

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Incentivizing Efficiency in Societal-Scale Cyber-Physical Systems

Permalink

<https://escholarship.org/uc/item/6ck1z3x3>

Author

Ratliff, Lillian Jane

Publication Date

2015

Peer reviewed|Thesis/dissertation

Incentivizing Efficiency in Societal-Scale Cyber-Physical Systems

by

Lillian Jane Ratliff

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Engineering—Electrical Engineering and Computer Sciences

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor S. Shankar Sastry, Chair

Professor Pravin Varaiya

Professor Lawrence C. Evans

Summer 2015

Incentivizing Efficiency in Societal-Scale Cyber-Physical Systems

Copyright 2015
by
Lillian Jane Ratliff

Abstract

Incentivizing Efficiency in Societal-Scale Cyber-Physical Systems

by

Lillian Jane Ratliff

Doctor of Philosophy in Engineering—Electrical Engineering and Computer Sciences

University of California, Berkeley

Professor S. Shankar Sastry, Chair

In the modernization of infrastructure systems such as energy, transportation, and health-care systems we are seeing the convergence of three research domains: Cyber-Physical Systems (CPS), Big Data, and the Internet of Things (IoT). Indeed, new CPS technologies are being deployed to create large sensor-actuator networks which produce massive quantities of data often in real-time which is, in turn, being used to inform everyday decision-making of the entities that engage with these large-scale infrastructure systems. As a consequence, such systems are quickly evolving into *societal-scale cyber-physical systems*.

The result of this increasing connectivity and interdependence is two-fold: more and more data is being collected, transmitted, and stored, and more and more actuation modalities are available, allowing new ways to influence the behavior of infrastructure systems. These new and pervasive sensing/actuation modalities present new opportunities for improving efficiency, yet they expose novel vulnerabilities. In energy CPS, for instance, smart metering technologies increase the availability of streaming data thereby enabling monetization of energy savings. Such savings can be realized by employing novel machine learning algorithms to customize offerings to consumers. On the other hand, the availability of this fine-grained consumer/system data and the increased number of access points to the broader system expose new privacy and security risks. Hence, there is an inherent efficiency-vulnerability tradeoff. This tradeoff is becoming more pronounced due to greater dependence on CPS technologies and the push towards more human-centric operations, i.e. integration of human decision-making and preferences into the closed-loop behavior of the system.

Beginning with the problem of modeling the non-cooperative agents that interact with these large-scale sociotechnical systems and thus, compete over scarce resources, we analyze the outcome of their strategic interactions. In particular, we create a characterization of Nash equilibria—termed *differential Nash equilibria*—in games on non-convex strategy spaces that is amenable to computation. We show that such non-degenerate differential Nash equilibria are structurally stable and generic thereby robust to small modeling errors and measurement noise. Introducing a planner tasked with coordinating these decision-makers, we leverage this characterization in the construction of a utility learning and incentive design

algorithm. We provide convergence results in both the case where agents play according to Nash and where they play using a myopic update rule.

Narrowing our focus to the demand-side of the smart grid, we consider that the planner will capitalize on new sensing/actuation modalities in the design incentives thereby exposing the efficiency-vulnerability tradeoff. We consider privacy risks introduced by smart metering technologies that produce streaming energy consumption data. On one hand the data has utility in the sense that it can help improve operations, yet on the other it exposes the consumer and the power company to greater privacy risk. We propose a solution that combines economic and statistics tools, i.e. privacy-aware service contracts in which service is differentiated according to privacy and consumers select based on their needs and wallet. We argue that the power company has an incentive to invest in security or purchase insurance because of inefficiencies that arise due to information asymmetries and we design insurance contracts accordingly. We provide a number of qualitative insights that have the potential to be useful for informing policy and regulations in the energy ecosystem. Finally, we conclude with an overview of the contributions and a discussion of future research directions for the near and far terms.

The contributions are the first steps towards an emerging systems theory of societal-scale cyber-physical systems in which there are many tightly coupled human-CPS decision-making loops and socioeconomic factors intricately woven into the fabric.

“At any street corner the feeling of absurdity
can strike any man in the face.”
— Albert Camus

Contents

Contents	ii
1 Introduction	1
1.1 Societal–Scale Cyber–Physical Systems	1
1.2 Efficiency–Vulnerability Tradeoff in Societal–Scale Cyber–Physical Systems	3
1.3 Transitioning to Human–Centric Systems	5
1.4 Systems Theory for Societal–Scale Cyber–Physical Systems	6
1.4.1 A Motivating Vignette	7
1.4.2 Contributions	7
2 Characterization and Computation of Local Nash	10
2.1 Game Formulation	12
2.2 Characterization of Local Nash Equilibria	16
2.3 Structural Stability	20
2.4 Genericity	23
2.5 Potential Games	26
2.6 Computation of Local Nash Equilibria	33
2.6.1 Examples	35
2.7 Inducing a Nash Equilibrium	38
2.8 Discussion	40
2.A Preliminaries	41
2.A.1 Algebraic Topology	44
2.A.2 Differential Topology	45
3 Utility Learning and Incentive Design	47
3.1 Problem Formulation	48
3.2 Utility Learning Formulation	51
3.2.1 Utility Learning Under Nash–Play	51
3.2.2 Utility Learning Under Myopic–Play	53
3.3 Incentive Design Formulation	54
3.3.1 Incentive Design Under Nash–Play	55
3.3.2 Incentive Design Under Myopic Play	57

3.4	Algorithm	58
3.5	Convergence Results	59
3.6	Uncertainty in Agent Play	68
3.7	Discussion	75
3.A	Preliminaries	78
4	Privacy–Aware Incentive Design	82
4.1	Privacy Contracts	84
4.2	Effects of Privacy Loss Risk on Contracts	91
4.2.1	Inferential Privacy Metrics	91
4.2.2	Contracts with Risk-Averse Consumers	94
4.2.3	Characterizing the Effects of Privacy Loss Risk	94
4.3	Insurance Contracts	101
4.3.1	Analysis of the Agent’s Decision	102
4.3.2	Analysis of the Insurer’s Decision	104
4.4	Discussion	107
5	Conclusion and Future Directions	109
5.1	Emerging Tools for Societal–Scale Cyber–Physical Systems	109
5.2	Future Plans and Frontiers	110
5.2.1	Merging Game–Theory and Statistical Learning	110
5.2.2	Vulnerability–Aware Incentives and the Emerging Data Market	111
5.2.3	Smart Urban Spaces	112
	Bibliography	114

Acknowledgments

I would like to acknowledge all of the individuals and their contributions in making this thesis possible. I have to express my extreme gratitude to Shankar Sastry for advising on my journey through graduate school. It has been a wonderful pleasure to work with him and with the many brilliant researchers that surround him. The advice I have received is invaluable and I hope it will serve as a foundation in my ongoing research efforts.

I would like to thank the members of my qualifying exam committee, Shankar Sastry, Pravin Varaiya, L. Craig Evans, and Claire Tomlin. The guidance they provided and discussions we have had helped to clarify the problems I focused on and helped to shape this thesis.

I cannot go without acknowledging Sam Burden, a wonderful and brilliant mentor, who has provided me with endless advice and has helped me navigate through graduate school and research. I have quite literally followed in his footsteps and I hope that he can continue to mentor me as we both take the next steps in our careers.

I would like to thank all of the wonderful people I have had the pleasure to discuss research with as well as *shoot the breeze* with—perhaps they are one in the same. It has been great to work with Dan Calderone, Sam Coogan, Roy Dong, Henrik Ohlsson, and Aaron Bestick. Dan has an uncanny ability to seek out small fractures in your work and ask very poignant questions that he somehow formulates using intuition that I could only dream of having. Sam C. has been a great sounding board for frustrations and concerns and has provided me with needed advice as we both forged through our time in graduate school. It is great to work with someone like Henrik, so positive and productive. My relationship with Roy has been the longest of all the collaborators I have had in graduate school. We have bonded on many things from our experiences as students to weird cartoons. I hope our collaboration as researchers and our friendship continues. Aaron was one of the first students I attempted to work with in a kind of advisor-type role. As it turned out, Aaron gave me more advice and direction than I could have ever given to him and I continue to use him as a resource in this capacity. I also would like to thank Dorsa Sadigh, Ram Vasudevan, Humberto Gonzalez, Walid Krichene, and Jaime Fisac for fruitful discussions. Given my stubborn and sometimes argumentative technique for working through a problem, I am glad they all were able to look past that and stick with me.

I would also like to acknowledge all the hard work, which often goes under the radar, that Larry Rohrbough, Jessica Gamble, Aimee Tabor, and Carolyn Winter do. I do not think any graduate student in our lab could survive without Jessica.

Without further adieu, I cannot express enough gratitude for my partner Dan Cook. He supports me in every capacity.

Chapter 1

Introduction

Spurred on by economic and technological changes, traditional infrastructure systems are evolving into societal-scale cyber-physical systems (S-CPS) in which accessible, easily deployable sensing/actuation devices are being integrated into everything from operations and management to everyday decision-making of the users of these systems. S-CPS are the backbone of modern society; our economy and daily lives depend on access to the services and products offered by these systems. Further, in the modernization of infrastructure systems such as energy, transportation, and healthcare systems we are seeing the convergence of three research domains: Cyber-Physical Systems (CPS), Big Data, and the Internet of Things (IoT). Indeed, societal-scale infrastructures are at an important inflection point in their operations due to increased interdependence on new CPS technologies such as wireless sensor/actuator networks, data-driven real-time learning techniques being implemented in the cloud, and ubiquitous mobile computing devices for intermediating between networks of wireless sensors and the cloud.

The result of this increasing connectivity and interdependence is two-fold: more and more data is being collected, transmitted, and stored, and more and more actuation modalities are available, allowing new ways to influence the behavior of our infrastructures. Both the information and the actuation commands are being transmitted across networks, such as the Internet, and these infrastructures are being operated in an increasingly decentralized manner yet are becoming more human-centric.

1.1 Societal-Scale Cyber-Physical Systems

In viewing infrastructure systems as S-CPS we must consider the various entities and decision makers as well as their roles, information exchanges, and motivations. At an abstract level, there is a *population of users* (e.g. vehicle drivers or electricity consumers), *providers* that offer goods or services to the user population (e.g. local Department of Transportation or power company) facilitated through a *cyber-physical network* often managed by the provider, a *regulation entity* (e.g. Department of Transportation, government, or utility commission)

that issues regulations and policies, and *third-party solution providers* (e.g. data aggregators, insurance companies, etc.) that offers goods and services to either some subset of the user population or providers.

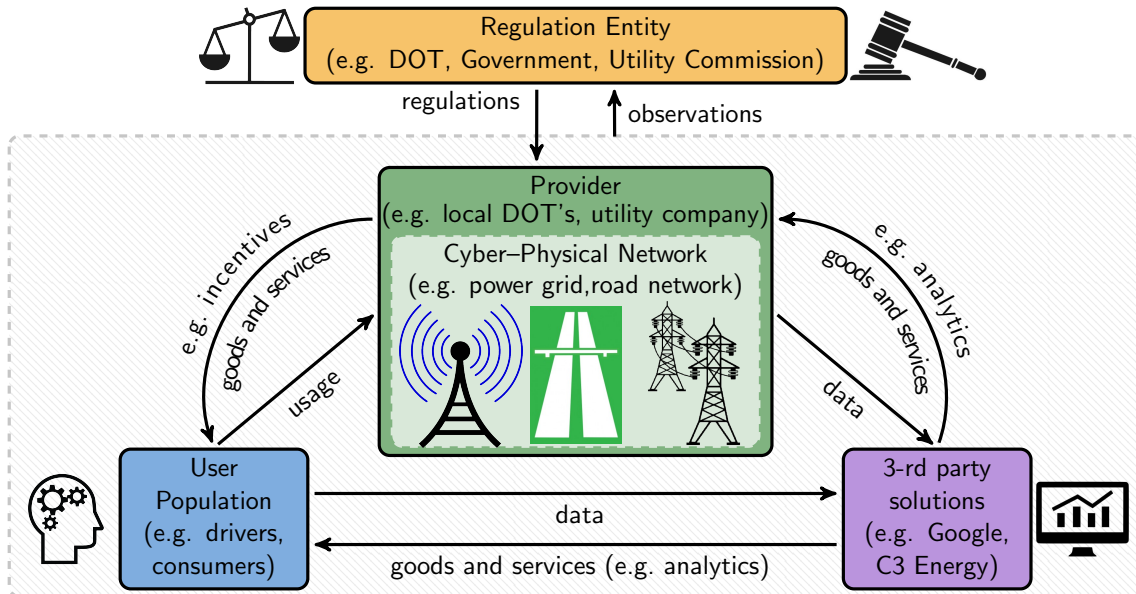


Figure 1.1: *An ontological view of S-CPS.* There are many agents in S-CPS having different capacities and interests in their interactions. At a granular level, we characterize these agents based on their roles in S-CPS. This includes regulation entities, service providers, users, and third-party solution providers.

Figure 1.1 provides an abstracted view of infrastructure systems as S-CPS. With this ontology in place, we can consider how these different agents interact, and the information and control authority given to each agent. It is necessary to understand and consider the information flows and information asymmetries of the different actors when trying to characterize even a subset of the larger S-CPS and, in particular, when designing new economic mechanisms or physical controls for the system. Such information asymmetries may come in the form of what is referred to as *adverse selection* (hidden information such as an agent’s preference over a particular good) or *moral hazard* (hidden actions such as an agent’s choice/decision that is unobserved) or both [LM02]. These information asymmetries lead naturally to inefficiencies in the way the system operates.

Within this framework, we can analyze the structure of the market these agents inhabit and identify whether or not these agents have an incentive to compromise the system. Furthermore, with these interconnections explicitly modeled, we can also consider where and what an outside adversary will attack. It is important to remark that malicious parties can be any one of the different actors/entities described or could be external to the system.

One way to categorize adversaries is based on their end goal. For instance, we can categorize adversaries as eavesdroppers, rational attackers, or anarchists. Eavesdroppers seek to collect information without affecting any part of the physical infrastructure. Eavesdroppers generally only present privacy risks, but this information can be used to compromise the security of systems other than just the infrastructure. Rational attackers seek to profit in a fashion which affects the physical dynamics. Although the intent of these attackers is not necessarily to damage the system, their actions may cause inefficiencies and instabilities. Rational attackers exist because of misaligned incentives. Finally, anarchists are adversaries who use all their available resources to maximize damage to the infrastructure. To model these adversaries, we must consider realistic models of what resources are available to them. In any of these categories their motivations could be political or financial.

Briefly connecting back to the types of information asymmetries, privacy issues relate nicely at a conceptual level to the class of information asymmetries known as adverse selection in which some information is hidden from another party whereas security problems relate nicely at a conceptual level to moral hazard in which some action is hidden or unobserved. For instance, in the former, consider a patient who wants to keep private that they smoke because their insurance may go up; however, their doctor could potentially use this information for diagnostics. Concerning the latter, consider a energy consumer who is stealing energy by spoofing their consumption; this can be viewed as their *action* being hidden from the power company.

This ontological view of S-CPS provides us with a framework and a set of conceptual models in which we can not only consider the traditional components such as the physical system and communication infrastructure, but also the many decision makers and their motivations and information they have access to in the context of the design, operation, and management of large-scale systems such as the power grid, transportation systems, and the interconnections between them. Moreover, we can utilize such a framework to develop tools which can lead to qualitative insights that will have the power to shape and inform policy development and regulation design.

1.2 Efficiency–Vulnerability Tradeoff in S-CPS

The advent of S-CPS brings with it new opportunities for improving efficiency while simultaneously exposing novel vulnerabilities. In transportation S-CPS, for instance, there is a heterogeneous set of users on the road network: bicyclists, pedestrians, and a whole continuum from traditional drivers to fully autonomous vehicles. Due to information asymmetries, user preferences and scarce resources such as the limited capacity of our road network, natural inefficiencies arise. There is a common goal perhaps of safety yet each individual wants to get to their destination as quickly as possible. Perhaps a planner is tasked with coordinating the users. This planner can take many forms from government agencies such as the department of transportation or local authority who implements new policies that affect operation to industry who introduces new service models and products such as shared car

models. Both government and industry are incorporating new CPS technologies such as embedded flow sensors into their operations and decision making. While at the same time this use of CPS technologies allow expanded services and potentially more efficient operations, they also expose security and privacy vulnerabilities. For example, recently we have seen hacks on such systems as Sensys Networks flow sensors—a wireless, embedded CPS technology—showing that through manipulation of the sensor data changes in traffic light control are possible [Zet14]. In addition, recent concerns over privacy around the Uber car sharing platform have been raised [BW14; WG14; TK14].

Transportation systems are not the only infrastructure system to see such changes. In energy S-CPS, smart metering technologies increase the availability of streaming data thereby enabling monetization of energy savings. Such savings can be realized by employing novel machine learning algorithms to customize offerings to consumers. On the other hand, the availability of this fine-grained consumer/system data and the increased number of access points to the broader system expose new privacy and security risks. Recently the US Department of Energy [Ene] and NIST [Ell14] have issued voluntary best practices for ensuring privacy and security of the smart grid. However, these policies are often not actionable, particularly when it comes to privacy. The reason for this is that privacy is subjective in its nature; it is interpreted by the individual. The emergence of such policies makes clear the fact that the integration of new CPS technologies, and ultimately the shift in the way large-scale infrastructure systems operate, is causing a shift in the needed types of regulation and policy. Furthermore, the existence of these best practices reinforces the need for a holistic, systems-theoretic understanding of vulnerabilities such as privacy and security risks in S-CPS.

In general, S-CPS often need to address two related problems concerning vulnerabilities. CPS technologies used for sensing/actuation may need to operate in exposed locations where tampering—sometimes by a user, e.g., electricity theft—may occur. On the other hand, sensor/actuator networks often collect data that is considered private by a user, or has the potential to reveal something private about a user, but is required to reach aggregate conclusions. This results in the user being exposed to privacy loss risk. Security investments may help protect against the former and the latter; however, they may not be the only solution. For instance, economic mechanisms such as fines, incentives, even insurance, can be used to augment security mechanisms that are in place.

One major concern of this modernization is that these security and privacy attacks can now occur on a larger scale. For example, in the past, electricity theft—viewed as a financial attack on the power company—usually required physically tampering with meters or power lines. Now, a hacker—or even a slightly sophisticated consumer [Law10a; Law10b]—can exploit software vulnerabilities and manipulate several smart meters at once. Furthermore, the large-scale, decentralized nature of these systems means that securing every node is not economically feasible. For this reason, recent research has investigated ways to make CPS more resilient. The aim is to design systems that recover from faults, failures, or attacks. If recovery is not possible, then the system is designed so that the performance degrades slowly, perhaps even gracefully, as an adversary attacks larger portions of the system. It

is also desirable that systems are able to detect attacks, even when the adversary is trying to intelligently design its attack to be discreet, for instance, by hiding behind the system dynamics. While there is a large body of literature addressing security attacks in various infrastructure systems (see, e.g., [Faw+14; Pas+12a; Pas+12b; Ami+13; Cár+09]), much of this work ignores the socioeconomic aspects that are often the driving force for attacks and the market structure that creates misaligned incentives.

As the above examples in the transportation and energy contexts show, there is an inherent efficiency–vulnerability tradeoff in S-CPS. Managing this tradeoff is key in the design, management, and operation of large–scale infrastructure systems. The efficiency–vulnerability tradeoff is becoming more pronounced due not only to greater interdependence on CPS technologies but also due to the fact that infrastructure systems are becoming more human–centric in that the user is being actively integrated into the system.

