to characterize mutual information under i.i.d. and Markovian input processes. This approach yields suboptimal solutions unless all past input and outputs are included in the state, which makes the problem computationally infeasible unless structural simplifications are applied. Alternatively, Markov decision processes yield numerical solutions for quadratic deviation, mutual information and hypothesis testing privacy metrics under Markov constraints on the input process. Furthermore, structural simplifications allow for simple expressions under i.i.d. inputs. Moreover, cost constraints and EHD are easily included into MDP models. The equivalence between rate distortion theory and EMUs with instantaneous power constraints and i.i.d. inputs enables the usage of efficient numerical algorithms (Blahut-Arimoto) and single letter lower bounds (Shannon Lower bound). Under Hypothesis Testing privacy metrics, and EHD with sufficiently large batteries, Chernoff-Stein Lemma allows single letter characterization of the asymptotic exponential decay rate of the Type II probability of error with i.i.d. inputs. Fisher information allows computation of bounds of the minimum variance of any unbiased estimator by assuming deterministic energy consumptions with or without cost constraints.

We propose the utilization of mutual information to measure privacy for two reason. First because its interpretation in terms of an adversary that minimizes log-loss with respect to an evolving soft-decision model [96] is well-matched to the evolving nature of energy distribution over time. Secondly, because mutual information provides a useful bridge to adjacent fields such as hypothesis testing [124], estimation [125], and statistical or machine learning [126].

In a nutshell, the simplicity and tractability of i.i.d. consumption models has captured the main focus in the literature, with some studies focusing on numerical solutions for Markovian energy consumptions. However, solutions for general input processes are only available numerically for quadratic deviation cost functions, and Fisher information, under either deterministic consumptions or learning algorithms (Q-Learning). Furthermore, in privacy and security settings it is typically interesting to characterize the worst-case performance. This interest is grounded on the need to provide guarantees that hold for every user, and is captured in the definition of privacy metrics such as differential privacy [91] or maximal leakage [96].

In the following, we focus on obtaining single letter guarantees that hold for a broad class of consumption processes. Specifically, we focus on bounded input processes, under no stationary or ergodic conditions. The results are also particularise

for processes with known expected average value. Moreover, we characterize the privacy-cost trade off in the presence of variable market prices. The provided results hold for mutual information, maximal leakage and maximal-alpha leakage.

Chapter 3

Universal privacy guarantees via the trapdoor channel

Within this chapter we focus on providing universal privacy guarantees for EMUs with access to a finite capacity battery. Our aim is to provide universal privacy guarantees that hold for a wide class of energy consumption models. In Section 3.3, we provide guarantees that hold for any bounded energy consumption. In Section 3.4, we particularize the guarantees by imposing a mean constraint on the energy consumption. It is important to remark that these bounds, i.e. on bounded input alphabets and on bounded input alphabets with average constraints, hold for any distribution on the consumption. We do not impose any stationary, ergodic or information stability constraints on the input source, i.e. the proposed bounds hold for every scenario in which the mutual information is defined.

The tightness of the upper bounds presented in Section 3.3 and Section 3.4 is demonstrated by constructing energy consumption processes that are tight with respect to the upper bounds for any feasible battery policy implemented by the EMU. In Section 3.5, we present numerical evaluations of the single letter upper and lower bounds derived on the previous sections.

The generality of the upper bounds, holding for any input distribution, is achieved thanks to inspiration borrowed from the trapdoor channel literature, and more generally, from the broader class of permuting channels. Therein, a combinatorial approach is employed to study the set of feasible transformations that the channel can implement, i.e. random permutations of the input. Thus, the key idea is understanding the combinatorial nature of the constraints imposed by battery channels. That is, given an input, a battery-based channel limits the set of feasible outputs, but do not impose constraints on the distribution of the output across the feasible set.

The term *universal* is borrowed from the source coding community, where universal data compression codes are those that achieve near optimal compression performance regardless of the statistics of the input source, e.g. LZ77 [127] and LZ78 [128]. The conceptual similarities between privacy and data compression studied in Section 2.3.5 and Section 5.4.1 further motivate this terminology.

3.1 System model

User energy consumption profiles tend to exhibit non-stationary features. Therefore, it is essential to employ consumption models that capture realistic temporal dynamics. To this end, we model user energy consumption as a discrete-time random process X^n with probability distribution P_{X^n} over the alphabet $\mathcal{X}^n = [\![0, \alpha]\!]^n$, where α denotes the maximum energy consumption of the user during one time step. We focus on integer random variables for presentation purposes but the results generalize to arbitrary discrete time random processes over any discrete alphabet. The random variable X_i describes the energy consumed by a user at time step i with $i \in \{0, 1, ..., n-1\}$. This model accommodates the non-stationary statistics observed in user energy consumption profiles [58, 2] and Figure 3.1.

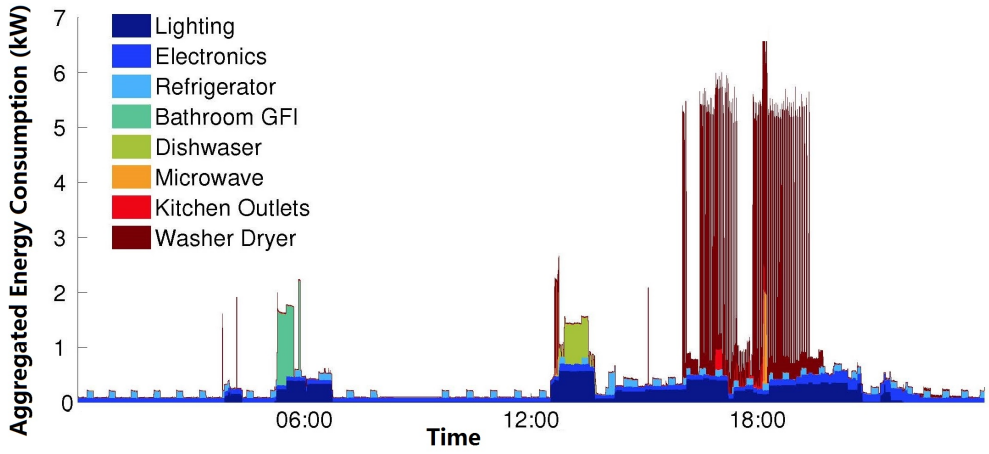

Figure 3.1. An example of energy consumption over the course of a day for one of the houses in REDD [2].

The energy requested from the UP is modelled by a discrete-time random process Y^n where the random variable Y_i describes the energy requested from the UP at time step $i \in \{0, 1, ..., n-1\}$. The energy request alphabet $\mathcal{Y} \subseteq \mathbb{Z}$ is larger than $\mathcal{X}$ since we consider UPs that are able to satisfy the energy consumption of the user even when no battery is available; in addition, $\mathcal{Y}$ contains negative values to capture the possibility that the user sells energy back to the grid.

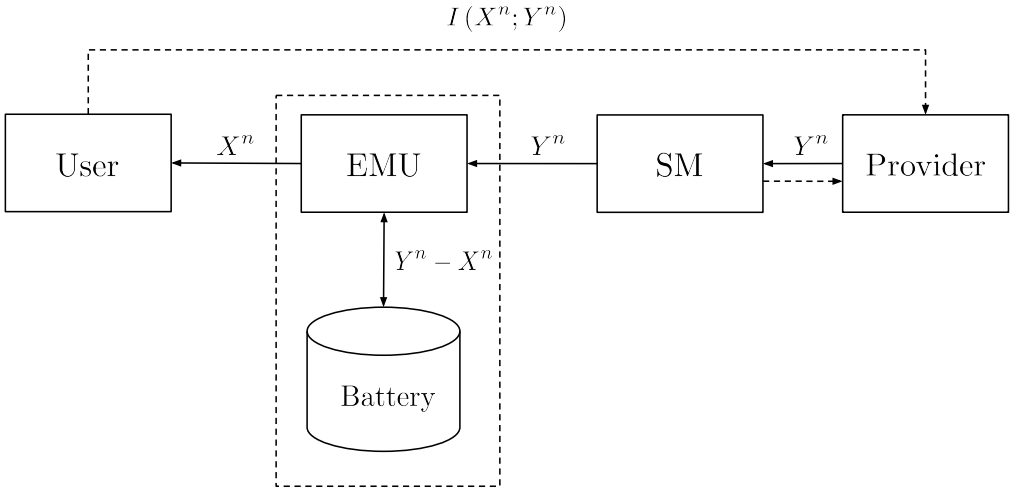

Figure 3.2. Energy Management System with Finite Battery Model

To characterize the level of privacy that the user is guaranteed when the realization of X^n is known to the provider only via Y^n, we use Shannon's mutual information as a measure of privacy. The fundamental properties of mutual information are summarized in Section 2.2.6. To that end, the EMU implements a policy that maps consumption sequences X^n to energy request sequences Y^n. These policies aim to minimize the mutual information between X^n and Y^n while satisfying some operational constraints.

Definition 3.1. *Given a random process X^n modelling the energy consumption of the user and an EMU implementing a battery policy $\pi \in \mathcal{P}_{Y^n|X^n}$, the information leakage is defined as*

$$\mathcal{I}(X^n, \pi) \triangleq \frac{1}{n} I(X^n; Y^n). \tag{3.1}$$

Remark 3.1. *Note that the information leakage is normalized over n in order to preserve consistence with the literature. However, our results hold in the finite block length regime. Thus, our results characterize the total information leakage as well as the asymptotic leakage rate.*

Figure 3.2 depicts the energy management system and the random processes therein. The privacy guarantee is defined in terms of the information leakage from the user to the provider, and the task of the EMU is to choose a battery policy that minimizes the information leakage while satisfying the operational constraints.

The battery can store up to $\beta \in \mathbb{N}$ energy units. Thus, the battery state at time step i, denoted by S_i, takes values in $\mathcal{S} = [\![0, \beta]\!]$ and is governed by the charging

dynamics:

$$S_i = s_0 + \sum_{k=0}^{i-1} Y_k - \sum_{k=0}^{i-1} X_k, \tag{3.2}$$

where $s_0 \in \mathcal{S} = [\![0, \beta]\!]$ is the initial battery state. At this point, we need to clarify the limitations of the battery model that we assume. First, we do not impose any charge or discharge rate limits on the battery, i.e. the battery stores or provides as much energy as the EMU requests at any given time step. Secondly, we do not consider the strain induced on the health of the battery by the charge and discharge cycles. This strain is an important factor in real battery systems and, indeed, practical EMU policies should consider the damage caused by rapid charge and discharge strategies. Within the proposed setting, a power outage occurs when $S_i + Y_i - X_i < 0$; energy is wasted when $S_i + Y_i - X_i > \beta$. We focus on EMUs that do not allow power outages nor energy waste, these ensures the energy demand of the user is always satisfied and only the necessary energy is bought. Definition 3.2 captures the set of energy requests that the EMU can implement.

Definition 3.2. *Given an initial state $s_0 \in \mathcal{S}$, an $n \in \mathbb{N}$, and an energy consumption sequence $\mathbf{x} \in \mathcal{X}^n$ as the inputs of an EMU with battery capacity β, the set of feasible energy requests is*

$$\mathcal{Y}^n(s_0, \mathbf{x}) \triangleq \{\mathbf{y} \in \mathcal{Y}^n : s_i \in [\![0, \beta]\!] \text{ for all } i \in [\![0, n]\!]\}. \tag{3.3}$$

Given the initial state $s_0 \in \mathcal{S}$ and the energy consumption $\mathbf{x} \in \mathcal{X}^n$, the EMU must select an energy request sequence $\mathbf{y}$ that satisfies the feasibility constraint $\mathbf{y} \in \mathcal{Y}^n(s_0, \mathbf{x})$. The following definition describes the selection process.

Definition 3.3. *Given an EMU with battery capacity β and initial state s_0, the set of feasible battery policies over the output alphabet $\mathcal{Y}$ is given by*

$$\Omega(s_0, \beta) \triangleq \{P_{Y^n|X^n} : \operatorname{supp}(P_{Y^n|X^n=\mathbf{x}}) \subseteq \mathcal{Y}^n(s_0, \mathbf{x}) \text{ for all } \mathbf{x} \in \mathcal{X}^n\}. \tag{3.4}$$

Remark 3.2. *It is important to note that the proposed definition for the set of feasible battery policies $\Omega(s_0, \beta)$ contains non-causal battery policies. Implementation of non-causal battery policies rely on precise forecasting capabilities.*

Remark 3.3. *Note the dependency of the set of feasible request sequences $\mathcal{Y}^n(s_0, \mathbf{x})$ and the set of feasible battery policies $\Omega(s_0, \beta)$ on the initial state s_0. However, the notation $P_{Y^n|X^n=\mathbf{x},S_0=s_0}$ is avoided to emphasize that the initial state is considered a known system parameter and not a random variable.*

Remark 3.4. *Note that within this setting, the conditional probability distribution $P(Y^n|X^n)$ models the battery charge/discharge process. The set of charg-*

ing/discharging models that the battery is able to implement given the feasibility constraints imposed in Definition 3.2 is characterized by the set of feasible battery policies $\Omega(s_0, \beta)$.

3.2 Challenges and methodology

In this section, we address the main challenges posed by the proposed system model, in particular we highlight the challenges posed by the memory introduced by the finite capacity battery and the generality of the input process. Therein, we aim to understand the difficulty posed by the setting and limitations of previous approaches. Subsequently, we propose a new approach that sheds light on the problem. This new approach is inspired by the results of Ahlswede and Kaspi [129] on the trapdoor channel. Our proposal hinges on the idea that, under certain conditions, the SMs privacy problem and the trapdoor channel are equivalent.

3.2.1 Challenges introduced by this model

Our objective is to characterize the minimum information leakage that a finite capacity battery can guarantee, i.e. we aim to obtain a bound holding for every input distribution. Minimizing mutual information over a set of random transformations is a hard problem in general that often arises in information theory. The vast majority of the literature is devoted to subadditive distortion constraints and stationary sources, with a strong focus on additive constraints and stationary ergodic sources [130]. The challenges introduced by more general sources and fidelity criteria are portrayed in the lack of a coding theorem proving the operational lossy compression meaning of this minimization. Unlike systems with instantaneous power constraints (Section 2.3.5), battery policies are not defined by additive constraints. Although some work has been done on more general fidelity criteria, such as context-dependent fidelity criteria [131], minimization of the mutual information subject to general, non additive or subadditive constraints is still an open problem [130]. Furthermore, the input processes considered here do not posses the common assumptions of stationarity or ergodicity. This impedes the utilization of many of the tools typically employed in information theory and ergodic theory. This further hinders the tractability of the proposed system model.

In short, the memory introduced by the battery and the generality of the input process reduces the tractability of the problem. For these reasons, the solutions presented so far in the literature focus on simplifications of the above system model. Therein, the battery is removed or considered sufficiently large, thus avoiding the

issues raised by memory; or the input processes are assumed to be i.i.d. or Markovian, allowing structural simplifications and tools requiring ergodic or stationary properties. However, we argue that the need for these assumptions hinges around the probabilistic approaches employed in those studies. Probabilistic models are commonly employed in information theory since input sources, and channel models are governed by probabilistic laws, and in view of the great generality allowed by probabilistic models. However, we note that the constraints imposed on the energy consumption Y^n are constraints on its support and not on the probability distribution, i.e. Y^n is constraint to the set of feasible outputs $\mathcal{Y}(s_0, \mathbf{x})$ with any arbitrary probability, as shown in Definition 3.3. That is, battery policies are determined by combinatorial constraints. While probabilistic models provide great generality, and are able to model battery policies, we argue that a combinatorial analysis can provide the understanding and structural simplifications required to improve the tractability of the problem before a probabilistic model is employed.

3.2.2 Permuting and trapdoor channels

As discussed in Chapter 2, battery policies have been modelled in a variety of ways. The fundamental property governing battery policies is that, up to a constant β, the total amount of energy introduced into the battery equals the total amount of energy provided by the battery. Therein, batteries allow the energy request sequence to be a reordered or permuted version of the consumption sequence, but they do not allow energy to be generated or consumed. It is important to note that batteries do not only perform a permutation of the consumption sequences, e.g.

$$(4, 2, 3) \to (4, 3, 2), \tag{3.5}$$

but they are also able to aggregate symbols, e.g.

$$(4, 2, 3) \to (4, 5, 0). \tag{3.6}$$

Moreover, only those permutations that preserve the internal state of the battery, i.e. the difference between the sum of all previous inputs and outputs, within the operational limits are allowed. This idea, further developed in Section 3.2.3, brings our attention to the the general class of permuting channels [132].

Permuting channels inspired by DNA coding and transmission systems where symbols are inserted, deleted, permuted or substituted are studied from a geometric, combinatorial approach in [133, 134]. Similar asymptotically optimal codes are presented in [135]. Unlike on those system, batteries do not allow insertions or

deletions of symbols, nor they are limited by a maximum number of permutations. In fact, batteries allow permuting, adding or subtracting every input element, as long as the battery constraints are satisfied. In [136] a permuting channel able to permute inputs up to a k-distance and a k-buffer is presented. Therein, two players try to establish a stenographic communication by transmitting ordered packages. This scheme models covert communications via internet protocol (IP) or cache timing [137]. The k-buffer permuter is able to permute the channel inputs based on a buffer of capacity k, the aim of the jammer channel is to minimize the communication rate. Interestingly, in [136], a combinatorial approach with game theoretic tools is employed. This approach is closer to our system model, as no insertions or deletions are allowed.

A more similar problem is presented in the trapdoor channel, first introduced by D. Blackwell [138], and depicted on the front cover of the Dover edition of Robert B. Ash's classical book "Information Theory" [12]. In [138, 12], the trapdoor channel or chemical channel is described as a two-state channel, represented by a box with capacity for two balls. At every time step a new ball comes in (input) and one of the balls inside the box goes out (output). This behaviour is depicted in Figure 3.3. In the original work by Blackwell the balls are marked in a binary alphabet $\mathcal{X} = \mathcal{Y} = \{0, 1\}$ and the box has capacity for 2 balls, i.e. only one ball is left inside the box between time steps. However, interest generated by the channel has lead to generalizations to arbitrary ball alphabets and arbitrary box capacities [129]. Despite the attention gathered by the trapdoor channel and its simplicity, a closed form expression for its capacity is still unknown. The zero-error capacity was proved to be 0.5bits per channel use [139]. The feedback capacity was studied by [140]. Recurrent algorithms have been used to show that the capacity is strictly bigger that 0.5bits [141].

In 1987, Ahlswede and Kaspi [129] proposed two variations of the trapdoor channel: the permuting jammer and the permuting relay [142]. In the relay channel case, the trapdoor aims to communicate information by permuting a fixed sequence of input balls. In the jammer channel case the trapdoor acts as a jammer, aiming to minimize the communication rate by releasing one ball or another. Therein, at every time step a ball numbered $1, 2, ..., \hat{\alpha}$ is introduced into the box and one of the $\hat{\beta} + 1$ balls inside the box is extracted. Within the jammer model, the ball extracted from the box is selected in order to minimize the mutual information between the input and the output. Note that the extraction criteria is not probabilistic and is instead analysed using combinatorial tools. Thus, by following a combinatorial approach, it is shown [129, Proposition 1] that the Shannon capacity C_j of such channels is lower

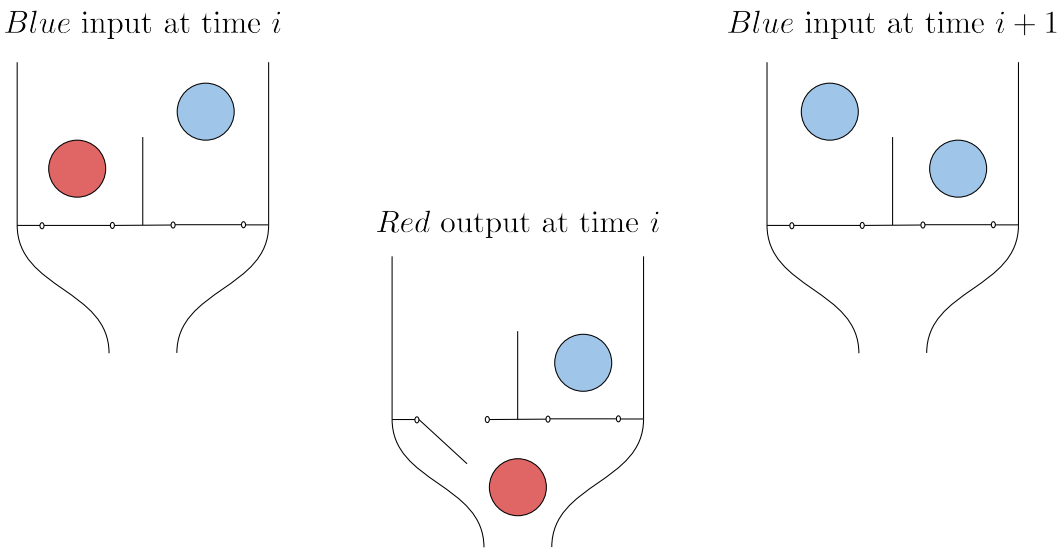

Figure 3.3. Diagram depicting the functioning of a trapdoor channel with $\beta = 1$, as illustrated on the cover of Robert B. Ash's classical book [12]

bounded by

$$C_j \geq \frac{\log \hat{\alpha}}{\hat{\beta}+1}. \tag{3.7}$$

Moreover, when $\hat{\alpha} = 2$ the capacity is upper bounded by

$$C_j \leq \frac{1}{\hat{\beta}+1}. \tag{3.8}$$

Despite the difficulty of characterizing single letter expressions for the capacity of the original model of the trapdoor channel [138], the combinatorial approach followed in [129] allowed for the characterization of single letter bounds holding for general input processes in the case of the jamming trapdoor channel. These results, and the ideas therein, inspire our work on battery policies. In the following, this inspiration is further motivated by showing that, under certain circumstances, the battery and trapdoor channels are equivalent.

3.2.3 Equivalence between trapdoor and battery channel

In [129], Ahlswede and Kaspi define the trapdoor channel as a box containing b_0 *blue* balls and $\beta - b_0$ *red* balls. The operation of the channel is depicted in Figure 3.3. At time i a new ball X_i coloured *blue* or *red* is thrown into the box. Immediately after, one of the $\beta + 1$ balls inside the box is selected and taken out of the box. Let Y_i denote the ball extracted at time i. Following this model, the number of *blue* balls

inside the box at time i is given by

$$b_i = b_0 + \sum_{k=0}^{i-1} bl(X_k) - \sum_{k=0}^{i-1} bl(Y_k), \tag{3.9}$$

where the indicator function $bl(\cdot)$ equals 1 when its argument is coloured *blue* and 0 otherwise. Similarly, when $\mathcal{X} = \mathcal{Y} = \{red,\ blue\}$ the number of *red* balls inside the box at time i is given by $r_i = \beta - b_i$. Replacing $b_i = \beta - r_i$ and $b_0 = \beta - r_0$ into (3.9) yields:

$$r_i = r_0 + \sum_{k=0}^{i-1} bl(Y_k) - \sum_{k=0}^{i-1} bl(X_k). \tag{3.10}$$

For a box of capacity β, the set of feasible output sequences $\mathcal{Y}^n(r_0, \mathbf{x})$ is defined as the set of outputs $\mathbf{y} \in \mathcal{Y}^n$ than can be pulled out of the box given an initial state r_0 and an input sequence $\mathbf{x} \in \mathcal{X}^n$. The following definition captures the set of output sequences that the trapdoor can implement.

Definition 3.4. *Given an initial state $s_0 \in [\![0, \beta]\!]$, an $n \in \mathbb{N}$, and a sequence $\mathbf{x} \in [\![0, 1]\!]^n$ as the inputs of a trapdoor channel with capacity β, the set of feasible output sequences is*

$$\mathcal{Y}_T^n(r_0, \mathbf{x}) = \Big\{\mathbf{y} \in \mathcal{Y}^n : r_i + bl(y_i) - bl(x_i) \in [\![0, \beta]\!] \text{ for all } i\Big\}. \tag{3.11}$$

Following the same principle as in the battery channels, given the initial state $r_0 \in [\![0, \beta]\!]$ and the input sequence $\mathbf{x} \in \mathcal{X}^n$, the trapdoor channel must select an output sequence $\mathbf{y}$ that satisfies the feasibility constraint $\mathbf{y} \in \mathcal{Y}^n(s_0, \mathbf{x})$. This selection process, depicted in Figure 3.3, is captured by the following definition.

Definition 3.5. *Given a trapdoor with capacity β and initial state r_0, the set of feasible trapdoor policies over the output alphabet $\mathcal{Y}$ is given by*

$$\Omega_T(r_0, \beta) \triangleq \{P_{Y^n|X^n} : \text{supp}(P_{Y^n|X^n=\mathbf{x}}) \subseteq \mathcal{Y}_T^n(r_0, \mathbf{x}) \text{ for all } \mathbf{x} \in \mathcal{X}_T^n\}. \tag{3.12}$$

The following theorem shows that, for binary input and output alphabets, the battery and the trapdoor channels are equivalent since requesting energy from the grid corresponds to extracting a ball from the trapdoor channel. Similarly, replacing a ball from the trapdoor channel in (3.10) corresponds to charging the battery of the EMU in (3.2).

Theorem 3.1. *Consider an EMU with a battery of capacity β, initial state $s_0 \in \mathcal{S}$ and binary input and output alphabets i.e. $\mathcal{X}^n = \mathcal{Y}^n = \{0, 1\}^n$. Consider also a trapdoor channel with capacity for β balls, initial state s_0 and balls taking binary*

values, i.e. $\mathcal{X}_T^n = \mathcal{Y}_T^n = \{0,1\}^n$. *Then the two systems are equivalent, i.e. they have the same set of feasible input distributions* P_{X^n}*:*

$$\mathcal{P}_{X^n} = \mathcal{P}_{X_T^n} = \left\{P_{X^n} : \mathbf{x}^n \in \{0,1\}^n\right\}, \tag{3.13}$$

and the same set of feasible battery/trapdoor policies $P_{Y^n|X^n}$*:*

$$\Omega(s_0, \beta) = \Omega_T(s_0, \beta) \tag{3.14}$$

Proof: The proof for the input distributions (3.13) follows immediately by the theorem conditions. The proof for the set of feasible policies (3.14) follows by recalling the definition of feasible battery policies (3.4), i.e.

$$\Omega(s_0, \beta) \stackrel{\Delta}{=} \{P_{Y^n|X^n} : \text{supp}(P_{Y^n|X^n=\mathbf{x}}) \subseteq \mathcal{Y}^n(s_0, \mathbf{x}) \text{ for all } \mathbf{x} \in \mathcal{X}^n\} \tag{3.15}$$

and the definition of the set of feasible trapdoor policies (3.12), i.e.

$$\Omega_T(s_0, \beta) \stackrel{\Delta}{=} \{P_{Y_T^n|X_T^n} : \text{supp}(P_{Y_T^n|X_T^n=\mathbf{x}}) \subseteq \mathcal{Y}_T^n(s_0, \mathbf{x}) \text{ for all } \mathbf{x} \in \mathcal{X}_T^n\} \tag{3.16}$$

Thus, it suffices to show that the set of feasible request sequences $\mathcal{Y}^n(s_0, \mathbf{x})$ and $\mathcal{Y}_T^n(s_0, \mathbf{x})$ are equivalent. The proof is completed by noting that the set of feasible outputs $\mathcal{Y}_T^n(s_0, \mathbf{x})$ in the trapdoor channel (3.11) and the set of feasible energy requests $\mathcal{Y}^n(s_0, \mathbf{x})$ in the battery channel (3.3) are equivalent when both input and output alphabets are binary, i.e. $\mathcal{X}_T^n = \mathcal{X}^n = \mathcal{Y}_T^n = \mathcal{Y}^n = \{0,1\}^n$. ■

Remark 3.5. *It is important to note that the equivalence presented here for the binary case does not hold for other alphabets. This is because in the trapdoor channel [129] the output is constrained to permutations of the input sequence. However, in our setting, the output is not required to contain the same symbols as the input, as the consumed energy can be combined or divided across different time steps. For that reason, the combinatorial approach presented by Ahlswede and Kaspi [129] requires some modifications to properly model the battery channel but the main idea remains. In the following, we present a combinatorial approach, based on the one proposed by Ahlswede and Kaspi, that avoids the difficulties introduced by the probabilistic structures generally employed in the SMs privacy literature.*

This equivalence between the trapdoor and the battery channels immediately provides our first result.

Theorem 3.2. *Consider an EMU with a battery of capacity* β *and initial state* $s_0 \in \mathcal{S}$. *Let* X^n *be a random process taking values in* $\mathcal{X}^n = \mathcal{Y}^n = \{0,1\}^n$, *then there*

exists a battery policy $\hat{\pi} \in \Omega(s_0, \beta)$ *such that*

$$\mathcal{I}(X^n, \hat{\pi}) \leq \frac{1}{\beta + 1}. \tag{3.17}$$

Furthermore, there exist an energy consumption sequence $\hat{X}^n$*, such that for any battery policy* $\hat{\pi} \in \Omega(s_0, \beta)$ *the following holds*

$$\mathcal{I}(\hat{X}^n, \pi) = \frac{1}{\beta + 1}. \tag{3.18}$$

Proof: The proofs follow immediately from (3.7) and (3.8) derived in [129], and the equivalence between the trapdoor and the battery channels for binary alphabets proved in Theorem 3.1. ■

3.3 Privacy with arbitrary energy consumption

In this section, we provide bounds on the information leakage rate when no restrictions are imposed on the probability law of X^n. We first propose the construction of a feasible battery policy $\hat{\pi} \in \Omega(s_0, \beta)$ and characterize an upper bound on the information leakage rate $\mathcal{I}$ induced by $\hat{\pi}$ for any arbitrary random process X^n. Furthermore, we study the tightness of the upper bound by presenting the probability law of a random process $\hat{X}^n$ whose leakage is tight with respect to the upper bound. Moreover, the leakage rate induced by the random process $\hat{X}^n$ is shown to be independent of the employed battery policy $\pi \in \Omega(s_0, \beta)$.