1.3 Transitioning to Human–Centric Systems

S-CPS models that incorporate customers into solutions in a dynamic, bidirectional way are becoming commonplace. For example, this phenomena can be seen in the rise of new car–sharing models, third–party energy demand response aggregators, and the proliferation of personal health–monitoring devices. Companies are beginning to capitalize on access to consumer and system data. It is necessary to have a systematic understanding of the impact of these service models on standard operations as well as the resilience and sustainability of infrastructure systems.

It is important to consider not only human–CPS coupling at the individual level but also at the societal level. At the societal level we must consider many tightly coupled decision–making loops. What is the right way of understanding the interdependencies between these various closed–loop systems and what sorts of behaviors might emerge at the global level which are not readily observable at the local level? As the number of decision makers increases yet decisions are determined based on local information (or connections) we observe different emergent global behaviors. For example, in energy S-CPS, as we move towards a Distributed Energy Resource (DER) based system, i.e. microgrids and virtual power plants, operational decisions are being made based on local information but have an impact on the overall system behavior, e.g. frequency and voltage.

The same is true for transportation systems in which local decisions about what road to take, or even what lane to drive in, impact the global flow on the network. While this has always been the case in transportation systems, now these local decisions are being supported by CPS technologies such as cellphones which provide information about the aggregate system behavior to the individual thereby shaping the decisions they make. Conversely, observations of individual decision making are often used to inform decision making at the aggregate level (e.g. loop detectors that count individual cars are used to determine traffic light control policies which, in turn, affect traffic flow).

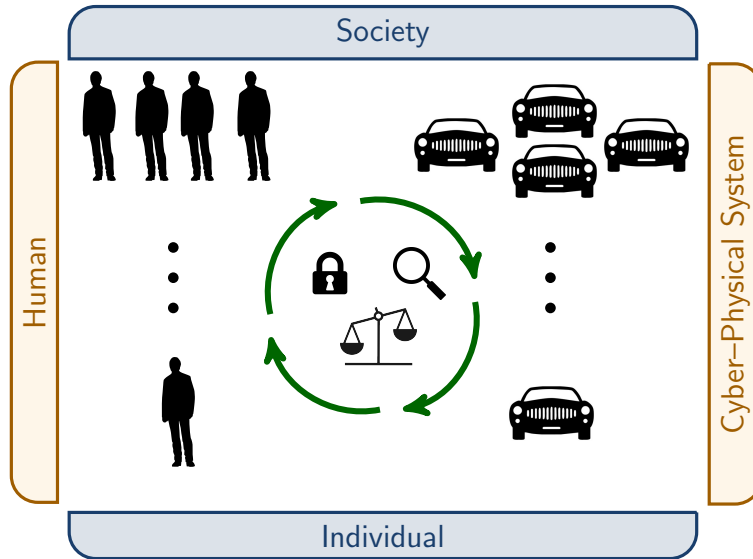


Figure 1.2: *Human-Enabled Cyber-Physical Systems* Societal-scale cyber-physical systems are becoming more distributed and even decentralized while at the same time becoming more human-centric. It is not enough to consider only individual (microscopic) or societal-scale (macroscopic) decision making in isolation. The coupling between the microscopic and macroscopic decision making presents unique and challenging problems in S-CPS for addressing system vulnerabilities and implementing policy and regulations.

Moreover, macroscopic data or aggregated data is often used to inform laws, regulations, and policies which then shape the way individuals interface with CPS infrastructure. Privacy and security laws and regulations are a prime example of this phenomena. Hence, there is a tight coupling between local decision making and global behavior (see Figure 1.2). A key observation in the move from traditional infrastructure systems and modern infrastructure systems as S-CPS is that *as constituents of S-CPS, particularly consumers, become more aware of the value of their data and vulnerabilities, they need to be compensated more to participate.*

1.4 Systems Theory for S-CPS

Realizing a systems theory for S-CPS in light of the ontological view we have constructed thus far is a broad agenda that will require contributions from a multi-disciplinary team of researchers over many years. In this dissertation, we focus on some fundamental problems that will likely contribute to this broader agenda. In particular, we develop a set of game-theoretic tools for addressing problems of adverse selection that arise in S-CPS. Referring back to Figure 1.1, we focus mainly on the interaction between the provider and its user

population. However, one can imagine the tools generalizing to other parts of the extended picture. While this is a small piece of this much larger S-CPS framework that needs to be developed, it is a foundational one and has the potential to inform the way we think about problems in the broader S-CPS context.

1.4.1 A Motivating Vignette

Consider a scenario in which there is a number of non-cooperative, self-interested agents that make up a society. They are all interested in sharing the consumption of some scarce resource. Imagine, for instance, that these agents are electricity consumers and the scarce resource is power or that these agents are drivers on a road network where capacity limits constrain the system. These agents strategically interact in competing over the scarce resource. The outcome may, for instance, be a Nash equilibrium. It is well known that Nash equilibria are not socially optimal (see, e.g. [Var04]) meaning that often the natural solution to which competing agents arrive is not efficient from a societal point of view.

Consider now a central planner who is tasked with coordinating these individuals around a more efficient outcome that is perhaps more desirable or even socially optimal in the sense that it maximizes social welfare. For instance, the central planner could be the power company or a local transportation authority. However, this planner does not know the underlying preferences of the agents (resulting in a problem of adverse selection). These preferences are what drive the agents' interactions to a particular outcome. Hence, in knowing them, the central planner can then shape them, resulting in changed behavior.

Often, in S-CPS, the central planner will lean on the underlying CPS infrastructure which may enable the efficiency gains that the central planner seeks. For instance, in energy systems the power provider utilizes streaming data from the smart meter to generate energy analytics which are in turn used to customize offerings to the consumer. In transportation systems, perhaps this CPS technology provides global positioning system (GPS) data from cellphones or flow data from embedded sensors in the road network. However, the reliance on this streaming data increases the potential for a privacy breach both on the part of the consumer and the system as well as increases security risk—the smart meter or flow sensor platform acts as an access point to the broader system.

Figure 1.3 is an abstraction of this vignette and it summarizes the contributions of this thesis. Keep this picture in mind as we walk through these contributions and as we move through the remainder of this thesis.

1.4.2 Contributions

Beginning with a class of games of *perfect information* on non-convex strategy spaces, in Chapter 2, we characterize the outcome of the players' strategic interaction by defining the *differential Nash equilibrium* concept. We show that this characterization is amenable to computation and that as such it has utility as a tool for analysis and synthesis in engineering problems with competitive agents. In Chapter 3, we consider a class of problems in which a

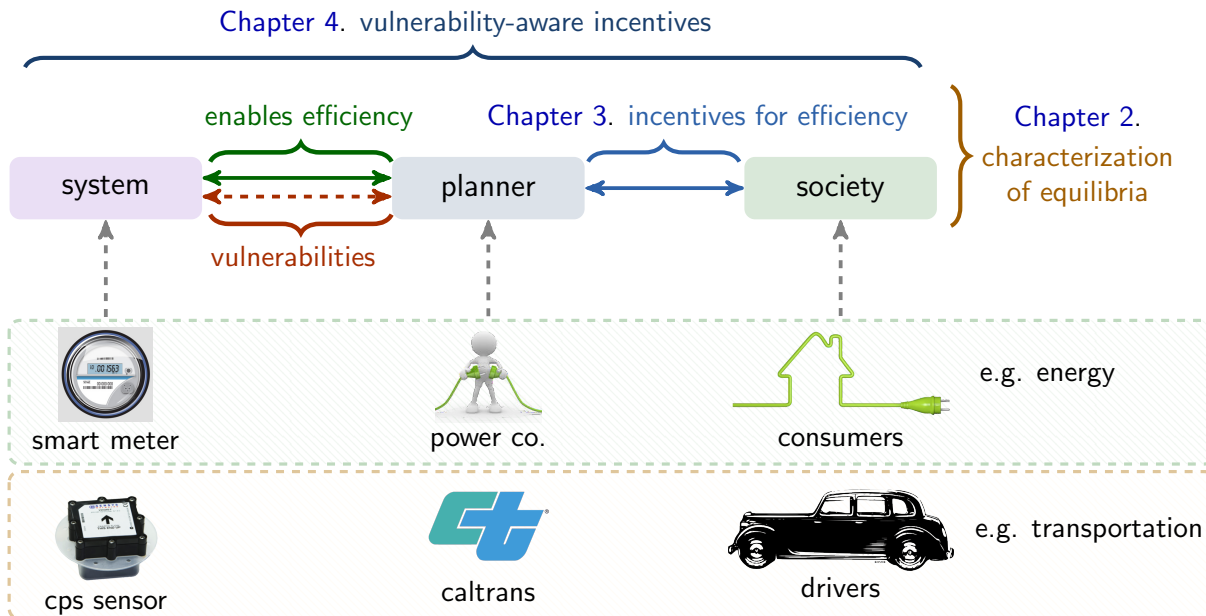


Figure 1.3: *Vignette and Contributions—Addressing the Efficiency–Vulnerability Tradeoff.* We depict the three main contributions of this dissertation numbered by the respective chapters and an abstraction of the motivating vignette.

central planner is tasked with coordinating possibly many non-cooperative agents but lacks knowledge of their preferences. We utilize the differential Nash equilibrium concept to define a utility learning and incentive design algorithm. Narrowing our focus to the demand-side of the smart grid in Chapter 4, we analyze consumer privacy, derive inference based privacy metrics, and design privacy-based service contracts.

Characterization and Computation of Local Nash Equilibria

Verifying that a strategy constitutes a Nash equilibrium in games with continuous strategy spaces requires testing that a non-convex inequality is satisfied on an open set, a generally intractable task. Further, it is often the case that strategy spaces are non-convex and agents are of bounded rationality both of which result in myopic play. We define a representation of local Nash equilibria—differential Nash equilibria—that is characterized by first- and second-order conditions on the players’ objective functions thereby allowing for the replacement of arguably computationally intractable inequalities defining Nash equilibria with conditions that only need to be checked at a single point. Additionally, we show that such equilibria are generic and structurally stable which implies that local Nash equilibria in an open-dense set of continuous games are non-degenerate—hence, isolated—differential Nash equilibria, and furthermore these equilibria persist under perturbations to player costs. As a

consequence, small modeling errors or environmental disturbances generally do not result in games with drastically different equilibrium behavior—a desirable property when considering a planner trying to induce such an equilibrium.

Utility Learning and Incentive Design Algorithm

Given that the Nash equilibrium agents reach playing according to their own selfish objectives is generally inefficient from a societal point of view, we consider a central planner who is tasked with coordinating the agents yet does not know their preferences. In this setting, we assume parameterized objective functions where the parameters are unknown to the planner. We design an algorithm for iteratively estimating these parameters and designing incentives. We consider both the case where the agents play according to a Nash equilibrium strategy and the case where the agents play according to a myopic update rule such as approximate best response. Both cases are cast in a unified framework and we formulate an online learning algorithm for estimating the parameters. Under reasonable assumptions we provide convergence results for the algorithm. The techniques used have strong ties to adaptive control and online learning or convex optimization.

Privacy–Aware Incentive Design

Finally, we focus on the demand–side of the smart grid in order to examine the kinds of vulnerabilities that arise when utilizing CPS infrastructure to support incentive design and physical control in S-CPS. In particular, we design privacy contracts and insurance for demand–side management. This work was in part motivated by our previous work on the fundamental limits of non–intrusive load monitoring in which we derived a bound on the ability of an adversary to successfully distinguish between hypotheses on consumer behavior [Don+13b; Don+13a]. This bound was then reinterpreted as an inferential metric for privacy. This led to work on quantification of the efficiency–privacy tradeoff [Don+14]. Given this fundamental tradeoff between efficiency and privacy, we design privacy–aware contracts to help manage it. The power company faces a problem of adverse selection; it does not know how consumers value privacy, but desires to have high-fidelity data for operations, e.g. to design incentives for energy efficient behavior or demand response programs. As a result, electricity service is offered as a product line differentiated according to privacy where consumers can self–select the level of privacy that fits their needs and wallet. We show the impact of viewing privacy as a good on the smart grid by studying fundamental issues such as the effects of risk and the distribution of types—the former has implications for security/insurance investment and the latter impacts privacy—on social welfare, and efficiency.

Chapter 2

Characterization and Computation of Local Nash Equilibria

Many engineering systems are complex networks in which intelligent actors make decisions regarding usage of shared, yet scarce, resources. Thinking back to the vignette we introduced in Section 1.4.1, in the transportation example, the drivers on a road network are vying over space as they navigate to their destination as quickly as possible while remaining safe. In the energy example, consumers are deciding when to use energy or even when to sell energy back to the grid and these decisions depend on the available energy supply which is naturally constrained and the price.

Game theory provides established techniques for modeling competitive interactions that have emerged as tools for analysis and synthesis of systems comprised of dynamically-coupled decision-making agents possessing diverse and oft-opposing interests (see, e.g. [Fri+12; SA05]). We focus on games with a finite number of agents where their strategy spaces are continuous, either a finite-dimensional differentiable manifold or an infinite-dimensional Banach manifold.

Previous work on continuous games with convex strategy spaces and player costs led to global characterization and computation of Nash equilibria [Baş87; Con+04; LB87]. Adding constraints led to extensions of nonlinear programming concepts, such as constraint qualification conditions, to games with generalized Nash equilibria [Dor+13; Fac+07; Ros65]. Imposing a differentiable structure on the strategy spaces yielded other global conditions ensuring existence and uniqueness of Nash equilibria and Pareto optima [Eke74; Sma75; Tho74]. In contrast, we aim to analytically characterize and numerically compute *local* Nash equilibria in continuous games on non-convex strategy spaces.

Bounding the rationality of agents can result in *myopic* behavior [Flå98; Flå99; Flå02], meaning that agents seek strategies that are optimal locally but not necessarily globally. Further, it is common in engineering applications for strategy spaces or player costs to be non-convex, for example when an agent's configuration space is a constrained set or a differentiable manifold [KK02; ME05]. These observations suggest that techniques for characterization and computation of local Nash equilibria have important practical applications.

Motivated by systems with myopic agents and non-convex strategy spaces, we seek an intrinsic characterization for local Nash equilibria that is structurally stable and amenable to computation. By generalizing derivative-based conditions for local optimality in nonlinear programming [Ber99] and optimal control [Pol97], we provide necessary first- and second-order conditions that local Nash equilibria must satisfy, and further develop a second-order sufficient-condition ensuring player strategies constitute a local Nash equilibrium. We term points satisfying this sufficient-condition *differential* Nash equilibria. In contrast to a pure optimization problem, this second-order condition is insufficient to guarantee a differential Nash equilibrium is isolated; in fact, games may possess a continuum of differential Nash equilibria. Hence, we introduce an additional second-order condition ensuring a differential Nash equilibrium is isolated.

Verifying that a strategy constitutes a Nash equilibrium in non-trivial strategy spaces requires testing that a non-convex inequality constraint is satisfied on an open set, a task we regard as generally intractable. In contrast, our sufficient conditions for local Nash equilibria require only the evaluation of player costs and their derivatives at single points. Further, our framework allows for numerical computations to be carried out when players' strategy spaces and cost functions are non-convex. Hence, we provide tractable tools for characterization and computation of differential Nash equilibria in continuous games.

We show that non-degenerate differential Nash equilibria are structurally stable. Consequently, model uncertainty or error that gives rise to a nearby game does not result in drastically different equilibrium behavior. This *structural stability* property is desirable in both the design of games as well as inverse modeling of agent behavior in competitive environments which will be discussed in detail in Chapter 3. We provide sufficient conditions ensuring that the flow generated by the gradient of each player's cost (*gradient play* [SA05] or *myopic tâtonnement* [Wal26]) converges locally to a differential Nash equilibrium. Structural stability ensures that following the flow generated by the gradient of each player's cost converges locally to a stable, non-degenerate differential Nash equilibrium.

Further, we show that non-degenerate differential Nash equilibria are generic among local Nash equilibria for games on finite-dimensional manifolds meaning that there is an open-dense set of games for which its local Nash equilibria are, in fact, non-degenerate differential Nash equilibria and therefore structurally stable and computable. Informally speaking, the fact that this set of games is *open* gives us a sense of stability and the fact that it is *dense* gives us a sense of *for-almost-all*.

The general game formulation is presented in Section 2.1. We follow with the characterization of local Nash equilibria in Section 2.2. In Section 2.3, we show that the characterization we provide—non-degenerate differential Nash equilibrium—is structurally stable. We show that non-degenerate differential Nash equilibria in finite-dimensional games are generic in Section 2.4 and we define potential games on non-convex strategy spaces as well as explore an interesting example of coupled oscillators—which is the preferred mathematical abstraction in many engineering applications—in Section 2.5. In Section 2.6, by taking a dynamical systems point of view, we show that non-degenerate differential Nash are amenable to computation using approximate myopic best response (*gradient play*). Further, we ex-

plore computation of non-degenerate differential Nash equilibria through several examples. Throughout the chapter we carry running examples that provide insight into the importance of the results and in Section 2.7 we return to the examples and highlight the importance of the results on incentive design. We remark that the results of this chapter provide the foundational tools and support for the incentive design and utility learning problem as presented in the sequel. Finally, we conclude with discussion in Section 2.8. The standard mathematical objects used throughout can be found in Appendix 2.A at the end of this chapter.

2.1 Game Formulation

The theory of games we consider concerns interaction between a finite number of rational agents that we refer to as *players*.

Consider a *complete information* game in which we have n selfish players with competing interests. The strategy spaces are topological spaces M_i for each $i \in \{1, \dots, n\}$. Note these can be finite-dimensional smooth manifolds or infinite-dimensional Banach manifolds. We denote the joint strategy space by $M = \prod_{i=1}^n M_i$. The players are each interested in minimizing a cost function representing their interests by choosing an element from their strategy space. We define player i 's cost to be a twice-differentiable function $f_i \in C^2(M, \mathbb{R})$.

The following definition describes the equilibrium behavior we are interested in:

Definition 2.1.1. *A strategy $(u_1, \dots, u_n) \in M$ is a **local Nash equilibrium** if there exist open sets $W_i \subset M_i$ such that $u_i \in W_i$ and for each $i \in \{1, \dots, n\}$,*

$$f_i(u_1, \dots, u_i, \dots, u_n) \leq f_i(u_1, \dots, u'_i, \dots, u_n), \quad \forall u'_i \in W_i \setminus \{u_i\}. \quad (2.1.1)$$

*If the above inequalities are strict, then we say (u_1, \dots, u_n) is a **strict local Nash equilibrium**. If $W_i = M_i$ for each i , then (u_1, \dots, u_n) is a **global Nash equilibrium**.*

Simply put, the above definition says that no player can unilaterally deviate from the Nash strategy and decrease her cost.

Prior to moving on to the characterization of local Nash equilibria, we describe the types of games the results apply to and why they are important in engineering applications.

Continuous games with finite-dimensional strategy spaces are described by the player strategy spaces M_1, \dots, M_n and their cost functions (f_1, \dots, f_n) . They arise in a number of engineering and economic applications, for instance, in modeling one-shot decision making problems arising in transportation, communication and power networks [Kri+14; Can+10; Par+01], or mixed strategies over discrete strategy spaces [SA05]. On the other hand, the consideration of mixed strategies in games with continuous finite-dimensional strategy spaces lead to games on infinite-dimensional strategy spaces. In particular, the mixed strategies are probability measures over the pure strategies [Gli52].

Continuous games with infinite-dimensional strategy spaces, regarded as open-loop differential games, are used in engineering applications in which there are agents coupled

through dynamics [Ba95]. They arise in problems such as energy management in buildings [Coo+13], travel-time optimization in transportation networks [BH12], and integration of renewables into energy systems [Zhu+12a].

Open-loop differential games often come in the following form. Let $L_2[0, T]$ denote the space of square integrable functions from $[0, T] \subset \mathbb{R}$ into \mathbb{R}^m . For an n -player game, strategy spaces are Banach manifolds, M_i for $i \in \{1, \dots, n\}$, modeled on $L_2[0, T]$. For each $t \in [0, T]$, let $x(t) \in \mathbb{R}^n$ denote the state of the game. The state evolves according to the dynamics

$$\dot{x}(t) = h(x(t), u_1(t), \dots, u_n(t)) \quad \forall t \in [0, T] \quad (2.1.2)$$

where $u_i \in M_i$ is player i 's strategy. We assume that $h(x, u_1, \dots, u_n)$ is continuously differentiable, globally Lipschitz continuous and all the derivatives in all its arguments are globally Lipschitz continuous. We denote by $f_i(u_1, \dots, u_n) = \hat{f}_i(x^{(x(0), u_1, \dots, u_n)}(T))$ player i 's cost function. The superscript notation on the state x indicates the dependence of the state on the initial state and the strategies of the players. Each \hat{f}_i is assumed twice continuously differentiable so that each f_i is C^2 -Fréchet-differentiable [Pol97, Thm. 5.6.10]. We pose each player's optimization problem as

$$\min_{u_i} \hat{f}_i(x^{(x(0), u_1, \dots, u_i, \dots, u_n)}(T)). \quad (2.1.3)$$

The co-state for player i evolves according to

$$\dot{p}_i(t) = -p_i(t) \frac{\partial h}{\partial x}(x(t), u_1(t), \dots, u_i(t), \dots, u_n(t)) \quad (2.1.4)$$

with final time condition

$$p_i(T) = D_x f_i(x^{(x(0), u_1, \dots, u_i, \dots, u_n)}(T)). \quad (2.1.5)$$

The derivative of the i -th player's cost function is given by

$$(D_i f_i)(t) = p_i(t) \frac{\partial h}{\partial u_i}(x(t), u_1, \dots, u_i(t), \dots, u_n(t)). \quad (2.1.6)$$

Remark 2.1.1. *We have formulated the open-loop differential game in terms of final-time cost optimization problems for each player. However, there is a transformation from running cost problems into final-time cost problems [Pol97, Chapter 4, §1].*

Before we dive into the details, let us consider a couple of simple examples that exhibit very interesting behavior. We return to these examples throughout the chapter as they highlight different aspects and the importance of our characterization of Nash equilibria.

Example 2.1 (Betty–Sue: Thermodynamic Coupling). *Consider a two player game between Betty and Sue. Let Betty's strategy space be $M_1 = \mathbb{R}$ and Sue's strategy space be $M_2 = \mathbb{R}$. Furthermore, let Betty's cost function be defined by*

$$f_1(u_1, u_2) = \frac{u_1^2}{2} - u_1 u_2$$

and let Sue’s cost function be defined by

$$f_2(u_1, u_2) = \frac{u_2^2}{2} - u_1 u_2.$$

This game can be thought of as an abstraction of two agents in a building occupying adjoining rooms. The first term in each of their costs represents an energy cost and the second term is a cost from thermodynamic coupling. The agents try to maintain the temperature at a desired set–point in thermodynamic equilibrium.

Definition 2.1.1 specifies that a point (μ_1, μ_2) is a Nash equilibrium if no player can unilaterally deviate and decrease their cost, i.e. $f_1(\mu_1, \mu_2) < f_1(u_1, \mu_2)$ for all $u_1 \in \mathbb{R}$ and $f_2(\mu_1, \mu_2) < f_2(\mu_1, u_2)$ for all $u_2 \in \mathbb{R}$.

Fix Sue’s strategy $u_2 = \mu_2$, and calculate

$$D_1 f_1 = \frac{\partial f_1}{\partial u_1} = u_1 - \mu_2 \tag{2.1.7}$$

Then, Betty’s optimal response to Sue playing $u_2 = \mu_2$ is $u_1 = \mu_2$. Similarly, if we fix $u_1 = \mu_1$, then Sue’s optimal response to Betty playing $u_1 = \mu_1$ is $u_2 = \mu_1$. For all $u_1 \in \mathbb{R} \setminus \{\mu_2\}$

$$-\frac{\mu_2^2}{2} < \frac{u_1^2}{2} - u_1 \mu_2 \tag{2.1.8}$$

so that $f_1(\mu_2, \mu_2) < f_1(u_1, \mu_2)$ for all $u_1 \in \mathbb{R} \setminus \{\mu_2\}$. Again, similarly, for all $u_2 \in \mathbb{R} \setminus \{\mu_1\}$

$$-\frac{\mu_1^2}{2} < \frac{u_2^2}{2} - u_2 \mu_1 \tag{2.1.9}$$

so that $f_2(\mu_1, \mu_1) < f_2(\mu_1, u_2)$ for all $u_2 \in \mathbb{R} \setminus \{\mu_1\}$. Hence, all the points on the line $u_1 = u_2$ in $M_1 \times M_2 = \mathbb{R}^2$ are strict local Nash equilibria—in fact, they are strict global Nash equilibria.

□

As the above example shows, continuous games can exhibit a continuum of equilibria. The following example will show this *pathology* is not limited to trivial strategy spaces.

The next example we consider is of coupled oscillators viewed as a game. Not only is the following an example of a game on non–trivial strategy spaces, as we will show, it exhibits a cadre of interesting characteristics including a continuum of global Nash equilibria. It is an important example because coupled oscillator models are the preferred mathematical abstraction for many engineering applications.

Coupled oscillator models—in particular, the Kuramoto oscillator model [Kur75]—are used widely for modeling attitude control and coordination (satellites, aircraft, etc.) [Cha+11; WKD91; Zou+12], generators synchronizing with the energy grid (microgrids) [Dör+13; DB12], traffic light control [Coo+15; AE05], robotics [Moi+10], biological networks [Wan+15; Wan+08], healthcare (pacemakers) [Pes75; DB78], and in a general engineering context [DB11; Jad+04; Sep+05a; Sep+07; Ge+08; Pal+07]. Coupled oscillator models are often viewed in

a game theoretic context in order to gain further insight into the system properties [CB13; Lee+08; Yin+12; Got+10; ZS01; Yin+10]. For example, both the problem of multiple satellites trying to synchronize and the problem of generators trying to synchronize their frequency with the larger power grid lend themselves naturally to a game theoretic framework. Indeed, consider each player to be a satellite (generator) that chooses its relative phase in an effort to synchronize but does so selfishly. Modeling the players as selfish may arise because they in fact have a desire to minimize their own effort regardless of other participants. On the other hand, it may arise through abstraction of information constraints or the myopic nature of the players.

We will consider two coupled oscillators and return to the general n -coupled oscillator problem in Section 2.5.

Example 2.2 (Jean–Paul: Coupled Oscillators). *Consider two coupled oscillators managed by Jean and Paul respectively. We denote the phase of oscillator i by $\theta_i \in \mathbb{S}^1$. Let Jean’s oscillator have phase θ_1 and his cost be given by*

$$f_1(\theta_1, \theta_2) = -\frac{1}{2} \cos(\theta_1 - \theta_2). \quad (2.1.10)$$

Similarly, let Paul’s oscillator have phase θ_2 and his cost be given by

$$f_2(\theta_1, \theta_2) = -\frac{1}{2} \cos(\theta_2 - \theta_1). \quad (2.1.11)$$

We show that this game has a continuum of global Nash equilibria. In particular, all points in the set

$$\{(\theta_1, \theta_2) \in \mathbb{S}^1 \times \mathbb{S}^1 \mid \theta_1 = \theta_2\}$$

are global Nash equilibria. Indeed, consider the points $\theta_1 - \theta_2 = 0$. First, since

$$D_1 f_1(\theta_1, \theta_2) = \frac{1}{2} \sin(\theta_1 - \theta_2) = 0,$$

$\theta_1 = \theta_2$ is a best-response to θ_2 . Similarly, since

$$D_2 f_2(\theta_1, \theta_2) = \frac{1}{2} \sin(\theta_2 - \theta_1) = 0,$$

$\theta_2 = \theta_1$ is a best-response to θ_1 . Let $\theta_2 = \beta$ be fixed. For $\theta_1 = \beta$ we have

$$-\frac{1}{2} \cos(\theta_1 - \beta) = -\frac{1}{2} \leq -\frac{1}{2} \cos(\theta'_1 - \beta), \quad \forall \theta'_1 \in \mathbb{S}^1 \setminus \{\beta\} \quad (2.1.12)$$

Similarly, let $\theta_1 = \alpha$ be fixed. For $\theta_2 = \alpha$, we have

$$-\frac{1}{2} \cos(\alpha - \theta_2) = -\frac{1}{2} \leq -\frac{1}{2} \cos(\alpha - \theta'_2), \quad \forall \theta'_2 \in \mathbb{S}^1 \setminus \{\alpha\} \quad (2.1.13)$$

Hence, the Nash equilibria of the game are exactly the points for which the coupled oscillators are synchronized in equilibrium.

On the other hand, if players are utility maximizers in the game (f_1, f_2) on $\mathbb{T}^2 = \mathbb{S}^1 \times \mathbb{S}^1$ where f_1 and f_2 are defined in (2.1.10) and (2.1.11), then it is straightforward to show (in a similar manner as above) that all points in the set

$$\{(\theta_1, \theta_2) \in \mathbb{S}^1 \times \mathbb{S}^1 \mid \theta_1 - \theta_2 = \pi\}$$

are global Nash equilibria. In other words, the set of Nash equilibria contains exactly the points for which the coupled oscillators are balanced in equilibrium.

We remark that this is quite a simple game; since cosine is an even function, $f_1 = f_2$. However, in the general n -coupled oscillator model, as we will see in Section 2.5, exhibits a continuum of Nash equilibria as well. We choose to carry this two-player example—as opposed to the n -player example—throughout the chapter because it is not only illustrative but simple enough for the reader to process while reading the text.

2.2 Characterization of Local Nash Equilibria

In this section, we characterize local Nash equilibria by paralleling results in nonlinear programming and optimal control that provide first- and second-order necessary and sufficient conditions for local optima.

The following definition of a differential game form is due to Stein [Ste10].

Definition 2.2.1. A **differential game form** is a differential 1-form $\omega : M_1 \times \cdots \times M_n \rightarrow T^*(M_1 \times \cdots \times M_n)$ defined by

$$\omega = \sum_{i=1}^n \psi_{M_i} \circ df_i. \tag{2.2.1}$$

where ψ_{M_i} are the natural bundle maps defined in (2.A.3) that annihilate those components of the covector field df_i not corresponding to M_i .

Remark 2.2.1. If each M_i is a finite-dimensional manifold of dimension m_i , then the differential game form has the following coordinate representation:

$$\omega_\varphi = \sum_{i=1}^n \sum_{j=1}^{m_i} \frac{\partial(f_i \circ \varphi^{-1})}{\partial y_i^j} dy_i^j \tag{2.2.2}$$

where (U, φ) is a product chart on M at $u = (u_1, \dots, u_n)$ with local coordinates given by $(v_1^1, \dots, v_1^{m_1}, \dots, v_n^1, \dots, v_n^{m_n})$ and where $U = \prod_{i=1}^n U_i$ and $\varphi = \times_{i=1}^n \varphi_i$. In addition, $f_i \circ \varphi^{-1}$ is the coordinate representation of f_i for $i \in \{1, \dots, n\}$. In particular, $\varphi_i(u_i) = (v_i^1, \dots, v_i^{m_i})$ where each $v_i^j : U_i \rightarrow \mathbb{R}$ is a coordinate function so that dv_i^j is its derivative. \square

The differential game form captures a differential view of the strategic interaction between the players. Indeed, ω indicates the direction in which the players can change their strategies to decrease their individual cost functions most rapidly. In particular, each player's cost function depends on its own choice variable as well as all the other players' choice variables. However, each player can only affect their payoff by adjusting their own strategy.

Definition 2.2.2. *A strategy $u = (u_1, \dots, u_n) \in M_1 \times \dots \times M_n$ is a **differential Nash equilibrium** if $\omega(u) = 0$ and $D_{ii}^2 f_i(u)$ is positive-definite for each $i \in \{1, \dots, n\}$.*

The second-order conditions used to define differential Nash equilibria are motivated by results in nonlinear programming that use first- and second-order conditions to assess whether a critical point is a local optima [Pol97], [Ber99].

The following proposition provides first- and second-order necessary conditions for local Nash equilibria. We remark that these conditions are reminiscent of those seen in nonlinear programming for optimality of critical points.

Proposition 2.2.1 ([Rat+13; Rat+14d]). *If $u = (u_1, \dots, u_n)$ is a local Nash equilibrium, then $\omega(u) = 0$ and $D_{ii}^2 f_i(u)$ is positive semi-definite for each $i \in \{1, \dots, n\}$.*

Proof. Suppose that $u = (u_1, \dots, u_n) \in M$ is a local Nash equilibrium. Then,

$$f_i(u) \leq f_i(u_1, \dots, u'_i, \dots, u_n), \quad \forall u'_i \in W_i \setminus \{u_i\} \quad (2.2.3)$$

for open $W_i \subset M_i$, $i \in \{1, \dots, n\}$. Suppose that we have a product chart (U, φ) , where $U = \prod_{i=1}^n U_i$ and $\varphi = \times_{i=1}^n \varphi_i$, such that $u \in U$.

Let $\varphi_i(u_i) = v_i$ for each i . Then, since φ is continuous, for each $i \in \{1, \dots, n\}$, we have that for all $v'_i \in \varphi_i(W_i \cap U_i) \setminus \{\varphi_i(u_i)\}$,

$$f_i \circ \varphi^{-1}(v_1, \dots, v_i, \dots, v_n) \leq f_i \circ \varphi^{-1}(v_1, \dots, v'_i, \dots, v_n). \quad (2.2.4)$$

Now, we apply Proposition 1.1.1 from [Ber99], if M_i is finite-dimensional, or Theorem 4.2.3(1) and Theorem 4.2.4(a) from [Pol97], if M_i is infinite-dimensional, to $f_i \circ \varphi^{-1}$. We conclude that for each $i \in \{1, \dots, n\}$, $D_i(f_i \circ \varphi^{-1})(v_1, \dots, v_n) = 0$ and for all $\nu \in \varphi_i(U_i \cap W_i)$,

$$D_{ii}^2(f_i \circ \varphi^{-1})(v_1, \dots, v_n)(\nu, \nu) \geq \alpha \|\nu\|^2, \quad (2.2.5)$$

i.e. it is a positive semi-definite bilinear form on $\varphi_i(U_i \cap W_i)$.

Invariance of the stationarity of critical points and the index of the Hessian with respect to coordinate change gives us $\omega(u) = 0$ and $D_{ii}^2 f_i(u)$ is a positive semi-definite for each $i \in \{1, \dots, n\}$. \square

We now show that the conditions defining a differential Nash equilibrium are sufficient to guarantee a strict local Nash equilibrium.

Theorem 2.2.1 ([Rat+13; Rat+14d]). *A differential Nash equilibrium is a strict local Nash equilibrium, i.e. if $\omega(u) = 0$ and $D_{ii}^2 f_i(u) > 0$, then u is a strict local Nash equilibrium.*

Proof. Suppose that $u = (u_1, \dots, u_n) \in M$ is a differential Nash equilibrium. Then, by the definition of differential Nash equilibrium, $\omega(u) = 0$ and $D_{ii}^2 f_i(u)$ is positive definite for each $i \in \{1, \dots, n\}$. The second-derivative conditions imply that $D_{ii}^2 (f_i \circ \varphi^{-1})(v_1, \dots, v_n)$ is a positive definite bilinear form where $v_i = \varphi_i(u_i)$ for any coordinate chart (U, φ) , with $\varphi = \times_i \varphi_i$, $U = \prod_i U_i$, and $u_i \in U_i$ for each $i \in \{1, \dots, n\}$. Using the isomorphism introduced in the appendix in (2.A.2), $\omega(u) = 0$ implies that for each $i \in \{1, \dots, n\}$, $D_i (f_i \circ \varphi^{-1})(v_1, \dots, v_n) = 0$. Let E_i be the model space, i.e. the underlying Banach space, in either the finite-dimensional or infinite-dimensional case. Applying either Proposition 1.1.3 from [Ber99] or Theorem 4.2.6 (a) from [Pol97] to to each $f_i \circ \varphi^{-1}$ with

$$(\varphi_1(u_1), \dots, \varphi_{i-1}(u_{i-1}), \varphi_{i+1}(u_{i+1}), \dots, \varphi_n(u_n))$$

fixed yields a neighborhood $W_i \subset E_i$ such that for all $v' \in W_i$,

$$f_i \circ \varphi^{-1}(v_1, \dots, v_i, \dots, v_n) < f_i \circ \varphi^{-1}(v_1, \dots, v', \dots, v_n). \quad (2.2.6)$$

Since φ is continuous, there exists a neighborhood $V_i \subset M_i$ of u_i such that for $V_i = \varphi_i^{-1}(W_i)$ and all $u'_i \in V_i \setminus \{u_i\}$,

$$f_i(u_1, \dots, u_i, \dots, u_n) < f_i(u_1, \dots, u'_i, \dots, u_n). \quad (2.2.7)$$

Therefore, differential Nash equilibria are strict local Nash equilibria. Due to the fact that both $\omega(u) = 0$ and definiteness of the Hessian are coordinate invariant, this is independent of choice of coordinate chart. \square

We remark that the conditions for differential Nash equilibria are not sufficient to guarantee that an equilibrium is isolated.

Example 2.1 (Betty–Sue: Continuum of Differential Nash). *Returning to the Betty–Sue example, we can check that at all the points such that $u_1 = u_2$, $\omega(u_1, u_2) = 0$ and $D_{ii}^2 f_i(u_1, u_2) = 1 > 0$ for each $i \in \{1, 2\}$. Hence, there is a continuum of differential Nash equilibria in this game.* \square

Example 2.2 (Jean–Paul: Continuum of Differential Nash). *Just as above, returning to the coupled oscillator example, we can check that at all the points in the set*

$$\{(\theta_1, \theta_2) \in \mathbb{S}^1 \times \mathbb{S}^1 \mid \theta_1 - \theta_2 = \pi\},$$

$\omega(\theta_1, \theta_2) = 0$ and $D_{ii}^2 f_i(\theta_1, \theta_2) = 1 > 0$ for each $i \in \{1, 2\}$. Hence, there is a continuum of differential Nash equilibria in this game. \square

We propose a sufficient condition to guarantee that differential Nash equilibria are isolated. We do so by combining ideas introduced by Rosen [Ros65] for convex games with concepts from Morse theory [Mil63], in particular second-order conditions on non-degenerate critical points of real-valued functions on manifolds.

At a differential Nash equilibrium $u = (u_1, \dots, u_n)$, consider the derivative of the differential game form

$$d\omega = \sum_{i=1}^n d(\psi_{M_i} \circ df_i). \tag{2.2.8}$$

Intrinsically, this derivative is a tensor field $d\omega \in T_2^0(M)$; at a point $u \in M$ where $\omega(u) = 0$ it is a bilinear form constructed from the uniquely determined continuous, symmetric, bilinear forms $\{d^2 f_i(u)\}_{i=1}^n$. We will refer to its local representation as the *Hessian* of the differential game form. Moreover, we want to emphasize that the derivative of the differential game form as defined above is distinct from the *exterior derivative* (see Appendix 2.A) of a differential form.