3.3.1 Upper bound on the information leakage rate

We propose a battery policy based on the code construction in [129]. The code proposed in the trapdoor channel context is the counterpart of the battery policy in the smart meter case. For the design of the random transformation that minimized the mutual information, Ahlswede and Kaspi show that, in the binary case, the trapdoor is always able to output a constant sequence of balls with length equal to the trapdoor size. In our case we are interested in larger alphabets, and batteries are not only able to permute, but also to aggregate energy. However, the idea of forcing a uniform, low entropy, behaviour in the output process inspired us in the design of our battery policies. The proposed policy structures the energy request sequences according to the output alphabet defined below.

Definition 3.6. *Consider the set of codewords of length* l *constructed by repetitions of* 0 *or* α *symbols, i.e.* $\mathcal{O}_l = \{(0, 0, \cdots, 0), (\alpha, \alpha, \cdots, \alpha)\} \subseteq \mathcal{X}^l$*. For* $n = lm$*, we*

define the block repetition alphabet as the set $\mathcal{O}_l^m$ of sequences obtained by the m-fold concatenation of codewords of length l, i.e.

$$\mathcal{O}_l^m = \mathcal{O}_l \times \mathcal{O}_l \times ... \times \mathcal{O}_l. \tag{3.19}$$

We now define a battery policy that maps the energy consumption of the user to output sequences constructed with the *block repetition alphabet* $\mathcal{O}_l^m$.

Definition 3.7. *A block battery policy $\hat{\pi} \in \Omega(s_0, \beta)$ is a mapping of the form*

$$\hat{\pi} : \mathcal{S} \times \mathcal{X}^n \to \mathcal{O}_l^m \cap \mathcal{Y}^n(s_0, \mathbf{x}). \tag{3.20}$$

Note that a block battery policy is nothing more than a strategy to assign to each input sequence a feasible energy request sequence constructed with a block repetition alphabet. With these definitions at hand we now provide the following privacy guarantee.

Theorem 3.3. *Consider an EMU with a battery of capacity β and initial state $s_0 \in \mathcal{S}$. Let X^n be a random process with X_i taking values in $\mathcal{X} = \{0, 1, \cdots, \alpha\}$ for $i = 1, 2, \cdots, n$, and let $\hat{\pi}$ denote a block battery policy with $l = \lfloor (\beta + 1)/\alpha \rfloor$ as described in Definition 3.7. Then*

$$\mathcal{I}(X^n, \hat{\pi}) \leq \frac{1}{\lfloor (\beta + 1)/\alpha \rfloor}. \tag{3.21}$$

Proof: Notice that the information leakage rate is upper bounded by

$$\mathcal{I}(X^n, \hat{\pi}) = \frac{1}{n} I(X^n; Y^n) \leq \frac{1}{n} H(Y^n). \tag{3.22}$$

Since Y^n takes values in $\mathcal{O}_l^m$ and $|\mathcal{O}_l^m| = 2^m$ we have that

$$\frac{1}{n} H(Y^n) \leq \frac{1}{n} \log |\mathcal{O}_l^m| = \frac{1}{n} \log(2^m) = \frac{m}{n} = \frac{1}{l}. \tag{3.23}$$

We now show that when $l \leq (\beta + 1)/\alpha$ there exists at least one block battery policy $\hat{\pi}$ that is feasible for every initial state s_0 and consumption $\mathbf{x}$. To prove this we show that for every realization $\mathbf{x}$ and initial state s_0 there exists an energy request sequence $\mathbf{y} \in \mathcal{O}_l^m$ such that $\mathbf{y}$ belongs to the set of feasible energy requests $\mathcal{Y}^n(s_0, \mathbf{x})$. The strategy is to notice that $\mathcal{O}_l^m \cap \mathcal{Y}^n(s_0, \mathbf{x}) \neq \emptyset$ for $m = 1$ and to then prove by induction that the non-emptiness of the intersection holds for $m \geq 1$.

For $m = 1$, the intersection $\{(0, 0, \cdots, 0), (\alpha, \alpha, \cdots, \alpha)\} \cap \mathcal{Y}_\beta^l(s_0, \mathbf{x}^l)$ is non-empty if and only if either the sequence $(0, 0, \cdots, 0)$ or $(\alpha, \alpha, \cdots, \alpha)$ belong to $\mathcal{Y}^n(s_0, \mathbf{x})$.

This implies that either

$$s_{i+1} = s_i + y_i - x_i = s_i + 0 - x_i \in \mathcal{S} = [\![0, \beta]\!] \tag{3.24}$$

or

$$s_{i+1} = s_i + y_i - x_i = s_i + \alpha - x_i \in \mathcal{S} = [\![0, \beta]\!] \tag{3.25}$$

hold for all $i \leq l$. In the first case, described in (3.24), we have that $x_i \geq 0$ for $i = 0, \cdots, l-1$. Hence, the energy stored in the battery, s_i, decreases monotonically. Therefore, since $s_0 \leq \beta$, all s_i belong to $\mathcal{S} = [\![0, \beta]\!]$ when $s_i \geq 0$ on the last time step, i.e.,

$$0 \leq s_0 - \sum_{i=0}^{l-1} x_i. \tag{3.26}$$

Similarly, in the case described by (3.25), we have that $x_i \leq \alpha$ and the energy stored increases monotonically. It is then sufficient to show that

$$s_0 - \sum_{i=0}^{l-1} x_i \leq \beta - \alpha l. \tag{3.27}$$

When $\beta - \alpha l \geq -1$ every integer $s_l \in [\![s_0 - l\alpha, s_0]\!]$ satisfies at least one of the inequalities given by (3.26) and (3.27). This ensures that either (3.26) or (3.27) hold for every $s_0 \in \mathcal{S}$ and $\mathbf{x} \in \mathcal{X}^l$, and therefore, the intersection $\mathcal{O}_l^m \cap \mathcal{Y}^n(s_0, \mathbf{x})$ is non-empty. This completes the proof for $m = 1$. The induction for $m \geq 1$ is straightforward as the proof for $m = 1$ holds for every initial state s_0. ■

Remark 3.6. *Theorem 3.3 provides an upper bound on the leakage induced by any energy consumption process taking values in a bounded alphabet $[\![0, \alpha]\!]$. This implies that when the EMU implements a block battery policy, no bounded energy consumption process leaks more than $\frac{1}{\lfloor(\beta+1)/\alpha\rfloor}$ bits per sample. Thus, Theorem 3.3 provides a worst case guarantee for any user with a bounded energy consumption.*

Note that in order to map input pairs $(s_0, \mathbf{x})$ to energy requests in $\mathcal{O}_l^m$ it suffices to forecast, at the start of each block of length $\lambda = (\beta + 1)/\alpha$, whether the battery will deplete during the current block, i.e. $s_0 - \sigma(\mathbf{x}^\lambda) \lneqq 0$. This shows the forecasting capabilities required to implement the battery policy described in Theorem 3.3.

3.3.2 Tightness of the upper bound

We now address the tightness of the upper bound presented in Theorem 3.3. To this end, we construct a random process modelling the energy consumption of the

user that is tight with respect to the result in Theorem 3.3 for every battery policy $\pi \in \Omega(s_0, \beta)$. The worst case input process presented by Ahlswede and Kaspi hinges on the same idea as the optimal battery policy, i.e. uniformity. The worst case input process is thus constructed by consecutively introducing a fixed number of balls of the same type into the trapdoor. When this number exceeds the capacity of the trapdoor, the trapdoor is forced to leak the type of ball that was introduced. Although the constraints imposed on batteries polices and on the trapdoor channel are not equivalent outside the binary case, the worst case input processes derived for the trapdoor channel provided inspiration for the design of our our worst case consumption processes.

Theorem 3.4. *Consider an EMU with a battery of capacity β and initial state s_0. Let $\hat{X}^n$ be a random process taking uniformly distributed values in $\mathcal{O}_l^m$ with $l = \lceil(\beta+1)/\alpha\rceil$. Let π be a feasible battery policy. Then*

$$\mathcal{I}(\hat{X}^n, \pi) = \frac{1}{\lceil(\beta+1)/\alpha\rceil}. \tag{3.28}$$

Proof: Expand $\mathcal{I}(\hat{X}^n, \pi)$ as

$$\mathcal{I}(\hat{X}^n, \pi) = \frac{1}{n} I(X^n; Y^n) = \frac{1}{n} H(X^n) - \frac{1}{n} H(X^n|Y^n). \tag{3.29}$$

When X^n is uniformly distributed over the alphabet $\mathcal{O}_l^m$, with $|\mathcal{O}_l^m| = 2^m$ and $n = ml$, we have that

$$\frac{1}{n} H(X^n) = \frac{1}{n} \log |\mathcal{O}_l^m| = \frac{m}{n} = \frac{1}{l}. \tag{3.30}$$

We now show that the equivocation rate $\frac{1}{n} H(X^n|Y^n)$ is 0 when X^n takes values in $\mathcal{O}_l^m$ with $l > \beta/\alpha$. To this aim, we prove by induction that when the input sequence $\mathbf{x}$ belongs to $\mathcal{O}_l^m$ with $l > \beta/\alpha$, the sets $\mathcal{Y}^n(s_0, \mathbf{x})$ of feasible output words generated by different consumption sequences are disjoint, i.e.

$$\mathcal{Y}^n(s_0, \mathbf{x}') \cap \mathcal{Y}^n(s_0, \mathbf{x}) = \emptyset \text{ for } \mathbf{x}' \neq \mathbf{x}. \tag{3.31}$$

As a result, any request sequence $\mathbf{y} \in \mathcal{Y}^n(s_0, \mathbf{x})$ unequivocally determines the generating input $\mathbf{x}$. In other words, given an output sequence $\mathbf{y}$ there is no uncertainty about the input $\mathbf{x}$, and therefore, the equivocation rate $\frac{1}{n} H(X^n|Y^n)$ is 0.

For $m = 1$ there are two possible inputs $(0, 0, \cdots, 0)$ and $(\alpha, \alpha, \cdots, \alpha)$. When $\mathbf{x} = (0, 0, \cdots, 0) \in \mathcal{O}_l^1$ the energy stored in the battery at time l is given by

$$s_l(0) = s_0 + \sum_{i=0}^{l-1} (y_i - 0). \tag{3.32}$$

Similarly, when $\mathbf{x} = (\alpha, \alpha, \cdots, \alpha) \in \mathcal{X}^l$ the energy stored in the battery at time l is given by

$$s_l(\alpha) = s_0 + \sum_{i=0}^{l-1}(y_i - \alpha). \tag{3.33}$$

Taking the difference between (3.32) and (3.33) yields:

$$s_l(0) - s_l(\alpha) = \sum_{i=0}^{l-1}\alpha = l\alpha. \tag{3.34}$$

When $s_l(\alpha) \in \mathcal{S} = [\![0, \beta]\!]$ we have that $s_l(0) = s_l(\alpha) + l\alpha \geq l\alpha$, showing that for $l\alpha > \beta$ the events $s_l(\alpha) \in \mathcal{S} = [\![0, \beta]\!]$ and $s_l(0) \in \mathcal{S} = [\![0, \beta]\!]$ do not occur simultaneously. This implies that the set of output sequences belonging to $\mathcal{Y}^n(s_0, (0, 0, \cdots, 0))$ and $\mathcal{Y}^n(s_0, (\alpha, \alpha, \cdots, \alpha))$ is empty for every initial state s_0. Therefore the sets are disjoint and $\frac{1}{n}H(X^n|Y^n) = 0$. The proof for $m > 1$ follows by induction and noticing that the proof above is valid for every initial state s_0. ■

Remark 3.7. *Theorem 3.4 shows that at least one consumption profile leaks* $\frac{1}{\lceil(\beta+1)/\alpha\rceil}$ *bits per sample, for any feasible battery policy. This implies that no upper bound with the generality of the one presented in Theorem 3.3 can guarantee a leakage smaller than* $\frac{1}{\lceil(\beta+1)/\alpha\rceil}$ *bits per sample.*

Note that the gap between both bounds is given by

$$G = \frac{1}{\lceil(\beta+1)/\alpha\rceil} - \frac{1}{\lfloor(\beta+1)/\alpha\rfloor} = \begin{cases} 0, & \text{when } (\beta+1)/\alpha \in \mathbb{Z} \\ \frac{1}{\lceil\frac{\beta+1}{\alpha}\rceil\lfloor\frac{\beta+1}{\alpha}\rfloor}, & \text{otherwise.} \end{cases} \tag{3.35}$$

Thus, the bounds are exact for integer values of $(\beta+1)/\alpha$, and their difference increase roughly with the square of $\lambda = (\beta+1)/\alpha$ otherwise. Note that this gap is a peculiarity introduced by the discrete time nature of the model. When the time needed to fully discharge the battery is an integer number of time steps, i.e. $(\beta+1)/\alpha \in \mathbb{Z}$ upper and lower bounds coincide. Otherwise, the feasibility of the upper bound is only guaranteed during $\lfloor(\beta+1)/\alpha\rfloor$ time steps, while the consumption sequence requires $\lceil(\beta+1)/\alpha\rceil$ time steps to leak one bit of information. It is important to note that the maximum energy consumption α is a function of the sampling interval, i.e. $\alpha = P_{\max}T$, where $P_{\max}$ is the peak power consumption of the user and T is the sampling interval. Thus, by selecting an adequate sampling interval, i.e. T such that $(\beta+1)/(P_{\max}T) \in \mathbb{Z}$, we can always make the bounds exact for arbitrary values of β and $P_{\max}$.

3.4 Privacy with an average energy constraint

The information leakage rate bounds provided in Section 3.3 do not impose any moment restriction on the random process modelling the energy consumption of the user. Indeed, they depend only on the peak energy consumption α and on the size of the battery β. However, one of the most widely used energy consumption metrics is the average energy consumption over an arbitrary time interval, being common for SMs to display this information to the user [29, 66]. In the following, we particularize the results in Theorem 3.3 and Theorem 3.4 to the case in which the average energy consumption of the user is specified. Therein, we analyse the impact of the average energy consumption on the privacy performance. We define the average energy consumption of the random process X^n as

$$\mu_n = \mathbb{E}\left[\frac{1}{n}\sum_{i=0}^{n-1} X_i\right]. \tag{3.36}$$

Note that since we do not impose any stationarity constraint on the random process X^n, the average energy consumption is a function of the time step n. This agrees with the non-stationary nature observed in energy consumption profiles of residential users [58].

Following the results on Section 3.3, the upper bound presented in this section relies on the availability of energy consumption forecasts up to $\lambda = (\beta+1)/\alpha$ time steps ahead. Moreover, the upper bound provides a worst case guarantee, i.e, it upper bounds the amount of information that any bounded consumption profile with a given average consumption leaks to the provider. Finally, the lower bound presented in this section shows the existence of a consumption profile tight to the upper bound. This implies that, for the same generality, there exist no significantly lower upper bound than the one here presented.

3.4.1 Upper bound on the information leakage rate

The following result provides an upper bound on the information leakage rate for random processes X^n with average energy consumption μ_n.

Theorem 3.5. *Consider a battery system with capacity β and initial state s_0. Let X^n be a random process taking values on $\mathcal{X}^n = [\![0,\alpha]\!]^n$ with average energy consumption μ_n. Let $\hat{\pi}$ be a block battery policy with $l = \lfloor(\beta+1)/\alpha\rfloor$, then*

$$\mathcal{I}(X^n,\hat{\pi}) \leq \frac{\max_{\epsilon\in[\frac{-s_0}{n\alpha},\frac{\beta-s_0}{n\alpha}]} H_2\left(\frac{\mu_n}{\alpha}+\epsilon\right)}{\lfloor(\beta+1)/\alpha\rfloor}, \tag{3.37}$$

where $H_2(p) = -p\log_2 p - (1-p)\log_2(1-p)$ *denotes the binary entropy.*

Proof: Let the output process Y^n take values in $\mathcal{O}_l^m$. Thus, the information leakage rate is upper bounded by

$$\begin{aligned}\mathcal{I}(X^n, \hat{\pi}) &\leq \frac{1}{n} H(Y^n) && (3.38)\\ &= \frac{1}{n}\sum_{i=0}^{m-1} H(Y_{il}, \ldots, Y_{(i+1)l-1} | Y_0, \ldots, Y_{il-1}) && (3.39)\\ &\leq \frac{1}{n}\sum_{i=0}^{m-1} H(Y_{il}, \ldots, Y_{(i+1)l-1}), && (3.40)\end{aligned}$$

where (3.39) follows by applying the chain rule and (3.40) follows from the fact that conditioning reduces entropy. Notice that (3.40) is the entropy of m sequences Y^l taking values on the binary alphabet $\mathcal{O}_l = \{(0, 0, \cdots, 0), (\alpha, \alpha, \cdots, \alpha)\}$. Therefore, the information leakage is upper bounded by

$$\begin{aligned}\mathcal{I}(X^n, \hat{\pi}) &\leq \frac{1}{n}\sum_{i=0}^{m-1} H(Y_{il}, \ldots, Y_{(i+1)l-1}) && (3.41)\\ &\leq \frac{1}{n}\sum_{i=0}^{m-1} H_2\Big(\mathbb{P}\left[Y^l = (\alpha, \alpha, \ldots, \alpha)\right]\Big) && (3.42)\\ &= \frac{1}{l} H_2\left(\frac{\mathbb{E}\left[\frac{1}{n}\sum_{i=0}^{n-1} Y_i\right]}{\alpha}\right), && (3.43)\end{aligned}$$

with equality for the case in which each sequence Y^l is independent and identically distributed, i.e. with distribution

$$\mathbb{P}\left[Y^l = (\alpha, \alpha, \ldots, \alpha)\right] = \frac{\mathbb{E}\left[\frac{1}{n}\sum_{i=0}^{n-1} Y_i\right]}{\alpha}, \tag{3.44}$$

and

$$\mathbb{P}\left[Y^l = (0, 0, \ldots, 0)\right] = 1 - \mathbb{P}\left[Y^l = (\alpha, \alpha, \ldots, \alpha)\right]. \tag{3.45}$$

We now bound the average energy requested from the grid as a function of the average energy consumption of the user and the battery size. Dividing (3.2) by n and taking the expected value w.r.t P_{X^n, Y^n} yields

$$\mathbb{E}\left[\frac{1}{n}\sum_{i=0}^{n-1} Y_i\right] = \mathbb{E}\left[\frac{1}{n}\sum_{i=0}^{n-1} X_i\right] + \mathbb{E}\left[\frac{S_n - s_0}{n}\right], \tag{3.46}$$

or equivalently, noting that S_n takes values in $[\![0, \beta]\!]$ and recalling (3.36),

$$\mu_n - \frac{s_0}{n} \leq \mathbb{E}\left[\frac{1}{n}\sum_{i=0}^{n-1} Y_i\right] \leq \mu_n + \frac{\beta - s_0}{n}. \tag{3.47}$$

This implies that

$$\mathbb{P}\left[Y^l = (\alpha, \alpha, \ldots, \alpha)\right] = \frac{\mathbb{E}\left[\frac{1}{n}\sum_{i=0}^{n-1} Y_i\right]}{\alpha} = \frac{\mu_n}{\alpha} - \epsilon \tag{3.48}$$

for some $\epsilon \in [\frac{-s_0}{n\alpha}, \frac{\beta - s_0}{n\alpha}]$. Recall that for $l \leq (\beta+1)/\alpha$, and for every initial state $s_0 \in \mathcal{S}$ and input realization $\mathbf{x} \in \mathcal{X}^n$ there exists a sequence $\mathbf{y} \in \mathcal{O}_l^m$ such that $\mathbf{y}$ belongs to the set of feasible energy requests $\mathcal{Y}^n(s_0, \mathbf{x})$. Thus, the above strategy is always feasible, completing the proof. ■

3.4.2 Tightness of the upper bound

Proceeding in a similar fashion as in Section 3.3, we now prove that the upper bound in Theorem 3.5 is tight for a certain class of random processes modelling the energy consumption.

Theorem 3.6. *Consider a battery system with capacity β and initial state s_0. Let $\hat{X}^n$ be a random process with average energy consumption μ_n and taking values in $\mathcal{O}_l^m$ with $l = \lceil(\beta+1)/\alpha\rceil$. Let π be a feasible battery policy, then*

$$\mathcal{I}(\hat{X}^n, \pi) = \frac{H_2\left(\frac{\mu_n}{\alpha}\right)}{\lceil(\beta+1)/\alpha\rceil}. \tag{3.49}$$

Proof: It follows from (3.43) that the entropy rate of the random process X^n taking values in $\mathcal{O}_l^m$ is upper bounded by

$$\frac{1}{n}H(X^n) \leq \frac{1}{l}H_2\left(\frac{\mathbb{E}\left[\frac{1}{n}\sum_{i=0}^{n-1} X_i\right]}{\alpha}\right) = \frac{1}{l}H_2\left(\frac{\mu_n}{\alpha}\right), \tag{3.50}$$

with equality when the m symbols X^l forming X^n are i.i.d.. We now recall that when X^n takes values in $\mathcal{O}_l^m$ with $l > \beta/\alpha$ the input $\mathbf{x}$ can be uniquely determined from the output sequence $\mathbf{y}$ and $H(X^n|Y^n) = 0$. Thus, by selecting $l = \lceil(\beta+1)/\alpha\rceil$, we have that

$$\mathcal{I}(\hat{X}^n, \pi) = \frac{1}{n}H(X^n) - \frac{1}{n}H(Y^n|X^n) = \frac{1}{n}H(X^n) \leq \frac{H_2\left(\frac{\mu_n}{\alpha}\right)}{\lceil(\beta+1)/\alpha\rceil}. \tag{3.51}$$

with equality when when the m symbols X^l forming X^n are i.i.d.. This concludes de proof. ■

3.5 Numerical results

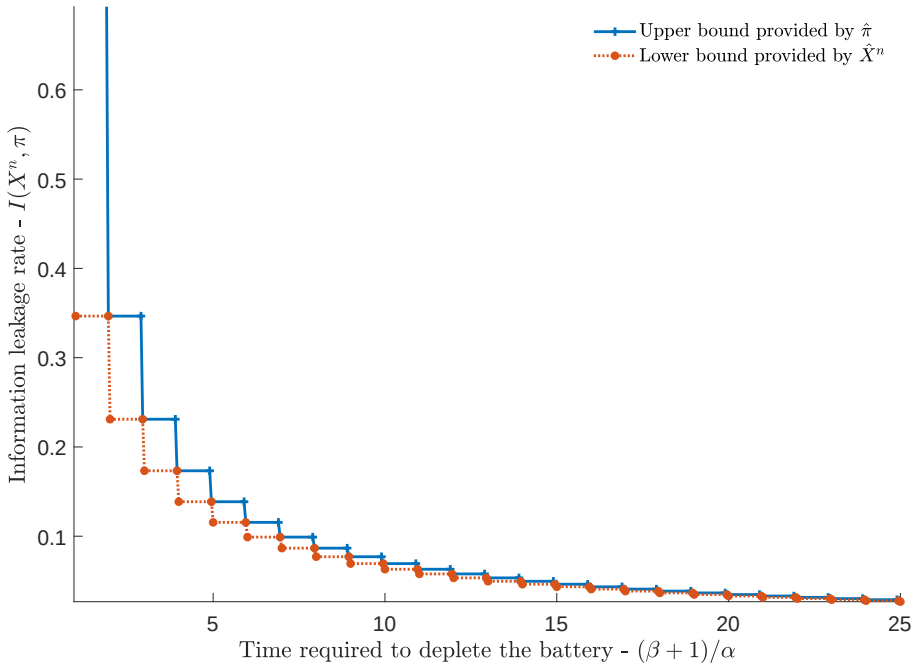

Figure 3.4. Upper bound on the information leakage rate of an EMU as a function of the ratio between the battery size and the peak power consumption.

In this section we numerically evaluate the bounds derived in the previous sections. The upper bound derived in Theorem 3.3 is depicted for different battery sizes in Figure 3.4. It is interesting to note that the privacy guarantees increase significantly for small values of $(\beta + 1)/\alpha$ but the benefit vanishes as the size of the battery increases.

The upper bound on the information leakage rate of infinitely large sequences when the average energy consumption of the user is known is illustrated in Figure 3.5 and Figure 3.6. As expected, the binary entropy term in Theorem 3.5 introduces concavity in the upper bound as shown in Figure 3.5. Interestingly, Figure 3.6 shows that the information leakage rate reduction as the size of the battery increases is less significant for extreme values of the average energy consumption. For representation purposes we evaluate the upper bounds when n goes to ∞, and defined the infinity length average energy consumption $\mu = \lim_{n\to\infty} \mu_n$ which we assume to exits.

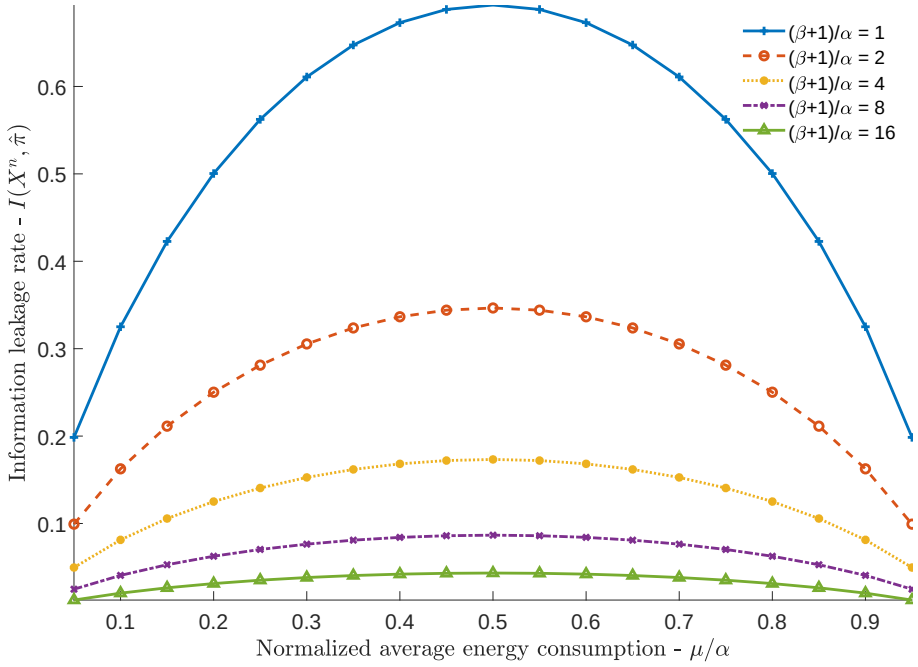

Figure 3.5. Upper bound on the information leakage rate of an EMU with infinitely large sequences as a function of the average energy consumption of the user for different values of the ratio between the battery size and the peak power consumption.

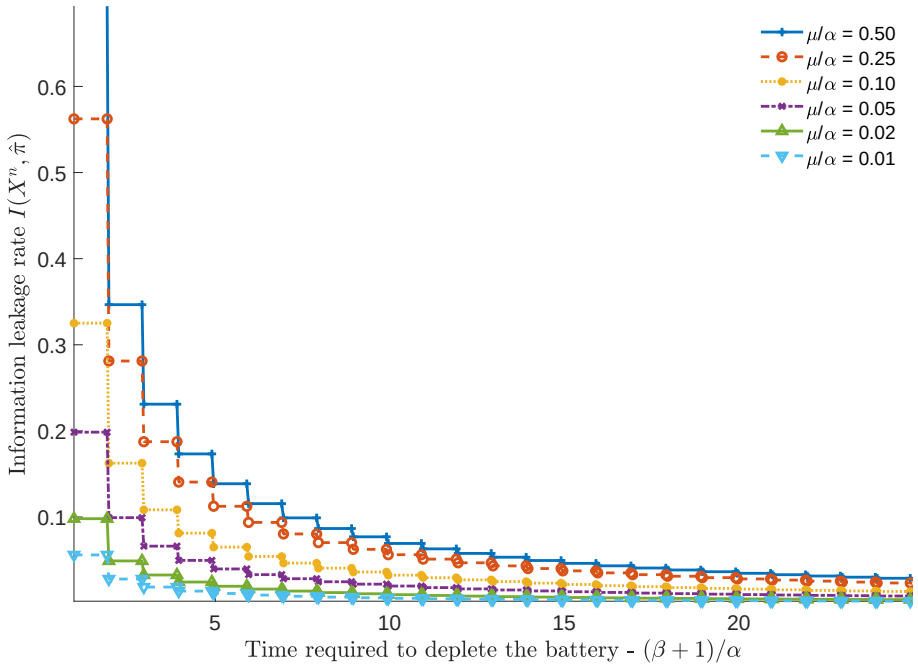

Figure 3.6. Upper bound on the information leakage rate of an EMU with infinitely large sequences, as a function of the ratio between the battery size and the peak power consumption for different values of the average energy consumption of the user

3.6 Conclusion

In this chapter we have studied the information leakage rate of EMUs with finite battery capacity for general random processes modelling the energy consumption of the user. Our results are inspired by previous works on permuting channels, and in particular by the jamming case of the trapdoor channel. Following the coding strategies proposed for the trapdoor channel, where the output is kept constant for a fixed number of time steps, we defined the family of block battery policies. These policies map energy consumption sequences to a block repetition alphabet, constant in blocks of length $\lambda = (\beta + 1)/\alpha$. We have shown that these policies are feasible for arbitrary random processes taking values on a bounded alphabet, thus, providing an upper bound on the information leakage holding for a wide class of input processes. While the resulting policy is non-causal, detailed analysis shows that knowing $(\beta + 1)/\alpha$ time steps ahead suffices to achieve optimality. Thus, we envision practical implementations that rely on consumption forecasting. This approach also provides insight on what forecasting capabilities are needed.

Furthermore, we have make specific the analysis to the case in which the average energy consumption of the user is known and we have concluded that extreme values of the average energy consumption provide lower values of information leakage to the utility provider. The tightness of the upper bound was studied by presenting the probability law of a consumption process tight to the upper bound. The construction of this consumption process, inspired by the work on the trapdoor channel, shows that energy consumption process that consume either no energy or maximum power in cycles of length λ leak the maximum amount of information to the provider. This further emphasizes the key role played by the parameter $\lambda = (\beta + 1)/\alpha$, i.e. by the time needed to deplete a fully charged battery.