Theorem 2.2.2 ([Rat+13; Rat+14d]). *If $u = (u_1, \dots, u_n)$ is a differential Nash equilibrium and $d\omega(u)$ is non-degenerate, then u is an isolated strict local Nash equilibrium.*

Proof. Since u is a differential Nash equilibrium, Theorem 2.2.1 gives us that it is a strict local Nash equilibrium. The following argument shows that it is isolated. Non-degeneracy of $d\omega(u)$ at a critical point is invariant with respect to the choice of coordinates. It suffices to choose a coordinate chart (U, φ) with $\varphi = \times_{i=1}^n \varphi_i$ and $U = \prod_{i=1}^n U_i$ and show the result with respect to φ . Let E denote the underlying model space of the manifold $M_1 \times \dots \times M_n$. Define the map $g : E \rightarrow E$ by

$$g(\varphi(u)) = \sum_{i=1}^n D_i(f_i \circ \varphi^{-1})(\varphi(u)) \tag{2.2.9}$$

Note that g is the coordinate representation of the differential game form ω . Zeros of the function g define critical points of the game and its derivative at critical points is $d\omega_\varphi$. Since u is a differential Nash equilibrium, $\omega(u) = 0$. Further, since $d\omega_\varphi(u)$ is non-degenerate—the map $A(v)(w) = d\omega_\varphi(u)(v, w)$ is a linear isomorphism—we can apply the Inverse Function Theorem [Abr+88, Thm. 2.5.2] to get that g is a local diffeomorphism at u , i.e. there exists an open neighborhood V of u such that the restriction of g to V establishes a diffeomorphism between V and an open subset of E . Thus, only $\varphi(u)$ could be mapped to zero near $\varphi(u)$. Non-degeneracy of $d\omega(u)$ is invariant with respect to choice of coordinates. Therefore, independent of the choice of φ , u is isolated. \square

Definition 2.2.3. *Differential Nash equilibria $u = (u_1, \dots, u_n)$ such that $d\omega(u)$ is non-degenerate are termed **non-degenerate differential Nash equilibria**.*

Example 2.1 (Betty–Sue: Degeneracy and Breaking Symmetry). *Return again to the Betty–Sue example in which we showed that there is a continuum of Nash equilibria; in fact, all the points on the line $u_1 = u_2$ are differential Nash equilibria and at each of these points we have*

$$d\omega(u_1, u_2) = \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix} \tag{2.2.10}$$

so that $\det(d\omega(u_1, u_2)) = 0$. Hence, all of the equilibria are degenerate. By breaking the symmetry in the game, we can make $(0, 0)$ a non-degenerate differential Nash equilibrium; i.e. we can remove all but one of the equilibria. Indeed, let Betty's cost be given by $\tilde{f}_1(u_1, u_2) = \frac{u_1^2}{2} - au_1u_2$ and let Sue's cost remain unchanged. Then the local representation of the derivative of the differential game form $\tilde{\omega}$ of the game (\tilde{f}_1, f_2) is

$$d\tilde{\omega}(u_1, u_2) = \begin{bmatrix} 1 & -a \\ -1 & 1 \end{bmatrix} \tag{2.2.11}$$

Thus for any value of $a \neq 1$, $(0, 0)$ is a non-degenerate differential Nash equilibrium. This shows that small modeling errors can remove degenerate differential Nash equilibria. \square

Example 2.2 (Jean–Paul: Degeneracy and Breaking Symmetry). *Returning to the coupled oscillator example, recall that we showed there was a continuum of differential Nash equilibria. However, they are not non-degenerate since $\det(d\omega(\theta_1, \theta_1)) = 0$. Just as in the Betty–Sue example, by breaking the symmetry in the game, we can make $(0, 0)$ a non-degenerate differential Nash equilibrium. Indeed, consider the game (\tilde{f}_1, f_2) where $\tilde{f}_1 = \frac{1}{2} \cos(b\theta_1 - \theta_2)$. Then the local representation of the differential game form $\tilde{\omega}$ of the game (\tilde{f}_1, f_2) is*

$$\begin{bmatrix} -\frac{1}{2}b^2 \cos(b\theta_1 - \theta_2) & \frac{1}{2}b \cos(b\theta_1 - \theta_2) \\ \frac{1}{2} \cos(\theta_1 - \theta_2) & -\frac{1}{2} \cos(\theta_1 - \theta_2) \end{bmatrix} \tag{2.2.12}$$

Hence, for any value of $b \notin \{0, 1\}$, $(0, 0)$ is a non-degenerate differential Nash equilibrium. We remark that if players are utility maximizers in the same game, it is straightforward to show that $(0, \pi)$ is a non-degenerate differential Nash using the same perturbation. Why we make this distinction between utility maximization and cost minimization in this example will become clear when we return to this example and show that minimizing (maximizing) corresponds to synchronizing (balancing) the phases of the coupled oscillators.

In a neighborhood of a non-degenerate differential Nash equilibrium there are no other Nash equilibria. This property is desirable particularly in applications where a central planner is designing incentives to induce a socially optimal or otherwise desirable equilibrium that optimizes the central planner's cost; if the desired equilibrium resides on a continuum of equilibria, then due to measurement noise or myopic play, agents may be induced to play a nearby equilibrium that is suboptimal for the central planner. In Section 2.7, we extend Example 2.1 and 2.2 by introducing a central planner. But first, we show that non-degenerate differential Nash equilibria are structurally stable.

2.3 Structural Stability

Examples demonstrate that global Nash equilibria may fail to persist under arbitrarily small changes in player costs [Eke74]. A natural question arises: do local Nash equilibria persist under perturbations? Applying structural stability analysis from dynamical systems theory,

we answer this question affirmatively for non-degenerate differential Nash equilibria subject to smooth perturbations in player costs.

Let $M = M_1 \times \cdots \times M_n$ and $f_1, \dots, f_n : M \rightarrow \mathbb{R}$ be C^2 player cost functions, $\omega : M \rightarrow T^*M$ the associated differential game form (2.2.1), and suppose $u \in M$ is a non-degenerate differential Nash equilibrium, i.e. $\omega(u) = 0$ and $d\omega(u)$ is non-degenerate. We show that for all $\tilde{f}_i \in C^\infty(M, \mathbb{R})$ sufficiently close to f_i there exists a unique non-degenerate differential Nash equilibrium $\tilde{u} \in M$ for $(\tilde{f}_1, \dots, \tilde{f}_n)$ near u .

Proposition 2.3.1 (Parameterized Structural Stability [Rat+14b]). *Non-degenerate differential Nash equilibria are parametrically structurally stable: given $f_1, \dots, f_n \in C^2(M, \mathbb{R})$, $\zeta_1, \dots, \zeta_n \in C^2(M, \mathbb{R})$, and a non-degenerate differential Nash equilibrium $u \in M$ for (f_1, \dots, f_n) , there exist neighborhoods $U \subset \mathbb{R}$ of 0 and $W \subset M$ of u such that for all $s \in U$ there exists a unique non-degenerate differential Nash equilibrium $\tilde{u}(s) \in W$ for $(f_1 + s\zeta_1, \dots, f_n + s\zeta_n)$.*

Proof. Define $\tilde{f}_j : M_1 \times \cdots \times M_n \times \mathbb{R} \rightarrow \mathbb{R}$ by

$$\tilde{f}_j(u, s) = f_j(u) + s\zeta_j(u)$$

and $\tilde{\omega} : M_1 \times \cdots \times M_n \times \mathbb{R} \rightarrow T^*(M_1 \times \cdots \times M_n)$ by

$$\tilde{\omega}(u, s) = \sum_{i=1}^n \tilde{\psi}_{M_i} \circ d\tilde{f}_i(u, s)$$

for all $s \in \mathbb{R}$ and $u \in M_1 \times \cdots \times M_n$ and where $\tilde{\psi}_{M_i} : T^*(M_1 \times \cdots \times M_n \times \mathbb{R}) \rightarrow T^*(M_1 \times \cdots \times M_n \times \mathbb{R})$. Observe that $D_1\tilde{\omega}((u_1, \dots, u_n), 0)$ is invertible since u is a non-degenerate differential Nash equilibrium for (f_1, \dots, f_n) . Therefore by the Implicit Function Theorem [Abr+88, Prop. 3.3.13 (iii)], there exist neighborhoods $V \subset \mathbb{R}$ of 0 and $W \subset M$ of u and a smooth function $\sigma \in C^\infty(V, W)$ such that

$$\forall s \in V, u \in W : \tilde{\omega}(u, s) = 0 \iff u = \sigma(s).$$

Furthermore, since $\tilde{\omega}$ is continuously differentiable, there exists a neighborhood $U \subset V$ of 0 such that $d\tilde{\omega}(\sigma(s), s)$ is invertible for all $s \in U$. We conclude for all $s \in U$ that $\sigma(s) \in M$ is the unique Nash equilibrium for $((f_1 + s\zeta_1)|_W, \dots, (f_n + s\zeta_n)|_W)$, and furthermore that $\sigma(s)$ is a non-degenerate differential Nash equilibrium. \square

We remark that the preceding analysis extends directly to any finitely-parameterized perturbation. For an arbitrary perturbation, we have the following.

Theorem 2.3.1 (Structural Stability [Rat+14b; Rat+14d]). *Non-degenerate differential Nash equilibria are structurally stable: let $u \in M$ be a non-degenerate differential Nash equilibrium for $(f_1, \dots, f_n) \in C^2(M, \mathbb{R}^n)$. Then there exist neighborhoods $U \subset C^2(M, \mathbb{R}^n)$ of (f_1, \dots, f_n) and $W \subset M$ of u and a C^2 Fréchet-differentiable function $\sigma \in C^2(U, W)$ such that for all $(\tilde{f}_1, \dots, \tilde{f}_n) \in U$ the point $\sigma(\tilde{f}_1, \dots, \tilde{f}_n)$ is the unique non-degenerate differential Nash equilibrium for $(\tilde{f}_1, \dots, \tilde{f}_n)$ in W .*

Proof. Consider the operator $\Omega \in C^1(C^1(M, \mathbb{R}^n) \times M, \mathbb{R}^n)$ defined by

$$\Omega((\tilde{f}_1, \dots, \tilde{f}_n), (u_1, \dots, u_n)) = \sum_{i=1}^n \psi_{M_i} \circ d\tilde{f}_i(u_1, \dots, u_n). \quad (2.3.1)$$

Note that the right-hand side is the differential game form $\tilde{\omega}(u_1, \dots, u_n)$ for the game $(\tilde{f}_1, \dots, \tilde{f}_n)$. Suppose that $u = (u_1, \dots, u_n)$ is a non-degenerate differential Nash equilibrium. A straightforward application of Proposition 2.4.20 [Abr+88] implies that the operator Ω is C^1 Fréchet-differentiable. In addition,

$$D_2\Omega((f_1, \dots, f_n), (u_1, \dots, u_n)) = d\omega(u_1, \dots, u_n). \quad (2.3.2)$$

Since $d\omega(u)$ is an isomorphism by assumption, we can apply the Implicit Function Theorem [Abr+88, Prop. 3.3.13 (iii)] to Ω to get an open neighborhood $W \subset M$ of u and $V \subset C^2(M, \mathbb{R}^n)$ of (f_1, \dots, f_n) and a smooth function $\sigma \in C^2(V, W)$ such that

$$\forall \tilde{f} \in V, v \in W : \Omega(\tilde{f}, v) = 0 \iff v = \sigma(\tilde{f})$$

where $\tilde{f} = (\tilde{f}_1, \dots, \tilde{f}_n)$. Furthermore, since Ω is continuously differentiable, there exists a neighborhood $U \subset V$ of (f_1, \dots, f_n) such that $d\Omega(\tilde{f}, \sigma(\tilde{f}))$ is invertible for all $\tilde{f} \in U$. Thus, for all $\tilde{f} \in U$, $\sigma(\tilde{f}) \in M$ is the unique non-degenerate differential Nash equilibrium in W . \square

Let us return to Example 2.1 and examine what can happen in the degenerate case.

Example 2.1 (Betty–Sue: Structural Instability). *Let us recall again the Betty–Sue example in which we have a game admitting a continuum of differential Nash equilibria. An arbitrarily small perturbation can make all the equilibria disappear. Ineed, let $\varepsilon \neq 0$ be arbitrarily small and consider Betty’s perturbed cost function*

$$\tilde{f}_1(u_1, u_2) = \frac{u_1^2}{2} - u_1u_2 + \varepsilon u_1. \quad (2.3.3)$$

Let Sue’s cost function remain unchanged. A necessary condition that a Nash equilibrium $(u_1, u_2) \in M_1 \times M_2$ must satisfy is $\omega(u_1, u_2) = 0$ (see Proposition 2.2.1) thereby implying $D_1\tilde{f}_1(u_1, u_2) = u_1 - u_2 + \varepsilon = 0$ and $D_2f_2(u_1, u_2) = u_2 - u_1 = 0$. This can only happen for $\varepsilon = 0$. Hence, any perturbation εu_1 with $\varepsilon \neq 0$ will remove all the Nash equilibria. \square

Example 2.2 (Jean–Paul: Structural Instability). *Just as in the Betty–Sue example, we can show that an arbitrary smooth perturbation to the game can remove all the Nash equilibria. Consider the game (\tilde{f}_1, f_2) where $\tilde{f}_1 = \frac{1}{2} \cos(\theta_1 - \theta_2) + \varepsilon \theta_1$ for $\varepsilon \neq 0$. Then, since $\omega(\theta_1, \theta_2) = 0$ is a necessary condition for Nash, it is easy to see that for any point (θ_1, θ_2) to be a Nash equilibrium, $\varepsilon = 0$.*

2.4 Genericity

In this section, we show local Nash equilibria are generically non-degenerate differential Nash equilibria in games on finite-dimensional manifolds; there is an open-dense set of games whose local Nash equilibria are non-degenerate differential Nash equilibria. Throughout this section we will assume the strategy spaces of players are finite-dimensional unless otherwise stated.

As we showed in the previous section, non-degenerate differential Nash equilibria are amenable to computation since they satisfy first- and second-order conditions reminiscent of those from nonlinear programming. Continuing with the dynamical systems theory perspective, the following result is analogous to the fact that non-degenerate singularities are generic [BT10].

Theorem 2.4.1 (Genericity [Rat+14b]). *Non-degenerate differential Nash equilibria are generic among local Nash equilibria: for any smooth boundaryless manifolds M_1, \dots, M_n there exists an open-dense subset $G \subset C^\infty(M_1 \times \dots \times M_n, \mathbb{R}^n)$ such that for all $(f_1, \dots, f_n) \in G$, if $(u_1, \dots, u_n) \in M_1 \times \dots \times M_n$ is a local Nash equilibrium for (f_1, \dots, f_n) , then (u_1, \dots, u_n) is a non-degenerate differential Nash equilibrium for (f_1, \dots, f_n) .*

Proof. For the sake of ease of presentation, we present the proof only for the two player case and remark that extended it is straightforward after several algebraic manipulations.

Consider a two-player game where player i 's cost function is $f_i \in C^\infty(M_1 \times M_2, \mathbb{R})$. Let $J^2(M_1 \times M_2, \mathbb{R}^2)$ denote the second order jet bundle containing 2-jets $j^2 f$ such that $f = (f_1, f_2) : M_1 \times M_2 \rightarrow \mathbb{R}^2$. Let (U, φ) be a product chart on $M_1 \times M_2$ that contains (μ_1, μ_2) . The dimensions of M_1 and M_2 are m_1 and m_2 respectively and we define $m = m_1 + m_2$. We define $S(m)$ to be the symmetric $m \times m$ matrices as follows

$$S(m) = \{A \in \mathbb{R}^{m \times m} \mid A = A^T\}. \quad (2.4.1)$$

For $(A_1, A_2) \in S(m)^2$, we can partition each A_i as follows:

$$A_i = \begin{bmatrix} A_i^{11} & A_i^{12} \\ A_i^{21} & A_i^{22} \end{bmatrix} \quad (2.4.2)$$

where $A_i^{kj} \in \mathbb{R}^{m_k \times m_j}$ for $j, k \in \{1, 2\}$. The non-degeneracy of a differential Nash equilibrium is determined by the determinant of $D\omega$. Recall that $D\omega$ is constructed from components of the symmetric matrices $D^2 f_1$ and $D^2 f_2$, i.e. the Hessians of f_1 and f_2 . Hence, we partition the space $S(m)^2$ into two subsets $S_1(m)$ and $S_2(m)$ defined as follows:

$$S_1(m) = \left\{ \begin{bmatrix} A_1^{11} & A_1^{21} \\ A_2^{12} & A_2^{22} \end{bmatrix} \in \mathbb{R}^{m \times m} \mid A_1, A_2 \in S(m) \right\} \quad (2.4.3)$$

and

$$S_2(m) = \left\{ \begin{bmatrix} A_2^{11} & A_2^{21} \\ A_1^{12} & A_1^{22} \end{bmatrix} \in \mathbb{R}^{m \times m} \mid A_1, A_2 \in S(m) \right\} \quad (2.4.4)$$

where $S_1(m)$ is the space corresponding to $D\omega$ and $S_2(m)$ is the space in which matrices constructed from the other pieces of the player Hessians that were excluded in the construction of $D\omega$. Then $J^2(M_1 \times M_2, \mathbb{R}^2)$ is locally diffeomorphic to

$$\mathbb{R}^m \times \mathbb{R}^2 \times \mathbb{R}^{m_1+m_2} \times \mathbb{R}^{m_1+m_2} \times \mathbb{R}^{\frac{m(m+1)}{2}} \times \mathbb{R}^{\frac{m(m+1)}{2}} \quad (2.4.5)$$

and the 2-jet extension of $f = (f_1, f_2)$ at a point $(\mu_1, \mu_2) \in M_1 \times M_2$, namely $j^2 f(\mu_1, \mu_2)$, in coordinates is given by

$$\begin{aligned} &(\varphi(\mu_1, \mu_2), ((f_1 \circ \varphi^{-1})(\varphi(\mu_1, \mu_2)), (f_2 \circ \varphi^{-1})(\varphi(\mu_1, \mu_2))), Df_1(\mu_1, \mu_2), \\ &Df_2(\mu_1, \mu_2), D^2 f_1(\mu_1, \mu_2), D^2 f_2(\mu_1, \mu_2)). \end{aligned} \quad (2.4.6)$$

Define

$$Z(m) = \{A \in S_1(m) \mid \det(A) = 0\}. \quad (2.4.7)$$

The set $Z(m)$ is algebraic and hence, admits a canonical Whitney stratification having finitely many algebraic strata (see [Gib+76, Chapter 1, Theorem 2.7]), i.e. it is the finite union of submanifolds. By its construction, $Z(m)$ has no interior points. Hence, it has co-dimension at least one. Now, we consider the subset of the jet bundle $J^2(M_1 \times M_2, \mathbb{R}^2)$ defined by

$$G_1 = \mathbb{R}^m \times \mathbb{R}^2 \times \{0_{\mathbb{R}^{m_1}}\} \times \mathbb{R}^{m_2} \times \mathbb{R}^{m_1} \times \{0_{\mathbb{R}^{m_2}}\} \times Z(m) \times S_2(m) \quad (2.4.8)$$

where $0_{\mathbb{R}^{m_i}}$ is the zero vector in \mathbb{R}^{m_i} . Note that $\{0_{\mathbb{R}^{m_i}}\}$ has co-dimension m_i . Hence, G_1 is the union of submanifolds of co-dimension at least $m_1 + m_2 + 1$. By the Jet Transversality Theorem (see Section 2.A.2, Theorem 2.A.3 or [Hir76, Theorem 2.8]) and Proposition 2.A.1, since $m_1 + m_2 + 1 > m_1 + m_2$, for generic $f = (f_1, f_2)$, the image of the 2-jet extension $j^2 f$ is disjoint from G_1 . Hence, there is an open-dense set of functions $f = (f_1, f_2)$ such that for each $(\mu_1, \mu_2) \in M_1 \times M_2$, whenever $D_1 f_1(\mu_1, \mu_2) = 0$ and $D_2 f_2(\mu_1, \mu_2) = 0$ (i.e. $\omega(\mu_1, \mu_2) = 0$), the derivative of the differential game form has non-zero determinant (i.e. $\det d\omega(\mu_1, \mu_2) \neq 0$). Note that the conditions $\omega(\mu_1, \mu_2) = 0$ and $\det(d\omega(\mu_1, \mu_2)) \neq 0$ are coordinate-invariant. Hence, this result is independent of the choice of chart.

Similarly, consider another subset of $J^2(M_1 \times M_2, \mathbb{R}^2)$ defined by

$$\begin{aligned} G_2 = &\mathbb{R}^m \times \mathbb{R}^2 \times \{0_{\mathbb{R}^{m_1}}\} \times \mathbb{R}^{m_2} \times \mathbb{R}^{m_1} \times \{0_{\mathbb{R}^{m_2}}\} \times Z(m_1) \\ &\times \mathbb{R}^{m_1 \times m_2} \times S(m_2) \times S(m_1) \times \mathbb{R}^{m_1 \times m_2} \times Z(m_2) \end{aligned} \quad (2.4.9)$$

where $Z(m_i)$ is the subset of symmetric matrices $S(m_i)$ such that for $A \in Z(m_i)$, $\det(A) = 0$. Since $Z(m_i)$ are algebraic and have no interior points, we may again use the Whitney stratification theorem [Gib+76, Chapter 1, Theorem 2.7] to get that each $Z(m_i)$ is the union of submanifolds of co-dimension at least 1. Hence, G_2 is the union of submanifolds and has co-dimension at least $m_1 + m_2 + 2$. Application of the Jet Transversality Theorem 2.A.3 and Proposition 2.A.1 yields an open-dense set of functions $f = (f_1, f_2)$ such that when $\omega(\mu_1, \mu_2) = 0$ we have $\det(D_{ii}^2 f_i(\mu_1, \mu_2)) \neq 0$ for each $i \in \{1, 2\}$.

Since the intersection of two open–dense sets is open–dense, we have an open–dense set of functions $f = (f_1, f_2)$ such that for each $(\mu_1, \mu_2) \in M_1 \times M_2$ whenever $\omega(\mu_1, \mu_2) = 0$, $\det(D_{ii}^2 f_i(\mu_1, \mu_2)) \neq 0$ for each $i \in \{1, 2\}$ and $\det(d\omega(\mu_1, \mu_2)) \neq 0$ independent of the choice of chart.

Thus, there exists an open–dense set $G \subset C^\infty(M_1 \times M_2, \mathbb{R}^2)$ such that for all $f = (f_1, f_2) \in G$, if $(\mu_1, \mu_2) \in M_1 \times M_2$ is a local Nash equilibrium, then (μ_1, μ_2) is a non–degenerate differential Nash equilibrium. Indeed, suppose $(f_1, f_2) \in G$ and $(\mu_1, \mu_2) \in M_1 \times M_2$ is a local Nash equilibrium. Then, by Proposition 2.2.1, (μ_1, μ_2) necessarily satisfies $\omega(\mu_1, \mu_2) = 0$ and $D_{ii}^2 f_i(\mu_1, \mu_2) \geq 0$ for each $i \in \{1, 2\}$. However, since $(f_1, f_2) \in G$, $\det(D_{ii}^2 f_i(\mu_1, \mu_2)) \neq 0$ so that $D_{ii}^2 f_i(\mu_1, \mu_2) > 0$. Hence, (μ_1, μ_2) is a differential Nash equilibrium. Further, $(f_1, f_2) \in G$ implies that $\det(d\omega(\mu_1, \mu_2)) \neq 0$; hence, (μ_1, μ_2) is non–degenerate. \square

Again, let $M = M_1 \times \dots \times M_n$ and denote the set of players $\mathcal{I} = \{1, \dots, n\}$. Given $f_1, \dots, f_n \in C^\infty(M, \mathbb{R})$, we define the set of local Nash equilibria as

$$\begin{aligned} \text{LN}(f_1, \dots, f_n) = \{u \in M \mid & \text{for each } i \in \mathcal{I}, W_i \subset M_i, \\ & f_i(u) \leq f_i(u_1, \dots, u_{i-1}, u'_i, u_{i+1}, \dots, u_n) \ \forall u'_i \in W_i \setminus \{u_i\}\} \end{aligned} \quad (2.4.10)$$

and the set of non–degenerate differential Nash equilibria as

$$\text{DN}(f_1, \dots, f_n) = \{u \in M \mid \omega(u) = 0, D_{ii}^2 f_i(u) > 0 \ \forall i \in \mathcal{I}, \det(d\omega(u)) \neq 0\}. \quad (2.4.11)$$

Remark 2.4.1. *As remarked in the proof of Theorem 2.4.1, we only show the details for the two player case as the proof readily generalizes and its derivation would merely obfuscate the core tools needed for the proof.*

In Section 2.2, we showed that $\text{DN}(f_1, \dots, f_2) \subset \text{LN}(f_1, \dots, f_n)$ for all $f_1, \dots, f_n \in C^\infty(M_1 \times \dots \times M_n, \mathbb{R})$. Theorem 2.4.1 shows that $\text{LN}(f_1, \dots, f_n) = \text{DN}(f_1, \dots, f_n)$ for all (f_1, \dots, f_n) in an open–dense subset $G \subset C^\infty(M_1 \times \dots \times M_n, \mathbb{R}^n)$. In other words, the set of local Nash equilibria is generically equivalent to the set of non–degenerate differential Nash equilibria.

Another interpretation of Theorem 2.4.1 is that degenerate local Nash equilibria can become non–degenerate differential Nash equilibria under an arbitrarily small perturbation to the game. In Section 2.3, we showed the converse is not true: non–degenerate differential Nash equilibria persist under small smooth perturbations to player costs.

The results of Theorem 2.4.1 specialize to the case of zero–sum games.

Corollary 2.4.1 (Genericity for Zero–Sum Games). *Non–degenerate differential Nash equilibria are generic among local Nash equilibria in zero–sum games: for any smooth boundaryless manifolds M_1, M_2 there exists an open dense subset $G \subset C^\infty(M_1 \times M_2, \mathbb{R}^2)$ such that for all $(f, -f) \in G$, if $(p, q) \in M_1 \times M_2$ is a local Nash equilibrium for $(f, -f)$, then (p, q) is a non–degenerate differential Nash equilibrium for $(f, -f)$.*

The proof of the above corollary is omitted since it follows the exact reasoning of the proof for Theorem 2.4.1.

2.5 Potential Games

In this section, we extend the notion of *potential game* as introduced by in [MS96] to games on manifolds with the goal of building the proper tools needed for expounding upon the coupled-oscillator game introduced in Example 2.2. It is an example on a non-trivial strategy space which is based on a fundamental model that is used in a wide variety of engineering applications.

First, let us define a potential game. We will recall the results for continuous games defined by functions (f_1, \dots, f_n) on strategy spaces $E_1 \times \dots \times E_n$ where each E_i is an interval of real numbers. We use the notation $E = E_1 \times \dots \times E_n$ and $E_{-i} = E_1 \times \dots \times E_{i-1} \times E_{i+1} \times \dots \times E_n$.

Definition 2.5.1 (Potential Game [MS96]). *A function $\Phi : E \rightarrow \mathbb{R}$ is an **exact potential** for the game (f_1, \dots, f_n) on E if for each $i \in \{1, \dots, n\}$ and for every $u_{-i} \in E_{-i}$ we have*

$$f_i(v, u_{-i}) - f_i(w, u_{-i}) = \Phi(v, u_{-i}) - \Phi(w, u_{-i}), \quad \text{for every } v, w \in E_i. \quad (2.5.1)$$

*The game (f_1, \dots, f_n) on E is an **exact potential game** (or just **potential game**) if it admits an exact potential.*

The following results on characterizing continuous potential games were introduced by Monderer and Shapley.

Lemma 2.5.1 ([MS96]). *Consider a game (f_1, \dots, f_n) where for each $i \in \{1, \dots, n\}$, player i has strategy space E_i which is an interval of real numbers. Suppose each f_i is continuously differentiable with respect to the choice variable u_i and let $\phi : E_1 \times \dots \times E_n \rightarrow \mathbb{R}$. Then ϕ is a potential for the game if and only if ϕ is continuously differentiable and*

$$\frac{\partial f_i}{\partial u_i} = \frac{\partial \phi}{\partial u_i}, \quad \text{for each } i, j \in \{1, \dots, n\}. \quad (2.5.2)$$

The next theorem is a classical result from physics and it boils down to an application of *de Rham cohomology* [Lee12].

Theorem 2.5.1 ([MS96]). *Suppose we have a game defined by cost functions (f_1, \dots, f_n) which are continuously differentiable where player i has strategy space $E_i \subset \mathbb{R}$ which is an interval of real numbers. The game is a potential game if and only if*

$$\frac{\partial^2 f_i}{\partial u_i \partial u_j} = \frac{\partial^2 f_j}{\partial u_i \partial u_j}, \quad \text{for every } i, j \in \{1, \dots, n\}. \quad (2.5.3)$$

Further, if the payoff functions satisfy (2.5.3) and $v = (v_1, \dots, v_n) \in E_1 \times \dots \times E_n$ is an arbitrary, fixed strategy profile, then a potential function for the game is given by

$$\phi(u_1, \dots, u_n) = \sum_{i=1}^n \int_0^1 \frac{\partial f_i}{\partial u_i}(u(t)) \frac{du_i}{dt}(t) dt \quad (2.5.4)$$

Potential games are an interesting class of games because they possess nice properties. Indeed, if the strategy spaces are compact, then every potential game possess at least one pure strategy Nash equilibrium [MS96]. Further, if the potential function for the game is (jointly) convex, then natural dynamics such as gradient descent converge (see, e.g., [Kri+15]).

It is interesting to note that Monderer and Shapley provide a local condition—in both Lemma 2.5.1 and Theorem 2.5.1—that can be checked to determine if a game is a potential game. We will see that, while not directly obvious, this result can be extended to continuous games on non-trivial strategy spaces by employing tools from differential geometry.

Let $\pi_{M_{-i}} : M_1 \times \cdots \times M_n \rightarrow M_{-i}$ where $M_{-i} = M_1 \times \cdots \times M_{i-1} \times M_{i+1} \times \cdots \times M_n$ for each $i \in \{1, \dots, n\}$ be natural projection maps. Let $\pi_{M_{-i}}^*$ denote its *pullback*, i.e. for a function $g : M_{-i} \rightarrow \mathbb{R}$, $\pi_{M_{-i}}^*g(u) = g(\pi_{M_{-i}}(u))$.

Definition 2.5.2 (Potential Game on Manifold). *A game (f_1, \dots, f_n) on $M_1 \times \cdots \times M_n$ where each M_i is a smooth, connected manifold (without boundary) is a **potential game** if there exists a potential function $\phi \in C^\infty(M_1 \times \cdots \times M_n)$ such that $f_i - \phi \in \text{im } \pi_{M_{-i}}^*$ for each $i \in \{1, \dots, n\}$.*

This definition says that f_i and ϕ differ by a function of u_{-i} only and u_{-i} are the decision variables that selected by players in $-i$, i.e. they are not controlled or determined by player i . The following two propositions, due to Stein, appeared in an unpublished technical report [Ste10] and hence, we provide the statements and the proofs here with some minor modifications.

Proposition 2.5.1 ([Ste10]). *A game is a potential game if and only if its differential game form is exact.*

Proof. Suppose the game is a potential game so that there exists a potential function ϕ such that $f_i - \phi \in \text{im } \pi_{M_{-i}}^*$ for each $i \in \{1, \dots, n\}$. Note that this is to say that f_i and ϕ differ by a function of u_{-i} only. Then it is straightforward to show that

$$\sum_{i=1}^n \frac{\partial(f_i - \phi)}{\partial u_i} du_i = 0 \tag{2.5.5}$$

so that we have

$$\omega = \sum_{i=1}^n \frac{\partial f_i}{\partial u_i} du_i = \sum_{i=1}^n \frac{\partial \phi}{\partial u_i} du_i = d\phi. \tag{2.5.6}$$

Hence, ω is exact. On the other hand, if the differential game form ω is exact, then there exists some function $\phi \in C^\infty(M_1 \times \cdots \times M_n)$ such that $\omega = d\phi$. Since $\psi_{M_i} \circ \psi_{M_i} = \text{id}$ (where id is identity) and $\psi_{M_i} \circ \psi_{M_j} = 0$, we have

$$\psi_{M_i}(df_i) = \psi_{M_i}(\omega) = \psi_{M_i}(d\phi).$$

Since d is linear, $\psi_{M_i} \circ d(f_i - \phi) = 0$. Furthermore, since M_i is connected, $\ker(\psi_{M_i} \circ d) = \text{im } \pi_{M_{-i}}^*$. Thus, $f_i - \phi \in \text{im } \pi_{M_{-i}}^*$. \square

Proposition 2.5.2 ([Ste10]). *Suppose we have a game (f_1, \dots, f_n) on $M = M_1 \times \dots \times M_n$ where each M_i is a smooth, connected, compact manifold without boundary. The differential game form ω associated with (f_1, \dots, f_n) is exact if and only if it is closed.*

Proof. If the differential game form ω is exact, then it is closed since all exact forms are closed (since $d \circ d = 0$ by Theorem 2.A.1). Hence, we only need to prove that if $d\omega = 0$, then it is exact.

From de Rham's Theorem [Lee12, Theorem 18.14], we know that if $\int_c \omega = 0$ for all smooth 1-cycles c , then ω is exact. Hence, we need to show that $\int_c \omega = 0$. First, fix $(v_1, \dots, v_n) \in M$. Let $\iota_{v_{-i}} : M_i \rightarrow M$ be the natural inclusion map such that $\iota_{v_{-i}}(v) = (v_1, \dots, v_{i-1}, v, v_{i+1}, \dots, v_n)$. Then the map $H_1^\infty(M_1) \times \dots \times H_1^\infty(M_n) \rightarrow H_1^\infty(M)$ given by

$$([\alpha_1], \dots, [\alpha_n]) \mapsto \left[\sum_{i=1}^n \iota_{v_{-i}*} \alpha_i \right]$$

is an isomorphism. Indeed, by Künneth's formula (see, e.g. [Spa81, Chapter 5, §3], [Hat02, Chapter 3]), the map $H_1(M_1) \times \dots \times H_1(M_n) \rightarrow H_1(M)$, i.e. for singular homology, is an isomorphism. Furthermore, for any smooth manifold M , the map from smooth singular homology to singular homology $\iota_* : H_p^\infty(M) \rightarrow H_p(M)$ induced by inclusion is an isomorphism (see [Lee12, Theorem 18.7] and also Appendix 2.A). Thus, without loss of generality, we can write the smooth 1-cycle c as $c = \sum_{i=1}^n \iota_{v_{-i}*} \alpha_i$ where each α_i is a smooth singular 1-cycle on M_i .

Now, we can compute the following for arbitrary smooth 1-cycle c :

$$\int_c \omega = \int_{\sum_{i=1}^n \iota_{v_{-i}*} \alpha_i} \omega = \sum_{i=1}^n \int_{\iota_{v_{-i}*} \alpha_i} \omega = \sum_{i=1}^n \int_{\alpha_i} \iota_{v_{-i}}^* \omega = \sum_{i=1}^n \int_{\alpha_i} df_{i,v_{-i}} = 0 \quad (2.5.7)$$

where we use the notation $df_{i,v_{-i}}$ for the derivative of f_i given that v_{-i} is fixed. The pullback of the form ω by $\iota_{v_{-i}}$ is equal to $df_{i,v_{-i}}$, i.e. $\iota_{v_{-i}}^* \omega = df_{i,v_{-i}}$, thereby justifying the second to last equality. The last equality holds by Stokes' theorem [Lee12, Theorem 16.10]. \square

Corollary 2.5.1. *Suppose we have a game (f_1, \dots, f_n) on $M = M_1 \times \dots \times M_n$ where each M_i is a smooth, connected, compact manifold without boundary and let ω be its differential game form. The game is a potential game if and only if its differential game form is closed and a potential function for the game is given by*

$$\phi(u_1, \dots, u_n) = \int_\gamma \omega \quad (2.5.8)$$

where γ is any piecewise differentiable path such that $(v_1, \dots, v_n) \mapsto (u_1, \dots, u_n)$ where we have fixed (v_1, \dots, v_n) .

The proof is straightforward and thus, it is omitted.

Essentially Proposition 2.5.2 is a generalization of Theorem 2.5.1 and it allows us to extend the local conditions (Equations (2.5.2) and (2.5.3)) to continuous games on manifolds that are smooth, compact, and connected. To make the connection between (2.5.3) and the differential game form more clear, let us consider a two player game (f_1, f_2) on $M_1 \times M_2$. The condition that $d\omega = 0$, which by Corollary 2.5.1 ensures the game is a potential game, can be expressed in local coordinates as

$$d\omega = d\left(\frac{\partial f_1}{\partial u_1} du_1 + \frac{\partial f_2}{\partial u_2} du_2\right) = \left(\frac{\partial^2 f_2}{\partial u_2 \partial u_1} - \frac{\partial^2 f_1}{\partial u_2 \partial u_1}\right) du_1 \wedge du_2 \quad (2.5.9)$$

Hence, $d\omega = 0$ is equivalent to

$$\frac{\partial^2 f_2}{\partial u_2 \partial u_1} = \frac{\partial^2 f_1}{\partial u_2 \partial u_1}. \quad (2.5.10)$$

Note that d is the *exterior derivative* and not the usual differential (see Appendix 2.A).

Remark 2.5.1. Recall the definition of $d\omega$; for this two-player game, in coordinates we can express $d\omega$ as

$$\begin{bmatrix} \frac{\partial^2 f_1}{\partial u_1^2} & \frac{\partial^2 f_1}{\partial u_2 \partial u_1} \\ \frac{\partial^2 f_2}{\partial u_1 \partial u_2} & \frac{\partial^2 f_2}{\partial u_2^2} \end{bmatrix} \quad (2.5.11)$$

If the local representation of $d\omega$ was symmetric, then the condition in (2.5.10) would be satisfied. This exposes some interesting questions regarding computing the potential piece of the an arbitrary game. We can write any matrix as the sum of a symmetric matrix and an anti-symmetric matrix. Suppose we decompose $d\omega$ in to its symmetric and anti-symmetric parts and then, reconstruct the game corresponding to the symmetric piece, i.e. the potential part of the original game. Then, can we say anything about the relationship between the constructed potential game and the original game? Furthermore, will such insights lead to better methods for computing equilibria? Similar ideas have been explored for finite games, i.e. where the strategy spaces of the players are finite [Can+11; Can+13].

We remark that in the case the the strategy spaces are not compact and connected, we can no longer rely on this local condition to check if our game is a potential game. However, we can state the following.

Definition 2.5.3. For a game (f_1, \dots, f_n) on a smooth manifold without boundary $M = M_1 \times \dots \times M_n$, we say it is a **local potential game** if for point in M there exists a coordinate neighborhood U such that $\omega|_U = d\phi$ for a function $\phi \in C^\infty(U)$ where ω is the differential game form.

The above definition leads naturally to the following proposition.

Proposition 2.5.3. Suppose we have a game (f_1, \dots, f_n) on a smooth manifold without boundary $M = M_1 \times \dots \times M_n$. Then if $d\omega = 0$, the game is a local potential game.

The proof of the proposition is a direct application of the Poincaré Lemma [Lee12, Theorem 15.11].

We now return to Example 2.2 and show that it is in fact a potential game and furthermore, it admits a continuum of differential Nash equilibria. We remark that by restricting the strategy spaces to be *finite* subsets of \mathbb{S}^1 for each player, the coupled oscillator game was shown to be a potential game [Got+10]. Here we will extend this to the entire strategy space \mathbb{S}^1 using the notion of a potential game on a manifold. We will show that the n -coupled oscillator model admits a continuum of differential Nash equilibria and that a perturbed version of the two-player games has multiple non-degenerate differential Nash equilibria.