Chapter 4

Universal privacy guarantees under cost constraints

Variable market prices play a fundamental role in the smart grid, enabling new approaches to match energy generation and demand [27, 28]. However, variable market prices pose an extra constraint on privacy preserving battery policies, since the user is encouraged to charge and discharge the battery at particular times of the day in order to reduce the energy bill. Thus, it is important to understand how variable market prices impact the privacy guarantees studied in Chapter 3.

Within this chapter, we study the impact of variable market prices on the privacy guarantees. To that aim we model the market prices as a deterministic sequence known in advance by the user. This is in line with most modern billing systems for private consumers, where the different price zones and times are defined in advance. In Section 4.2, we address the main challenges introduced by the joint privacy-cost optimization. Therein we formulate the price constraint as a constraint on the battery state at market transition points. In Section 4.3, we derive upper and lower bounds on the privacy guarantee subject to feasibility and cost constraints. Finally, in Section 4.4, we numerically evaluate the presented bounds.

4.1 System model

Within this chapter we follow the energy management system depicted in Figure 4.1. The privacy guarantee is defined in terms of the information leakage from the user to the provider, and the task of the EMU is to choose a battery policy that minimizes the leakage while satisfying the operation and cost constraints.

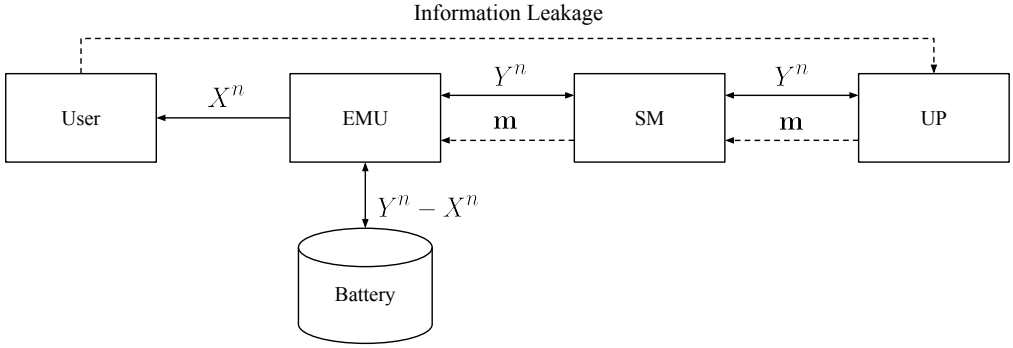

Figure 4.1. Energy Management System with finite battery model and variable market price $\mathbf{m}$.

Following with the set-up in Chapter 2, we model user energy consumption as a discrete-time random process X^n on alphabet $\mathcal{X}^n = [\![0, \alpha]\!]^n$, with probability $P_{X^n} \in \mathcal{P}_{X^n}$. Here $\mathcal{P}_{X^n}$ is a fixed family probability distributions that may contain non-stationary random processes. The EMU maps the consumption sequence $X^n \in \mathcal{X}^n$ to a request sequence $Y^n \in \mathcal{Y}^n$ using a battery policy $P_{Y^n|X^n}$. To be considered feasible, battery policy $P_{Y^n|X^n}$ must create a request sequence that meets the energy demands of the user and does not request energy it cannot use or store. Formal definitions of the sets of feasible energy request and feasible battery policies are given in Definition 3.2 and 3.3.

Our aim in feasible policy design is to minimize privacy subject to a constraint on policy cost. Towards this end, we define our measures of information leakage (where privacy is high when information leakage is low) and cost. We measure a battery policy's information leakage by its worst-case performance.

Definition 4.1. *The information leakage of policy $P_{Y^n|X^n}$ is*

$$\bar{\mathcal{I}}(P_{Y^n|X^n}) = \max_{P_{X^n} \in \mathcal{P}_{X^n}} \frac{1}{n} I(X^n; Y^n). \tag{4.1}$$

We measure the cost of a policy $P_{Y^n|X^n}$ as the difference between the user's energy bill under that policy and the user's energy bill under the feasible battery policy that minimizes the energy bill. Under this definition, cost can be negative only for infeasible policies. To calculate energy bills, we model the energy market price as a deterministic sequence, $\mathbf{m} \in \mathbb{R}^n$. Under this definition, the cost of an energy request sequence $\mathbf{y}$ is $\mathbf{m}^T\mathbf{y}$. We assume that the market price is constant over each of K blocks of time. The price and duration of the k-th block, $k = 0, 1, \ldots, K-1$, are

m_k and l_k, respectively, giving

$$\mathbf{m} = (\underbrace{m_0, \ldots, m_0}_{l_0}, \underbrace{m_1, \ldots, m_1}_{l_1}, \ldots, \underbrace{m_{K-1}, \ldots, m_{K-1}}_{l_{K-1}}). \tag{4.2}$$

Definition 4.2. *Consider an EMU with battery capacity β, initial state $s_0 \in \mathcal{S}$, and market price $\mathbf{m}$. The system cost of energy consumption sequence $\mathbf{x} \in \mathcal{X}^n$ under battery policy $P_{Y^n|X^n}$ is*

$$g(Y^n, \mathbf{x}) = \mathbb{E}_{P_{Y^n|X^n=\mathbf{x}}}[\mathbf{m}^T Y^n - \mathbf{m}^T \mathbf{y}^*(\mathbf{x})], \tag{4.3}$$

where $\mathbf{y}^*(\mathbf{x}) = \operatorname{argmin}_{\mathbf{y}\in\mathcal{Y}^n(s_0,\mathbf{x})} \mathbf{m}^T\mathbf{y}$ is the sequence that induces the minimum feasible cost. For any $\Delta \geq 0$, the set of feasible Δ-affordable battery policies is

$$\Gamma(\Delta) \triangleq \left\{P_{Y^n|X^n} \in \Omega(s_0, \beta) : g(Y^n, \mathbf{x}) \leq \Delta \ \text{ for all } \mathbf{x} \in \mathcal{X}^n\right\}. \tag{4.4}$$

Remark 4.1. *Note the implicit dependency of the systems cost function $g(Y^n, \mathbf{x})$ and of the set of feasible battery policies $\Gamma(\Delta)$ on the output alphabet $\mathcal{Y}$, the initial battery state s_0 and the capacity of the battery β. However, here the simplified notation $g(Y^n, \mathbf{x})$ and $\Gamma(\Delta)$ is preferred to $g(Y^n, \mathbf{x}, \mathcal{Y}, s_0, \beta)$ and $\Gamma(\Delta, \mathcal{Y}, s_0, \beta)$ for simplicity, and because the following focuses on the influence of the budget constraint Δ while the influence of $\mathcal{Y}, s_0$ and β is studied in Chapter 3 and Section 4.2.2.*

Finally, the privacy-cost function defines the optimal tradeoff between privacy and cost achievable by feasible battery policies.

Definition 4.3. *Given an EMU with battery capacity β, initial state s_0 and market price $\mathbf{m}$, the privacy cost function is defined, for each $\Delta \geq 0$, as*

$$\mathcal{I}(\Delta) \triangleq \min_{P_{Y^n|X^n} \in \Gamma(\Delta)} \bar{\mathcal{I}}(P_{Y^n|X^n}). \tag{4.5}$$

The following lemma shows that the privacy cost function $\mathcal{I}(\Delta)$ is invariant to the order of optimization.

Lemma 4.1. *The privacy cost function $\mathcal{I}(\Delta)$ is invariant with respect to the optimization order, i.e.*

$$\mathcal{I}(\Delta) = \min_{\Gamma(\Delta)} \max_{\mathcal{P}_{X^n}} \frac{1}{n} I(X^n; Y^n) = \max_{\mathcal{P}_{X^n}} \min_{\Gamma(\Delta)} \frac{1}{n} I(X^n; Y^n). \tag{4.6}$$

Proof: Let $\pi_1, \pi_2 \in \Gamma(\Delta)$ be Δ-affordable battery policies. Note that any linear combination $\pi_3 = \lambda\pi_1 + (1-\lambda)\pi_2$, with $\lambda \in [0,1]$, satisfies the cost constraint:

$$g(Y_3^n, \mathbf{x}) = \mathbf{m}^T \mathbb{E}[\lambda Y_1^n + (1-\lambda) Y_2^n] - \mathbf{m}^T \mathbf{y}^*(\mathbf{x}) \leq \Delta, \tag{4.7}$$

and the support constraint:

$$\text{supp}(Y_3^n) = \text{supp}(Y_1^n) \cup \text{supp}(Y_2^n) \subseteq \mathcal{Y}^n(s_0, \mathbf{x}). \tag{4.8}$$

Thus $\Gamma(\Delta)$ is convex. The minimax theorem [143], shows that

$$\max_{a\in\mathcal{A}} \min_{b\in\mathcal{B}} f(a,b) = \min_{b\in\mathcal{B}} \max_{a\in\mathcal{A}} f(a,b) \tag{4.9}$$

when $\mathcal{A}$, $\mathcal{B}$ are compact convex sets and $f : \mathcal{A} \times \mathcal{B} \to \mathbb{R}$ is a continuous function concave for fixed b and convex for fixed a. Thus, by concave-convex properties of the mutual information, and the compactness and convexity of the sets $\mathcal{P}_{X^n}$ and $\Gamma(\Delta)$ the proof is completed. ■

4.2 Challenges and methodology

We now focus on the main challenges introduced by variable market prices. Specifically, we notice the utility of modelling cost constraints as a constraint on the battery state at market transition points. These characterizations allow us to combine feasibility and cost constraints into a single, more tractable, optimization problem. Therein, at every time step, the battery state is in $[\![0, \beta]\!]$ as a result of the feasibility constraints, while at market transition points, the battery state is further constraint to a smaller interval by the budget constraints. Two main approaches to tackle this problem are then proposed, the first approach relies on solving the optimization problem one market block at a time, while the second one considers the whole time interval as a single block.

4.2.1 Analysis of market price constraints

In Chapter 3, feasible battery policies that guarantee that the battery state belongs to the interval $[\![0, \beta]\!]$ were presented. Thus, in order to improve the tractability of the problem, it is useful to present the cost constraints as a constraint on the set of battery states. Recall that the the energy bill is given by

$$B = \mathbf{m}^T Y^n = \sum_{k=0}^{K-1} m_k \left(\sum_{i=i_k}^{i_{k+1}-1} Y_i \right) \tag{4.10}$$

where i_k denotes the time step at which the k-th market change takes place, i.e. $i_{k+1} = i_k + l_k$, with $i_0 = 0$. By the battery charging dynamics, this implies that

$$B = \sum_{k=0}^{K-1} m_k \left(\sum_{i=i_k}^{i_{k+1}-1} (S_{i+1} - S_i + X_i) \right) \tag{4.11}$$

$$= \mathbf{m}^T X^n + \sum_{k=0}^{K-1} m_k \left(S_{i_{k+1}} - S_{i_k} \right) \tag{4.12}$$

$$= \mathbf{m}^T X^n + \sum_{k=0}^{K-1} m_k (T_{k+1} - T_k) \tag{4.13}$$

where T_k denotes the battery state at the k-th market transition point, i.e. $T_k = S_{i_k}$. Regrouping terms we have that

$$B = \mathbf{m}^T X^n - m_0 T_0 + (m_0 - m_1) T_1 + \cdots + (m_{K-2} - m_{K-1}) T_{K-1} + m_{K-1} T_K, \tag{4.14}$$

$$= \mathbf{m}^T X^n - m_0 s_0 + \sum_{k=1}^{K} \delta_k T_k, \tag{4.15}$$

where $\boldsymbol{\delta}$ denotes the vector of market price differences, i.e. $\delta_k = m_{k-1} - m_k$ for $k = 1, 2, \ldots, K-1$ and $\delta_K = m_{K-1}$. That is, the price paid for the energy depends on the initial battery state s_0, the energy consumption X^n and the battery state at transition points. Therefore, the additional price paid for the privacy of sequence $\mathbf{x}$ is given by

$$g(Y^n, \mathbf{x}) = \mathbb{E}_{P_{Y^n|X^n=\mathbf{x}}}[\mathbf{m}^T Y^n - \mathbf{m}^T \mathbf{y}^*(\mathbf{x})] \tag{4.16}$$

$$= \mathbb{E}_{P_{Y^n|X^n=\mathbf{x}}} \left[\sum_{k=1}^{K} \delta_k (T_k - T_k^*) \right]. \tag{4.17}$$

Thus, the additional price paid for the energy depends solely on the battery states at transition points. It is worth noting that in general, the optimal battery states at transition points T_k^* depends on the market $\mathbf{m}$, the energy consumption of the user $\mathbf{x}$ and the output alphabet $\mathcal{Y}$. Interestingly, for large output alphabets, i.e. $\mathcal{Y} = \mathbb{Z}$ we have that

$$\mathbf{m}^T \mathbf{y}^*(\mathbf{x}) = \mathbf{m}^T X^n - m_0 s_0 + \min_{T^K \in [\![0,\beta]\!]^K} \sum_{k=1}^{K} \delta_k T_k \tag{4.18}$$

$$= \mathbf{m}^T X^n - m_0 s_0 + \sum_{k=1}^{K} \delta_k \beta \mathbb{1}\{\delta_k < 0\}. \tag{4.19}$$

That is, the strategy to minimize the cost is to fully charge the battery, i.e. $T_k^* = \beta$, when the future price is more expensive than the current one, i.e. when $m_{k-1} - m_k = \delta_k < 0$. Similarly, the optimal strategy is to fully discharge the battery, i.e. $T_k^* = 0$, when the future price is cheaper than the current one, i.e. when $m_{k-1} - m_k = \delta_k > 0$.

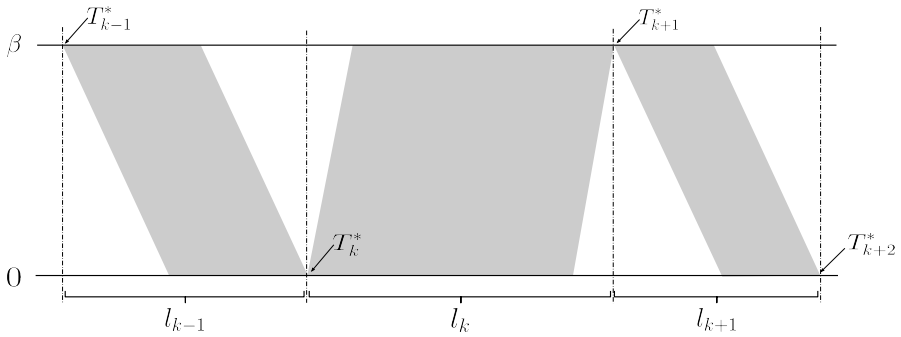

Figure 4.2. Battery state evolution over time. In grey the region containing all possible sequences of battery states that achieve the minimum price, where T^*_k denotes the battery states at transition points that yield the minimum feasible price. Note the different charging and discharging slopes, as the extremes of $\mathcal{Y}$ are allowed different values.

Note that the feasibility of this strategy is not guaranteed when the request alphabet is constrained, as there exist sequences $\mathbf{x}$ such that the battery can not be fully charged/discharged during one market block. The impact of constrained request alphabets is studied in the following section. When $\mathcal{Y} = \mathbb{Z}$, the additional price paid for the privacy is given by

$$g(Y^n, \mathbf{x}) = \mathbb{E}_{P_{Y^n|X^n=\mathbf{x}}} \left[\sum_{k=1}^{K} \delta_k (T_k - T^*_k) \right]. \tag{4.20}$$

Furthermore, when $\mathcal{Y} = \mathbb{Z}$, the additional price paid at each time step is non-negative, and thus

$$g(Y^n, \mathbf{x}) = \mathbb{E}_{P_{Y^n|X^n=\mathbf{x}}} \left[\sum_{k=1}^{K} |\delta_k| |T_k - T^*_k| \right] = \mathbb{E}_{P_{T^K|X^n=\mathbf{x}}} \left[\sum_{k=1}^{K} |\delta_k| T^\Delta_k \right], \tag{4.21}$$

where $T^\Delta_k = |T_k - T^*_k|$ denotes the deviation from the optimal battery state at market transition point k. That is, deviating T^Δ_k units from the optimal market price at transition k induces an extra cost δ_k.

Interestingly, when $\Delta = 0$ and $\mathcal{Y} = \mathbb{Z}$, the battery states at transition points are constrained to a single value, i.e. $T^\Delta_k = 0$ for all $k = 0, 1, \cdots, K-1$, as depicted in Figure 4.2. This simplifies the probabilistic structure of the problem, introduced by the usage of the expected value in the cost constraint, enabling the utilization of some the combinatorial arguments developed in Chapter 3. Furthermore, it enables the independent optimization of each of the K market blocks, since the initial and final battery state of each market block are fixed, Figure 4.2. This allows us to construct the worst case consumption process when $\Delta = 0$, while relaxation of the single value constraints generalizes the result to arbitrary alphabets. Similarly, when

$\Delta = \infty$, the problem reduces to that solved in Chapter 3 using a combinatorial argument. Increasing or reducing the energy requested at market transition points allows us to fine tune the battery state at transition points. Thus, reducing the price paid for the energy and modelling scenarios with arbitrary price constraints.

4.2.2 Impact of the output alphabet on information leakage

In the following, we characterize the impact of the output alphabet on the information leakage $\mathcal{I}(\Delta)$. In particular, we show that the information leakage does not increase when the policy operates with a constrained output alphabet $\mathcal{Y}_c$. Lemma 4.2 shows that it is possible to remove extreme values, i.e. $y_i \notin [\![0, \alpha]\!]$, while retaining the feasibility of the sequence $\mathbf{y} \in \mathcal{Y}^n(s_0, \mathbf{x})$. This is depicted in Figure 4.3.

Lemma 4.2. *Let two output alphabets $\mathcal{Y}_c^n$ and $\mathcal{Y}^n$ be such that $[\![0, \alpha]\!]^n \subseteq \mathcal{Y}_c^n \subseteq \mathcal{Y}^n \subseteq \mathbb{Z}^n$. Then there exists a function $F_n : \mathcal{Y}^n \to \mathcal{Y}_c^n$ such that for any $(s_0, \mathbf{x}) \in \mathcal{S} \times \mathcal{X}^n$ and $\mathbf{y} \in \mathcal{Y}^n(s_0, \mathbf{x})$ it holds that*

$$F_n(\mathbf{y}) \in \mathcal{Y}_c^n(s_0, \mathbf{x}), \tag{4.22}$$

where $\mathcal{Y}_c^n(s_0, \mathbf{x})$ denotes the feasible set with output alphabet $\mathcal{Y}_c$.

Proof: We first define the set of functions $\{h_i\}_{i=1}^n$ that allows us to define F_n. For each function h_i with $i \in [\![1, n]\!]$ set $d_i \in [\![0, (y_i - \alpha)^+]\!]$ and define $h_i : \mathcal{Y}^n \to \mathcal{Y}_c^n$ as

$$h_i(\mathbf{y}) = \begin{cases} \mathbf{y} + d_i(\mathbf{e}_{i+1} - \mathbf{e}_i), & \text{when } i \in [\![1, n-1]\!] \\ \mathbf{y} - d_i \mathbf{e}_i, & \text{otherwise,} \end{cases} \tag{4.23}$$

where $\mathbf{e}_i \in \mathbb{N}^n$ denotes the ith unit vector of the standard base, i.e. $(\mathbf{e}_i)_k = \mathbb{1}\{k = i\}$. That is, the function h_i reallocates the purchase of d_i units of energy from time step i to the next time step $i + 1$. Note that when this occurs on the last time step, i.e. when $i = n$, the excess energy request is not reallocated but removed from the sequence. Let $\mathbf{s} \in \mathcal{S}^{n+1}$ be the sequence of battery states induced by the feasible sequence $\mathbf{y}$. By the battery charging dynamics (3.2), the sequence of battery states induced by $\tilde{\mathbf{y}} = h_i(\mathbf{y})$ is given by

$$\tilde{\mathbf{s}} = \mathbf{s} - d_i \mathbf{e}_{i+1}, \tag{4.24}$$

with $d_i \in [\![0, (y_i - \alpha)^+]\!]$. Note that

$$\tilde{s}_{i+1} = s_{i+1} - d_i \leq s_{i+1} \leq \beta, \tag{4.25}$$

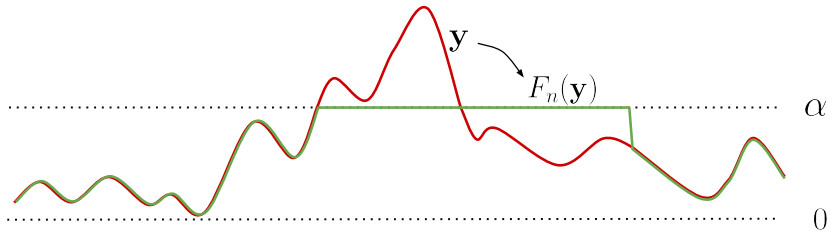

Figure 4.3. Original energy request sequence $\mathbf{y}$ taking values outside $[\![0, \alpha]\!]$ (red) and final energy request sequence $F_n(\mathbf{y})$ taking values in $[\![0, \alpha]\!]$ (green).

and since $x_i \leq \alpha \leq y_i - d_i$

$$\tilde{s}_{i+1} = \tilde{s}_i + (y_i - d_i) - x_i \geq \tilde{s}_i \geq 0. \tag{4.26}$$

Thus $\tilde{\mathbf{s}} \in \mathcal{S}^{n+1}$, and $h_i(\mathbf{y}) \in \mathcal{Y}^n(s_0, \mathbf{x})$. The above argument shows that any excess energy request, i.e. $y_i \geq \alpha$, can be reallocated to the next time step. A similar argument shows that any excess can be reallocated to the previous time step. Furthermore any excess energy selling, i.e. $y_i < 0$, can be reallocated to the next and previous time steps without impacting the feasibility of the energy request. A recursive application of the arguments above yields the existence of the function F_n constructed as

$$F_n(\mathbf{y}) = h_n \circ h_{n-1} \cdots \circ h_1(\mathbf{y}), \tag{4.27}$$

so that $F_n(\mathbf{y}) \in \mathcal{Y}_c^n(s_0, \mathbf{x})$. This completes the proof. ■

The lemma above shows that battery policies that operate over an output alphabet with a maximum energy request that matches the peak energy consumption of the user, i.e. $\mathcal{Y} = [\![0, \alpha]\!]$, are sufficient to satisfy the feasibility. However, in general the function F_n does not preserve the price paid for the energy, as $\mathbf{y}$ and $F_n(\mathbf{y})$ may yield different energy bills. The following lemma identifies the conditions that guarantee that the energy bill does not change after the application of F_n.

Lemma 4.3. *Let the output alphabet be $\mathcal{Y}_c^n \subseteq [\![-\beta/\underline{l}, \beta/\underline{l} + \alpha]\!]^n$ where $\underline{l} = \min_k l_k$ and l_k is the length of the k-th market price period as defined in (4.2). Consider a Δ-feasible battery policy $P_{Y^n|X^n} \in \Gamma(\Delta)$. Then there exists a function $\widehat{F}_{n,\mathbf{m}} : \mathcal{Y}^n \to \mathcal{Y}_c^n$ such that $P_{\widehat{F}_{n,\mathbf{m}}(Y^n)|X^n}$ is a Δ-feasible battery policy for output alphabet $\mathcal{Y}_c$.*

Proof: A sufficient condition for any mapping $\widehat{F}_{n,\mathbf{m}}$ reallocating energy consumption to preserve the energy bill, is to preserve the total amount of energy requested during each market block k. By the battery charging dynamics (3.2), the total amount of energy requested during any block of length l_k, starting at $i = i_k$, is given

by

$$\sum_{i=i_k}^{i_k+l_k} Y_i = T_{k+1} - T_k + \sum_{i=i_k}^{i_k+l_k} X_i, \tag{4.28}$$

Where we recall that T_{k+1} and T_k denote the battery state at transition points between market blocks, as described in (4.13). Note that as T_{k+1}, T_k take values in $[\![0, \beta]\!]$ and X_i take values in $[\![0, \alpha]\!]$ for all i, we have that

$$\sum_{i=i_k}^{i_k+l_k} Y_i \in [\![-\beta, \beta + l_k\alpha]\!] \tag{4.29}$$

That is, the total energy requested during any market block takes values in $[\![-\beta, \beta + l_k\alpha]\!]$. When the output alphabet $\mathcal{Y}_c$ allows such requests i.e. when $\mathcal{Y}_c^n \supseteq [\![-\beta/l_k, \beta/l_k + \alpha]\!]^n$, there is no need to reallocated energy between market blocks. No energy reallocation between markets blocks implies no difference on the price paid for the energy. ■

The lemmas above show that the resulting output sequences $F_n(\mathbf{y})$ and $\widehat{F}_n(\mathbf{y})$ do not depend on the input pair $(s_0, \mathbf{x})$ and instead depend only on the original output sequence $\mathbf{y}$. This insight leads to the following result.

Lemma 4.4. *Let $\mathcal{I}_{\mathcal{Y}}(\Delta)$ represent the privacy-cost function under output alphabet $\mathcal{Y}^n$, then for any $\mathcal{Y} \supset \mathcal{Y}_{\mathcal{X}} = \mathcal{X}$:*

$$\mathcal{I}_{\mathcal{Y}_{\mathcal{X}}}(\infty) = \mathcal{I}_{\mathcal{Y}}(\infty). \tag{4.30}$$

Furthermore for any $\Delta \geq 0$, $\mathcal{Y} \supset \mathcal{Y}_c = [\![-\beta/\underline{l}, \beta/\underline{l} + \alpha]\!]^n$ as defined in Lemma 4.3, it holds that:

$$\mathcal{I}_{\mathcal{Y}_c}(\Delta) = \mathcal{I}_{\mathcal{Y}}(\Delta). \tag{4.31}$$

Proof: Let $\Gamma_{\mathcal{Y}}(\Delta)$ denote the set of feasible Δ-affordable battery policies under output alphabet $\mathcal{Y}$. Lemma 4.2 states the existence of a function $F_n : \mathcal{Y}^n \to \mathcal{Y}_X^n$ such that if $P_{Y^n|X^n} \in \Gamma(\infty)$ then $P_{F_n(Y^n)|X^n} \in \Gamma(\infty)$. The function F_n induces the Markov chain

$$X^n \to Y^n \to F_n(Y^n). \tag{4.32}$$

Therefore $I(X^n; F_n(Y^n)) \leq I(X^n; Y^n)$ by the data processing inequality. The converse follows by noting that $\Gamma_{\mathcal{Y}_X}(\infty) \subseteq \Gamma_{\mathcal{Y}}(\infty)$, and therefore

$$\min_{\Gamma_{\mathcal{Y}}(\infty)} I(X^n; Y^n) \leq \min_{\Gamma_{\mathcal{Y}_X}(\infty)} I(X^n; Y^n). \tag{4.33}$$

For the second statement, Lemma 4.3 shows the existence of a function $\widehat{F}_{n,\mathbf{m}} : \mathcal{Y}^n \to \mathcal{Y}_c^n$ such that if $P_{Y^n|X^n} \in \Gamma(\Delta)$ then $P_{\widehat{F}_n(Y^n)|X^n} \in \Gamma_c(\Delta)$. Noting that the function $\widehat{F}_{n,\mathbf{m}}$ induces the Markov chain

$$X^n \to Y^n \to \widehat{F}_{n,\mathbf{m}}(Y^n) \tag{4.34}$$

yields $I(X^n; \widehat{F}_{n,\mathbf{m}}(Y^n)) \leq I(X^n; Y^n)$ by the data processing inequality. The converse follows by noting that $\Gamma_c(\Delta) \subseteq \Gamma(\Delta)$, and therefore

$$\min_{\Gamma(\Delta)} I(X^n; Y^n) \leq \min_{\Gamma_c(\Delta)} I(X^n; Y^n). \tag{4.35}$$

■

Lemma 4.4 shows that under certain assumptions, the privacy cost function does not vary when the EMU operates with a constrained output alphabet. This result is consistent with prior results reported for privacy based on hypothesis testing [87] and for i.i.d. energy consumptions with instantaneous power constraints [120, 10, 144] We note that the proof for the existence of the function $\widehat{F}_n$ presented in Lemma 4.3 requires forecasting of $\underline{l}$ time steps ahead.

We now improve the upper bounds presented on Chapter 3. Intuitively, smaller sampling periods, i.e. the time elapsed between time steps, lead to larger privacy leakages. This intuition leads to the following result.

Lemma 4.5. *Let n be a positive integer and $\lambda = (\beta + 1)/\alpha$. Then the privacy cost function under no cost constraints, is bounded by*

$$\mathcal{I}(\infty) \leq \frac{1}{n} \left\lceil \frac{n}{\lambda} \right\rceil. \tag{4.36}$$

Proof: Let T_0 denote the time elapsed between two time steps, and $P_{\max} = \alpha/T_0$ denote the maximum power consumption of the user. Let $n' = n\alpha$ denote the corresponding number of samples if the sampling period is reduced to $T_0' = T_0/\alpha = 1/P_{\max}$, i.e. $T = T_0 n = T_0' n'$ with T denoting the total time. Then $\alpha' = P_{\max} T_0' = 1$. Moreover, the total information leakage satisfies

$$n\mathcal{I}(\infty) \leq n'\mathcal{I}'(\infty). \tag{4.37}$$

Which follows the intuition that reducing the sampling period can not decrease the total information leakage. To prove this, let $\mathbf{y}'$ be a feasible sequence for T_0', inducing battery states $\mathbf{s}'$. Then, for T_0, the battery states $\mathbf{s}$, induced by the downsampled

sequence $f^n(\mathbf{y})$ with

$$f_k(\mathbf{y}') = \sum_{i=k\alpha}^{(k+1)\alpha-1} y'_i, \tag{4.38}$$

are a subsequence of the original battery states $\mathbf{s}'$, in particular

$$s_i = s_0 + \sum_{k=0}^{i} \left(f_k(\mathbf{y}') - f_k(\mathbf{x}')\right) = s'_{i\alpha}. \tag{4.39}$$

Thus, $f_k(\mathbf{y}')$ is feasible if $\mathbf{y}'$ is feasible, and inducing a smaller leakage by the data processing inequality. Recall that Theorem 3.3 shows that

$$n'\mathcal{I}'(\infty) \leq \left\lceil \frac{n'}{\lfloor(\beta+1)/\alpha'\rfloor} \right\rceil = \left\lceil \frac{n\alpha}{\lfloor(\beta+1)/1\rfloor} \right\rceil = \left\lceil \frac{n\alpha}{\beta+1} \right\rceil = \left\lceil \frac{n}{\lambda} \right\rceil. \tag{4.40}$$

This completes the proof. ■

Note that for a sampling period T_0 and a maximum power consumption $P_{\max} = \alpha/T_0$, the total amount of information leaked during a time interval $T = nT_0$ is bounded by

$$n\mathcal{I}(\infty) \leq \left\lfloor \frac{n}{\lambda} \right\rfloor = \left\lfloor \frac{T/T_0}{(\beta+1)/(P_{\max}T_0)} \right\rfloor = \left\lfloor \frac{TP_{\max}}{\beta+1} \right\rfloor. \tag{4.41}$$

Thus, for sampling periods satisfying $T_0 \leq (\beta+1)/P_{\max}$, the total amount of information leaked to the provider is independent of the sampling period T_0. This implies that, when block battery policies are employed, increasing the sampling frequency of SMs over $(\beta+1)/P_{\max}$ does not increase the information leakage. This enables the UP to measure the energy consumption of the user arbitrarily often without increasing the information about the activities performed by the user. This property of block battery policies is of particular interest as the sampling frequency of smart meters is predicted to increase [66].