Before we dive into the example, let us recall the basic model for n coupled oscillators (see [Pal+07; Sep+07] and references therein for a more thorough presentation). Each oscillator will have a position $r_k = x_k + iy_k \in \mathbb{C}$ and a phase $\theta_k \in \mathbb{S}^1$. We will denote the collection of positions by $r = (r_1, \dots, r_n)$ and the collection of phases by $\theta = (\theta_1, \dots, \theta_n)$. The collection of all the phases θ evolves on the n -torus, denoted

$$\mathbb{T}^n = \underbrace{S^1 \times \dots \times S^1}_{n\text{-times}}.$$

The coupled oscillator model is given by

$$\dot{r}_j = e^{i\theta_j} \tag{2.5.12}$$

$$\dot{\theta}_j = u_j(r, \theta), \quad j = 1, \dots, n \tag{2.5.13}$$

where $u_j(r, \theta)$ is some control input.

We can split the control input u_j into three terms,

$$u_j = \theta_0 + u_j^{\text{spac}}(r, \theta) + u_j^{\text{ori}}(\theta), \quad j = 1, \dots, n \tag{2.5.14}$$

where $\theta_0 \in \mathbb{R}$ is a constant, $u_j^{\text{spac}}(r, \theta)$ is the spacing control, and $u_j^{\text{ori}}(\theta)$ is the orientation control. If we ignore the spacing control, i.e. $u_j^{\text{spac}} = 0$, then we obtain what is referred to as the *phase model*:

$$\dot{\theta}_j = \theta_0 + u_j^{\text{ori}}(\theta), \quad j = 1, \dots, N \tag{2.5.15}$$

which is a system of coupled-phase oscillators with identical natural frequency θ_0 .

We say the phases θ_k and θ_j are *phase locked* if $\dot{\theta}_{kj} = 0$ where we use the abbreviated notation $\theta_{kj} = \theta_k - \theta_j$. A *synchronized phase arrangement* θ is a phase-locked arrangement for which $\theta_k = \theta_j$ for all pairs j and k .

Let us consider n coupled oscillators whose coupling is prescribed by a complete graph $\mathcal{G} = (E, V)$ where E is the set of edges and V is the set of vertices, i.e. if oscillator $v_k \in V$ is connected to oscillator $v_j \in V$ then there is an edge $e_{kj} \in E$. We denote the phase of oscillator v_j by θ_j . Let L be the Laplacian matrix of the graph \mathcal{G} (see, e.g., [GR01, Chapter 13]) which is defined to have entries

$$L_{k,j} = \begin{cases} b_k, & \text{if } k = j \\ -1, & \text{if } e_{kj} \in E \\ 0, & \text{otherwise} \end{cases} \tag{2.5.16}$$

where b_k is the degree of vertex v_k . We can define the *Laplacian phase potential* by

$$\phi(\theta) = \frac{1}{2n} \langle e^{i\theta}, L e^{i\theta} \rangle = \frac{1}{2n} \sum_{k=1}^n \langle e^{i\theta_k}, L_k e^{i\theta} \rangle = \frac{1}{2n} \sum_{i=1}^n \left(b_k - \sum_{j \in N_k} \langle e^{i\theta_k}, e^{i\theta_j} \rangle \right) \quad (2.5.17)$$

where L_k is the k -th row of L and N_k is the index set of vertices connected to v_k . Assuming \mathcal{G} is undirected, the gradient of $\phi(\theta)$ is $\frac{\partial \phi}{\partial \theta_k} = \frac{1}{N} \langle i e^{i\theta_k}, L_k e^{i\theta} \rangle$. In the phase model, if we choose the gradient-based control

$$u_j^{\text{ori}} = \frac{K_1}{n} \langle i e^{i\theta_j}, L_j e^{i\theta} \rangle, \quad j \in \{1, \dots, n\} \quad (2.5.18)$$

with $K_1 \neq 0$, then the potential $\Phi(\theta)$ evolves monotonically since

$$\dot{\phi}(\theta) = \frac{\partial \phi^T}{\partial \theta} \dot{\theta} = \frac{K_1}{n^2} \sum_{j=1}^n \langle i e^{i\theta_j}, L_j e^{i\theta} \rangle^2. \quad (2.5.19)$$

Furthermore, the phase model with gradient-based orientation control is given by

$$\dot{\theta}_j = \theta_0 + \frac{K_1}{n} \sum_{l \in N_j} \sin \theta_{jl}. \quad (2.5.20)$$

If $K_1 < 0$, then the phases are synchronized under the feedback law; conversely, if $K_1 > 0$, then the phases are balanced under the feedback law. Details of these results are given in [Sep+05b]. As shown in [Pal+07], maximizing the potential function ϕ leads to balanced phases and minimizing the potential function ϕ leads to phase synchronization.

The interesting thing to note here is that the above system is a simplified Kuramoto model of identical coupled-phase oscillators with limited interaction as specified through the graph structure. As we have pointed out in Section 2.1, coupled oscillator models of this type are commonly used in many engineering application domains.

Example 2.3 (Coupled Oscillators is a Potential Game). *Consider n -coupled oscillators with an interaction structure specified by a undirected, complete graph where the nodes represent the oscillators and the edges indicate a connection between oscillators. As before, we denote the phase of the j -th oscillator by θ_j and we use the notation $\theta = (\theta_1, \dots, \theta_n)$ for the collection of phases of all the oscillators. Each phase θ_j represents a point on the unit circle S^1 .*

Recall the Laplacian phase potential ϕ defined in (2.5.17). We claim the game (f_1, \dots, f_n) on \mathbb{T}^n where

$$f_j(\theta) = -\frac{1}{n} \sum_{l \in N_j} \cos \theta_{jl} \quad (2.5.21)$$

is a potential game in which the potential function it admits is ϕ . This is straightforward to check. By Proposition 2.5.1, we need only verify that the differential game form ω is exact and satisfies $\omega = d\phi$. Indeed,

$$d\phi = \frac{1}{n} \sum_{k=1}^n \left(\sum_{j \in N_k} \sin \theta_{kj} \right) d\theta_k \quad (2.5.22)$$

which is exactly the differential game form ω for the game (f_1, \dots, f_n) in which the players are cost minimizers. On the other hand, if each player is a utility maximizer with

$$\tilde{f}_j(\theta) = -f_j(\theta) = \frac{1}{n} \sum_{l \in N_j} \cos \theta_{jl}, \quad (2.5.23)$$

then it is straightforward to show that the game is a potential game with potential function $-\phi$. \square

Example 2.4 (Coupled Oscillator Game Admits a Continuum of Nash). *Again, consider n -coupled oscillators with an interaction structure specified by a undirected, complete graph where the nodes represent the oscillators and the edges indicate a connection between oscillators. Let the phase of oscillator j be denoted by $\theta_j \in \mathbb{S}^1$ and let its cost be given by*

$$f_j = -\frac{1}{n} \sum_{l \in N_j} \cos(\theta_j - \theta_l) \quad (2.5.24)$$

where N_j is the index set of oscillators that are coupled to oscillator j . Consider the potential function

$$\phi(\theta_1, \dots, \theta_n) = -\frac{1}{2n} \sum_{i=1}^n \left(\sum_{j \in N_i} \cos(\theta_i - \theta_j) \right). \quad (2.5.25)$$

Just as in the previous example, the differential game form for the oscillator game satisfies $\omega = d\phi$. Note that this indicates that the potential function for a game is not necessarily unique—in particular, here we just dropped the constants b_k from the potential function in (2.5.17).

We claim that all points in the set

$$\{(\theta_1, \dots, \theta_n) \in \mathbb{S}^1 \times \dots \times \mathbb{S}^1 \mid \theta_l = \theta_j, \forall l, j \in \{1, \dots, n\}\} \quad (2.5.26)$$

are global Nash equilibria of the game. Indeed, consider player l and fix $\theta_j = \beta$ for all $j \neq l$. Then $\theta_l = \beta$ is a best response—i.e. in the set of optimizers—by oscillator l to all other oscillators playing $\theta_j = \beta$. In particular,

$$f_l(\beta, \dots, \beta) = -\frac{|N_l|}{n} < -\frac{1}{n} \sum_{j \in N_l} \cos(\theta'_l - \beta), \quad \forall \theta'_l \in \mathbb{S}^1 \setminus \{\beta\}, \quad (2.5.27)$$

where $|\cdot|$ is the cardinality. Thus there is a continuum of Nash equilibria for which the oscillators are synchronized. In fact there is a continuum of differential Nash equilibria; this is easily seen by checking that $D_{ii}^2 f_l(\theta_1, \dots, \theta_n) > 0$ when $\theta_l = \theta_j$ for all l, j . \square

It is interesting to note that if we considered the same game with the modification that (f_1, \dots, f_n) are now utility functions and the oscillators are utility maximizers, then there is a continuum of Nash equilibria now at all $(\theta_1, \dots, \theta_n)$ such that the oscillators are *balanced*.

Connecting the above example back to Example 2.2, we have shown that for games on non-trivial strategy spaces—even potential games which have inherently nice properties in terms of existence and uniqueness of Nash equilibria and computation—pathologies and undesirable properties can arise. Further, the example demonstrates the results in the previous section for an example on a non-trivial strategy space that is common mathematical abstraction in many engineering applications.

While one may notice the symmetry in the game described in Example 2.2, breaking that symmetry may still result in multiple Nash equilibria.

Example 2.2 (Jean–Paul: Preferred Phase). *Return again to the Jean–Paul oscillator example, i.e. $n = 2$. We will perturb Jean’s cost and leave Paul’s unchanged. Let Jean’s cost be*

$$\tilde{f}_1 = -\frac{1}{2} \cos(\gamma\theta_1 - \theta_2) \tag{2.5.28}$$

and Paul’s cost be

$$f_2 = -\frac{1}{2} \cos(\theta_2 - \theta_1) \tag{2.5.29}$$

where in this example Jean and Paul have different preferences for their phase. Allowing γ to take values in $\mathbb{N} \setminus \{1\}$, there are at least $\gamma - 1$ non-degenerate differential Nash equilibria:

$$\left\{ (\theta_1, \theta_2) \in \mathbb{S}^1 \times \mathbb{S}^1 \mid \theta_1 = \theta_2 = \frac{2(k-1)\pi}{\gamma-1}, k \in \{1, \dots, \gamma-1\} \right\}. \tag{2.5.30}$$

The above set contains only stable, non-degenerate differential Nash equilibria of the game (f_1, f_2) since points in this set satisfy $\omega(\theta_1, \theta_2) = 0$, $D_{ii}^2 f_i(\theta_1, \theta_2) > 0$, and $\det(d\omega(\theta_1, \theta_2)) \neq 0$. In fact, they are (non-strict) global Nash equilibria. ■

In our framework, stable equilibria attract nearby trajectories under the gradient flow of the game, a fact which can be leveraged by a central planner. Indeed, the n -coupled oscillator game can be thought of as an abstraction of generators or inverters—perhaps even microgrids—coupling to the grid [DB12; Dör+13] where each of them is individually managed. Due to the existence of a continuum of Nash equilibria, it is possible that the players will equilibrate on a socially undesirable outcome. A central planner vying to coordinate the individuals would therefore benefit from considering these second-order conditions when designing incentives.

2.6 Computation of Local Nash Equilibria

Our sufficient conditions for local Nash equilibria based on first- and second-order properties of player costs closely parallel theoretical developments in nonlinear programming [Ber99] and optimal control [Pol97]. In this section we further exploit this analogy by proposing an iterative *steepest descent* algorithm for computation of differential Nash equilibria.

We adopt a dynamical systems perspective of an n -player game over the strategy space $U = \prod_{i=1}^n U_i$ with player costs $f_i : U \rightarrow \mathbb{R}$. Specifically, we consider the continuous-time dynamical system generated by the negative of the player's individual gradients:

$$\dot{u} = -\omega(u). \tag{2.6.1}$$

Gradient play may be viewed as a *better response* strategy instead of a *best response* strategy; in particular, this is a myopic tâtonnement process in which players adjust their current strategy in a gradient direction [SA05]. If $\mu \in U$ is a differential Nash equilibrium, then $\omega(\mu) = 0$. These dynamics are *uncoupled* in the sense the dynamics \dot{u}_i for each player do not depend on the cost function of the other player. It is known that such uncoupled dynamics need not converge to local Nash equilibria [HMC03]. However, we provide the following result on convergence of these dynamics.

Proposition 2.6.1 ([Rat+13; Rat+14d]). *If μ is a differential Nash equilibrium and the spectrum of $d\omega$ is strictly in the right-half plane, then μ is an exponentially stable stationary point of the continuous-time dynamical system (2.6.1).*

The above result was stated for the finite-dimensional case in [Rat+13, Prop. 4] and the proof of the stated result is an application of [Abr+88, Thm. 4.3.4].

We say a non-degenerate differential Nash equilibrium is *stable* if the spectrum of $d\omega$ is strictly in the right-half plane. Equilibria that are stable—thereby attracting using decoupled myopic approximate best response—persist under small perturbations [Rat+13, Section IV]. Furthermore, Theorem 2.3.1 shows that convergence of uncoupled gradient play to such *stable* non-degenerate differential Nash equilibria persists under small smooth perturbations to player costs since the spectrum varies continuously [DS67, Lemma 6.3].

Remark 2.6.1. *Theorem 2.3.1 shows that convergence of uncoupled gradient play to such stable non-degenerate differential Nash equilibria persists under small smooth perturbations to player costs.*

Toward developing a numerical algorithm that approximates Nash equilibria, we study the forward-Euler approximation to (2.6.1). If each U_i is finite-dimensional, then by fixing a step size $h > 0$, we obtain the discrete-time dynamical system

$$u^{k+1} = u^k - h\omega(u^k). \tag{2.6.2}$$

Note that a differential Nash equilibrium is a fixed point of (2.6.2). Linearizing around such an equilibrium, we obtain the following sufficient condition ensuring nearby strategies converge to the Nash equilibrium under iteration of (2.6.2).

Proposition 2.6.2. *For each $i \in \{1, \dots, n\}$, let U_i be finite-dimensional. If μ is a differential Nash equilibrium and all eigenvalues of $-d\omega(\mu)$ are in the open left-half plane, then there exists $\eta > 0$ such that for all $h \in (0, \eta)$, μ is an exponentially stable fixed point of the discrete-time dynamical system (2.6.2).*

We interpret iteration of (2.6.2) as a *steepest-descent* algorithm analogous to techniques in nonlinear programming [Ber99], and terminate the iteration when $\|\omega(u^k)\|$ becomes sufficiently small. In fact, if the players are identical so that $f_1 = f_2 = f$, the algorithm exactly reduces to gradient descent on f with constant stepsize. A less trivial case where (2.6.2) reduces to gradient descent arises when $f_1 \neq f_2$ yet ω is an *exact* 1-form, i.e. it is a potential game. In this case, there exists a smooth function ϕ such that $\omega = d\phi$, and hence (2.6.2) is again equivalent to gradient descent on ϕ .

The analogy between gradient descent algorithms for nonlinear programming and the formula in (2.6.2) suggests a technique to numerically approximate differential Nash equilibria in the class of open-loop differential games described in 2.1. In particular, the derivative (2.1.4) can be approximated using techniques from numerical optimal control [Pol97], and hence the formula in (2.6.2) may be iterated to approximate differential Nash equilibria in the game.

Note that Proposition 2.6.2 only ensures *local* convergence of iterates of (2.6.2) to stable, non-degenerate differential Nash equilibria. However, we have observed empirically in the examples described in the next section that our proposed algorithm converges to a stationary point of (2.6.2) when initialized from almost every randomly-sampled initial condition.

Existing methods for iterative approximation of Nash equilibria generally employ the *relaxation technique*, where players alternately update their strategies by averaging their current strategy with the best response to the other player's current strategy,

$$u_i^{k+1} = \alpha u_i^k + (1 - \alpha) \arg \min_{\mu_i \in U_i} f_i(\mu_i, u_{-i}^k), \tag{2.6.3}$$

where $\alpha \in (0, 1)$ is a parameter and we again use the notation u_{-i} to denote the strategies of all players other than i . Assuming convexity in the strategy space and cost functions to ensure there exists a unique Nash equilibrium, it is known that iterating (2.6.3) converges to the Nash equilibrium [Baş87; Con+04; UR94]. Each iteration of (2.6.3) requires the solution of a (generally non-convex) optimization problem at every iteration; in contrast, our scheme requires only the evaluation of derivatives of the player costs at a single point.

2.6.1 Examples

In this section we demonstrate the preceding theoretical and algorithmic developments in examples with (i) nonlinear and (ii) infinite-dimensional strategy spaces. Our proposed method applies broadly, but we present examples where Nash equilibria are known explicitly so that we may evaluate the scalability and accuracy of our algorithm.

Location Game

Here we consider two-player game on the unit circle, \mathbb{S}^1 . The player costs $f_i : \mathbb{S}^1 \times \mathbb{S}^1 \rightarrow \mathbb{R}$ are given by

$$\begin{aligned} f_1(\theta_1, \theta_2) &= -\cos \theta_1 + \alpha_1 \cos(\theta_1 - \theta_2) \\ f_2(\theta_1, \theta_2) &= -\cos \theta_2 + \alpha_2 \cos(\theta_2 - \theta_1) \end{aligned}$$




Figure 2.1: Visualization of a two–player game played on a nonlinear strategy space (the torus, $\mathbb{S}^1 \times \mathbb{S}^1$) as described in Section 2.6.1. Curves indicate the set of points where $d_i f_i = 0$ for each $i \in \{1, 2\}$; we have $\omega = 0$ wherever the curves intersect. There are two differential Nash equilibria, indicated by a dark circle; most initial conditions converge to one of these via the steepest descent algorithm of Section 2.6. The empirical basin of attraction for each Nash equilibria is illustrated by the filled region containing the point.

where $\alpha_1, \alpha_2 \in \mathbb{R}$ are parameters. An interpretation of these costs is that both players wish to be near zero but far from each other. This game is a location game that is an abstraction of a game that has many applications. In coordinates, the game form $\omega(\theta_1, \theta_2)$ is

$$\left[\sin \theta_1 - \alpha_1 \cos(\theta_1 - \theta_2) \quad \sin \theta_2 - \alpha_2 \cos(\theta_2 - \theta_1) \right]$$

and the Hessian $d\omega(\theta_1, \theta_2)$ is

$$\begin{bmatrix} \cos \theta_1 - \alpha_1 \cos(\theta_1 - \theta_2) & \alpha_1 \cos(\theta_1 - \theta_2) \\ \alpha_2 \cos(\theta_2 - \theta_1) & \cos \theta_2 - \alpha_2 \cos(\theta_2 - \theta_1) \end{bmatrix}.$$

Theorem 2.2.2 implies that any point (θ_1, θ_2) for which $\omega(\theta_1, \theta_2) = 0$ and $d\omega(\theta_1, \theta_2) > 0$ is an isolated local Nash equilibrium. Numerically, with $\alpha_1 = 1, \alpha_2 = 1.05$ we find two such equilibria situated symmetrically around the zero angle: one near $(1, -1.1)$ and the other near $(-1, 1.1)$. Points where $\omega = 0$ but $d\omega$ is not positive–definite are located at $(0, \pi)$ and (π, π) . Applying the steepest descent algorithm of Section 2.6 with constant step–size 0.1 and termination tolerance 1×10^{-3} , we find empirically that most initial conditions converge to one of the two stable Nash equilibria within a few hundred iterations. See Figure 2.1 for a visualization of this example.