4.3 Privacy with cost constraints

Theorem 4.1 bounds the information leakage for arbitrary cost constraints Δ. The proof proceeds by constructing a battery policy that combines two components for every request sequence. One of the components guarantees the feasibility constraint, while the other guarantees the cost constraint.

4.3.1 Upper bound on the information leakage rate

Theorem 4.1. *Consider an EMU with battery capacity β, initial state s_0, market price $\mathbf{m}$, and output alphabet $\mathcal{Y}^n$ satisfying $\mathcal{Y}_c^n \subseteq \mathcal{Y}^n$ with $\mathcal{Y}_c^n$ defined in Lemma 4.4, then*

$$\mathcal{I}(\Delta) \leq \mathcal{I}(\infty) + \mathcal{I}_\Gamma(\Delta), \tag{4.42}$$

where

$$\mathcal{I}_\Gamma(\Delta) = \min_{P_{\tilde{T}^K|T^K} \in \Gamma_\omega(\Delta)} \max_{P_{T^K} \in \mathcal{P}_{T^K}} \frac{1}{n} I(\tilde{T}^K - T^K; T^K). \tag{4.43}$$

Here T^K and $\tilde{T}^K$ are random processes in $[\![0, \beta]\!]^K$, $\mathcal{P}_{T^K}$ denotes the set of possible distributions over $[\![0, \beta]\!]^K$, and the joint distribution between T^K and $\tilde{T}^K$ is determined by

$$\Gamma_\omega(\Delta) = \left\{ P_{\tilde{T}^K|T^K} : \mathbb{E}\left[\sum_{k=1}^{K} \delta_k (\tilde{T}_k - T_k^*) \right] \leq \Delta \right\}, \tag{4.44}$$

where $T_k^* = \beta \mathbb{1}\{\delta_k < 0\}$ is the battery state at transition times achieving the minimum cost, $\delta_k \in \mathbb{R}$ denotes the vector of market price differences, with entries given by $\delta_k = m_{k-1} - m_k$ for $k = 1, 2, \ldots, K-1$ and $\delta_K = m_{K-1}$.

Proof: We prove the result for $\mathcal{Y}^n = \mathbb{Z}^n$; Lemma 4.4 generalizes the proof for every $\mathcal{Y}^n$ satisfying $\mathcal{Y}_c^n \subseteq \mathcal{Y}^n$. The proof follows by dividing the optimization process into two steps. In the first step, we present a battery policy ω such that the resulting request sequence V_ω^n satisfies the power outage and energy waste constraints, i.e., $\omega \in \Omega(s_0)$ as defined in (3.4). These policies are discussed in Chapter 3. In the second step, we define a random vector V_γ^n such that $Y^n = V_\omega^n + V_\gamma^n$ also satisfies the cost constraints. Specifically, for T_k denoting the battery state induced by V_ω^n and X^n at market transition time k, we set

$$V_\gamma^n = \sum_{k=1}^{K} \Big((\mathbf{e}_{i_k - 1} - \mathbf{e}_{i_k})(\tilde{T}_k - T_k) \Big), \tag{4.45}$$

where i_k denotes the time step of the k-th market transition. This implies that the battery state induced by $V_\omega^n + V_\gamma^n$ and X^n at market transition time k, is $\tilde{T}_k$. Thus,

$$g(Y^n, \mathbf{x}) = \mathbb{E}_{P_{T^K|X^n = \mathbf{x}}} \left[\sum_{k=1}^{K} \delta_k (\tilde{T}_k - T_k^*) \right], \tag{4.46}$$

where (4.46) follows by (4.20) and recalling that $\mathcal{Y}^n = \mathbb{Z}^n$. Selecting the transformation γ determining $\tilde{T}^K$ from the set described in (4.44) yields

$$\begin{aligned} I(X^n;Y^n) &\leq I(X^n;V_\omega^n) + I(X^n;V_\gamma^n|V_\omega^n) && (4.47)\\ &= I(X^n;V_\omega^n) + H(V_\gamma^n|V_\omega^n) - H(V_\gamma^n|V_\omega^n,X^n,T^K) && (4.48)\\ &= I(X^n;V_\omega^n) + H(\tilde{T}^K - T^K|V_\omega^n) - H(\tilde{T}^K - T^K|T^K) && (4.49)\\ &\leq I(X^n;V_\omega^n) + I(\tilde{T}^K - T^K;T^K), && (4.50) \end{aligned}$$

where (4.48) follows as X^n and V_ω^n determine T^K by the battery charging dynamics (3.2); (4.49) follows by (4.45) and noting that $\tilde{T}^K - T^K$ is independent of V_ω^n and X^n given T^K. Thus

$$\begin{aligned} n\mathcal{I}(\Delta) &= \min_{P_{Y^n|X^n}\in\Gamma(\Delta)} \max_{\mathcal{P}_{X^n}} I(X^n;Y^n) && (4.51)\\ &\leq \min_{\gamma\in\Gamma_\omega(\Delta)} \min_{\omega\in\Omega(s_0)} \max_{\mathcal{P}_{X^n}} \Big(I(X^n;V_\omega^n) + I(\tilde{T}^K - T^K;T^K) \Big) && (4.52)\\ &\leq \min_{\omega\in\Omega(s_0)} \max_{\mathcal{P}_{X^n}} I(X^n;V_\omega^n) + \min_{\gamma\in\Gamma_\omega(\Delta)} \max_{\mathcal{P}_{T^K}} I(\tilde{T}^K - T^K;T^K). && (4.53) \end{aligned}$$

This completes the proof. ■

Corollary 4.1. *Consider an EMU with battery capacity β, initial state s_0, market price $\mathbf{m}$, and output alphabet $\mathcal{Y}^n$ satisfying $\mathcal{Y}_c^n \subseteq \mathcal{Y}^n$ with $\mathcal{Y}_c^n$ defined in Lemma 4.4, then for any input process X^n:*

$$\mathcal{I}(\Delta) \leq \frac{1}{n}\left\lceil \frac{n}{\lambda} \right\rceil + \mathcal{I}_\Gamma(\Delta), \tag{4.54}$$

where $\mathcal{I}_\Gamma(\Delta)$ is defined in (4.43). Furthermore, for any input process X^n with expected energy consumption μ_n:

$$\mathcal{I}(\Delta) \leq \frac{1}{n}\left\lceil \frac{n}{\lambda} \right\rceil \max_{\epsilon\in[\frac{-s_0}{n\alpha},\frac{\beta-s_0}{n\alpha}]} H_2\left(\frac{\mu_n}{\alpha} + \epsilon\right) + \mathcal{I}_\Gamma(\Delta), \tag{4.55}$$

where $\mathcal{I}_\Gamma(\Delta)$ is defined in (4.43).

Proof: The proof of (4.54) follows by Theorem 3.3 and Lemma 4.5. The proof of (4.55) follows by Theorem 3.5, Lemma 4.5, and using that $\sum_{i=0}^{n-1}(V_\gamma)_i = 0$. ■

While direct computation of the information leakage in (4.5) relies on finding an n-dimensional joint distribution satisfying $\Gamma(\Delta)$, the bound presented in (4.42) relies on a K-dimensional distribution and the simplified version of $\Gamma(\Delta)$ defined in (4.44). This significantly eases the computation of the information leakage as described in Section 4.4. The following corollary provides a single letter bound for the additional information leakage rate induced by the market.

Corollary 4.2. *Consider an EMU with battery capacity β, initial state s_0, market price $\mathbf{m}$, and output alphabet $\mathcal{Y}^n$ satisfying $\mathcal{Y}_c^n \subseteq \mathcal{Y}^n$ with $\mathcal{Y}_c^n$ defined in Lemma 4.4, then for any input process X^n:*

$$\mathcal{I}_\Gamma(\Delta) \leq \left(1 - \frac{\Delta}{\Delta_{\max}}\right)^+ \frac{K}{n} \log(\beta + 1). \tag{4.56}$$

Proof: Note that by (4.43) and cardinality bounds on the mutual information, we have that

$$n\mathcal{I}_\Gamma(0) \leq \log |T^K| = K \log(\beta + 1). \tag{4.57}$$

Moreover, by (4.44), for any feasible battery state $\tilde{T}_k$ in $[\![0, \beta]\!]$ and optimal battery state T_k^* the system cost function is bounded by

$$\mathbb{E}\left[\sum_{k=1}^{K} \delta_k(\tilde{T}_k - T_k^*)\right] \leq \beta \|\boldsymbol{\delta}\|_1 = \Delta_{\max}. \tag{4.58}$$

Therefore, for $\Delta \geq \Delta_{\max}$, any feasible battery state $\tilde{T}_k$ satisfies the cost function. Thus setting $\tilde{T}^K = T^K$ satisfies the cost constraint, and yields that $\mathcal{I}_\Gamma(\Delta) = 0$. A time-sharing argument, based on the convexity of $\Gamma(\Delta)$, proved in Lemma 4.1, completes the proof. ■

4.3.2 Tightness of the upper bound

We now address the tightness of the upper bound presented in Theorem 4.1. To that end, we construct a random process modelling the energy consumption of the user that is tight with respect to the result in Theorem 4.1 for every battery policy in $\Gamma(\Delta)$.

Theorem 4.2. *The privacy cost function $\mathcal{I}(\Delta)$ is lower bounded by*

$$\frac{1}{n} \sum_{k=0}^{K-1} \left(\kappa_k + \log(l_k \alpha - \kappa_k \lceil \lambda \rceil \alpha) - \min_{\Gamma_T(\Delta)} H(T_k)\right) \leq \mathcal{I}(\Delta), \tag{4.59}$$

where $\kappa_k = \lfloor l_k / \lceil \lambda \rceil - 1 \rfloor^+$ and

$$\Gamma_T(\Delta) = \left\{ P_{T^K} : \mathbb{E}\left[\sum_{k=1}^{K} \delta_k(T_k - T_k^*)\right] \leq \Delta \right\}. \tag{4.60}$$

Proof: We proceed to prove by presenting the probability law P_{W^n} of a random process W^n achieving the lower bound. Let the input alphabet $\mathcal{W}^n$ be divided

according to the market price partitioning, i.e.

$$\mathcal{W}^n = \mathcal{W}^{l_1} \times \mathcal{W}^{l_2} \times ... \times \mathcal{W}^{l_K}. \tag{4.61}$$

Let each market partition be further divided in two, i.e. $\mathcal{W}^{l_k} = \mathcal{A}_k \times \mathcal{B}_k$. Let

$$\mathcal{A}_k = \mathcal{O}^{\kappa_k}_{\lceil\lambda\rceil} = \left\{ \underbrace{(0, 0, \cdots, 0)}_{\lceil\lambda\rceil}, \underbrace{(\alpha, \alpha, \cdots, \alpha)}_{\lceil\lambda\rceil} \right\}^{\kappa_k}, \tag{4.62}$$

with $\kappa_k = \lfloor l_k / \lceil\lambda\rceil - 1 \rfloor^+$ and where $\mathcal{O}^{\kappa_k}_{\lceil\lambda\rceil}$ is the block repetition alphabet presented in Definition 3.6. Moreover let

$$\mathcal{B}_k = \{\mathbf{w}_0, \mathbf{w}_1, \cdots, \mathbf{w}_N\} \subset \mathcal{X}^{l_k - \lceil\lambda\rceil\kappa_k}, \tag{4.63}$$

with $\sigma(\mathbf{w}_i) = i$ and $N = \alpha(l_k - \lceil\lambda\rceil\,\kappa_k)$, i.e. no two elements in $\mathcal{B}_k$ share the same total consumption. Let W^n take uniformly distribution values over $\mathcal{W}^n$, then for any feasible battery policy it holds that

$$\begin{aligned} n\mathcal{I}(\Delta) &\geq \min_{\Gamma(\Delta)} I(W^n; Y^n) && (4.64)\\ &= \log|\mathcal{W}^n| - \max_{\Gamma(\Delta)} H(W^n|Y^n) && (4.65)\\ &= \log|\mathcal{W}^n| - \max_{\Gamma(\Delta)} H(T^K|Y^n) && (4.66)\\ &\geq \log|\mathcal{W}^n| - \max_{\Gamma(\Delta)} \sum_{k=0}^{K-1} H(T_k), && (4.67) \end{aligned}$$

where inequality (4.64) follows as $P_{W^n} \in \mathcal{P}_{X^n}$. Equality (4.65) follows as W^n take uniformly distributed values over $\mathcal{W}^n$. Equality (4.66) holds by the Bayes' rule

$$H(W^n|Y^n) = H(W^n|T^K, Y^n) + H(T^K|Y^n) - H(T^K|W^n, Y^n), \tag{4.68}$$

and using that $H(W^n|T^K, Y^n) = 0$ by Theorem 3.3, and $H(T^K|W^n, Y^n) = 0$ by the battery charging dynamics. Finally, inequality (4.67) follows by the chain rule and the fact that conditioning reduces entropy. This completes the proof. ■

4.4 Numerical results

In this section, we numerically assess the upper bounds on the privacy cost described in Theorem 4.2 and Theorem 4.1. We model the market price after the UK Economy 7 tariff, where users are charged an off-peak price of 0.071 £/kWh within a 7 hour block and a peak price of 0.152 £/kWh otherwise [81]. We assume the user has an LG Chem RESU 6.5 battery with a capacity of 4.2 kWh and a peak power of 4.2 kW [145].

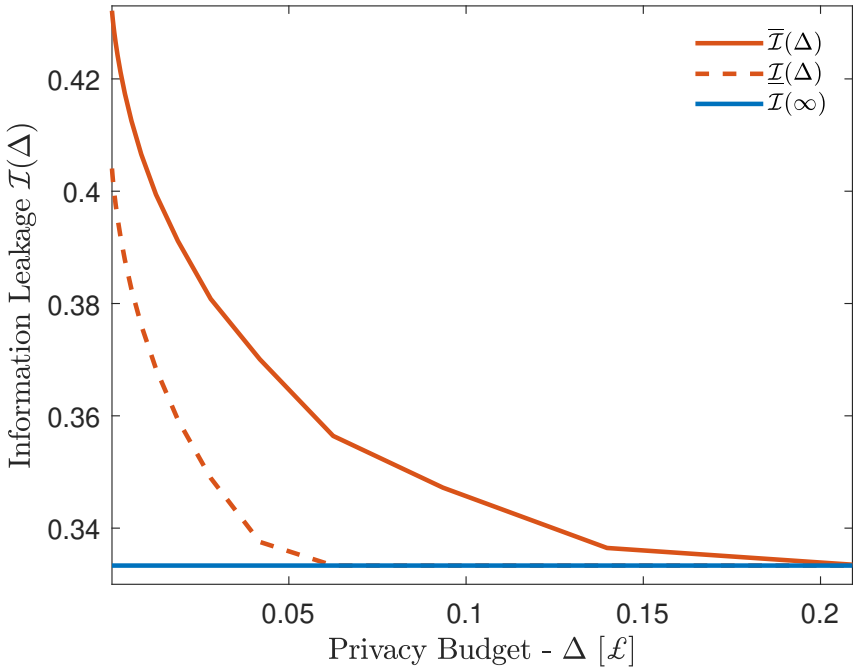

Figure 4.4. Information leakage $\mathcal{I}(\Delta)$ against privacy budget Δ with $\underline{\mathcal{I}}(\Delta)$, $\mathcal{I}(\Delta)$ and $\bar{\mathcal{I}}(\Delta)$ denoting lower bound, exact value, and upper bound respectively.

For simplicity we match the users' maximum power consumption to the peak power of the battery, i.e., 4.2 kW [65]. The SM sends the UP integrated energy readings every 10 min. Thus, we set the time elapsed between time steps i and $i+1$ to 10 min. Defining 175W as 1 unit of energy yields the following parameters in our system model: market lengths $l_0 = 7\text{h} \times 3\,\text{samples/hour} = 21$ and $l_1 = 17\text{h} \times 3\,\text{samples/hour} = 51$; corresponding market prices of $m_0 = 0.152£/\text{kWh} \times 3\,\text{samples/hour} \times 375\,\text{W/unit} = 1.90\,\text{cents}$ per sample per energy unit; and $m_1 = 0.071£/\text{kWh} \times 3\,\text{samples/hour} \times 375\,\text{W/unit} = 0.88\,\text{cents}$ per sample per energy unit; maximum consumption between time steps $\alpha = 4.5\,\text{kW}/(3\,\text{samples/hour})/(375\,\text{W/unit}) = 4$; battery capacity $\beta = 4.125\,\text{kW}/(375\,\text{W/unit}) = 11$.

Figure 4.7 depicts the bounds on the privacy cost $\mathcal{I}(\Delta)$ for different values of the system cost Δ and initial battery state $s_0 = 0$ during a one day period, i.e. $n = 24\,\text{h}/0.3\,\text{h} = 72$. It can be seen how the information leakage decreases with the increase of the privacy budget. The convexity of the set of feasible battery policies on Δ can be appreciated on the upper bound. For large values of the system cost Δ the cost constraint is always satisfied, i.e. $\mathcal{I}_\Gamma(\Delta) = 0$, and the privacy leakage is governed by the feasibility constraints. Figure 4.4 shows the impact of the privacy budget Δ on the information leakage. Note how the information leakage decreases logarithmically with the privacy budget.

Figure 4.5 depicts the impact of the maximum energy consumption α on the information leakage. It can be seen how larger values of α increase the information

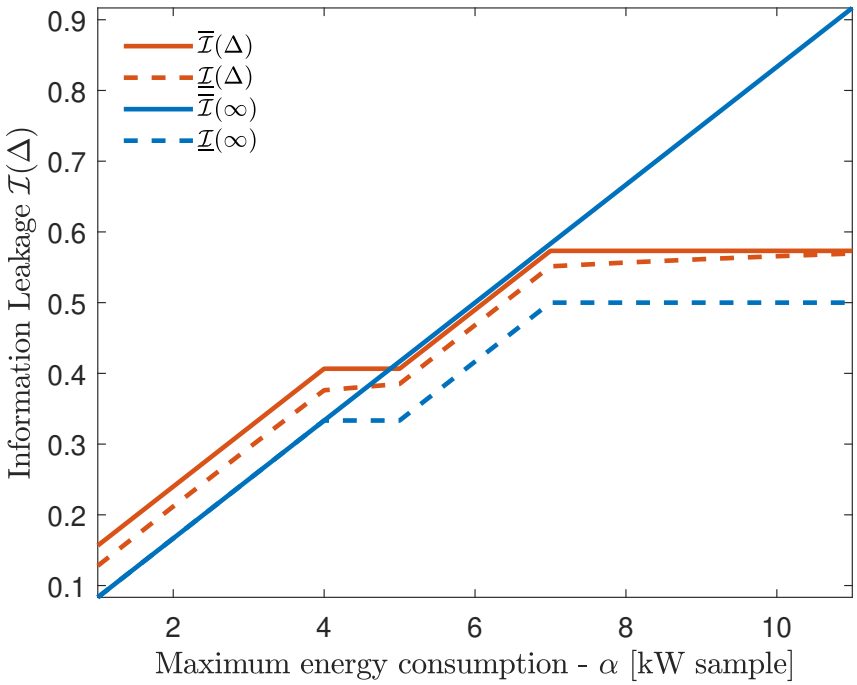

Figure 4.5. Information leakage $\mathcal{I}(\Delta)$ against maximum energy consumption α with $\underline{\mathcal{I}}(\Delta)$ and $\bar{\mathcal{I}}(\Delta)$ denoting lower bound and upper bound respectively.

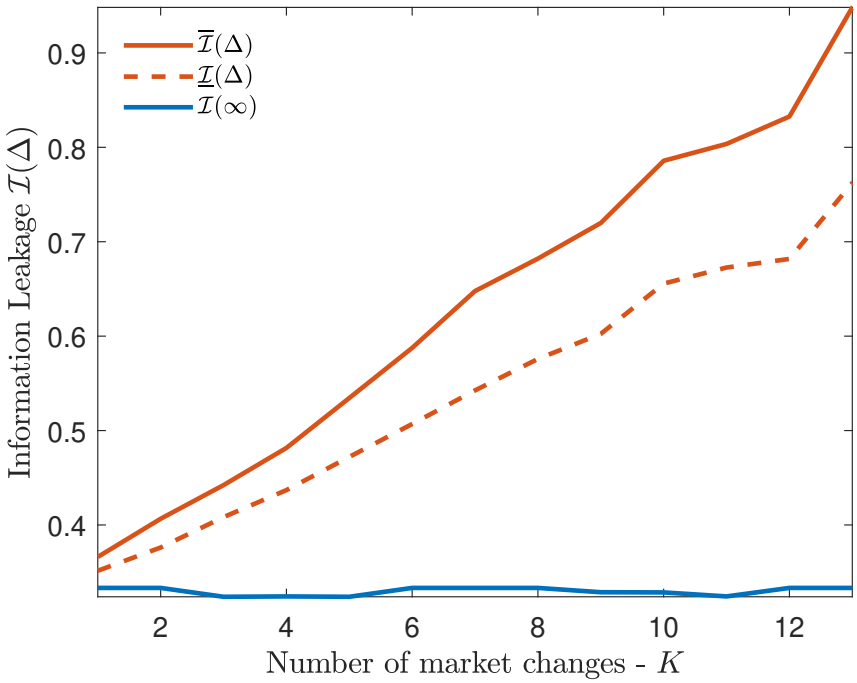

Figure 4.6. Information leakage against number of market changes K for a fixed total length n with $\underline{\mathcal{I}}(\Delta)$, $\mathcal{I}(\Delta)$ and $\bar{\mathcal{I}}(\Delta)$ denoting lower bound, exact value, and upper bound respectively.

leakage. The step like behaviour is due to the discrete time approximation of the system and the sampling period. Figure 4.6 shows the impact of the number of market changes K on the privacy guarantee. Therein the length of the consumption

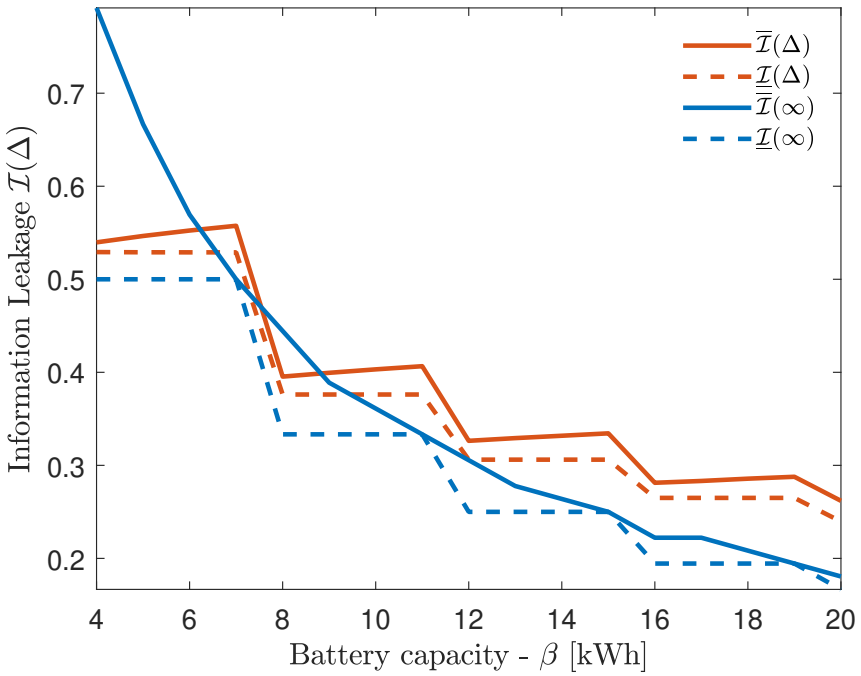

Figure 4.7. Information leakage $\mathcal{I}(\Delta)$ against battery capacity β with $\underline{\mathcal{I}}(\Delta)$ and $\bar{\mathcal{I}}(\Delta)$ denoting lower bound and upper bound respectively.

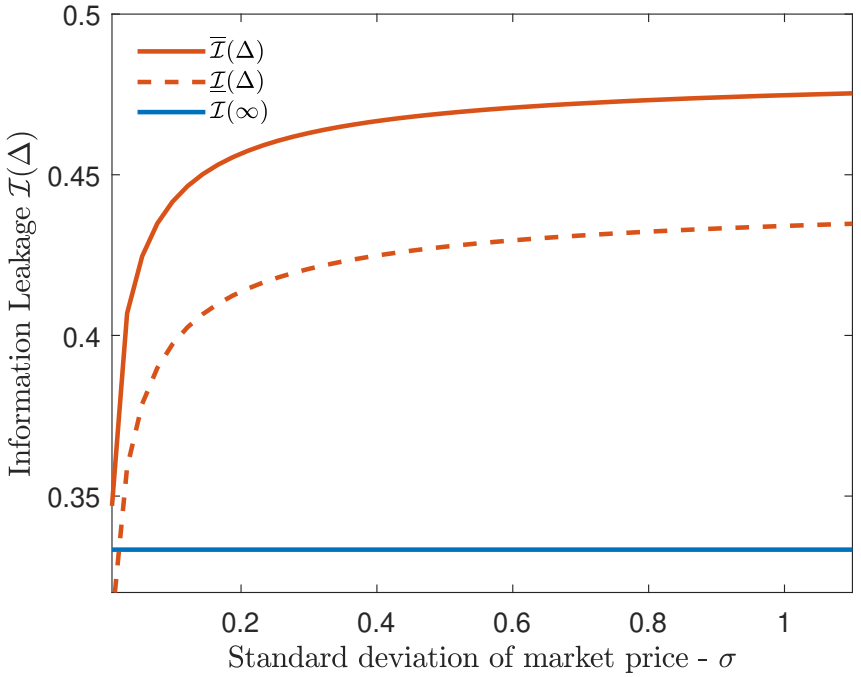

Figure 4.8. Information leakage $\mathcal{I}(\Delta)$ against standard deviation of the market σ with $\underline{\mathcal{I}}(\Delta)$, $\mathcal{I}(\Delta)$ and $\bar{\mathcal{I}}(\Delta)$ denoting lower bound, exact value, and upper bound respectively.

sequence, n, is preserved as K increases. It can be seen that the information leakage increases linearly with the number the market changes. It can also be appreciated how the gap between upper and lower bound increases as the number of market

changes increases. Figure 4.7 shows the impact of the battery capacity on the privacy guarantee. As described on the previous chapter, the privacy guarantee decreases with the inverse of the battery capacity. Figure 4.8 shows the impact of the standard deviation of the market price prices on the privacy guarantee. Therein the UK 7 tariff is employed and the corresponding prices are modified to match the desired standard deviation σ. The information leakage increases rapidly with σ for small values of σ, to stabilize later on the large σ regime.

4.5 Conclusion

In this chapter we have studied the tradeoff between privacy and energy cost in the presence of variable market prices. Therein, we studied the market price constraints, noticing that the energy bill is determined by the battery state at transition points. In Section 4.3 an upper bound holding for any distribution on the consumption is derived for the case when a maximum privacy budget is available. The upper bound is derived in two steps, on the first step the feasibility constraints are satisfied as studied in Chapter 3. On the second step, the battery states at transition points are optimized in order to satisfy the budget constraints. The tightness of the bound is studied by deriving the probability law of an energy consumption process inducing a fixed information leakage for any feasible battery policy. Numerical simulations were presented in Section 4.4. Analysis of the numerical simulations shows that leakage decreases with the privacy budget, while it increases for larger deviations of the market price and number of market changes.

Chapter 5

Single-letter bounds of universal privacy guarantees

In this section, we provide single letter bounds to the information leakage and extend these bounds and the results obtained in previous chapters to other privacy metrics. Section 5.1 studies the necessary and sufficient conditions for the existence of shared feasible requests. Therein, the cardinalities of the minimal covering and packing sets are bounded. This enables the exact characterization of the information leakage when no privacy budget is allocated. Section 5.2 provides single letter upper and lower bounds for any privacy budget. Numerical simulations evaluating the obtained bounds are then presented in Section 5.3. Finally, Section 5.4 extends the validity of the results obtained through out this thesis to maximal leakage and maximal α-leakage. The extension is accomplished via a general result showing that under certain conditions, maximal leakage, maximal alpha leakage and mutual information coincide.

In continuity with Chapter 4, in this chapter, we follow the system model defined in Section 4.1.

5.1 Challenges and methodology: Geometry of the set of feasible request

In this section, we study the geometry of the set of feasible energy request sequences. Therein, we characterize the fundamental properties determining the existence of shared feasible sequences, and provide cardinality bounds to the set covering and packing of the input and output alphabets.

5.1.1 Shared output sequences

For a given β, Lemma 5.1 characterizes a necessary and sufficient condition under which all input pairs $(s_0, \mathbf{x})$ in a set $\mathcal{A}$ share a common feasible output $\mathbf{y}_\mathcal{A}$. Such shared output sequences are good for privacy since a UP that sees $\mathbf{y}_\mathcal{A}$ cannot distinguish which input pair $(s_0, \mathbf{x}) \in \mathcal{A}$ caused it. Conversely, when two inputs $(s_0, \mathbf{x}), (\hat{s}_0, \hat{\mathbf{x}})$ share no feasible output $\mathbf{y}_\mathcal{A}$, the EMU cannot conceal from the UP which pair caused the request. The following measure of distance is useful for that analysis.

Definition 5.1. *The distance between two input pairs* $(s_0, \mathbf{x}), (\hat{s}_0, \hat{\mathbf{x}}) \in \mathcal{S} \times \mathcal{X}^n$ *is defined as*

$$d_n\Big((s_0, \mathbf{x}), (\hat{s}_0, \hat{\mathbf{x}})\Big) = \max_{i \in [\![0, n-1]\!]} \Big| \Big(s_0 - \sigma(\mathbf{x}^i)\Big) - \Big(\hat{s}_0 - \sigma(\hat{\mathbf{x}}^i)\Big) \Big|. \tag{5.1}$$

Where we recall that $\sigma(\cdot)$ denotes the sum over all the elements of a vector, e.g. $\sigma(\mathbf{x}^i)$ denotes the total accumulated energy consumption up to time $i-1$, i.e. $\sigma(\mathbf{x}^i) = \sum_{k=0}^{i-1} x_k$. Let z_i denote $s_0 - \sigma(\mathbf{x}^i)$ and note that the definition above satisfies the conditions of a distance, i.e. it satisfies the triangle inequality

$$\begin{aligned} d_n((s_0, \mathbf{x}), (s_0', \mathbf{x}')) &= \max_{i \in [\![0, n-1]\!]} |z_i - z_i'| && (5.2)\\ &\leq \max_{i \in [\![0, n-1]\!]} (|z_i - z_i''| + |z_i'' - z_i'|) && (5.3)\\ &\leq \max_{i \in [\![0, n-1]\!]} |z_i - z_i''| + \max_{j \in [\![0, l]\!]} |z_j'' - z_j'|, && (5.4) \end{aligned}$$

the symmetry $d_n((s_0, \mathbf{x}), (s_0', \mathbf{x}')) = d_n((s_0', \mathbf{x}'), (s_0, \mathbf{x}))$, and the identity of indiscernibles $d_n((s_0, \mathbf{x}), (s_0, \mathbf{x})) = 0$ conditions.

Lemma 5.1 shows that the distance between input pairs determines the existence of a shared feasible output $\mathbf{y}$. The result emphasizes the central role that battery capacity β plays in privacy.

Lemma 5.1. *Let* $\mathcal{A}$ *denote a subset of the input pair alphabet* $\mathcal{S}_0 \times \mathcal{X}^n$*. Then, the following two statements are equivalent.*

a) The distance between every two pairs $(s_0, \mathbf{x}), (\hat{s}_0, \hat{\mathbf{x}}) \in \mathcal{A}$ *is less than or equal to the capacity of the battery, i.e.*

$$d_n\Big((s_0, \mathbf{x}), (\hat{s}_0, \hat{\mathbf{x}})\Big) \leq \beta \text{ for all } (s_0, \mathbf{x}), (\hat{s}_0, \hat{\mathbf{x}}) \in \mathcal{A}. \tag{5.5}$$

b) All sequences in $\mathcal{A}$ *share a feasible request* $\mathbf{y}_\mathcal{A}$*, i.e.*

$$\mathbf{y}_\mathcal{A} \in \bigcap_{(s_0, \mathbf{x}) \in \mathcal{A}} \mathcal{Y}^n(s_0, \mathbf{x}). \tag{5.6}$$

Proof: Let the sequence $\mathbf{y}_{\mathcal{A}}$ be such that for all i:

$$\sigma(\mathbf{y}_{\mathcal{A}}^i) \triangleq -\min_{(s_0,\mathbf{x})\in\mathcal{A}}(s_0 - \sigma(\mathbf{x}^i)), \tag{5.7}$$

Thus, for any $(\hat{s}_0, \hat{\mathbf{x}}) \in \mathcal{A}$, the battery state at time $i+1$ is

$$s_{i+1} = (\hat{s}_0 - \sigma(\hat{\mathbf{x}}^i)) - \min_{(s_0,\mathbf{x})\in\mathcal{A}}(s_0 - \sigma(\mathbf{x}^i)). \tag{5.8}$$

Since $d_n\Big((s_0,\mathbf{x}),(\hat{s}_0,\hat{\mathbf{x}})\Big) \leq \beta$ implies that $s_{i+1} \in [\![0,\beta]\!]$ for all i, it follows that $\mathbf{y}_{\mathcal{A}}$ is a feasible sequence. The converse follows since for any sequence $\mathbf{y}$ and any input pairs $(s_0,\mathbf{x}),(\hat{s}_0,\hat{\mathbf{x}}) \in \mathcal{A}$ such that $d_n\Big((s_0,\mathbf{x}),(\hat{s}_0,\hat{\mathbf{x}})\Big) > \beta$, the absolute difference between the corresponding battery states at some time step i satisfies

$$\Big|s_{i+1} - \hat{s}_{i+1}\Big| = \Big|(s_0 - \sigma(\mathbf{x}^i)) - (\hat{s}_0 - \sigma(\hat{\mathbf{x}}^i))\Big| > \beta. \tag{5.9}$$

Thus s_{i+1} and $\hat{s}_{i+1}$ cannot both belong to $\mathcal{S} = [\![0,\beta]\!]$. ■

In the following we particularize the previous result to the case in which the battery state at time n is in a given range.

Lemma 5.2. *Let $\mathcal{A}_n(\mathcal{S}_0,\mathcal{Z}_n)$ denote a subset of the input alphabet $\mathcal{S}\times\mathcal{X}^n$ satisfying*

$$\mathcal{A}_n(\mathcal{S}_0,\mathcal{Z}_n) = \Big\{(s_0,\mathbf{x}) \in \mathcal{S}_0 \times \mathcal{X}^n : s_0 - \sigma(\mathbf{x}) \in \mathcal{Z}_n\Big\}, \tag{5.10}$$

with $\mathcal{S}_0 = [\![\underline{\mathcal{S}_0}, \overline{\mathcal{S}_0}]\!] \subseteq \mathcal{S}$ and $\mathcal{Z}_n = [\![\underline{\mathcal{Z}_n}, \overline{\mathcal{Z}_n}]\!] \subseteq [\![\underline{\mathcal{S}_0} - n\alpha, \overline{\mathcal{S}_0}]\!]$, where $\bar{\cdot}$ and $\underline{\cdot}$ denote the maximum and minimum element of an ordered set. Then the following two statements are equivalent.

a) All sequences in $\mathcal{A}$ share a feasible request $\mathbf{y}_{\mathcal{A}}$, i.e.

$$\mathbf{y}_{\mathcal{A}} \in \bigcap_{(s_0,\mathbf{x})\in\mathcal{A}} \mathcal{Y}^n(s_0,\mathbf{x}). \tag{5.11}$$

b) One or more of the following holds

$$n\alpha + \overline{\mathcal{Z}_n} - \underline{\mathcal{S}_0} \leq \beta, \tag{5.12}$$

$$\overline{\mathcal{S}_0} - \underline{\mathcal{Z}_n} \leq \beta, \tag{5.13}$$

$$\overline{\mathcal{S}_0} - \underline{\mathcal{S}_0} + i^*\alpha \leq \beta + d, \tag{5.14}$$

where $i^* = ((\overline{\mathcal{Z}_n} - \underline{\mathcal{Z}_n}) - (\overline{\mathcal{S}_0} - \underline{\mathcal{S}_0}) + n\alpha)/2\alpha$ and $d = \alpha\min(i^* - \lfloor i^* \rfloor, \lceil i^* \rceil - i^*)$.

Proof: By Lemma 5.1, it suffices to show that (5.12) to (5.14) are equivalent to (5.5). Note that (5.5) holds if and only if

$$\max_{(s_0,\mathbf{x})\in\mathcal{A}} \Big(s_0 - \sigma(\mathbf{x}^i)\Big) - \min_{(s_0,\mathbf{x})\in\mathcal{A}} \Big(s_0 - \sigma(\mathbf{x}^i)\Big) \leq \beta \tag{5.15}$$

for all $i \in [\![0, n-1]\!]$. The geometrical analysis shown in Figure 5.1 reveals (5.15) is equivalent to

$$\min\{\overline{\mathcal{S}_0}, \overline{\mathcal{Z}_n} + (n-i)\alpha\} - \max\{\underline{\mathcal{S}_0} - i\alpha, \underline{\mathcal{Z}_n}\} \leq \beta \tag{5.16}$$

for all $i \in [\![0, n-1]\!]$. This holds if and only if, for all i, one of the following holds

$$\overline{\mathcal{S}_0} - \underline{\mathcal{S}_0} + i\alpha \leq \beta, \tag{5.17}$$
$$\overline{\mathcal{S}_0} - \underline{\mathcal{Z}_n} \leq \beta, \tag{5.18}$$
$$\overline{\mathcal{Z}_n} + (n-i)\alpha - \underline{\mathcal{S}_0} + i\alpha \leq \beta, \tag{5.19}$$
$$\overline{\mathcal{Z}_n} + (n-i)\alpha - \underline{\mathcal{Z}_n} \leq \beta. \tag{5.20}$$

Algebraic manipulation of equations (5.17) and (5.20) completes the proof. ■

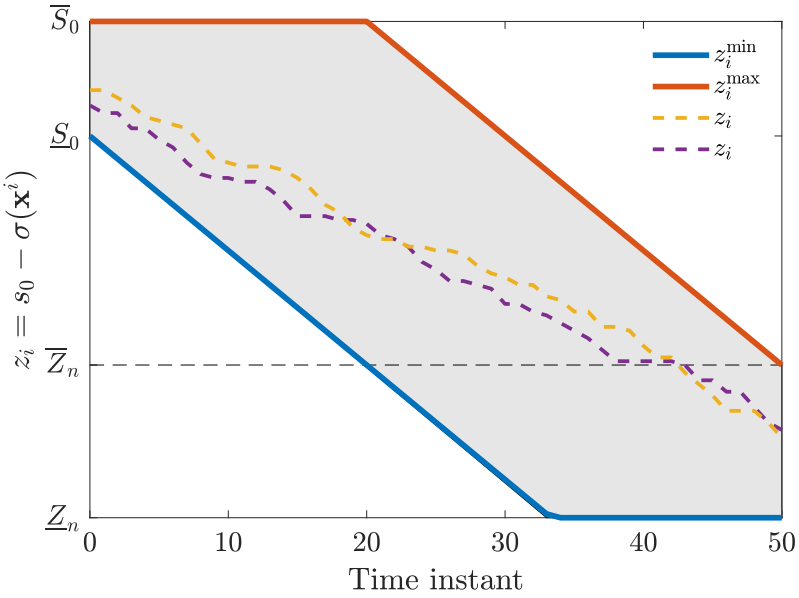

Figure 5.1. Evolution of the battery state when no energy is introduced into the battery, where $z_i = s_0 - \sigma(\mathbf{x}^i)$ takes values in the grey area.

5.1.2 Cardinality bounds

Building on Lemma 5.1, Theorem 5.1 gives an upper bound on the number of distinguishable input pairs $(s_0, \mathbf{x}^n) \in \mathcal{S}_0 \times \mathcal{X}^n$, where $\mathcal{S}_0 \subseteq \mathcal{S}$ is the set of possible

initial battery states. The result is obtained by building a covering $\{\mathcal{A}_i\}$ of $\mathcal{S}_0 \times \mathcal{X}$ such that all input pairs in each $\mathcal{A}_i$ share a common feasible request $\mathbf{y}_i$. The result shows that the minimum time $\lambda \triangleq \lfloor(\beta+1)/\alpha\rfloor$ needed to fully discharge a battery of capacity β under maximal consumption $\alpha \triangleq \max \mathcal{X}$ is a central parameter in the construction of privacy preserving battery policies.

Theorem 5.1. *Let $\mathcal{S}_0 = [\![\underline{\mathcal{S}_0}, \overline{\mathcal{S}_0}]\!]$ and $\mathcal{S}_n = [\![\underline{\mathcal{S}_n}, \overline{\mathcal{S}_n}]\!]$ be two intervals denoting the possible states of the battery at times 0 and n. There exists a set of request sequences $\mathcal{V}^n(\mathcal{S}_0, \mathcal{S}_n) \subseteq \mathcal{Y}^n$ such that*

$$\log\left|\mathcal{V}^n(\mathcal{S}_0,\mathcal{S}_n)\right| \leq \left\lfloor \frac{n-\lambda+d_0}{\lambda} \right\rfloor^+ + \log\left\lceil \frac{(n+d_0) \bmod \lambda + \lambda \mathbb{1}\{n \geq l_z\}}{|\mathcal{S}_n|\alpha^{-1}} + \mathbb{1}_{\mathcal{B}} \right\rceil. \tag{5.21}$$

with $d_0 = (\overline{\mathcal{S}_0} - \underline{\mathcal{S}_0})/\alpha \leq \lambda$, and $\mathbb{1}_{\mathcal{B}}$ the indicator function

$$\mathbb{1}_{\mathcal{B}} = \mathbb{1}\left\{ \frac{\beta + l_r\alpha}{|\mathcal{S}_n|} \leq 1 + 2\left\lfloor \frac{l_r\alpha}{|\mathcal{S}_n|} \right\rfloor \right\}. \tag{5.22}$$

Moreover, for every input pair $(s_0, \mathbf{x}) \in \mathcal{S}_0 \times \mathcal{X}^n$, at least one sequence $\mathbf{v} \in \mathcal{V}^n(\mathcal{S}_0)$ is feasible, i.e.

$$\mathcal{Y}^n(s_0, \mathbf{x}) \cap \mathcal{V}^n(\mathcal{S}_0) \neq \emptyset, \tag{5.23}$$

and guarantees that $s_n \in \mathcal{S}_n$.

Proof: Consider the first $l_z = \beta/\alpha - d_0$ time steps, and note that all possible input sequences during those l_z time steps, i.e. $\mathcal{S}_0 \times \mathcal{X}^{l_z}$, are contained in the set

$$\mathcal{A}_0 = \{(s_0, \mathbf{x}) \in \mathcal{S}_0 \times \mathcal{X}^{l_z} : s_0 - \sigma(\mathbf{x}) \in [\![\underline{\mathcal{S}_0} - l_z\alpha, \overline{\mathcal{S}_0}]\!]\}, \tag{5.24}$$

as the condition $s_0 - \sigma(\mathbf{x}) \in [\![\underline{\mathcal{S}_0} - l_z\alpha, \overline{\mathcal{S}_0}]\!]$ is always satisfied. By (5.12) in Lemma 5.2, there exists request sequence $\mathbf{y}_z$ feasible for every input pair in $\mathcal{A}_0$.

Following a similar reasoning, consider the set of possible input pairs during the subsequent λ times steps, i.e. $\mathcal{S} \times \mathcal{X}^\lambda$. Define a cover of the input alphabet, $\mathcal{S} \times \mathcal{X}^\lambda \subseteq (\mathcal{A}_1 \bigcup \mathcal{A}_2)$, with subsets given by

$$\mathcal{A}_1 = \left\{(s_0, \mathbf{x}) \in \mathcal{S} \times \mathcal{X}^\lambda : s_0 - \sigma(\mathbf{x}) \in [\![0, \beta]\!]\right\}, \tag{5.25}$$

and

$$\mathcal{A}_2 = \left\{(s_0, \mathbf{x}) \in \mathcal{S} \times \mathcal{X}^\lambda : s_0 - \sigma(\mathbf{x}) \in [\![-\lambda\alpha, -1]\!]\right\}, \tag{5.26}$$

where the union $\mathcal{A}_1 \bigcup \mathcal{A}_2$ contains all sequences in $\mathcal{S} \times \mathcal{X}^\lambda$ as $s_0 - \sigma(\mathbf{x}) \in [\![-\lambda\alpha, -1]\!] \cup [\![0, \beta]\!]$. By Lemma 5.2, there exists a shared feasible sequence $\mathbf{y}_i$ for all pairs in $\mathcal{A}_i$ with $i = 1, 2$. The above argument can be recursively applied $\kappa = \lfloor (n - l_z)/\lambda \rfloor^+$ times.

For the remaining $l_r = l - l_z - \kappa\lambda$ time steps, define a cover of the input alphabet, $\mathcal{S} \times \mathcal{X}^{l_r} \subseteq \bigcup_i \mathcal{A}'_i$, with subsets:

$$\mathcal{A}'_i = \left\{ (s_0, \mathbf{x}) \in \mathcal{S} \times \mathcal{X}^{l_r} : s_0 - \sigma(\mathbf{x}) \in \mathcal{Z}_i \right\}, \tag{5.27}$$

and the $\mathcal{Z}_i$ intervals satisfying the conditions of Lemma 5.2, having the same cardinality as $\mathcal{S}_n$, and such that $s_0 - \sigma(\mathbf{x}) \in \bigcup_{i=0}^{N-1} \mathcal{Z}_i$ for all $(s, \mathbf{x}) \in \mathcal{S} \times \mathcal{X}^{l_r}$. The required number of intervals N is thus given by

$$N \leq \left\lceil \frac{\beta + l_r\alpha}{|\mathcal{S}_n|} \right\rceil + \mathbb{1}\left\{ \frac{\beta + l_r\alpha}{|\mathcal{S}_n|} \leq 1 + 2\left\lfloor \frac{l_r\alpha}{|\mathcal{S}_n|} \right\rfloor \right\}, \tag{5.28}$$

where the additional one comes from the constraints imposed by Lemma 5.2. By Lemma 5.2, for each i there exists a sequence $\mathbf{y}'_i$ feasible for all input pairs in $\mathcal{A}_{l_r}(\mathcal{S}, \mathcal{Z}_i)$. Note that for each i, the final state of the battery for all sequences in $\mathcal{A}_{l_r}(\mathcal{S}, \mathcal{Z}_i)$ resulting from $\mathbf{y}'_i$ lies on the interval

$$s_0 - \sigma(\mathbf{x}) + \sigma(\mathbf{y}'_i) \in [\![\underline{\mathcal{Z}_i} + \sigma(\mathbf{y}'_i), \overline{\mathcal{Z}_i} + \sigma(\mathbf{y}'_i)]\!], \tag{5.29}$$

adding $\underline{\mathcal{Z}_i} + \sigma(\mathbf{y}'_i) - \underline{\mathcal{S}_n}$ to the last time step of the request sequence $(\mathbf{y}'_i)_{l_r}$ yields

$$[\![\underline{\mathcal{Z}_i} + \sigma(\mathbf{y}'_i), \overline{\mathcal{Z}_i} + \sigma(\mathbf{y}'_i)]\!] = \mathcal{S}_n, \tag{5.30}$$

while preserving the feasibility constraint, as only the final state is modified. The proof is completed by setting

$$\mathcal{V}^n(\mathcal{S}_0, \mathcal{S}_n) = \{\mathbf{y}_z\} \times \underbrace{\{\mathbf{y}_1, \mathbf{y}_2\} \times \ldots \times \{\mathbf{y}_1, \mathbf{y}_2\}}_{\kappa} \times \{\mathbf{y}'_1, \cdots, \mathbf{y}'_N\}. \tag{5.31}$$

■

The construction of the set of request sequences given by (5.31) shows the forecasting capabilities required to implement optimal battery policies. Note that in order to map input pairs $(s_0, \mathbf{x})$ to energy requests in $\mathcal{V}^n(\mathcal{S}_0, \mathcal{S}_n)$ it suffices to forecast, at the start of each block of length λ, whether the battery will deplete during the current block, i.e. $s_0 - \sigma(\mathbf{x}^\lambda) \lesseqgtr 0$. Moreover, in order to achieve the minimum feasible price, the EMU requires knowledge of the total amount of energy consumed during the last $l_r \leq \lambda$ time steps of each market block. The following theorem shows the previous result is tight.

Theorem 5.2. *Let $\mathcal{S}_0 = \{s_0\}$ and $\mathcal{S}_n = [\![\underline{\mathcal{S}_n}, \overline{\mathcal{S}_n}]\!]$ denote the possible states of the battery at times 0 and n. Then there exists a set $\mathcal{W}^n(\mathcal{S}_0, \mathcal{S}_n) \subseteq \mathcal{X}^n$ with cardinality*

$$\log|\mathcal{W}^n(\mathcal{S}_0, \mathcal{S}_n)| \geq \left\lfloor \frac{n - \lceil\lambda\rceil}{\lceil\lambda\rceil} \right\rfloor^+ + \log \left\lceil \frac{n \bmod \lceil\lambda\rceil + \lceil\lambda\rceil \mathbb{1}\{n \geq \lceil\lambda\rceil\}}{|\mathcal{S}_n|\alpha^{-1}} \right\rceil, \quad (5.32)$$

for $\lambda = (\beta+1)/\alpha$. Moreover, no sequence $\mathbf{y} \in \mathbb{Z}^n$ satisfies $s_n \in \mathcal{S}_n$ and is feasible, i.e.

$$\mathcal{Y}^n(s_0, \mathbf{w}) \cap \mathcal{Y}^n(s_0', \mathbf{w}') = \emptyset, \quad (5.33)$$

for any distinct $(s_0, \mathbf{w})$, $(\hat{s}_0, \hat{\mathbf{w}})$ in $\mathcal{W}^n(\mathcal{S}_0, \mathcal{S}_n)$.

Proof: Consider the first $\lceil\lambda\rceil$ time steps. Let the input alphabet during those $\lceil\lambda\rceil$ time steps be defined by $\mathbf{w}_1, \mathbf{w}_2 \in \mathcal{X}^{\lceil\lambda\rceil}$, with $\mathbf{w}_1$ and $\mathbf{w}_2$ such that $\sigma(\mathbf{w}_1) = 0$ and $\sigma(\mathbf{w}_2) \geq \lceil\lambda\rceil\,\alpha$. This implies that

$$d\Big((s_0, \mathbf{w}_1), (s_0, \mathbf{w}_2)\Big) = |\sigma(\mathbf{w}_1) - \sigma(\mathbf{w}_2)| = \lceil\lambda\rceil\,\alpha > \beta. \quad (5.34)$$

Therefore, by Lemma 5.1, no output sequence is shared between $(s_0, \mathbf{w}_1)$ and $(s_0, \mathbf{w}_2)$, i.e

$$\mathcal{Y}^{\lceil\lambda\rceil}(s_0, \mathbf{w}_1) \cap \mathcal{Y}^{\lceil\lambda\rceil}(s_0, \mathbf{w}_2) = \emptyset. \quad (5.35)$$

In view of this, the input sequence $\mathbf{w}_{\mathbf{y}} \in \{\mathbf{w}_1, \mathbf{w}_2\}$, and the initial battery state of the second block $s_{\lceil\lambda\rceil} = s_0 - \sigma(\mathbf{w}_{\mathbf{y}}) + \sigma(\mathbf{y})$ are uniquely determined by the output sequence $\mathbf{y}$. The argument above can be applied recursively for the subsequent $\hat{\kappa} = \lfloor (n - \lceil\lambda\rceil)/\lceil\lambda\rceil \rfloor^+$ blocks.

Following a similar reasoning, let the alphabet defining the remaining $l_r = n - \hat{\kappa}\lceil\lambda\rceil$ time steps be given by $\{\mathbf{w}_1', \mathbf{w}_2', ..., \mathbf{w}_N'\} \subseteq \mathcal{X}^{l_r}$ with $\mathbf{w}_i$ such that $\sigma(\mathbf{w}_i) = i|\mathcal{S}_n|$ and

$$N = \left\lfloor \frac{l_r \alpha}{|\mathcal{S}_n|} \right\rfloor = \left\lceil \frac{n \bmod \lceil\lambda\rceil + \lceil\lambda\rceil \mathbb{1}\{n \geq \lceil\lambda\rceil\}}{|\mathcal{S}_n|\alpha^{-1}} \right\rceil. \quad (5.36)$$

Consequently, for any given $\mathbf{y}$, only one sequence $\mathbf{w}_i'$ satisfies the constraint $s_n = s_{\hat{\kappa}\lceil\lambda\rceil} - \sigma(\mathbf{w}_i') + \sigma(\mathbf{y}) \in \mathcal{S}_n$ simultaneously. Setting

$$\mathcal{W}^n(\mathcal{S}_0, \mathcal{S}_n) = \underbrace{\{\mathbf{w}_1, \mathbf{w}_2\} \times ... \times \{\mathbf{w}_1, \mathbf{w}_2\}}_{\kappa} \times \{\mathbf{w}_1', \cdots, \mathbf{w}_N'\}. \quad (5.37)$$

completes the proof. ■

5.2 Single-letter universal privacy bounds under cost constraints

In the following, we bound the information leakage given in Definition 3.1. We first recover the results presented in Theorem 3.3 and Theorem 3.4, i.e. when $\Delta = \infty$, using the tools developed in this chapter. Moreover, we provide tight, single letter characterizations for the case in which no extra money is allocated for privacy purposes, i.e. $\Delta = 0$. In Section 5.2.1 and Section 5.2.2, we provide single letter upper and lower bounds on the information leakage for arbitrary values of Δ.

Theorem 5.3. *The privacy cost function under no energy bill constraints $\mathcal{I}(\infty)$ is bounded by*

$$\frac{1}{n}\left\lfloor \frac{n}{\lceil \lambda \rceil} \right\rfloor \leq \mathcal{I}(\infty) \leq \frac{1}{n}\left\lfloor \frac{n}{\lambda} \right\rfloor, \tag{5.38}$$

where $\lambda = (\beta + 1)/\alpha$.

Proof: Upper bound. Theorem 5.1 shows the existence of the set $\mathcal{V}^n(\{s_0\}, \mathcal{S}_n)$ with cardinality bounded by

$$\log |\mathcal{V}^n(\{s_0\}, \mathcal{S}_n)| \leq \left\lfloor \frac{n-\lambda}{\lambda} \right\rfloor^+ + \log \left\lceil \frac{n \bmod \lambda + \lambda \mathbb{1}\{n \geq \lambda\}}{\lambda} \right\rceil = \left\lfloor \frac{n}{\lambda} \right\rfloor, \tag{5.39}$$

such that the intersection $\mathcal{V}^n(\{s_0\}) \cap \mathcal{Y}^n(s_0, \mathbf{x})$ is not empty for every input pair $(s_0, \mathbf{x})$. Letting the output Y^n take values in $\mathcal{V}^n(\{s_0\}) \cap \mathcal{Y}^n(s_0, \mathbf{x})$ completes the proof.

Lower bound. Theorem 5.2 states the existence of the set $\mathcal{W}^n = \mathcal{W}^n(\{s_0\}, \mathcal{S})$ with cardinality bounded by

$$\log |\mathcal{W}| \geq \left\lfloor \frac{n-\lceil \lambda \rceil}{\lceil \lambda \rceil} \right\rfloor^+ + \log \left\lceil \frac{n \bmod \lceil \lambda \rceil + \lceil \lambda \rceil \, \mathbb{1}\{n \geq \lambda\}}{\lceil \lambda \rceil} \right\rceil = \left\lfloor \frac{n}{\lceil \lambda \rceil} \right\rfloor, \tag{5.40}$$

such that no two sequences in $\mathcal{W}^n$ share a common output sequence, i.e. $H(W^n|Y^n) = 0$. Letting W^n take uniformly distributed values over $\mathcal{W}^n$ completes the proof. ■

For integer values of λ, lower an upper bounds on Lemma 5.3 coincide, providing the exact value of the information leakage $n\mathcal{I}(\infty) = \lfloor n/\lambda \rfloor$. Consequently, the step behaviour of the privacy guarantee when n increases, is not a peculiarity introduced by the tools used in this paper, but the real behaviour of the system.

Theorem 5.4. *The privacy cost function when no deviation from the minimum feasible price is allowed $\mathcal{I}(0)$, is bounded by*

$$\frac{1}{n}\sum_{k=0}^{K-1}\left(\left\lfloor\frac{l_k-\lceil\lambda\rceil}{\lceil\lambda\rceil}\right\rfloor^{+}+\log\frac{l_k \bmod \lceil\lambda\rceil+\lceil\lambda\rceil\,\mathbb{1}\{(l_k\geq\lceil\lambda\rceil\}}{\alpha^{-1}}\right) \tag{5.41}$$

$$\leq\mathcal{I}(0)\leq\frac{1}{n}\sum_{k=0}^{K-1}\left(\left\lfloor\frac{l_k-\lambda}{\lambda}\right\rfloor^{+}+\log\frac{l_k \bmod \lambda+\lambda\mathbb{1}\{l_k\geq\lambda\}}{\alpha^{-1}}\right), \tag{5.42}$$

where $\lambda=(\beta+1)/\alpha$.

Proof: Upper bound. Let T_k denote the state of the battery at the k-th market transition point, and let T_k^* denote the value of T_k that achieves the minimum cost. Theorem 5.1 shows the existence of a set $\mathcal{V}_k=\mathcal{V}_{l_k}(\{T_{k-1}^*\},\{T_k^*\})$ with cardinality bounded by

$$\log|\mathcal{V}_k|\leq\left\lfloor\frac{l_k-\lambda}{\lambda}\right\rfloor^{+}+\log\left\lceil\frac{l_k \bmod \lambda+\lambda\mathbb{1}\{l_k\geq\lambda\}}{\alpha^{-1}}\right\rceil \tag{5.43}$$

such that at least one sequence $\mathbf{v}\in\mathcal{V}_k$ is feasible, i.e. $\mathcal{V}_k\cap\mathcal{Y}^{l_k}(T_{k-1},\mathbf{x}^{l_k})\neq\emptyset$, and guarantees that $T_k=T_k^*$. Letting the output Y^n take values in

$$\mathcal{V}_0\times\mathcal{V}_1\times\cdots\times\mathcal{V}_{K-1}\cap\mathcal{Y}^n(s_0,\mathbf{x}) \tag{5.44}$$

completes the proof.