Open-loop Linear Quadratic Differential Game

As an illustration of the generality and scalability of our proposed algorithm, we numerically determine open-loop Nash equilibrium inputs in a linear-quadratic (LQ) game played between m players. We consider the linear time-invariant differential equation

$$\dot{x} = Ax + \sum_{j=1}^m B_j u_j, \quad x(0) = x_0$$

with player costs

$$f_i = \int_0^T x^T(t) Q_i x(t) + \sum_{j=1}^m u_j^T(t) R_{ij} u_j(t) dt$$

for each $i \in \{1, \dots, m\}$. It is known that, under non-degeneracy conditions on the game parameters [Eis82], the unique open-loop Nash equilibria strategy is given by the linear state feedback $u_i(t) = -R_{ii}^{-1} B_i^T P_i(t) x(t)$ for each $i \in \{1, \dots, m\}$ where $P_i \in \mathbb{R}^{n \times n}$ satisfies $P_i(T) = 0$ and

$$-\dot{P}_i = P_i A + A^T P_i + Q_i - P_i \sum_{j=1}^m B_j R_{jj}^{-1} B_j^T P_j.$$

Using the discretization scheme for optimal control problems described in Chapter 4 of [Pol97], we numerically approximate differential Nash equilibria for this game using the steepest descent algorithm of Section 2.6 and compare the result with the corresponding time-discretized approximation of this closed-form expression. By considering randomly generated examples where the entries of A , B_i , Q_i , and R_{ij} are chosen from a standard normal distribution (and subsequently positive semi-definiteness is enforced for Q_i and R_{ij}) for each $i, j \in \{1, \dots, m\}$, we find empirically that the limit point of our algorithm is insensitive to initialization and yields strategies that are quantitatively similar to the time-discretized analytical formula. Figure 2.2 shows how the relative error between the two solutions decreases as the number N of time samples increases in a typical randomly generated example; specifically,

$$A = \begin{bmatrix} -2.28 & 0.96 \\ 0.69 & 0.23 \end{bmatrix}, B_1 = \begin{bmatrix} -0.53 \\ 0.39 \end{bmatrix}, B_2 = \begin{bmatrix} -0.04 \\ -0.49 \end{bmatrix},$$

$$Q_1 = \begin{bmatrix} 1.69 & -0.64 \\ -0.64 & 3.02 \end{bmatrix}, Q_2 = \begin{bmatrix} 1.36 & -0.08 \\ -0.08 & 4.39 \end{bmatrix},$$

$$R_{11} = 1.85, R_{12} = 0.06, R_{21} = 1.1, R_{22} = 1.38.$$

We use a stepsize of 1 and a termination tolerance of 10^{-4} .

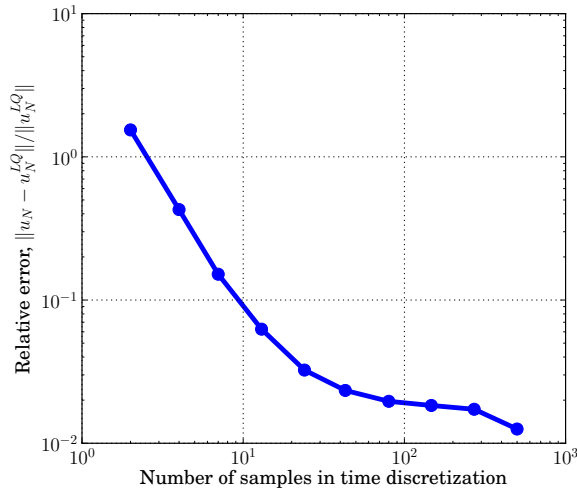


Figure 2.2: Relative error between open-loop Nash equilibria obtained from discretized analytical formula, u_N^{LQ} , and from steepest descent algorithm of Section 2.6, u_N , in the infinite-dimensional (linear-quadratic) game described in Section 2.6.1, with N samples in the time discretization. The relative error between the two solutions decreases as N increases.

2.7 Inducing a Nash Equilibrium

The problem of inducing Nash equilibria through incentive mechanisms appears in engineering applications including energy management [Coo+13] and network security [Rat+12; Zhu+12b]. The central planner aims to shift the Nash equilibrium of the agents’ game to one that is desirable from its perspective. Thus the central planner optimizes its cost subject to constraints given by the inequalities that define a Nash equilibrium. This requires verifying non-convex inequalities on an open set—a generally intractable task. A natural solution is to replace these inequalities with first- and second-order sufficient conditions on each agent’s optimization problem. As the Betty-Sue and Jean-Paul examples show (Example 2.1 and 2.2 respectively), these necessary conditions are not enough to guarantee the desired Nash is isolated; the additional constraint that $d\omega$ be non-degenerate must be enforced.

Example 2.1 (Betty-Sue: Inducing Nash). *Consider a central planner who desires to optimize the cost of deviating from the temperature τ ,*

$$f_p(u_1, u_2) = (u_1 - \tau)^2 + (u_2 - \tau)^2,$$

by inducing the agents to play $(u_1, u_2) = (\tau, \tau)$. The planner does so by selecting $a \in \mathbb{R}$ and augmenting Betty’s and Sue’s costs:

$$\tilde{f}_1^a(u_1, u_2) = f_1(u_2, u_2) + \frac{a}{2}(u_1 - \tau)^2,$$

$$\tilde{f}_2^a(u_1, u_2) = f_2(u_1, u_2) + \frac{a}{2}(u_2 - \tau)^2.$$

Recall that Betty's nominal cost is

$$f_1(u_1, u_2) = \frac{u_1^2}{2} - u_1 u_2$$

and Sue's nominal cost is

$$f_2(u_1, u_2) = \frac{u_2^2}{2} - u_1 u_2.$$

The differential game form of the augmented game $(\tilde{f}_1^a, \tilde{f}_2^a)$ is

$$\tilde{\omega}(u_1, u_2) = (u_1 - u_2 + a(u_1 - \tau))du_1 + (u_2 - u_1 + a(u_2 - \tau))du_2$$

and the Hessian of the differential game form at the differential Nash equilibrium is

$$d\tilde{\omega}(u_1, u_2) = \begin{bmatrix} 1 + a & -1 \\ -1 & 1 + a \end{bmatrix}.$$

For any $a \in (-1, \infty)$, (τ, τ) is a differential Nash equilibrium of $(\tilde{f}_1^a, \tilde{f}_2^a)$ since $\tilde{\omega}(\tau, \tau) = 0$ and $d_{ii}^2 \tilde{f}_i^a(\tau, \tau) > 0$. For any $a \in (-1, 0]$, the game $(\tilde{f}_1^a, \tilde{f}_2^a)$ undesirable behavior. Indeed, recall Section 2.6 in which we consider the gradient dynamics for a two player game. For values of $a \in (-1, 0)$, $d\tilde{\omega}$ is indefinite so that the equilibrium of the gradient system is a saddle point. Hence, if agents perform gradient play and happen to initialize on the unstable manifold, then they will not converge to any equilibrium. Further, while $a = 0$ seems like a natural choice since it means not augmenting the players costs at all, it in fact gives rise to a continuum of equilibria. However, for $a > 0$, $d\tilde{\omega}$ is positive definite so that, as the results of Section 2.6 points out, (τ, τ) is a stable, non-degenerate differential Nash equilibrium and as such, the gradient dynamics will converge and the value of a determines the contraction rate. \square

We can consider a similar scenario for Jean and Paul.

Example 2.2 (Jean–Paul: Inducing Nash). Consider a central planner who desires to optimize the cost of deviating from the phase τ :

$$f_p(u_1, u_2) = (u_1 - \tau)^2 + (u_2 - \tau)^2.$$

Perhaps Jean and Paul are managers of generators or inverters and they need to connect back to the grid in order to transfer power. Let the planner be the manager of the larger power grid and suppose it can induce Jean and Paul to adjust their phases to $(\theta_1, \theta_2) = (\tau, \tau)$ by selecting $b \in \mathbb{R}$ and augmenting Jean's and Pauls costs:

$$\tilde{f}_1^b(\theta_1, \theta_2) = f_1(\theta_1, \theta_2) + \frac{b}{2} \cos(\theta_1 - \tau)$$

$$\tilde{f}_2^b(\theta_1, \theta_2) = f_2(\theta_1, \theta_2) + \frac{b}{2} \cos(\theta_2 - \tau).$$

Recall that Jean's nominal cost is

$$f_1(\theta_1, \theta_2) = -\frac{1}{2} \cos(\theta_1 - \theta_2)$$

and Paul's nominal cost is

$$f_2(\theta_1, \theta_2) = -\frac{1}{2} \cos(\theta_2 - \theta_2).$$

The differential game form of the augmented game $(\tilde{f}_1^b, \tilde{f}_2^b)$ is

$$\tilde{\omega}(\theta_1, \theta_2) = \frac{1}{2} (\sin(\theta_1 - \theta_2) - b \sin(\tau - \theta_1)) d\theta_1 + \frac{1}{2} (\sin(\theta_2 - \theta_2) - b \sin(\tau - \theta_2)) d\theta_2$$

and the Hessian of the differential game form at the differential Nash equilibrium is

$$d\tilde{\omega}(\theta_1, \theta_2) = \begin{bmatrix} \frac{1}{2} (\cos(\theta_1 - \theta_2) + b \cos(\tau - \theta_1)) & -\frac{1}{2} \cos(\theta_1 - \theta_2) \\ -\frac{1}{2} \cos(\theta_1 - \theta_2) & \frac{1}{2} (\cos(\theta_1 - \theta_2) + b \cos(\tau - \theta_2)) \end{bmatrix}.$$

For any $b \in (-1, \infty)$, (τ, τ) is a differential Nash equilibrium of $(\tilde{f}_1^b, \tilde{f}_2^b)$ since $\tilde{\omega}(\tau, \tau) = 0$ and $d_{ii}^2 \tilde{f}_i^b(\tau, \tau) > 0$. For any $b \in (-1, 0]$, the game $(\tilde{f}_1^b, \tilde{f}_2^b)$ undesirable behavior. Indeed, just as we described for the above example, for values of $b \in (-1, 0)$, $d\tilde{\omega}$ is indefinite so that the equilibrium of the gradient system is a saddle point. Further, while $b = 0$ seems like a natural choice since it means not augmenting the players costs at all, it in fact gives rise to a continuum of equilibria. However, for $b > 0$, $d\tilde{\omega}$ is positive definite so that (τ, τ) is a stable, non-degenerate differential Nash equilibrium and as such, the gradient dynamics will converge and the value of b determines the contraction rate. \square

These examples indicates how undesirable behavior can arise when the operator $d\omega$ is degenerate. Further, if the goal is to induce a particular Nash equilibrium amongst competitive agents, then it is not enough to consider only necessary and sufficient conditions for Nash equilibria; inducing stable, non-degenerate differential Nash equilibria leads to desirable and structurally stable behavior. We will utilize $d\omega$ to enforce desirable convergence properties in the construction of utility learning and incentive design algorithms in the sequel.

2.8 Discussion

By paralleling results in non-linear programming and optimal control, we developed first- and second-order necessary and sufficient conditions that characterize local Nash equilibria in continuous games on both finite- and infinite-dimensional strategy spaces. We further provided a second-order sufficient condition guaranteeing differential Nash equilibria are non-degenerate and, hence, isolated. We showed that non-degenerate differential Nash

equilibria are structurally stable and thus small modeling errors or environmental disturbances generally will not result in games with drastically different equilibrium behavior. Further, as a result of structural stability, our characterization of non-degenerate differential Nash equilibria is amenable to computation—in particular, gradient play converges to stable, non-degenerate differential Nash equilibria. We illustrated that such a characterization has value for the design of incentives to induce a desired equilibria. By enforcing not only non-degeneracy but also stability of a differential Nash equilibrium, the central planner can ensure that the desired equilibrium is isolated and that gradient play will converge locally. We will expound upon this concept in greater detail in the sequel.

There are a number of interesting directions for future research, some of which we touched upon throughout the chapter. For instance, we mentioned decomposing the differential game form in such a way that we extract the piece of the game that corresponds to a potential game. In particular, an application of Hodge’s Decomposition [Abr+88, Chapter 8, §5] to the differential game form ω tells us that it is the sum of an exact form, a co-exact form, and a harmonic form. Using tools such as this from differential geometry, we can recover the potential piece of a game which, in a sense, is the *cooperative* piece of the game since it corresponds to the part of each player’s objective function which is common across players. One interesting direction for future research is in trying to utilize this information to inform methods of computing Nash equilibria. We are actively pursuing generalizations of *sufficient descent* techniques from nonlinear programming [Ber99] to develop algorithms which provably converge to differential Nash equilibria over larger regions of attraction. We are investigating using the potential piece to select sufficient descent directions. Moreover, another natural extension could be in the consideration of strategy spaces that are manifolds with boundary including regular constraint sets and manifolds with generalized boundaries [Jon+01]. Nash equilibria in games with constraints are typically called generalized Nash equilibria and there has been extensive research in the class of convex games going all the way back to the classical paper by Rosen [Ros65]. What has received little attention is non-convex games, where either the objective functions are non-convex or even strategy spaces themselves are non-convex. These are difficult problems, yet many tools from non-linear programming have the potential to extend to *non-cooperative programming*. The contributions in this chapter are certainly first steps toward this end.

Appendix 2.A Preliminaries

This appendix contains the standard mathematical objects used throughout this section in particular and document in general (see [Lee12; Abr+88] for a more detailed introduction).

Suppose that M is second-countable and a Hausdorff topological space. Then a *chart* on M is a homeomorphism φ from an open subset U of M to an open subset of a Banach space. We sometimes denote a chart by the pair (U, φ) . Two charts (U_1, φ_1) and (U_2, φ_2) are *C^r -compatible* if and only if the composition $\varphi_2 \circ \varphi_1^{-1} : \varphi_1(U_1 \cap U_2) \rightarrow \varphi_2(U_1 \cap U_2)$ is a C^r -diffeomorphism. A *C^r -atlas* on M is a collection of charts $\{(U_\alpha, \varphi_\alpha)\}_{\alpha \in A}$ any two of which

are C^r -compatible and such that the U_α 's cover M . A *smooth manifold* is a topological manifold with a smooth atlas. We use the term *manifold* generally; we specify whether it is a finite- or infinite-dimensional manifold only when it is not clear from context. If a covering by charts takes their values in a Banach space E , then E is called the *model space* and we say that M is a *C^r -Banach manifold*. We remark that one can form a manifold modeled on any linear space in which one has theory of differential calculus; we use Banach manifolds so that we can utilize the inverse function theorem.

Suppose that $f : M \rightarrow N$ where M, N are C^k -manifolds. We say f is of *class C^r* with $0 \leq r \leq k$, and we write $f \in C^r(M, N)$, if for each $u \in M$ and a chart (V, ψ) of N with $f(u) \in V$, there is a chart (U, φ) of M satisfying $u \in U$, $f(U) \subset V$, and such that the local representation of f , namely $\psi \circ f \circ \varphi^{-1}$, is of class C^r . If $N = \mathbb{R}$, then ψ can be taken to be the identity map so that the local representation is given by $f \circ \varphi^{-1}$.

Each $u \in M$ has an associated *tangent space* T_uM , and the disjoint union of the tangent spaces is the *tangent bundle* $TM = \coprod_{u \in M} T_uM$. The *co-tangent space* to M at $u \in M$, denoted T_u^*M , is the set of all real-valued linear functionals—or, simply, the dual—on the tangent space T_uM , and the disjoint union of the co-tangent spaces is the *co-tangent bundle* $T^*M = \coprod_{u \in M} T_u^*M$. Both TM and T^*M are naturally smooth manifolds [Abr+88, Thm. 3.3.10 and Ch. 5.2 resp.].

For a vector space E we define the vector space of continuous $(r + s)$ -multilinear maps $T_s^r(E) = L^{r+s}(E^*, \dots, E^*, E, \dots, E; \mathbb{R})$ with s copies of E and r copies of E^* and where E^* denotes the dual. We say elements of $T_s^r(E)$ are *tensors* on E , *contravariant* of order r and *covariant* of order s . Further, we use the notation $T_s^r(M)$ to denote the *vector bundle of tensors contravariant of order r and covariant of order s* [Abr+88, Def. 5.2.9]. In this notation, $T_0^1(M)$ is identified with the tangent bundle TM and $T_1^0(M)$ with the cotangent bundle T^*M .

Suppose $f : M \rightarrow N$ is a mapping of one manifold into another, and $u \in M$, then by means of charts we can interpret the derivative of f on each chart at u as a linear mapping $df(u) : T_uM \rightarrow T_{f(u)}N$. When $N = \mathbb{R}$, the collection of such maps defines a 1-form $df : M \rightarrow T^*M$. More generally, a 1-form is a continuous map $\omega : M \rightarrow T^*M$ satisfying $\pi \circ \omega = \text{Id}_M$ where $\pi : T^*M \rightarrow M$ is the natural projection mapping $\omega(p) \in T_p^*M$ to $p \in M$.

We use the notation $\Omega^k(M)$ to denote the k -forms on M (for more details, please see [Abr+88, Chapter 7, §3]). In particular, $\Omega^1(M) = T_1^0(M)$. We define the *exterior derivative* $d : \Omega^k(M) \rightarrow \Omega^{k+1}(M)$ as for finite dimensional manifolds as follows:

Theorem 2.A.1 ([Abr+88, Theorem 7.4.1]). *For every smooth manifold M of dimension n , there is unique family of linear maps $d^k(U) : \Omega^k(U) \rightarrow \Omega^{k+1}(U)$ where U is open in M and that we merely denote by d , called the **exterior derivative**, such that*

1. d is \mathbb{R} -linear and for $\alpha \in \Omega^k(U)$ and $\beta \in \Omega^\ell(U)$, $d(\alpha \wedge \beta) = d\alpha \wedge \beta + (-1)^k \alpha \wedge d\beta$,
2. If $f : U \rightarrow \mathbb{R}$ is a real-valued function, then df is just the differential of f , df ,
3. $d^2 = d \circ d = 0$,

4. d is natural with respect to restrictions, i.e. if $U \subset V \subset M$ are open and $\alpha \in \Omega^k(V)$, then $d(\alpha|_U) = (d\alpha)|_U$.

We remark that the extension of the definition of exterior derivative for infinite dimensional manifolds can be found in [Abr+88, Supplement 7.4A]. We will say a 1-form ω is closed if its *exterior derivative* is zero, i.e. $d\omega = 0$.

A point $u \in M$ is said to be a *critical point* of a map $f \in C^r(M, \mathbb{R})$, $r \geq 2$ if $df(u) = 0$. At a critical point $u \in M$, there is a uniquely determined continuous, symmetric, bilinear form (termed the *Hessian*) $d^2f(u) \in T_2^0(M)$ such that $d^2f(u)$ is defined for all $v, w \in T_uM$ by $d^2(f \circ \varphi^{-1})(\varphi(u))(v_\varphi, w_\varphi)$ where φ is any product chart at u and v_φ, w_φ are the local representations of v, w respectively [Pal63, Prop. in §7]. We say $d^2f(u)$ is *positive semi-definite* if there exists $\alpha \geq 0$ such that for any chart φ ,

$$d^2(f \circ \varphi^{-1})(\varphi(u))(v, v) \geq \alpha \|v\|^2, \quad \forall v \in T_{\varphi(u)}E. \quad (2.A.1)$$

If $\alpha > 0$, then we say $d^2f(u)$ is *positive-definite*. Both $\omega(u) = 0$ and positive definiteness are invariant with respect to the choice of coordinate chart.

Given a Banach space E and a bounded, symmetric bilinear form B on E , we say that B is *non-degenerate* if the linear map $A : E \rightarrow E^*$ defined by $A(v)(w) = B(v, w)$ is a linear isomorphism of E onto E^* , otherwise B is *degenerate*. A critical point u of f is called *non-degenerate* if the Hessian of f at u is non-degenerate [Pal63, Def. in §7]. Degeneracy is independent of the choice of coordinate chart.