Lower bound. Theorem 5.2 states the existence of the set of input sequences $\mathcal{W}_k=\mathcal{W}_{l_k}(\{T_{k-1}^*\},\{T_k^*\})$ with cardinality bounded by

$$\log|\mathcal{W}_k|\geq\left\lfloor\frac{l_k-\lceil\lambda\rceil}{\lceil\lambda\rceil}\right\rfloor^{+}+\log\left\lceil\frac{l_k \bmod \lceil\lambda\rceil+\lceil\lambda\rceil\,\mathbb{1}\{l_k\geq\lceil\lambda\rceil\}}{\alpha^{-1}}\right\rceil, \tag{5.45}$$

such that no two sequences in $\mathcal{W}_k$ share a common output sequence, i.e. $H(W_k|Y^n,W^{k-1})=0$. Letting W^n take uniformly distributed values over $\mathcal{W}_0\times\mathcal{W}_1\times\cdots\times\mathcal{W}_{K-1}$ completes the proof. ■

Similarly as with Lemma 5.3, for integer values of λ, lower an upper bounds on Lemma 5.4 coincide, providing the exact value of the privacy guarantee

$$\mathcal{I}(0)=\frac{1}{n}\sum_{k=0}^{K-1}\left(\left\lfloor\frac{l_k-\lambda}{\lambda}\right\rfloor^{+}+\log\frac{l_k \bmod \lambda+\lambda\mathbb{1}\{l_k\geq\lambda\}}{\alpha^{-1}}\right). \tag{5.46}$$

5.2.1 Upper bound the information leakage rate

Theorem 5.5 presents our main result for this section, where we provide a single letter bound on the information leakage for arbitrary cost constraints Δ. Following the approach of Theorem 4.1, the proof proceeds by constructing a battery policy that combines two components for every request sequence. One of the components guarantees the feasibility constraint, while the other guarantees the cost constraint. Single letter bounds, based on bounding the cardinality of the output, are then provided.

Theorem 5.5. *The privacy cost function is upper bounded by*

$$\mathcal{I}(\Delta) \leq \mathcal{I}(\infty) + \sum_{k=0}^{K} \log \left\lceil \frac{\beta+1}{\lfloor \Delta'/(K|\delta_k|) \rfloor} \right\rceil. \tag{5.47}$$

with $\Delta' = \Delta + \sum_{k=1}^{K} |\delta_k|$.

Proof: The proof follows by providing a bound on $\mathcal{I}_\Gamma(\Delta)$ presented on Theorem 4.1. In particular, for each market block k we propose the deterministic mapping $T_k \to \tilde{T}_k$:

$$\tilde{T}_k = \begin{cases} T_k \bmod d_k, & \text{when } T_k^* = 0 \\ (\beta - T_k) \bmod d_k, & \text{otherwise (i.e. when } T_k^* = \beta), \end{cases} \tag{5.48}$$

for positive integers $d_k > 0$. Therein, at any market transition point k, the extra price paid due to the deviation of $\tilde{T}_k$ from its optimal value T_k^* is upper bounded by

$$\begin{aligned} \max_{P_{T_k}} \mathbb{E}\Big[\delta_k(\tilde{T}_k - T_k^*)\Big] &= \mathbb{1}\{T_k^* = 0\} \max_{P_{T_k}} \mathbb{E}\Big[\delta_k\Big(T_k \bmod d_k\Big)\Big] \\ &\quad + \mathbb{1}\{T_k^* = \beta\} \max_{P_{T_k}} \mathbb{E}\Big[\delta_k\Big((\beta - \tilde{T}_k) \bmod d_k - \beta\Big)\Big] && (5.49) \\ &= \mathbb{1}\{T_k^* = 0\}\Big(\delta_k(d_k - 1)\Big) + \mathbb{1}\{T_k^* = \beta\}\Big(-\delta_k(d_k - 1)\Big) && (5.50) \\ &= \Big|\delta_k\Big|(d_k - 1), && (5.51) \end{aligned}$$

Thus, the cost constraint is bounded by

$$\mathbb{E}\left[\sum_{k=1}^{K} \delta_k(\tilde{T}_k - T_k^*)\right] \leq \sum_{k=1}^{K} |\delta_k|(d_k - 1), \tag{5.52}$$

while the information leakage is bounded by

$$I(\tilde{T}^K - T^K; T^K) \leq \sum_{k=1}^{K} \log |\tilde{T}_k - T_k| \leq \sum_{k=1}^{K} \log \left\lceil \frac{\beta+1}{d_k} \right\rceil. \tag{5.53}$$

Thus, the information leakage is upper bounded by

$$\min_{\mathbf{d}} \sum_{k=1}^{K} \log \left\lceil \frac{\beta+1}{d_k} \right\rceil \qquad \text{s.t.} \qquad \sum_{k=1}^{K} |\delta_k|(d_k-1) \leq \Delta. \tag{5.54}$$

The approximate integer solution $d_k = \lfloor \Delta'/(K|\delta_k|) \rfloor$, with $\Delta' = \Delta + \sum_{k=1}^{K} |\delta_k|$, obtained by Lagrangian optimization of (5.54), satisfies the cost constraint

$$\sum_{k=0}^{K} |\delta_k|(d_k-1) = \sum_{k=0}^{K} |\delta_k| \left(\left\lfloor \frac{\Delta'}{K|\delta_k|} \right\rfloor - 1 \right) \tag{5.55}$$

$$\leq \sum_{k=0}^{K} |\delta_k| \left(\frac{\Delta'/K}{|\delta_k|} - 1 \right) = \Delta. \tag{5.56}$$

Moreover, the resulting information leakage is given by

$$I(\tilde{T}^K - T^K; T^K) \leq \sum_{k=0}^{K} \log \left\lceil \frac{(\beta+1)}{\lfloor \Delta'/(K|\delta_k|) \rfloor} \right\rceil. \tag{5.57}$$

This completes the proof. ■

5.2.2 Tightness of the upper bound

We now provide a lower bound on the privacy cost function $\mathcal{I}(\Delta)$. The lower bound is derived by constructing an energy consumption profile P_{W^n} that achieves the lower bound for any feasible policy in $\Gamma(\Delta)$.

Theorem 5.6. *The privacy guarantee $\mathcal{I}(\Delta)$ as is lower bounded by*

$$\mathcal{I}(\Delta) \geq \frac{1}{\lceil \lambda \rceil} + \frac{1}{n} \sum_{k=0}^{K-1} \log \left(\frac{(\beta+1)}{\Delta'/(K|\delta_k|)} \right) + \gamma, \tag{5.58}$$

with $\Delta' = \Delta + \sum_{k=1}^{K} |\delta_k|$ and

$$\gamma = \frac{1}{n} \sum_{k=0}^{K-1} \Big(\log c_k - c_k - 1/\ln(2) \Big), \tag{5.59}$$

where $c_k = (l_k \bmod \lceil \lambda \rceil)/\lceil \lambda \rceil + \mathbb{1}\{l_k \geq \lceil \lambda \rceil\}$.

Proof: Let the random process W^n take uniformly distributed values over $\mathcal{W}_0 \times \mathcal{W}_1 \times \cdots \times \mathcal{W}_{K-1}$, with $\mathcal{W}_k = \mathcal{W}^{l_k}(\{T^*_{k-1}\}, \{T^*_k\})$, as defined in Lemma 5.4 and Theorem 5.2. Thus

$$H(W^n) = \sum_{k=0}^{K-1} \left(\left\lfloor \frac{l_k - \lceil \lambda \rceil}{\lceil \lambda \rceil} \right\rfloor^+ + \log \frac{l_k \bmod \lceil \lambda \rceil + \lceil \lambda \rceil \, \mathbb{1}\{l_k \geq \lceil \lambda \rceil\}}{\alpha^{-1}} \right), \tag{5.60}$$

while

$$H(W^n|Y^n) = H(T^K|Y^n) - H(T^K|W^n, Y^n) + H(W^n|Y^n, T^K) \leq H(T^K), \quad (5.61)$$

as $H(W^n|Y^n, T^K)$ by construction. The maximization of $H(T^K)$ over $\Gamma(\Delta)$ can be relaxed to maximizing

$$\max_{\Gamma(\Delta)} H(T^K) \leq \max_{\Gamma(\Delta)} \sum_{k=0}^{K-1} H(T_k) \leq \max_{\Pi(\Delta)} \sum_{k=0}^{K-1} H(T_k^\Delta), \quad (5.62)$$

with

$$\Pi(\Delta) \triangleq \left\{ P_{Y^n|X^n} : \mathbb{E}\left[\sum_{k=1}^{K} |\delta_k| T_k^\Delta\right] \leq \Delta \text{ and } T_k^\Delta \geq 0 \right\}, \quad (5.63)$$

where we now allow $T_k^\Delta = |T_k - T_k^*|$ to expand outside $[\![0, \beta]\!]$ into $\mathbb{Z}_+^K$, taking values over the positive integers, while keeping the cost constraints. The distribution that maximizes the entropy of $T_k^\Delta \geq 0$ for a given expected value $\mu_k = \mathbb{E}[T_k^\Delta]$ is the geometric distribution, i.e.

$$H(T_k^\Delta) = \log(1+\mu_k) + \log\left(1 + \frac{1}{\mu_k}\right)^{\mu_k} \leq \log(1+\mu_k) + 1/\ln(2). \quad (5.64)$$

Constructing the Lagrangian multiplier, we derive that the maximum is achieved for $\mu_k = \Delta'/(K|\delta_k|) - 1$. Therefore

$$\max_{\Pi(\Delta)} \sum_{k=0}^{K-1} H(T_k^\Delta) = \sum_{k=0}^{K-1} \left(\log \frac{\Delta'}{K|\delta_k|} + 1/\ln(2) \right). \quad (5.65)$$

This yields

$$\begin{aligned}
n\mathcal{I}(\Delta) &\geq H(W^n) - \max_{\Pi(\Delta)} \sum_{k=0}^{K-1} H(T_k^\Delta) && (5.66)\\
&= \sum_{k=0}^{K-1} \left(\left\lfloor \frac{l_k - \lceil\lambda\rceil}{\lceil\lambda\rceil} \right\rfloor^+ + \log \frac{l_k \bmod \lceil\lambda\rceil + \lceil\lambda\rceil \mathbb{1}\{l_k \geq \lceil\lambda\rceil\}}{\alpha^{-1}} - \log \frac{\Delta'}{K|\delta_k|} - \frac{1}{\ln(2)} \right) && (5.67)\\
&= \sum_{k=0}^{K-1} \left(\frac{l_k}{\lceil\lambda\rceil} + \log \frac{(\beta+1)}{\Delta'/(K|\delta_k|)} + \gamma_k \right) && (5.68)\\
&= \frac{n}{\lceil\lambda\rceil} + \sum_{k=0}^{K-1} \log \frac{(\beta+1)}{\Delta'/(K|\delta_k|)} + \sum_{k=0}^{K-1} \gamma_k, && (5.69)
\end{aligned}$$

where

$$\gamma_k = \left\lfloor \frac{l_k - \lceil\lambda\rceil}{\lceil\lambda\rceil} \right\rfloor^+ - \frac{l_k}{\lceil\lambda\rceil} + \log \frac{l_k \bmod \lceil\lambda\rceil + \lceil\lambda\rceil \mathbb{1}\{l_k \geq \lceil\lambda\rceil\}}{(\beta+1)\alpha^{-1}} - \frac{1}{\ln(2)}, \quad (5.70)$$

and therefore, noting that

$$\frac{l_k}{\lceil\lambda\rceil} - \left\lfloor \frac{l_k}{\lceil\lambda\rceil} \right\rfloor = \frac{l_k \bmod \lceil\lambda\rceil}{\lceil\lambda\rceil}, \tag{5.71}$$

and that $\log(x/\lambda) \geq \log(x/\lceil\lambda\rceil)$ for every $x \in \mathbb{R}$, we have that

$$\gamma_k \geq \log c_k - c_k - 1/\ln(2), \tag{5.72}$$

with

$$c_k = \frac{l_k \bmod \lceil\lambda\rceil + \lceil\lambda\rceil \, \mathbb{1}\{l_k \geq \lceil\lambda\rceil\}}{\lceil\lambda\rceil}. \tag{5.73}$$

This completes the proof. ■

Note that the gap between the single letter upper and lower bounds presented on theorems 5.5 and 5.6 is given by

$$G = \mathcal{I}(\infty) + \frac{1}{n}\sum_{k=0}^{K} \log\left\lceil \frac{\beta+1}{\lfloor \Delta'/(K|\delta_k|)\rfloor} \right\rceil - \left(\frac{1}{\lceil\lambda\rceil} + \frac{1}{n}\sum_{k=0}^{K-1} \log\left(\frac{(\beta+1)|}{\Delta'/(K|\delta_k|)} \right) + \gamma \right), \tag{5.74}$$

with $\Delta' = \Delta + \sum_{k=1}^{K} |\delta_k|$, $\gamma = \frac{1}{n}\sum_{k=0}^{K-1} \Big(\log c_k - c_k - 1/\ln(2) \Big)$, and $c_k = (l_k \bmod \lceil\lambda\rceil)/\lceil\lambda\rceil + \mathbb{1}\{l_k \geq \lceil\lambda\rceil\}$. Recall that by Theorem 4.5, we have that

$$\mathcal{I}(\infty) \leq \frac{1}{n}\left\lceil \frac{n}{\lambda} \right\rceil. \tag{5.75}$$

Therefore

$$G = \frac{1}{n}\left(\left\lceil \frac{n}{\lambda} \right\rceil - \frac{n}{\lceil\lambda\rceil}\right) + \frac{1}{n}\sum_{k=0}^{K} \log\left(\left\lceil \frac{\beta+1}{\lfloor \Delta'/(K|\delta_k|)\rfloor} \right\rceil \frac{\Delta'/(K|\delta_k|)}{(\beta+1)} \right) - \gamma. \tag{5.76}$$

Moreover, note that $\log(c_k) - c_k \geq -1$ for $c_k \in (1,2)$, and therefore

$$\gamma_k \geq \begin{cases} -1 - \frac{1}{\ln(2)} \approx -2.44, & \text{when } l_k \geq \lceil\lambda\rceil \\ -\frac{l_k \bmod \lceil\lambda\rceil}{\lceil\lambda\rceil} + \log \frac{l_k \bmod \lceil\lambda\rceil}{\lceil\lambda\rceil} - \frac{1}{\ln(2)} & \text{otherwise.} \end{cases} \tag{5.77}$$

This implies that, for slow changing markets, i.e. $l_k \geq \lceil\lambda\rceil$, the gap between upper and lower bound increases by $\approx 2.4/n$ bits per market change, plus the rounding errors described in (5.76).

5.3 Numerical results

In this section, we numerically assess the upper bounds on the privacy cost derived in this chapter. We model the market price after the UK Economy 7 tariff, where users are charged an off-peak price of 0.071 £/kWh within a 7 hour block and a peak price of 0.152 £/kWh otherwise [81]. We assume the user has a Tesla Powerwall battery with a capacity of 13.5kWh and a peak power of 5kW. We let the maximum power consumption of the user be $4.5kW$ [65]. The SM sends the UP integrated energy readings every 15 minutes following EU specifications for SMs [65]. Thus, we set the time elapsed between time steps i and $i+1$ to 15min. These yields the following values in our system model: market lengths $l_0 = 7\text{h} \times 4\,\text{samples/hour} = 28$ and $l_1 = 17\text{h} \times 4\,\text{samples/hour} = 68$; corresponding market prices of $m_0 = 0.152£/\text{kWh} \times 4\,\text{samples/hour} = 0.38\,\text{cents per sample per watt}$; and $m_1 = 0.071£/\text{kWh} \times 4\,\text{samples/hour} = 0.1775\,\text{cents per sample per watt}$; maximum consumption between time steps $\alpha = 4500\,\text{W}/4\,\text{samples/hour} = 1125$; battery capacity $\beta = 13499$.

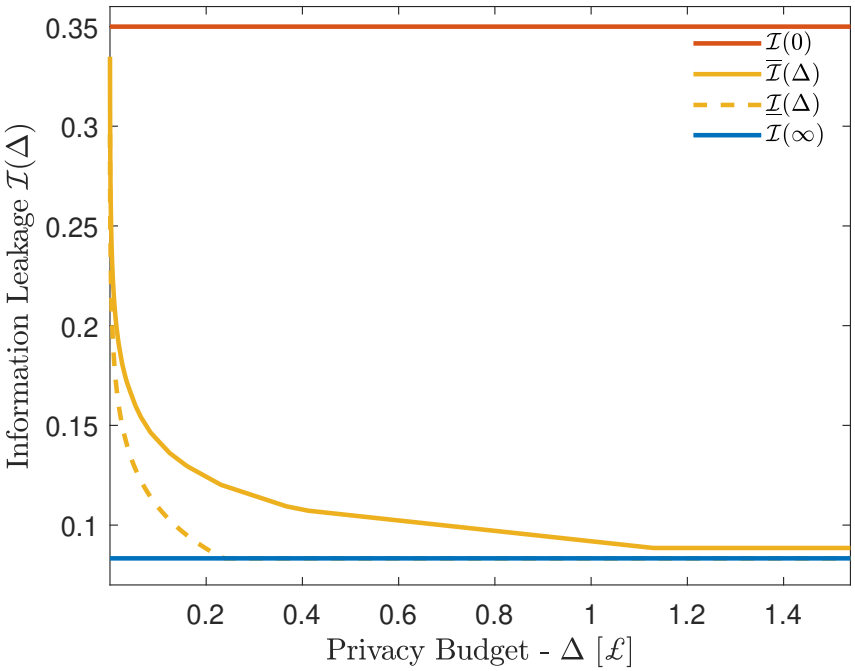

Figure 5.2. Information leakage $\mathcal{I}(\Delta)$ against privacy budget with $\underline{\mathcal{I}}$, $\mathcal{I}$ and $\bar{\mathcal{I}}$ denoting lower bound, exact value, and upper bound respectively.

Figure 5.2 shows the impact of the privacy budget Δ. It can be seen how the information leakage decreases logarithmically with the privacy budget. Figure 5.3 shows the impact of the maximum energy consumption α on the information leakage. It can be seen how the privacy guarantee increases linearly with the α. The bounds for different privacy budget are parallel to each other. Figure 5.4 shows the impact of the number of market changes K on the privacy guarantee. Therein the total length

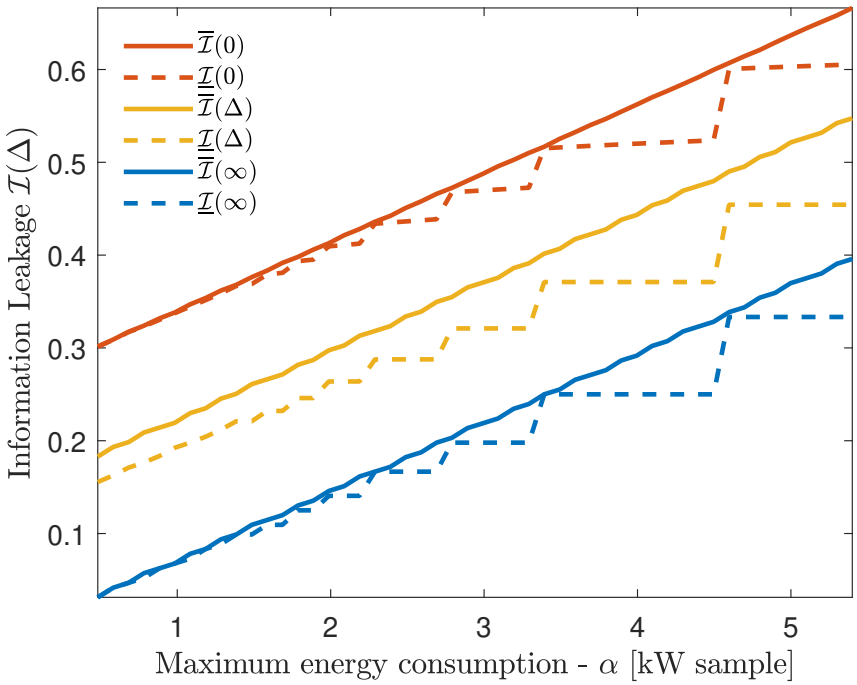

Figure 5.3. Information leakage $\mathcal{I}(\Delta)$ against maximum energy consumption α with $\underline{\mathcal{I}}$ and $\bar{\mathcal{I}}$ denoting lower bound and upper bound respectively.

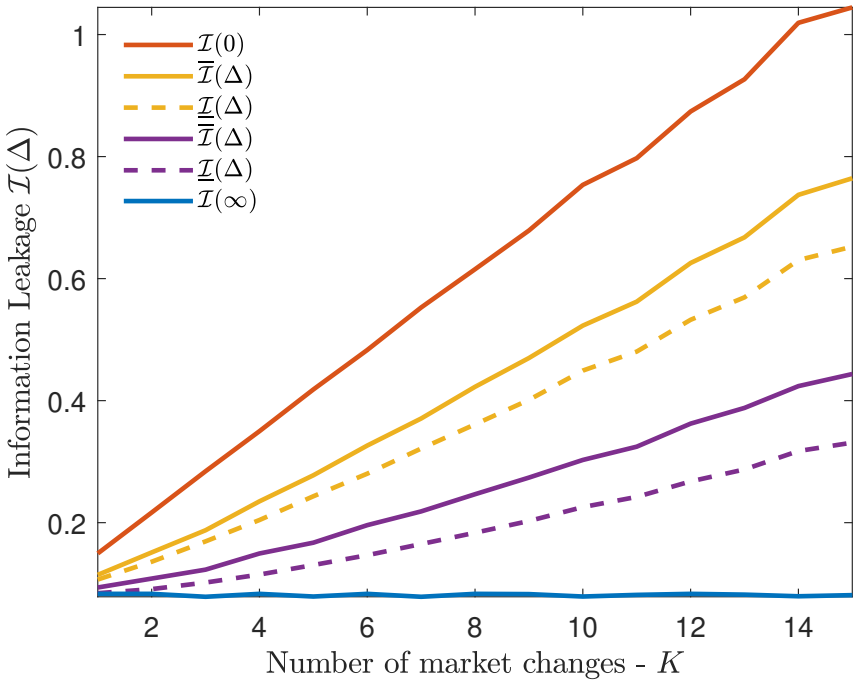

Figure 5.4. Information leakage against number of market changes K for a fixed total length n with $\underline{\mathcal{I}}$, $\mathcal{I}$ and $\bar{\mathcal{I}}$ denoting lower bound, exact value, and upper bound respectively.

of the sequence n is preserved as K increases. It can be seen that the information leakage increases linearly with the number the market changes and interestingly, the privacy budget defines the slope of the increase. It is also interesting to note

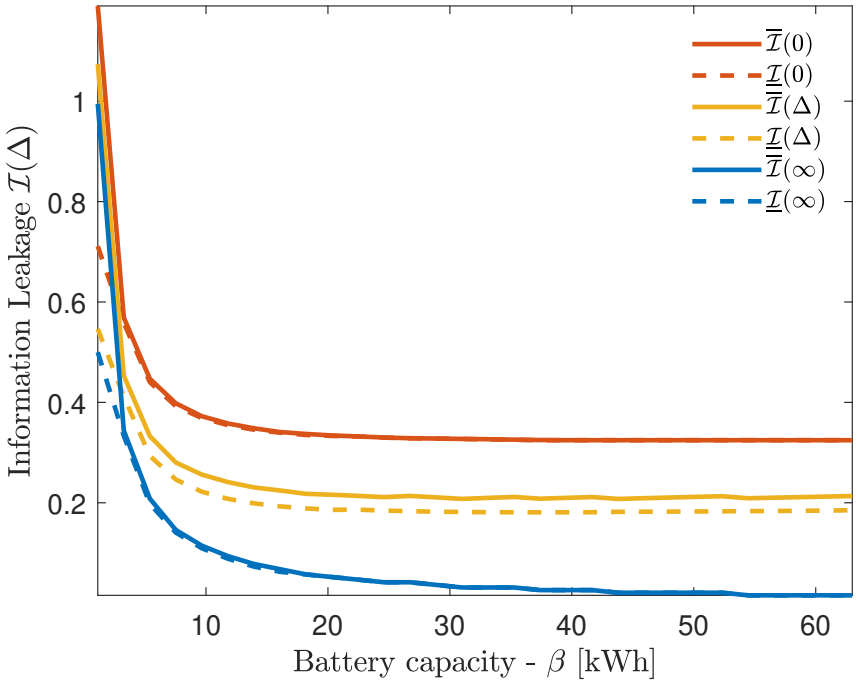

Figure 5.5. Information leakage $\mathcal{I}(\Delta)$ against battery capacity β with $\underline{\mathcal{I}}$ and $\bar{\mathcal{I}}$ denoting lower bound and upper bound respectively.

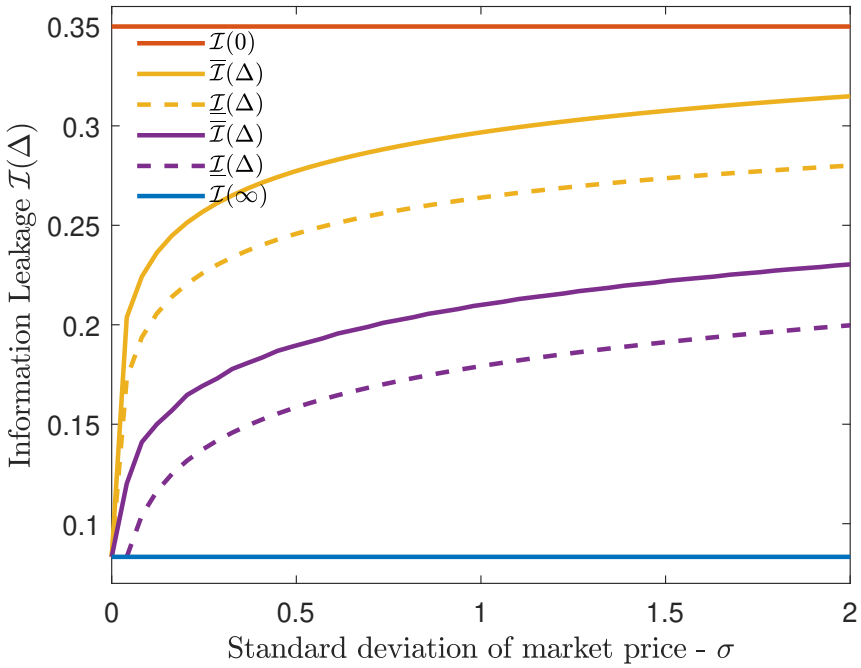

Figure 5.6. Information leakage $\mathcal{I}(\Delta)$ against standard deviation of the market σ with $\underline{\mathcal{I}}$, $\mathcal{I}$ and $\bar{\mathcal{I}}$ denoting lower bound, exact value, and upper bound respectively.

that the gap between upper and lower bound increases as the number of market changes increases. Figure 5.5 shows the impact of the battery capacity on the privacy guarantee. Interestingly, for small values of β the leakage is decreases inversely with the battery size. However, for larger values of β increasing the battery size does not

reduce the leakage. Thus, the privacy leakage is governed by the market component for large values of β, and by the no market component $\mathcal{I}(\infty)$ in the small β regime. Figure 5.6 shows the impact of the standard deviation of the market price prices on the privacy guarantee. Therein the UK 7 tariff is employed and the corresponding prices are modified to match the desired standard deviation σ. The information leakage increases logarithmically with σ for small values of σ, to stabilize later on the large σ regime.

5.3.1 Comparison with previous results

The diversity of modelling approaches and the novelty of this work hinders the direct comparison of the result obtained in this thesis with those available in the literature. Firstly, this thesis measures the privacy leakage in terms of mutual information, hindering the direct comparison with those works based on other metrics such as hypothesis testing and heuristic metrics. This issue is partly solved by the extension of this work to other metrics derived in Section 5.4 and by the links previously established between mutual information and other privacy metrics described in Section 2.2. Moreover, there exist a large diversity across the literature on the capabilities available to the EMU, in this work, for instance, we model a finite capacity battery, while other works consider alternative energy sources with average or maximum power generation constraints [122, 53]. Furthermore, this thesis is focused on the derivation of worst case universal privacy guarantees, a completely novel approach in the literature, while the main focus of the existing literature is on obtaining bounds valid for specific i.i.d. or Markovian energy consumption models. Finally, the battery policies proposed in this thesis allow for non-causal behaviours, i.e. they require precise forecasting of future energy consumptions, while the main body of literature focuses on causal battery policies.

However, it is possible to compare some of the results derived in this thesis with some of the previous works available in the literature. In our work, for instance, Theorem 5.3 shows that the privacy cost function under no energy bill constraints $\mathcal{I}(\infty)$ is bounded by

$$\frac{1}{n}\left\lfloor \frac{n}{\lceil \lambda \rceil} \right\rfloor \leq \mathcal{I}(\infty) \leq \frac{1}{n}\left\lfloor \frac{n}{\lambda} \right\rfloor, \tag{5.78}$$

where $\lambda = (\beta+1)/\alpha$. Figure 5.7, shows the similarity between the results obtain on Theorem 5.3 and those obtained in [57, section 4.1] and [13] for i.i.d. inputs on binary alphabets, i.e. for $\alpha = 1$. It is important to note that the bounds provided on Theorem 5.3 hold for any input distribution with offline battery policies. The numerically calculated values depicted in Figure 5.7 are based on binary i.i.d. input

loads with causal policies. For larger input and output alphabets, the non-causal approach, i.e. the available forecasting capabilities, can provide an advantage over the causal approach. Similarly, for larger alphabets, the i.i.d. assumptions can lead to a significant reduction on the achievable information leakage rate.

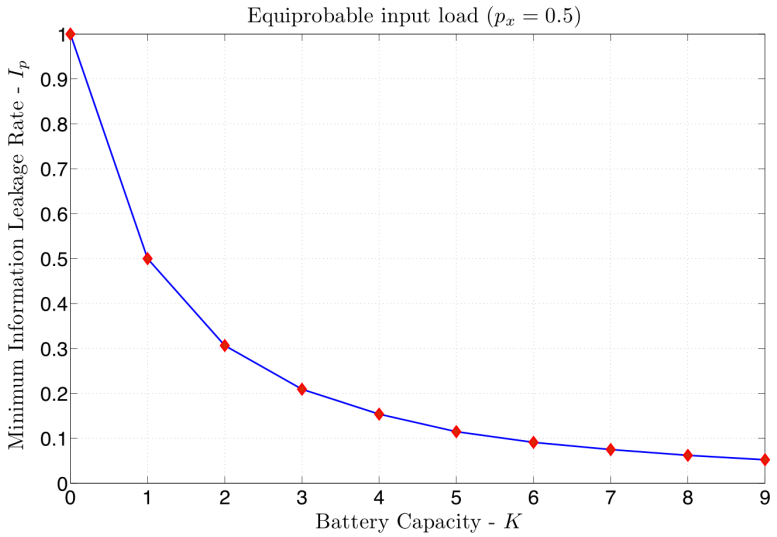

Figure 5.7. Numerical simulations for minimum information leakage rate $\mathcal{I}(\infty) = \mathcal{I}_p$, versus battery capacity $\beta = K$ for binary i.i.d. inputs, i.e. $\alpha = 1$ [13, Fig. 7].

Moreover, Figure 5.8 shows the results obtained in [14] and developed over [111–114]. Therein, it is shown that for i.i.d. binomially distributed energy consumptions, battery size $\beta = 50$, and maximum consumption $\alpha = m_x = 20$, the leakage rate equals ≈ 0.045, while Theorem 5.3 shows that for worst case consumptions models, the information leakage rate equals $\approx 1/\lambda = \alpha/(\beta + 1) = 20/51 \approx 0.39$, showing a factor of ≈ 8.67 between the worst case guarantee presented in this thesis, and the i.i.d. binomially distributed case shown in Figure 5.8. At the same time, for $\alpha = m_x = 5$, Figure 5.8 shows a leakage rate of ≈ 0.015 while Theorem 5.3 shows that for worst case consumptions models, the information leakage rate equals $\approx 1/\lambda = \alpha/(\beta + 1) = 5/51 = 0.098$, giving a factor of ≈ 6.5. This difference is due the fact that the binomial distribution does not model a worst case scenario. Moreover, the growing difference between worst and binomial i.i.d cases as α increases highlights that the worst case scenario where $X_i \in \{0, \alpha\}$ departs from the binomially distributed $X_i \sim Bi(\alpha, 0.5)$ as α increases.

5.4 Generalization to other metrics

In the following we generalize some of our previous results to other information leakage metrics. Lemma 5.3 generalizes Lemma 4.4, showing that for any distribution

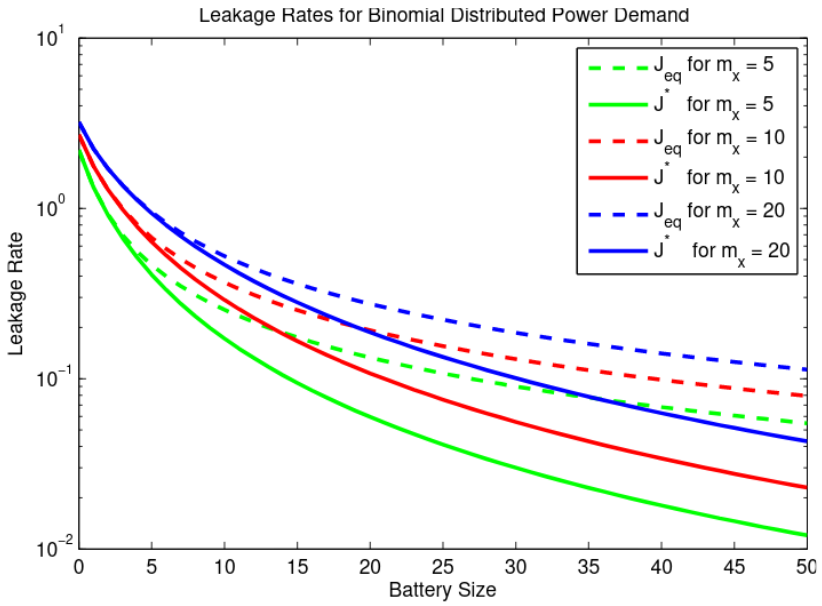

Figure 5.8. Suboptimal (dashed line) and optimal (solid line) leakage rate for i.i.d. Binomially distributed demand $X_i \sim Bi(\alpha, 0.5)$ for $\alpha = m_x = \{5, 10, 20\}$ [14].

modelling the energy consumption of the user, and for any metric satisfying the data processing inequality, constraining the output alphabet does not increase the information leakage.

Lemma 5.3. *Let the energy consumption of the user X^n be modelled by P_{X^n} taking values over $\mathcal{X}^n \subset \mathbb{R}^n$. Let $\Gamma(\Delta)$ denote the set of Δ-affordable battery policies. Let $\mathcal{L}(X^n, Y^n)$ denote any information metric satisfying the data processing inequality. Then for any $\mathcal{Y} \supset \mathcal{Y}_{\mathcal{X}} = \mathcal{X}$:*

$$\min_{\Gamma(\infty)} \mathcal{L}(X^n, Y^n) = \min_{\Gamma(\infty)} \mathcal{L}(X^n, Y^n_{\mathcal{X}}) \tag{5.79}$$

Furthermore for any $\Delta \geq 0$, $\mathcal{Y} \supset \mathcal{Y}_c = [\![-\beta/\underline{l}, \beta/\underline{l} + \alpha]\!]^n$ (as defined in Lemma 4.3) it holds that:

$$\min_{\Gamma(\Delta)} \mathcal{L}(X^n, Y^n) = \min_{\Gamma(\Delta)} \mathcal{L}(X^n, Y^n_c). \tag{5.80}$$

Proof: The proof follows by noting that Lemma 4.4 holds for any metric satisfying the data processing inequality and any input distribution. ■

The following lemma presents an ordering of different privacy metrics.

Lemma 5.4. *Let P_{X^n,Y^n} denote a joint distribution, then for any $\alpha \in [0,\infty]$ it holds that [96]:*

$$I(X^n;Y^n) \leq \mathcal{L}_\alpha(X^n \to Y^n) \tag{5.81}$$

$$= \sup_{\mathcal{P}_{X^n}} I_\alpha^S(X^n;Y^n) = \sup_{\mathcal{P}_{X^n}} I_\alpha^A(X^n;Y^n) \tag{5.82}$$

$$\leq \mathcal{L}(X^n \to Y^n) \leq \log|\operatorname{supp}(Y^n)|. \tag{5.83}$$

where I_α^S and I_α^A denote the Sibson and the Arimoto mutual information respectively, $\mathcal{P}_{X^n}$ is the set of all distributions with support on $\mathcal{X}^n$, $\mathcal{L}_\alpha(X^n \to Y^n)$ denotes the maximal α-leakage and $\mathcal{L}(X^n \to Y^n)$ denotes the maximal leakage as defined in Section 2.2.5.

The ordering established by Lemma 5.4 implies that all the lower bounds on the mutual information presented in this thesis hold for maximal leakage, maximal α-leakage, Sibson and Arimoto channel capacity (with a support constrained input distribution). Furthermore, note that the upper bounds on the mutual information derived by bounding the cardinality of the output also hold. This leads to Theorem 5.7.

Theorem 5.7. *Let $\mathcal{P}_{X^n}$ denote the family of probability distributions with support on $[\![0,\alpha]\!]^n$, let $\Gamma(\Delta)$ denote the set of Δ-affordable feasible battery policies, then*

$$\min_{\Gamma(\Delta)} \max_{\mathcal{P}_{X^n}} \mathcal{L}(X^n \to Y^n) \leq \frac{1}{n}\left\lfloor \frac{n}{\lambda} \right\rfloor + \sum_{k=0}^{K} \log\left\lceil \frac{\beta+1}{\lceil 2\Delta/(K|\delta_k|)\rceil} + 0.5 \right\rceil, \tag{5.84}$$

furthermore, for $\Delta = 0$ we have that

$$\min_{\Gamma(0)} \max_{\mathcal{P}_{X^n}} \mathcal{L}(X^n \to Y^n) \leq \frac{1}{n}\sum_{k=0}^{K-1}\left(\left\lfloor \frac{l_k - \lambda}{\lambda} \right\rfloor^+ + \log \frac{l_k \bmod \lambda + \lambda \mathbb{1}\{l_k \geq \lambda\}}{\alpha^{-1}} \right), \tag{5.85}$$

and for $\Delta = \infty$ we have that

$$\min_{\Gamma(\infty)} \max_{\mathcal{P}_{X^n}} \mathcal{L}(X^n \to Y^n) \leq \frac{1}{n}\left\lfloor \frac{n}{\lambda} \right\rfloor. \tag{5.86}$$

Proof: The proof follows by noticing that Theorems 5.4 and 5.5 provide a bound on the cardinality of the output alphabet $\mathcal{Y}^n$, and Lemma 5.4. ■

These results suggest a stronger equivalence between mutual information and maximal leakage presented on the following section.

5.4.1 Equivalence between privacy metrics

The following theorem provides an upper bound on the maximal leakage that holds for general input processes. For simplicity, all the upper bounds are proved for maximal leakage. However, we note that by Lemma 5.4 upper bounds on the maximal leakage also upper bound maximal α-leakage, Sibson and Arimoto channel capacity.

Theorem 5.8. *Let P_{X^n,Y^n} denote any joint probability distribution on finite alphabets $\mathcal{X}^n$ and $\mathcal{Y}^n$. Then for any $\epsilon > 0$ and sufficiently large n, there exist a random process Y_ϵ^n such that $\mathbb{P}[Y_\epsilon^n \neq Y^n | X^n] \leq \epsilon$, and*

$$\mathcal{L}(X^n \to Y_\epsilon^n) \leq \bar{H}(Y^n), \tag{5.87}$$

where $\bar{H}(Y^n)$ denotes the sup-entropy, i.e. the smallest real number γ such that for all $\epsilon' > 0$:

$$\lim_{n\to\infty} \mathbb{P}\left[\log \frac{1}{P_{Y^n}(Y^n)} \geq \gamma + \epsilon'\right] = 0. \tag{5.88}$$

Furthermore, when Y^n is stationary ergodic, and for any $\epsilon > 0$ and sufficiently large n, there exist a random process Y_ϵ^n such that $\mathbb{P}[Y_\epsilon^n \neq Y^n | X^n] \leq \epsilon$ and

$$\mathcal{L}(X^n \to Y_\epsilon^n) \leq H(Y^n). \tag{5.89}$$

Proof: For any $\epsilon > 0$, sufficiently large n, and any general process Y^n, [146, Theorem 3] shows the existence of a collection of M n-tuples $\{\mathbf{y}_1^n, \cdots, \mathbf{y}_M^n\}$ such that $\log M \leq \bar{H}(Y^n)$ and

$$\mathbb{P}[Y^n \notin \{\mathbf{y}_1^n, \cdots, \mathbf{y}_M^n\}] \leq \epsilon. \tag{5.90}$$

Let Y_ϵ^n equal Y^n when $Y^n \in \{\mathbf{y}_1^n, \cdots, \mathbf{y}_M^n\}$ and $\mathbf{y}_1^n$ otherwise. Thus, the cardinality bounds imply that

$$\mathcal{L}(X^n \to Y_\epsilon^n) \leq \log M \leq \bar{H}(Y^n). \tag{5.91}$$

This completes the proof of the first part of the theorem. The second part follows as for stationary ergodic sources the required number of n-tuples satisfies $\log M \leq H(Y^n)$. ■

We now show that under certain mild assumptions, all our previous results (except Theorem 4.1) generalize to maximal leakage and maximal α-leakage, the following definition is useful for that purpose.

Definition 5.2. *For any $\epsilon > 0$, the set of ϵ-feasible Δ-affordable battery policies is given by*

$$\Gamma_\epsilon(\Delta) = \left\{ P_{Y^n|X^n} : \mathbb{P}[Y^n \notin \mathcal{Y}^n(s_0, \mathbf{x})] \leq \epsilon \text{ and } g(Y^n, \mathbf{x}) \leq \Delta, \ \forall \mathbf{x} \in \mathcal{X}^n \right\}. \quad (5.92)$$

That is, we allow a non zero, arbitrarily small, probability of wasting energy or leaving the energy consumption of the user unsatisfied. Note that in real scenarios, where there exist a non zero probability of energy waste and unsatisfied demand due to other sources, such as grid, battery or appliance failure, an arbitrary small probability is acceptable.

Lemma 5.5. *For any process X^n modelling the energy consumption of the user, any $\epsilon > 0$ and sufficiently large n it holds that*

$$\min_{\Gamma_\epsilon(\Delta)} \mathcal{L}(X^n \to Y^n) \leq \min_{\Gamma(\Delta)} \bar{H}(Y^n), \quad (5.93)$$

where $\bar{H}(Y^n)$ denotes the sup-entropy rate. Furthermore, for stationary ergodic processes X^n modelling the energy consumption of the user, any $\epsilon > 0$ and sufficiently large n:

$$\min_{\Gamma_\epsilon(\Delta)} \mathcal{L}(X^n \to Y^n) \leq \min_{\hat{\Gamma}(\Delta)} H(Y^n), \quad (5.94)$$

where $\hat{\Gamma}(\Delta)$ denotes all policies in $\Gamma(\Delta)$ such that Y^n is ergodic.

Proof: Denote by $\hat{Y}^n$ the process achieving the minimum entropy in (5.94). Theorem 5.8 states the existence of a random process Y_ϵ^n such that $\mathbb{P}[Y_\epsilon^n \neq \hat{Y}^n | X^n] \leq \epsilon$ and

$$\mathcal{L}(X^n \to Y_\epsilon^n) \leq \bar{H}(\hat{Y}^n) = \min_{\Gamma(\Delta)} \bar{H}(Y^n). \quad (5.95)$$

Note that since $\hat{Y}^n \in \mathcal{Y}^n(s_0, \mathbf{x})$ it holds that:

$$\mathbb{P}[Y_\epsilon^n \notin \mathcal{Y}^n(s_0, \mathbf{x})] = \mathbb{P}[Y_\epsilon^n \neq \hat{Y}^n | X^n] \leq \epsilon. \quad (5.96)$$

Furthermore, since $g(\hat{Y}^n, \mathbf{x}) \leq \Delta$, letting $\mathbf{y}_1 = 0^n$ shows that

$$g(Y_\epsilon^n, \mathbf{x}) \leq g(\hat{Y}^n, \mathbf{x}) + g(\mathbf{y}_1, \mathbf{x}) \mathbb{P}[Y_\epsilon^n \neq \hat{Y}^n | X^n] \leq \Delta. \quad (5.97)$$

Therefore, $P_{Y_\epsilon^n|X^n} \in \Gamma_\epsilon(\Delta)$ and

$$\min_{\Gamma_\epsilon(\Delta)} \mathcal{L}(X^n \to Y^n) \leq \mathcal{L}(X^n \to Y_\epsilon^n) \leq \min_{\Gamma(\Delta)} \bar{H}(Y^n). \quad (5.98)$$

This completes the proof for general output processes. The proof for stationary ergodic processes follows by (5.89). ■

The following theorem provides an immediate consequence of this result.

Theorem 5.9. *Let the energy consumption of the user be modelled by a discrete stationary ergodic stochastic process X^n with average energy consumption μ_n, then for any $\epsilon > 0$ and sufficiently large n:*

$$\min_{\Gamma_\epsilon(\Delta)} \max_{\mathcal{P}_{X^n}} \frac{1}{n}\mathcal{L}(X^n \to Y^n) \leq H_2\left(\frac{\mu_n}{\alpha}\right)\frac{1}{\lambda} + \frac{1}{n}\sum_{k=0}^{K} \log\left\lceil \frac{\beta+1}{\lceil 2\Delta/(K|\delta_k|)\rceil} + 0.5\right\rceil. \quad (5.99)$$

for $\Delta = \infty$:

$$\min_{\Gamma_\epsilon(\infty)} \max_{\mathcal{P}_{X^n}} \frac{1}{n}\mathcal{L}(X^n \to Y^n) \leq H_2\left(\frac{\mu_n}{\alpha}\right)\frac{1}{\lambda}. \quad (5.100)$$

Proof: The proofs follows from Theorem 5.5 and Lemma 5.5. ■

Lemma 5.5 shows that the minimum achievable entropy, subject to a fidelity criteria or distortion constrain $\Gamma(\Delta)$, upper bounds the minimum achievable maximal leakage, subject to an ϵ relaxation of the fidelity criteria, i.e. $\Gamma_\epsilon(\Delta)$. The following theorem shows that under certain conditions on the fidelity criteria, the minimum achievable maximal leakage is upper bounded by the minimum achievable mutual information, under the same fidelity criteria. Since maximal leakage upper bounds mutual information, this shows that both metrics share the minimum subject to a certain class of fidelity criteria.

Theorem 5.10. *Let X^n and $d_n(\mathbf{x}, \mathbf{y})$ denote an input source and a fidelity criteria $d_n : \mathcal{X}^n \times \hat{\mathcal{X}}^n \to \mathbb{R}$, such that the rate distortion theorem holds, e.g. let X^n be stationary ergodic and d_n be a subadditive fidelity criteria with a reference letter, i.e. $\exists a : \mathbb{E}[d_1(x, a)] = \rho < \infty$ [130, Theorem 11.5.1], then*

$$\lim_{n\to\infty} \min_{\mathbb{E}\left[\frac{1}{n}d_n(X^n,Y^n)\right]\leq\Delta} \frac{1}{n}\mathcal{L}(X^n \to Y^n) = \lim_{n\to\infty} \min_{\mathbb{E}\left[\frac{1}{n}d_n(X^n,Y^n)\right]\leq\Delta} \frac{1}{n} I(X^n; Y^n). \quad (5.101)$$

Proof: Let R denote the smallest real number for which there exist a mapping

$$X^n \to M \to \hat{X}^n, \quad (5.102)$$

such that M takes values in $\{1, 2, \cdots, 2^{nR}\}$ and $\mathbb{E}\left[\frac{1}{n}d_n(X^n, \hat{X}^n)\right] \leq \Delta$. Note that by cardinality bounds and the data processing inequality

$$nR \geq \mathcal{L}(X^n \to M) \geq \mathcal{L}(X^n \to \hat{X}^n). \quad (5.103)$$

Furthermore, as $X^n \to \hat{X}^n$ satisfies the distortion constraint

$$\mathcal{L}(X^n \to \hat{X}^n) \geq \min_{\mathbb{E}\left[\frac{1}{n}d_n(X^n,Y^n)\right]\leq\Delta} \mathcal{L}(X^n \to Y^n). \tag{5.104}$$

Finally, for sources and fidelity criteria satisfying some version of the rate distortion theorem, and for large enough n, we have that

$$\min_{\mathbb{E}\left[\frac{1}{n}d_n(X^n,Y^n)\right]\leq\Delta} \frac{1}{n} I(X^n;Y^n) \geq R \geq \min_{\mathbb{E}\left[\frac{1}{n}d_n(X^n,Y^n)\right]\leq\Delta} \frac{1}{n}\mathcal{L}(X^n \to Y^n). \tag{5.105}$$

Noticing that the maximal leakage upper bounds the mutual information shows that the above chain of inequalities holds with equality. ■

This shows that although the mutual information and the maximal leakage are not generally equal, their infima subject to certain class of fidelity criteria are equal. Furthermore, by Lemma 5.4 this also implies the that maximal α-leakage, Sibson and Arimoto capacity for any $\alpha \in [0,\infty]$ share the same infima when the source and fidelity criteria satisfy the aforementioned constraints. This provides a new method to characterize rate distortion functions a hard problem in general. In the privacy optimization context, and for sources and fidelity criteria satisfying the rate distortion theorem, this settles the discussion about whether mutual information or maximal leakage should be employed as a privacy measure. Therein, this equivalence gives a new operational meaning to mutual information minimization, and a more fundamental, well studied framework for maximal leakage minimization. Finally, the following lemma provides a simple example of this.

Lemma 5.6. *Let X^n denote a bounded, asymptotically mean stationary (AMS) process modelling the energy consumption of the user. Let the EMU have access to an AES with average power constraint $\bar{P}$, then*

$$\lim_{n\to\infty} \min_{\Pi(\bar{P})} \frac{1}{n}\mathcal{L}(X^n \to Y^n) = \lim_{n\to\infty} \min_{\Pi(\bar{P})} \frac{1}{n} I(X^n;Y^n), \tag{5.106}$$

where

$$\Pi(\bar{P}) = \left\{ P_{Y^n|X^n} : \mathbb{E}\left[\frac{1}{n}\sum_{i=0}^{n-1}(X_i - Y_i)\right] \leq \bar{P} \text{ and } Y_i \leq X_i \text{ for all } i \right\}. \tag{5.107}$$

Proof: By [130] the rate distortion equals the Shannon's rate distortion for AMS sources X^n and additive fidelity criteria with a reference letter, i.e. $\exists a : \mathbb{E}[d_1(x,a)] = \rho < \infty$. Thus, by Theorem 5.10 it suffices to show that $\Pi(\bar{P})$ is an additive fidelity

criteria with a reference letter. Let the fidelity criteria be defined by

$$d_1(x,y) = \begin{cases} x-y, & \text{when } y \leq x \\ \infty & \text{otherwise .} \end{cases} \tag{5.108}$$

Let the n-th extension have the additive formulation

$$d_n(\mathbf{x},\mathbf{y}) = \frac{1}{n}\sum_{i=0}^{n-1} d_1(x_i,y_i). \tag{5.109}$$

Finally, for bounded consumptions, the reference letter $a=0$, yields the expected distortion $\mathbb{E}[d_1(x,a)] = \mathbb{E}[x] < \infty$. This completes the proof. ■

5.5 Conclusions

In this chapter, we have presented the necessary and sufficient conditions of the existence of shared feasible request. This enabled the characterization of the cardinalities of the minimal covering and packing sets, which yield the single letter characterization of the information leakage when no privacy budget is available. Furthermore, we have introduced structural simplifications to the battery policy and the worst case consumption presented in Chapter 4. The proposed battery policy and worst case consumption provide single letter upper and lower bounds on the information leakage rate for arbitrary privacy budgets. This showed the impact of the different parameters in the privacy-cost function, providing interesting operational insights to the smart meter privacy problem.

Finally, in Section 5.4 we have showed that the results derived in this thesis hold for maximal leakage, maximal α-leakage, and Sibson and Arimoto channel capacities. Moreover, we have shown that, under certain conditions, the aforementioned privacy metrics are equivalent.

Chapter 6

Conclusions and future work

6.1 Conclusions

In Chapter 1, we motivated this work by describing the new challenges faced by the electricity grid and highlighting the importance of advanced sensing and communication infrastructure, and in particular of smart meters, to tackle this new challenges. Subsequently, we noted the growing privacy concerns raised by recent privacy scandals and gave a few examples on how the data collected by smart meters can be used to infer private information about the users. Thus, emphasizing the need to develop privacy preserving mechanisms within the smart grid.

In Chapter 2, we reviewed the existing literature on smart meter privacy when access to energy storage or energy harvesting devices is available. We noticed that the simplicity and tractability of i.i.d. consumption models has captured the main focus in the literature, with some studies focusing on numerical solutions for Markovian energy consumptions. However, in privacy and security settings one is typically interested in the worst-case performance. This interest hinges on the need to provide guarantees that hold for every user and is captured in the definition of privacy metrics such as differential privacy or maximal leakage. For these reasons, it is important to characterize privacy guarantees that hold for a wide range of random processes that capture the diversity of energy consumption patterns and user profiles.

In Chapter 3, we focused on providing universal privacy guarantees that hold for a wide class of energy consumption models. However, we noticed that this model greatly limits the tractability of the problem. In particular, we discussed the difficulties posed by the memory introduced by the battery and the generality of the input process. These difficulties hinder the characterization of mutual information bounds as they impede the utilization of many of the tools typically used in both information theory and ergodic theory. To tackle these difficulties we noted that battery policies

are subject to constraints on their support, but they are not subject to constraints on the probability assigned to each realization. This idea inspired us to model battery policies from a combinatorial perspective, which provided understanding and structural simplifications of the system constraints.

In particular, the code construction employed by Ahlswede and Kaspi in the analysis of the trapdoor channel shed light on the structure of optimal battery policies and worst case energy consumptions. Therein, the construction of the optimal jamming strategy (battery policy) and the worst case input process (worst case consumption) are based on the idea of repetition and uniformity. For the policy design, Ahlswede and Kaspi show that in the binary case, the trapdoor is always able to output a constant sequence of balls with length equal to the trapdoor size. For the worst case consumption, consecutively introducing more balls of the same type than the trapdoor can store, forces the trapdoor to leak the type of ball introduced. Finally, these inspiration was grounded by showing the equivalence between battery policies and the jammer case of the trapdoor channel for binary input and output alphabets. However, this equivalence does not hold for arbitrary alphabets, since the battery is able to aggregate energy.

These ideas inspired the construction of block battery policies, providing mutual information privacy guarantees that hold for any energy consumption taking values on a bounded alphabet. The optimality of block battery policies was characterized by presenting an energy consumption process whose leakage is tight with respect to the upper bound provided by the proposed battery policy. Following the construction by Ahlswede and Kaspi, both battery policy and worst case consumption relied on the idea of uniformity, with both repeating a sequence of either no consumption or maximum consumption for the length required to deplete a fully charged battery. Thus, one bit of information is leaked every λ time steps, where λ denotes the time needed to deplete a fully charged battery. Furthermore, we studied the case in which the average energy consumption of the user is specified, showing that the information leakage is governed by the binary entropy of the average energy consumption. Moreover, this showed that the aforementioned battery policy and worst case consumption preserve their optimality when the average energy consumption is specified. It is important to note that the proposed battery policy requires non-causal information about whether the battery will be depleted in the next λ time steps. Although not far-fetched, these forecasting capabilities are not always available to the EMU.

In Chapter 4 we recalled the importance of variable market prices, and their fundamental role in matching energy demand and generation. Consequently, we focused on understanding the fundamental tradeoff between privacy optimization and

the price paid for the energy. Our key insight to solve this multi-objective optimization problem relied on expressing the price paid for the energy as a function of the battery state at market transition points. This allowed us to combine feasibility and cost constraints into a single more tractable optimization problem with constrained battery states.

For the construction of the battery policy, we started with the block battery policy derived in the previous chapter. The energy requested by this market unaware policy was then increased or reduced at market transition points. This allows the fine tuning of the battery state at transition points, reducing the price paid for the energy, and modelling scenarios with arbitrary price constraints. The analysis was also particularized to consider scenarios with known average energy consumptions. As in the previous chapter, this policy requires the forecasting of whether the battery will deplete in the next λ time steps. At the same time, the requirements to guarantee the price constraints depend on the maximum amount of energy that the user can sell or request from the provider at each instant. For the construction of the worst case consumption, we started with an EMU constrained to the minimum feasible cost. This effectively forces the battery states at the start and end of each market block to a single value, allowing the study of each market block independently of each other. This constraint was then relaxed to model arbitrary privacy budgets.

In Chapter 5 we noted the importance of single letter expressions characterizing the upper and lower bounds on the information leakage. Single letter expressions provide more fundamental insights than numerical results, helping understand the dependence on each parameter, and increasing the tractability of the analysis when integrated into a larger framework. Subsequently, we derived the necessary and sufficient conditions under which all the consumption sequences in a given set share a feasible request. This characterization is fundamental in the design of privacy preserving battery policies since the existence of shared request sequences determines whether the EMU is able conceal the energy consumption sequence that induced the energy request. This allowed us to bound the cardinality of the minimal covering and packing sets, and to obtain a tight single letter characterization of the information leakage when no privacy budget is allocated.

Moreover, we proposed structural simplifications on the battery policy and worst case consumptions derived in Chapter 4 that yield single letter upper and lower bounds on the information leakage. This single letter characterization showed that the cost constraint adds an additional component to the leakage induced by the feasibility constraints. This additional component, induced by the market, increases linearly with the number of market changes while it increases logarithmically with the battery size and the variance of the market prices. Finally, doubling the privacy

budget reduces the information leakage in one bit. These show how the different energy tariffs impact the information leakage, help the users decide on a privacy-cost tradeoff point, and help understand the impact of the battery size on the information leakage, among others. Finally, we extended our previous results to other privacy metrics. In particular we showed that the upper and lower bounds derived in this thesis hold for maximal leakage, maximal α-leakage, and Sibson and Arimoto channel capacities. These results were obtained by showing that, under certain conditions, these privacy metrics are equivalent.

6.2 Future work

We now discuss some possible future lines of research.

- **Characterize the impact of causality on the privacy guarantees**. Implementation of the battery policies presented in this thesis requires non-causal information on whether the battery will deplete in the following λ time steps. Thus, the work here presented can be further developed by determining what forecasting capabilities are typically available in real scenarios. Furthermore, it is interesting to characterize the privacy-forecasting tradeoff, therein describing the achievable privacy as a function of the forecasting capability of the EMU.

- **Introduce more complex battery models.** In our study, we do not consider the wear and tear of the battery, and we do not limit the maximum and minimum discharge rates. As discussed in Section 2.1.3, the characteristics of commercial batteries are typically within the requirements of block battery polices. Although our requirements are met by some batteries, and are typical assumptions in the SM privacy literature, certain batteries do not match the requirements of block battery policies. This work can thus be further developed by characterizing the privacy loss induced by more complex battery models.

- **Introduce more constraints on the consumption.** Within this thesis we have characterized bounds on the information leakage that hold for any bounded energy consumption. Furthermore we have made specific the analysis to consider scenarios where the average energy consumption of the user is specified. Another possible way forward is to further make specific the analysis for scenarios in which other properties of the energy consumption are specified, e.g. variance or autocorrelation.

- **Exploit the link between mutual information and maximal leakage**. In Chapter 5 we showed the equivalence between maximal leakage and mutual

information under certain scenarios. This equivalence may be useful in the characterization of rate distortion bounds, as it provides an alternative characterization method. The equivalence also provides a new operational meaning to the minimization of mutual information, thus exploring this link can shed some light on other scenarios such as the definition of new privacy measures.

Bibliography

[1] Vivopower. Rooftop: Why behind the smart meters? [Online]. Available: https://vivopower.com/rooftop-solar-installation-why-behind-the-meter-works/

[2] J. Z. Kolter and M. J. Johnson, "Redd: A public data set for energy disaggregation research," in *Workshop Data Mining Appl. Sustainability (SIGKDD), San Diego, CA*, vol. 25, no. Citeseer, 2011, pp. 59–62.

[3] A. J. Conejo, J. M. Morales, and L. Baringo, "Real-time demand response model," *IEEE Trans. Smart Grid*, vol. 1, no. 3, pp. 236–242, 2010.

[4] S. McLaughlin, P. McDaniel, and W. Aiello, "Protecting consumer privacy from electric load monitoring," in *Proc. ACM Conf. Computer Commun. security*, 2011, pp. 87–98.

[5] W. Yang, N. Li, Y. Qi, W. Qardaji, S. McLaughlin, and P. McDaniel, "Minimizing private data disclosures in the smart grid," in *Proc. ACM Conf. Computer Commun. security*, 2012, pp. 415–427.

[6] O. Tan, D. Gündüz, and J. G. Vilardebó, "Optimal privacy-cost trade-off in demand-side management with storage," in *IEEE Int. Workshop Signal Process. Advances Wireless Commun. (SPAWC)*, 2015, pp. 370–374.

[7] Y. Sun, L. Lampe, and V. W. Wong, "Smart meter privacy: Exploiting the potential of household energy storage units," *IEEE Internet Things J.*, vol. 5, no. 1, pp. 69–78, 2017.

[8] Z. Li, T. J. Oechtering, and M. Skoglund, "Privacy-preserving energy flow control in smart grids," in *IEEE Int. Conf. Acoustics, Speech Signal Process. (ICASSP)*, Mar. 2016, pp. 2194–2198.

[9] J. Gómez-Vilardebó and D. Gündüz, "Privacy of smart meter systems with an alternative energy source," in *Proc. IEEE Int. Symp. Inf. Theory*, Istanbul, Turkey, Jul. 2013, pp. 2572 –2576.

[10] J. Gomez-Vilardebó and D. Gündüz, "Smart meter privacy for multiple users in the presence of an alternative energy source," *IEEE Trans. Inf. Forensics Security*, vol. 10, no. 1, pp. 132–141, 2014.

[11] G. Giaconi, D. Gündüz, and H. V. Poor, "Smart meter privacy with an energy harvesting device and instantaneous power constraints," in *IEEE Int. Conf. Commun. (ICC)*, 2015, pp. 7216–7221.

[12] R. B. Ash, *Information Theory*. Dover Publications Inc., New York, 1990.

[13] O. Tan, D. Gunduz, and H. V. Poor, "Increasing smart meter privacy through energy harvesting and storage devices," *IEEE J. Select. Areas Commun.*, vol. 31, no. 7, pp. 1331–1341, 2013.

[14] S. Li, A. Khisti, and A. Mahajan, "Information-theoretic privacy for smart metering systems with a rechargeable battery," *IEEE Trans. Inf. Theory*, vol. 64, no. 5, pp. 3679–3695, 2018.

[15] S. Cleemput, "Secure and privacy-friendly smart electricity metering," 2018.

[16] U.S.-Canada Power System Outage Task Force. Final report on the August 14, 2003 blackout in the United States and Canada: Causes andrecommendations. [Online]. Available: https://www.energy.gov/sites/prod/files/oeprod/DocumentsandMedia/BlackoutFinal-Web.pdf

[17] K. Vaillancourt, "Electricity transmission and distribution," *Energy Technol. Syst. Anal. programme*. [Online]. Available: https://iea-etsap.org/E-TechDS/PDF/E12_el-t&d_KV_Apr2014_GSOK.pdf

[18] (2015, Dec.) Paris agreement. Accessed 04-Sep-2019. [Online]. Available: https://treaties.un.org/pages/ViewDetails.aspx?src=TREATY&mtdsg_no=XXVII-7-d&chapter=27&clang=_en

[19] U. E. I. Administration. (2016) International energy outlook 2016. U.S. Energy Inform. Admin. Washington, D.C. Tech. Rep. DOE/EIA-0484(2016). [Online]. Available: https://www.eia.gov/outlooks/ieo/pdf/0484(2016).pdf

[20] BP. (2016) BP energy outlook. [Online]. Available: https://www.bp.com/en/global/corporate/energy-economics/energy-outlook.html

[21] European 2030 energy strategy. Accessed 04-Sep-2019. [Online]. Available: https://ec.europa.eu/energy/en/topics/energy-strategy-and-energy-union/2030-energy-strategy

[22] A. Ipakchi and F. Albuyeh, "Grid of the future," *IEEE Power Energy Mag.*, vol. 7, no. 2, pp. 52–62, Mar. 2009.

[23] CISA. Cyber-attack against ukrainian critical infrastructure. [Online]. Available: https://www.us-cert.gov/ics/alerts/IR-ALERT-H-16-056-01

[24] A. Rath. Indian blackouts of july 2012: What happened and why? [Online]. Available: https://medium.com/clean-energy-for-billions/indian-blackouts-of-july-2012-what-happened-and-why-639e31fb52ad

[25] E. C. T. F. for Smart Grids. Functionalities of smart grids and smart meters. [Online]. Available: http://bit.ly/2kNg2BI

[26] Voltage disturbances. standard en 50160. Accessed 23-Sep-2019. [Online]. Available: https://copperalliance.org.uk/uploads/2018/03/542-standard-en-50160-voltage-characteristics-in.pdf

[27] A.-H. Mohsenian-Rad, V. W. Wong, J. Jatskevich, and R. Schober, "Optimal and autonomous incentive-based energy consumption scheduling algorithm for smart grid," in *Innovative Smart Grid Technologies (ISGT)*, 2010, pp. 1–6.

[28] R. Hartway, S. Price, and C. Woo, "Smart meter, customer choice and profitable time-of-use rate option," *Energy*, vol. 24, no. 10, pp. 895–903, Oct. 1999. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0360544299000407

[29] G. Wood and M. Newborough, "Dynamic energy-consumption indicators for domestic appliances: Environment, behaviour and design," *Energy Buildings*, vol. 35, no. 8, pp. 821–841, 2003. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0378778802002414

[30] Directive 2009/72/ec of the European parliament and of the council. Accessed 23-Sep-2019. [Online]. Available: https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32009L0072

[31] E. . Young. Smart grid: A race worth winning? [Online]. Available: https://www.ourenergypolicy.org/wp-content/uploads/2016/03/EY-Smart-Grid-a-race-worth-winning.pdf

[32] C. Dictionary. Meaning of privacy in English. Accessed 23-Sep-2019. [Online]. Available: https://dictionary.cambridge.org/us/dictionary/english/privacy

[33] T. Guardian. The Cambridge analytica files. Accessed 23-Sep-2019. [Online]. Available: https://www.theguardian.com/news/series/cambridge-analytica-files

[34] ——. Apple contractors 'regularly hear confidential details' on siri recordings. Accessed 23-Sep-2019. [Online]. Available: https://www.theguardian.com/technology/2019/jul/26/apple-contractors-regularly-hear-confidential-details-on-siri-recordings

[35] E. F. Foundation. Are you infected by Sony-BMG's rootkit? Accessed 23-Sep-2019. [Online]. Available: https://www.eff.org/deeplinks/2005/11/are-you-infected-sony-bmgs-rootkit

[36] A. Narayanan and V. Shmatikov, "How to break anonymity of the netflix prize dataset," *arXiv preprint cs/0610105*, 2006.

[37] L. Sweeney, "Simple demographics often identify people uniquely," *Health (San Francisco)*, vol. 671, pp. 1–34, 2000.

[38] P. Golle, "Revisiting the uniqueness of simple demographics in the US population," in *Proc. ACM workshop Privacy Electron. Soc.*, 2006, pp. 77–80.

[39] E. L. Quinn, "Privacy and the new energy infrastructure," *Available at SSRN 1370731*, 2008.

[40] P. McDaniel and S. McLaughlin, "Security and privacy challenges in the smart grid," *IEEE Security Privacy*, vol. 7, no. 3, pp. 75–77, May. 2009.

[41] H. Y. Lam, G. S. K. Fung, and W. K. Lee, "A novel method to construct taxonomy electrical appliances based on load signatures," *IEEE Trans. Consumer Electronics*, vol. 53, no. 2, pp. 653–660, May. 2007.

[42] G. W. Hart, "Nonintrusive appliance load monitoring," *Proc. IEEE*, vol. 80, no. 12, Dec. 1992.

[43] I. Rouf, H. Mustafa, M. Xu, W. Xu, R. Miller, and M. Gruteser, "Neighborhood watch: Security and privacy analysis of automatic meter reading systems," in *Proc. Conf. Computer Commun. Security*, New York, NY, USA, 2012. [Online]. Available: http://doi.acm.org/10.1145/2382196.2382246

[44] M. A. Lisovich, D. K. Mulligan, and S. B. Wicker, "Inferring personal information from demand-response systems," *IEEE Security & Privacy*, vol. 8, no. 1, pp. 11–20, 2010.

[45] E. McKenna, I. Richardson, and M. Thomson, "Smart meter data: Balancing consumer privacy concerns with legitimate applications," *Energy Policy*, vol. 41, pp. 807–814, 2012.

[46] A. Prudenzi, "A neuron nets based procedure for identifying domestic appliances pattern-of-use from energy recordings at meter panel," in *Proc. IEEE Power Eng. Soc. Winter Meeting*, 2002.

[47] A. Molina-Markham, P. Shenoy, K. Fu, E. Cecchet, and D. Irwin, "Private memoirs of a smart meter," in *Proc. Workshop Embedded Sensing Syst. Energy-Efficiency Building*, 2010, pp. 61–66.

[48] M. Lisovich and S. Wicker, "Privacy concerns in upcoming residential and commercial demand-response systems," *IEEE Proc. Power Syst.*, vol. 1, no. 1, pp. 1–10, 2008.

[49] U. Greveler, B. Justus, and D. Loehr, "Forensic content detection through power consumption," in *Proc. IEEE Int. Conf. Commun.*, Ottawa, ON, Canada, Jun. 2012.

[50] M. Enev, S. Gupta, T. Kohno, and S. N. Patel, "Televisions, video privacy, and powerline electromagnetic interference," in *Proc. ACM Conf. Computer Commun. security*, 2011, pp. 537–550.

[51] C. Cuijpers and B. Koops, *Smart Metering and Privacy in Europe: Lessons from the Dutch Case*. Springer Netherlands, 2012, pp. 269–293.

[52] A. Bartoli, J. Hernandez-Serrano, M. Soriano, M. Dohler, A. Kountouris, and D. Barthel, "Secure lossless aggregation over fading and shadowing channels for smart grid M2M networks," *IEEE Trans. Smart Grid*, vol. 2, no. 4, pp. 844–864, Dec. 2011.

[53] S. R. Rajagopalan, L. Sankar, S. Mohajer, and H. V. Poor, "Smart meter privacy: A utility-privacy framework," in *Proc. IEEE Int. Conf. Smart Grid Commun.*, Brussels, Belgium, Oct. 2011, pp. 190–195.

[54] F. Li, B. Luo, and P. Liu, "Secure information aggregation for smart grids using homomorphic encryption," in *Proc. IEEE Int. Conf. Smart Grid Commun.*, Oct. 2010.

[55] The sense home energy monitor. Accessed 16-Aug-2019. [Online]. Available: https://sense.com/

[56] O. Tan, D. Gündüz, and H. V. Poor, "Smart meter privacy in the presence of energy harvesting and storage devices," in *Proc. IEEE Int. Conf. Smart Grid Commun.*, Tainan, Taiwan, Nov. 2012, pp. 664–669.

[57] D. Varodayan and A. Khisti, "Smart meter privacy using a rechargeable battery: Minimizing the rate of information leakage," in *Proc. IEEE Int. Conf. Acoust. Speech Sig. Proc.*, Prague, Czech Republic, May. 2011, pp. 1932–1935.

[58] G. Kalogridis, C. Efthymiou, S. Z. Denic, T. A. Lewis, and R. Cepeda, "Privacy for smart meters: Towards undetectable appliance load signatures," in *Proc. IEEE Int. Conf. Smart Grid Commun.*, Gaithersburg, MD, USA, Oct. 2010.

[59] G. Kalogridis, R. Cepeda, S. Z. Denic, T. Lewis, and C. Efthymiou, "Elecprivacy: Evaluating the privacy protection of electricity management algorithms," *IEEE Trans. Smart Grid*, vol. 2, no. 4, pp. 750–758, 2011.

[60] M. Arrieta and I. Esnaola, "Smart meter privacy via the trapdoor channel," in *Proc. IEEE Int. Conf. Smart Grid Commun.*, Dresden, Germany, Nov. 2017, pp. 277–282.

[61] M. Arrieta, I. Esnaola, and M. Effros, "Universal privacy guarantees for smart meters," *Proc. IEEE Int. Symp. Inf. Theory*, Jul. 2019.

[62] E. Liu, P. You, and P. Cheng, "Optimal privacy-preserving load scheduling in smart grid," in *Proc. IEEE Power Energy Soc. General Meeting (PESGM)*, 2016, pp. 1–5.

[63] J. Yang, G. Huang, and C. Wei, "Privacy-aware electricity scheduling for home energy management system," *Peer-to-Peer Networking Appl.*, vol. 11, no. 2, pp. 309–317, 2018.

[64] Y. H. Liu, S.-H. Lee, and A. Khisti, "Information-theoretic privacy in smart metering systems using cascaded rechargeable batteries," *IEEE Signal Process. Lett.*, vol. 24, no. 3, pp. 314–318, 2017.

[65] G. Giaconi, D. Gündüz, and H. V. Poor, "Privacy-aware smart metering: Progress and challenges," *IEEE Signal Process. Mag.*, vol. 35, no. 6, pp. 59–78, Nov. 2018.

[66] Department of Energy and Climate Change. (2014, Nov.) Smart metering equipment technical specifications: Version 2. [Online]. Available: https://www.gov.uk/government/consultations/smart-metering-equipment-technical-specifications-second-version

[67] O. Parson. (2012) Public data sets for NIALM. [Online]. Available: http://blog.oliverparson.co.uk/2012/06/public-data-sets-for-nialm.html

[68] H. Kim, M. Marwah, M. Arlitt, G. Lyon, and J. Han, "Unsupervised disaggregation of low frequency power measurements," in *Proc. SIAM Int. Conf. data mining*, 2011, pp. 747–758.

[69] Y. Sun, L. Lampe, and V. W. S. Wong, "EV-assisted battery load hiding: A Markov decision process approach," in *IEEE Int. Conf. Smart Grid Commun.*, Nov. 2016, pp. 160–166.

[70] Solar batteries: Who makes the best solar battery? [Online]. Available: https://www.solarguide.co.uk/solar-batteries#/

[71] L. Yang, X. Chen, J. Zhang, and H. V. Poor, "Cost-effective and privacy-preserving energy management for smart meters," *IEEE Trans. Smart Grid*, vol. 6, no. 1, pp. 486–495, Jan. 2015.

[72] J. Wu, J. Liu, X. S. Hu, and Y. Shi, "Privacy protection via appliance scheduling in smart homes," in *IEEE/ACM Int. Conf. Computer-Aided Design (ICCAD)*, 2016, pp. 1–6.

[73] Z. Chen and L. Wu, "Residential appliance DR energy management with electric privacy protection by online stochastic optimization," *IEEE Trans. Smart Grid*, vol. 4, no. 4, pp. 1861–1869, 2013.

[74] F. Farokhi and H. Sandberg, "Fisher information as a measure of privacy: Preserving privacy of households with smart meters using batteries," *IEEE Trans. Smart Grid*, vol. 9, no. 5, pp. 4726–4734, 2017.

[75] A. Lever, D. Sanders, N. Lehmann, M. Ravishankar, M. Ashcroft, G. Strbac, M. Aunedi, F. Teng, and D. Pudijanto, "Can storage help reduce the cost of a future UK electricity system?" *Carbon Trust, Feb*, 2016.

[76] Octopus energy tariff. Accessed 21-Jan-2019. [Online]. Available: https://octopus.energy/agile/

[77] G. Giaconi, D. Gündüz, and H. V. Poor, "Joint privacy-cost optimization in smart electricity metering systems," *arXiv preprint arXiv:1806.09715*, 2018.

[78] G. Giaconi, D. Gündüz, and H. V. Poor, "Optimal demand-side management for joint privacy-cost optimization with energy storage," in *IEEE Int. Conf. Smart Grid Commun.*, 2017, pp. 265–270.

[79] J. Yao and P. Venkitasubramaniam, "The privacy analysis of battery control mechanisms in demand response: Revealing state approach and rate distortion bounds," *IEEE Trans. Smart Grid*, vol. 6, no. 5, pp. 2417–2425, 2015.

[80] L. Yang, X. Chen, J. Zhang, and H. V. Poor, "Optimal privacy-preserving energy management for smart meters," in *IEEE INFOCOM 2014-IEEE Conf. Computer Commun.*, 2014, pp. 513–521.

[81] UK economy 7 energy tariff. Accessed 21-Jan-2019. [Online]. Available: https://www.moneysavingexpert.com/utilities/economy-7/

[82] I. Wagner and D. Eckhoff, "Technical privacy metrics: a systematic survey," *ACM Computing Surveys (CSUR)*, vol. 51, no. 3, p. 57, 2018.

[83] G. W. Hart, "Residential energy monitoring and computerized surveillance via utility power flows," *IEEE Technol. Soc. Mag.*, vol. 8, no. 2, pp. 12–16, 1989.

[84] O. Tan, J. Gómez-Vilardebó, and D. Gündüz, "Privacy-cost trade-offs in demand-side management with storage," *IEEE Trans. Inf. Forensics Security*, vol. 12, no. 6, pp. 1458–1469, 2017.

[85] C. R. Rao, "Information and the accuracy attainable in the estimation of statistical parameters." Springer, 1992, pp. 235–247.

[86] J. Neyman and E. S. Pearson, "On the problem of the most efficient tests of statistical hypotheses," *Philosoph. Trans. of Roy. Soc. of London.*, vol. 231, no. 694-706, pp. 289–337, 1933.

[87] Z. Li and T. J. Oechtering, "Privacy on hypothesis testing in smart grids," in *Proc. IEEE Inf. Theory Workshop (ITW)*, Jeju, South Korea, Oct. 2015, pp. 337–341.

[88] Z. Li, T. J. Oechtering, and D. Gündüz, "Smart meter privacy based on adversarial hypothesis testing," in *IEEE Int. Symp. Inf. Theory (ISIT)*, June 2017, pp. 774–778.

[89] S. Salehkalaibar, F. Aminifar, and M. Shahidehpour, "Hypothesis testing for privacy of smart meters with side information," *IEEE Trans. Smart Grid*, vol. 10, no. 2, pp. 2059–2067, 2017.

[90] J. Liao, L. Sankar, V. Y. F. Tan, and F. du Pin Calmon, "Hypothesis testing under mutual information privacy constraints in the high privacy regime," *IEEE Trans. Inf. Forensics Security*, vol. 13, no. 4, pp. 1058–1071, Apr. 2018.

[91] C. Dwork, "Differential privacy," *Encyclopedia Cryptography Security*, pp. 338–340, 2011.

[92] J. Zhao, T. Jung, Y. Wang, and X. Li, "Achieving differential privacy of data disclosure in the smart grid," in *IEEE INFOCOM 2014-IEEE Conf. Computer Commun.*, 2014, pp. 504–512.

[93] P. Cuff and L. Yu, "Differential privacy as a mutual information constraint," in *Proc. ACM SIGSAC Conf. Computer Commun. Security*, 2016, pp. 43–54.

[94] P. Kairouz, S. Oh, and P. Viswanath, "The composition theorem for differential privacy," *IEEE Trans. Inf. Theory*, vol. 63, no. 6, pp. 4037–4049, 2017.

[95] I. Issa, S. Kamath, and A. B. Wagner, "An operational measure of information leakage," in *Annu. Conf. Inf. Sci. Syst. (CISS)*, Mar. 2016, pp. 234–239.

[96] J. Liao, O. Kosut, L. Sankar, and F. P. Calmon, "A tunable measure for information leakage," in *IEEE Int. Symp. Inf. Theory (ISIT)*, 2018, pp. 701–705.

[97] I. Issa, A. B. Wagner, and S. Kamath, "An operational approach to information leakage," *arXiv preprint arXiv:1807.07878*, 2018.

[98] C. Braun, K. Chatzikokolakis, and C. Palamidessi, "Quantitative notions of leakage for one-try attacks," *Electron. Notes Theoretical Computer Sci.*, vol. 249, pp. 75–91, 2009.

[99] C. E. Shannon, "A mathematical theory of communication," *Bell system Tech. J.*, vol. 27, no. 3, pp. 379–423, 1948.

[100] T. M. Cover and J. A. Thomas, *Elements of information theory*. John Wiley & Sons, 2012.

[101] S. Kullback and R. A. Leibler, "On information and sufficiency," *Ann. Math. Statist.*, vol. 22, no. 1, pp. 79–86, 1951.

[102] S. Finster and I. Baumgart, "Privacy-aware smart metering: A survey," *IEEE Commun. Surveys & Tutorials*, vol. 16, no. 3, pp. 1732–1745, 2014.

[103] M. R. Asghar, G. Dán, D. Miorandi, and I. Chlamtac, "Smart meter data privacy: A survey," *IEEE Commun. Surveys & Tutorials*, vol. 19, no. 4, pp. 2820–2835, 2017.

[104] G. Kalogridis and S. Dave, "Pehems: Privacy enabled hems and load balancing prototype," in *IEEE Third Int. Conf. Smart Grid Commun.*, 2012, pp. 486–491.

[105] Y. Sun, L. Lampe, and V. W. Wong, "Combining electric vehicle and rechargeable battery for household load hiding," in *IEEE Int. Conf. Smart Grid Commun.*, 2015, pp. 611–616.

[106] D.-M. Arnold, H.-A. Loeliger, P. O. Vontobel, A. Kavcic, and W. Zeng, "Simulation-based computation of information rates for channels with memory," *IEEE Trans. Inf. Theory*, vol. 52, no. 8, pp. 3498–3508, 2006.

[107] L. Bahl, J. Cocke, F. Jelinek, and J. Raviv, "Optimal decoding of linear codes for minimizing symbol error," *IEEE Trans. Inf. Theory*, vol. 20, no. 2, pp. 284–287, 1974.

[108] S. Han, U. Topcu, and G. J. Pappas, "Event-based information-theoretic privacy: A case study of smart meters," in *Amer. Control Conf. (ACC)*, 2016, pp. 2074–2079.

[109] D. P. Bertsekas, *Dynamic programming and optimal control*. Athena scientific Belmont, MA, 1995, vol. 1, no. 2.

[110] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press, 2018.

[111] S. Li, A. Khisti, and A. Mahajan, "Structure of optimal privacy-preserving policies in smart-metered systems with a rechargeable battery," in *IEEE Int. Workshop Signal Process. Advances Wireless Commun. (SPAWC)*, 2015, pp. 375–379.

[112] S. Li, A. Khisti *et al.*, "Privacy preserving rechargeable battery policies for smart metering systems," in *Int. Zurich Seminar Communications-Proceedings*, 2016.

[113] S. Li, A. Khisti, and A. Mahajan, "Privacy-optimal strategies for smart metering systems with a rechargeable battery," in *Amer. Control Conf. (ACC)*, 2016, pp. 2080–2085.

[114] Y. B. Li, "Information theoretic privacy in smart metering systems using rechargeable batteries," Ph.D. dissertation, 2016.

[115] Y. Liu, A. Khisti, and A. Mahajan, "On privacy in smart metering systems with periodically time-varying input distribution," in *Proc. IEEE Global Conf. Signal Inf. Process. (GlobalSIP)*, 2017, pp. 513–517.

[116] G. Giaconi and D. Gündüz, "Smart meter privacy with renewable energy and a finite capacity battery," in *Proc. IEEE Int. Workshop Sig. Process. Advances Wireless Commun. (SPAWC)*, Edinburgh, UK, Jul. 2016, pp. 1–5.

[117] E. Erdemir, P. L. Dragotti, and D. Gündüz, "Privacy-cost trade-off in a smart meter system with a renewable energy source and a rechargeable battery," in *IEEE Int. Conf. Acoust. Speech Signal Process. (ICASSP)*, 2019, pp. 2687–2691.

[118] J.-X. Chin, T. T. De Rubira, and G. Hug, "Privacy-protecting energy management unit through model-distribution predictive control," *IEEE Trans. Smart Grid*, vol. 8, no. 6, pp. 3084–3093, 2017.

[119] J. Koo, X. Lin, and S. Bagchi, "Privatus: Wallet-friendly privacy protection for smart meters," in *Eur. Symp. Res. Computer Security*, 2012, pp. 343–360.

[120] D. Gündüz and J. Gómez-Vilardebó, "Smart meter privacy in the presence of an alternative energy source," in *IEEE Int. Conf. Commun. (ICC)*, 2013, pp. 2027–2031.

[121] D. Gündüz, J. Gomez-Vilardebó, O. Tan, and H. V. Poor, "Information theoretic privacy for smart meters," in *Inf. Theory Appl. Workshop (ITA),*, 2013, pp. 1–7.

[122] L. Sankar, S. R. Rajagopalan, S. Mohajer, and H. V. Poor, "Smart meter privacy: A theoretical framework," *IEEE Trans. Smart Grid*, vol. 4, no. 2, pp. 837–846, 2012.

[123] W. Karush, "Minima of functions of several variables with inequalities as side constraints," *M. Sc. Dissertation. Dept. of Mathematics, Univ. Chicago*, 1939.

[124] H. V. Poor and S. Verdú, "A lower bound on the probability of error in multihypothesis testing," *IEEE Trans. Inf. Theory*, vol. 41, no. 6, pp. 1992–1994, Nov. 1995.

[125] D. Guo, S. Shamai, and S. Verdú, "Mutual information and minimum mean-square error in gaussian channels," *IEEE Trans. Inf. Theory*, vol. 51, no. 4, pp. 1261–1282, Apr. 2005.

[126] A. Xu and M. Raginsky, "Information-theoretic analysis of generalization capability of learning algorithms," in *Proc. Advances Neural Inf. Process. Syst.*, 2017, pp. 2524–2533.

[127] J. Ziv and A. Lempel, "A universal algorithm for sequential data compression," *IEEE Trans. Inf. theory*, vol. 23, no. 3, pp. 337–343, 1977.

[128] ——, "Compression of individual sequences via variable-rate coding," *IEEE Trans. Inf. Theory*, vol. 24, no. 5, pp. 530–536, 1978.

[129] R. Ahlswede and A. Kaspi, "Optimal coding strategies for certain permuting channels," *IEEE Trans. Inf. Theory*, vol. 33, no. 3, pp. 310–314, May. 1987.

[130] R. M. Gray, *Entropy and information theory*. Springer Science & Business Media, 2011.

[131] B. Leiner and R. Gray, "Bounds on rate-distortion functions for stationary sources and context-dependent fidelity criteria," *IEEE Trans. Inf. Theory*, vol. 19, no. 5, pp. 706–708, 1973.

[132] T. Benjamin, "Coding for a noisy channel with permutation errors," *IEEE Trans. Inf. Theory*, vol. 21, no. 4, pp. 488–488, 1975.

[133] M. Kovačević and V. Y. Tan, "Codes in the space of multisets—coding for permutation channels with impairments," *IEEE Trans. Inf. Theory*, vol. 64, no. 7, pp. 5156–5169, 2018.

[134] ——, "Coding for the permutation channel with insertions, deletions, substitutions, and erasures," in *IEEE Int. Symp. Inf. Theory (ISIT)*, 2017, pp. 1933–1937.

[135] L. J. Schulman and D. Zuckerman, "Asymptotically good codes correcting insertions, deletions, and transpositions," *IEEE Trans. Inf. theory*, vol. 45, no. 7, pp. 2552–2557, 1999.

[136] R. Chakinala, A. Kumarasubramanian, R. Manokaran, G. Noubir, C. P. Rangan, and R. Sundaram, "Steganographic communication in ordered channels," in *Int. Workshop Inf. Hiding*, 2006, pp. 42–57.

[137] Z. Wang and R. B. Lee, "Covert and side channels due to processor architecture," in *22nd Annu. Computer Security Appl. Conf. (ACSAC'06)*, 2006, pp. 473–482.

[138] D. Blackwell, "Information theory," E. F. Beckenbach, Ed. New York: McGraw-Hill, 1961, pp. 182–193, mR:0129161.

[139] R. Ahlswede, N. Cai, and Z. Zhang, "Zero-error capacity for models with memory and the enlightened dictator channel," *IEEE Trans. Inf. Theory*, vol. 44, no. 3, 1998.

[140] H. H. Permuter, P. Cuff, B. V. Roy, and T. Weissman, "Capacity of the trapdoor channel with feedback," *CoRR*, vol. abs/cs/0610047, 2006.

[141] K. Kobayashi, "Capacity problem of trapdoor channel," in *General Theory Inf. Transfer Combinatorics*. Springer, 2006, pp. 1084–1087.

[142] ——, "Combinatorial structure and capacity of the permuting relay channel," *IEEE Trans. Inf. theory*, vol. 33, no. 6, pp. 813–826, 1987.

[143] J. v. Neumann, "Zur theorie der gesellschaftsspiele," *Mathematische annalen*, vol. 100, no. 1, pp. 295–320, 1928.

[144] J. Gomez-Vilardebó and D. Gündüz, "Smart meter privacy for multiple users in the presence of an alternative energy source," *IEEE Trans. Inf. Forensics Security*, vol. 10, no. 1, pp. 132–141, 2015.

[145] Accessed 21-Jan-2019. [Online]. Available: http://www.windandsun.co.uk/products/Batteries/Lithium-Ion-Batteries/LG-Chem-Lithium-Ion-Batteries

[146] T. S. Han and S. Verdú, "Approximation theory of output statistics," *IEEE Trans. Inf. Theory*, vol. 39, no. 3, pp. 752–772, 1993.