Consider smooth manifolds M_1, \dots, M_n . The product space $\prod_{i=1}^n M_i = M_1 \times \dots \times M_n$ is naturally a smooth manifold [Abr+88, Def. 3.2.4]. In particular, there is an atlas on $M_1 \times \dots \times M_n$ composed of *product charts* $(U_1 \times \dots \times U_n, \varphi_1 \times \dots \times \varphi_n)$ where (U_i, φ_i) is a chart on M_i for $i \in \{1, \dots, n\}$. We use the notation $\times_{i=1}^n \varphi_i = \varphi_1 \times \dots \times \varphi_n$ and $\prod_{i=1}^n U_i = U_1 \times \dots \times U_n$.

There is a canonical isomorphism at each point such that the cotangent bundle of the product manifold splits:

$$T_{(u_1, \dots, u_n)}^*(M_1 \times \dots \times M_n) \cong T_{u_1}^*M_1 \oplus \dots \oplus T_{u_n}^*M_n \quad (2.A.2)$$

where \oplus denotes the direct sum of vector spaces. There are natural bundle maps

$$\psi_{M_i} : T^*(M_1 \times \dots \times M_n) \rightarrow T^*(M_1 \times \dots \times M_n) \quad (2.A.3)$$

annihilating all the components other than those corresponding to M_i of an element in the cotangent bundle for each $i \in \{1, \dots, n\}$. In particular,

$$\psi_{M_i}(\omega_1, \dots, \omega_n) = (0, \dots, 0, \omega_i, 0, \dots, 0)$$

where $\omega = (\omega_1, \dots, \omega_n) \in T_u^*(M_1 \times \dots \times M_n)$ and 0 is the zero functional in $T_{u_j}^*M_j$ for each $j \neq i$.

Let $M = M_1 \times \cdots \times M_n$. Given a point $u = (u_1, \dots, u_n) \in M$, then $\iota_u^j : M_j \rightarrow M$ is the natural inclusion map where $\iota_u^j(\mu) = (u_1, \dots, u_{j-1}, \mu, u_{j+1}, \dots, u_n)$. Suppose we have a function $f : M \rightarrow \mathbb{R}$. Then the derivatives $D_i f(u)$ of the map $\mu_i \mapsto f(u_1, \dots, u_{i-1}, \mu_i, u_{i+1}, \dots, u_n)$ where $\mu_i \in M_i$ for each $i \in \{1, \dots, n\}$ are called the *partial derivatives* of f at $u \in M$ [Abr+88, Prop. 2.4.12]. They are given by $D_i f(u)(v_i) = df(u)(\bar{v}_i)$ where $v_i \in T_{u_i} M_i$ and $\bar{v}_i = (0, \dots, 0, v_i, 0, \dots, 0) \in T_u M$. Indeed, $d\iota_u^i : T_{u_i} M \rightarrow T_u M$ is a map such that $d\iota_u^i(u_i)(v_i) = \bar{v}_i$. Hence, by the chain rule, we have $D_i f(u) = d(f \circ \iota_u^i)(u_i) = df(u) \circ d\iota_u^i$. Further, we have that for $v = (v_1, \dots, v_n)$, $df(u)(v) = \sum_{i=1}^n D_i f(u)(v_i)$. For second-order partial derivatives, we use the notation $D_{ij}^2 f(u) = D_i(D_j f)(u)$.

Let $\phi : M \rightarrow N$ be a smooth map between smooth manifolds M and N and suppose $f : N \rightarrow \mathbb{R}$ is a smooth function on N . Then the *pullback* of f by ϕ is the smooth function $\phi^* f = f \circ \phi$ on M . Suppose that α is a 1-form on N . Then the pullback of α by ϕ is the 1-form $\phi^* \alpha$ on M defined by $(\phi^* \alpha)_p(X) = \alpha_{\phi(p)}(d\phi_p(X))$ for $p \in M$ and $X \in T_p M$. The *pushforward* of a vector $v \in T_p M$ is a vector $\phi_* v \in T_q N$ defined by $\phi_* v(f) = v(f \circ \phi)$ for all smooth functions $f : N \rightarrow \mathbb{R}$.

2.A.1 Algebraic Topology

In this section, we briefly review some concepts and notation from algebraic topology as needed for the main body of this chapter. A more detailed introduction can be found in [Lee12, Chapter 18].

We first introduce singular homology. The standard p -simplex is the simplex $\Delta_p = [e_0, \dots, e_p] \subset \mathbb{R}^p$ where $e_0 = 0$ and e_i is the i -th standard basis vector. Let M be a topological space. A continuous map $\sigma : \Delta_p \rightarrow M$ is called a *singular p -simplex in M* . The *singular chain group of M in degree p* , denoted $C_p(M)$, is the free abelian group generated by all singular p -simplices in M . An element of this group, called a *singular p -chain*, is a finite formal linear combination of singular p -simplices in M with integer coefficients. The *boundary* of a singular p -simplex $\sigma : \Delta_p \rightarrow M$ is the singular $(p-1)$ -chain $\partial\sigma$ defined by

$$\partial\sigma = \sum_{i=0}^p (-1)^i \sigma \circ F_{i,p}$$

where $F_{i,p} : \Delta_{p-1} \rightarrow \Delta_p$ maps Δ_{p-1} homeomorphically onto the boundary face $\partial_i \Delta_p$ by sending

$$e_0 \mapsto e_0, \dots, e_{i-1} \mapsto e_{i-1}, e_i \mapsto e_{i+1}, \dots, e_{p-1} \mapsto e_p.$$

The map $\partial : C_p(M) \rightarrow C_{p-1}(M)$ is called the *singular boundary operator*. A singular p -chain c is called a *cycle* if $\partial c = 0$ and a *boundary* if $c = \partial b$ for some singular $(p+1)$ -chain b . Let $Z_p(M)$ denote the set of singular p -cycles in M and $B_p(M)$ the set of singular p -boundaries. Note that since ∂ is a homomorphism, $Z_p(M)$ and $B_p(M)$ are subgroups of $C_p(M)$, and because $\partial \circ \partial = 0$, they satisfy $B_p(M) \subseteq Z_p(M)$. The p -th *singular homology group* of M is the quotient group

$$H_p(M) = \frac{Z_p(M)}{B_p(M)}.$$

The sequence of abelian groups and homomorphisms

$$\cdots \rightarrow C_{p+1}(M) \xrightarrow{\partial} C_p(M) \xrightarrow{\partial} C_{p-1}(M) \rightarrow \cdots$$

is the *singular chain complex* and $H_p(M)$ is the p -th homology group of this complex. The equivalence class in $H_p(M)$ of a singular p -cycle c is called its *homology class* and we denote it by $[c]$. We say two p -cycles are homologous if they differ by a boundary.

Now, we introduce smooth singular homology and we will state the connection between the two in a theorem. First, if M is now a smooth manifold, we call a map $\sigma : \Delta_p \rightarrow M$ a *smooth p -simplex in M* where σ is smooth in the sense that it has a smooth extension to a neighborhood of each point. The subgroup of $C_p(M)$ generated by smooth simplices is denoted by $C_p^\infty(M)$ and called the *smooth chain group in degree p* . We call elements in this group which are formal linear combinations of smooth simplices *smooth chains*. Hence, we define the *p -th smooth singular homology group of M* to be

$$H_p^\infty(M) = \frac{\text{Ker}(\partial : C_p^\infty(M) \rightarrow C_{p-1}^\infty(M))}{\text{Im}(\partial : C_{p+1}^\infty(M) \rightarrow C_p^\infty(M))}.$$

The inclusion map $\iota : C_p^\infty(M) \hookrightarrow C_p(M)$ commutes with the boundary operator and hence, induces a map on homology $\iota_* : H_p^\infty(M) \rightarrow H_p(M)$ defined by $\iota_*[c] = [\iota[c]]$.

Theorem 2.A.2 ([Lee12, Theorem 18.7]). *For any smooth manifold M , the map $\iota_* : H_p^\infty(M) \rightarrow H_p(M)$ induced by inclusion is an isomorphism.*

2.A.2 Differential Topology

Consider smooth manifolds M and N of dimension m and n respectively. An k -jet from M to N is an equivalence class $[x, f, U]_k$ of triples (x, f, U) where $U \subset M$ is an open set, $x \in U$, and $f : U \rightarrow N$ is a C^k map. The equivalence relation satisfies $[x, f, U]_k = [y, g, V]_k$ if $x = y$ and in some (and hence any) pair of charts adapted to f at x , f and g have the same derivatives up to order k . We use the notation $[x, f, U]_k = j^k f(x)$ to denote the k -jet of f at x . The set of all k -jets from M to N is denoted by $J^k(M, N)$. The jet bundle $J^k(M, N)$ is a smooth manifold (see [Hir76, Chapter 2] for the construction). For each C^k map $f : M \rightarrow N$ we define a map $j^k f : M \rightarrow J^k(M, N)$ by $x \mapsto j^k f(x)$ and refer to it as the *k -jet extension*.

Definition 2.A.1. *Let M, N be smooth manifolds and $f : M \rightarrow N$ be a smooth mapping. Let Z be a smooth submanifold of N and p a point in M . Then f intersects Z transversally at p (denoted $f \pitchfork Z$ at p) if either $f(p) \notin Z$ or $f(p) \in Z$ and $T_{f(p)}N = T_{f(p)}Z + (f_*)_p(T_pM)$.*

For $1 \leq k < s \leq \infty$ consider the jet map

$$j^k : C^s(M, N) \rightarrow C^{s-k}(M, J^k(M, N)) \tag{2.A.4}$$

and let $Z \subset J^k(M, N)$ be a submanifold. Define

$$\pitchfork^s(M, N; j^k, Z) = \{h \in C^s(M, N) \mid j^k h \pitchfork Z\}. \tag{2.A.5}$$

A subset of a topological space X is *residual* if it contains the intersection of countably many open-dense sets. We say a property is *generic* if the set of all points of X which possess this property is residual [BT10].

The following results will be used to prove genericity of non-degenerate differential Nash equilibria.

Theorem 2.A.3 (Jet Transversality Theorem [Hir76, Theorem 2.8]). *Let M, N be C^∞ manifolds without boundary, and let $Z \subset J^k(M, N)$ be a C^∞ submanifold. Suppose that $1 \leq k < s \leq \infty$. Then, $\bigcap^s(M, N; j^k, Z)$ is residual and thus dense in $C^s(M, N)$ endowed with the strong topology, and open if Z is closed.*

Proposition 2.A.1 ([GG73, Chapter II.4, Proposition 4.2]). *Let M, N be smooth manifolds and $Z \subset N$ a submanifold. Suppose that $\dim M < \text{codim} Z$. Let $f : M \rightarrow N$ be smooth and suppose that $f \pitchfork Z$. Then, $f(M) \cap Z = \emptyset$.*

The Jet Transversality Theorem and Proposition 2.A.1 can be used to show a subset of a jet bundle having a particular set of desired properties is generic. Indeed, consider the jet bundle $J^k(M, N)$ and recall that it is a manifold that contains jets $j^k f : M \rightarrow J^k(M, N)$ as its elements where $f \in C^k(M, N)$. Let $Z \subset J^k(M, N)$ be the submanifold of the jet bundle that *does not* possess the desired properties. If $\dim M < \text{codim} Z$, then for a generic function $f \in C^k(M, N)$ the image of the k -jet extension is disjoint from Z implying that there is an open-dense set of functions having the desired properties.

Chapter 3

Utility Learning and Incentive Design

We will pick up right where we left off in the last chapter: designing incentives to induce more desirable outcomes. Recall the vignette introduced in Section 1.4.1. Imagine now that not only do we have our drivers or energy consumers that make up *society* in our S-CPS, but we also have a planner who is tasked with coordinating these individuals. This can be the local transportation authority or the power company. It could also be a third-party solution provider such as a traffic-routing cellphone application (e.g. Waze, Google Maps) or a demand-response aggregation company (e.g. Ohmconnect).

We will consider a class incentive design problems in which the planner or coordinator, terms which we will use interchangeably, does not know the underlying preferences of the agents that it is trying to coordinate. In the economics literature these types of problems are known as problems of *asymmetric information*—meaning that the involved parties do not possess the same information sets and, as is often the case, one party possesses some information to which the other party is not privy. The particular type of information asymmetry which we consider, i.e. where the preferences of the agents are unknown to the planner, results in a problem of *adverse selection*. The classic example of adverse selection is the *market for lemons* [Ake70] in which the seller of a used car knows more about the car than the buyer. There are a number of components that are hidden from the buyer such as the maintenance upkeep history, engine health, etc. Hence, the buyer could end up with a *lemon* instead of a *cherry*.

We assume that agents, including the planner, are cost minimizers or alternatively, utility maximizers—while in this chapter we will formulate the entire problem given all agents are cost minimizers, the utility maximization formulation is completely analogous. When we say the planner does not know the underlying preferences of the agents, we are assuming that it does not know the value of the parameters of the agents' cost functions. In the following sections, by taking a non-Bayesian approach, we will formulate a utility learning and incentive design problems for both the case when the agents play according to a Nash equilibrium as well as the case when the agents play myopically—e.g. the agents play according to a myopic update rule. We formulate an algorithm for iteratively estimating preferences and designing incentives. By adopting tools from adaptive control and online learning, we show that the

algorithm converges under reasonable assumptions.

We remark that, in contrast, in Chapter 4 we will consider a similar incentive design problem, where we take a Bayesian approach to the estimation problem. In particular, we will assume the planner has a prior over the preferences of the agents.

The results in this chapter have strong ties to both the adaptive control literature [GS84; KV86; SB89] and the online learning literature [CBL06; Nem+09; Rag+10]. The former gives us tools to do tracking of both the observed output (agents' strategies) and the control input (incentive mechanism). It also allows us to go one step further and prove parameter convergence under some additional assumptions—*persistence of excitation*—on the problem formulation and, in particular, the utility learning and incentive design algorithm. The latter provides tools that allow us to generalize the algorithm and get faster convergence of the observed actions of the agents to a more desirable or even socially optimal outcome.

The remainder of this chapter is organized as follows. We first introduce the problem formulation in Section 3.1. We follow that with the utility learning problem for both the case when the agents play according to a Nash strategy and when they play myopically in Section 3.2. We formulate the incentive design problem in Section 3.3 again for both the case when agents play Nash and when they play myopically. In Section 3.4, we present the utility learning and incentive design algorithm and study its convergence in the case without noise in Section 3.5 and with noise in Section 3.6. Finally, we wrap up with discussion and future directions in Section 3.7.

3.1 Problem Formulation

Consider a scenario in which there are n non-cooperative, self-interested agents competing over some scarce resource. Let $U_i \subset \mathbb{R}^{p_i}$ denote the strategy space of agent i . We will use the notation $U = U_1 \times \cdots \times U_n$ to denote the joint strategy space of all the agents and $U_{-i} = U_1 \times \cdots \times U_{i-1} \times U_{i+1} \times \cdots \times U_n$ to denote the joint strategy space of all the agents excluding the i -th agent. We denote the i -th agent's cost function by $f_i(u_i, u_{-i})$ where $u_i \in U_i$ is its choice variable and $u_{-i} = (u_1, \dots, u_{i-1}, u_{i+1}, \dots, u_n) \in U_{-i}$ are the choice variables of all the agents excluding the i -th agent. We use the notation $f_i(u)$ with $u = (u_1, \dots, u_n)$ when more convenient.

Suppose there is planner tasked with inducing agents to play according to a more efficient or desirable equilibrium. In particular, this more desirable equilibrium is the minimizer of the planner's cost $f_c(u, v)$ where $v \in \mathbb{R}^n$ is the choice variable of the planner and we recall that u denotes the choice variables of the competitive agents. The planner's cost function could, in fact, be the *social cost*. That is, the sum of the costs of all the agents. Let (u^d, v^d) denote the minimizer of the planner's cost (which we assume is unique, otherwise the planner must choose amongst the set of minimizers).

The planner is tasked with finding mappings $\gamma_i \in \Gamma \subset C^2(U, \mathbb{R})$ for each $i \in \{1, \dots, n\}$ such that $\gamma_i(u^d) = v_i^d$ and $u^d = (u_1^d, \dots, u_n^d)$ is the collective response of the agents. For instance, if the agents are assumed to play according to a Nash equilibrium, then u^d must

be a Nash equilibrium in the game induced by $\gamma = (\gamma_1, \dots, \gamma_n)$, i.e.

$$u_i^d \in \arg \min_{u_i} \{f_i(u_i, u_{-i}^d) + \gamma_i(u_i, u_{-i}^d)\}. \quad (3.1.1)$$

so that u_i^d is a *best response* to u_{-i}^d for each $i \in \{1, \dots, n\}$. Formally, we define a Nash equilibrium as follows.

Definition 3.1.1 (Nash Equilibrium of the Incentivized Game). *A point $u \in U$ is a **Nash equilibrium** of the incentivized game $(f_1(u) + \gamma_1(u), \dots, f_n(u) + \gamma_n(u))$ if*

$$f_i(u) + \gamma_i(u) \leq f_i(u'_i, u_{-i}) + \gamma_i(u'_i, u_{-i}), \quad \forall u'_i \in U_i. \quad (3.1.2)$$

We remark that if for each i , the inequality in (3.1.2) holds only for a neighborhood $W_i \subset U_i$ of u_i , then u is a **local Nash equilibrium**.

On the other hand, the agents may play myopically according to some update rule—e.g. myopic approximate best response (see Section 2.6 in which we show convergence of gradient play to stable, non-degenerate differential Nash equilibria), approximate fictitious play, or one of many other update rules [FL98]. In this case, the agents' collective response under these dynamics must coincide with the desired response u^d .

We assume that the planner knows the parametric structure of the agents' cost functions and receives observations of the agents' choices over time; however, it does not know the parameters of the agents' cost functions. In particular, for each $i \in \{1, \dots, n\}$, we assume that the cost function of agent i has the parametric form given by

$$f_i(u) = \sum_{j=1}^{m_i} \phi_{i,j}(u) \theta_{i,j}^* \quad (3.1.3)$$

where each $\phi_{i,j} \in C^2(U, \mathbb{R})$ and $\{\phi_{i,1}, \dots, \phi_{i,m_i}\}$ is the set of basis functions for player i . These basis functions are known to the planner; however, the parameters $\theta_{i,j}^*$ are unknown. Let $\Phi_i(u) = [\phi_{i,1}(u) \cdots \phi_{i,m_i}(u)]^T$ and $\theta_i^* = [\theta_{i,1}^* \cdots \theta_{i,m_i}^*]^T$. The admissible set of parameters for player i is Θ_i , a compact subset of \mathbb{R}^{m_i} .

Assumption 3.1.1. *For each $i \in \{1, \dots, n\}$, $\theta_i^* \in \Theta_i$, i.e. each agent's true parameters are in the admissible set.*

In addition, we define the admissible set of incentive mappings $\Gamma \subset C^2(U, \mathbb{R})$ to be all such mappings

$$\gamma_i(u) = \sum_{j=1}^{s_i} \psi_{i,j}(u) \alpha_{i,j} \quad (3.1.4)$$

for a set of basis functions $\{\psi_{i,1}, \dots, \psi_{i,s_i}\}$ where each $\psi_{i,j} \in C^2(U, \mathbb{R})$ and $\alpha_i^T = [\alpha_{i,1} \cdots \alpha_{i,s_i}]$ are parameters. We use the notation $\Psi_i(u) = [\psi_{i,1}(u) \cdots \psi_{i,s_i}(u)]^T$ and we denote the collection of all parameters by $\alpha^T = (\alpha_1^T, \dots, \alpha_n^T)$.

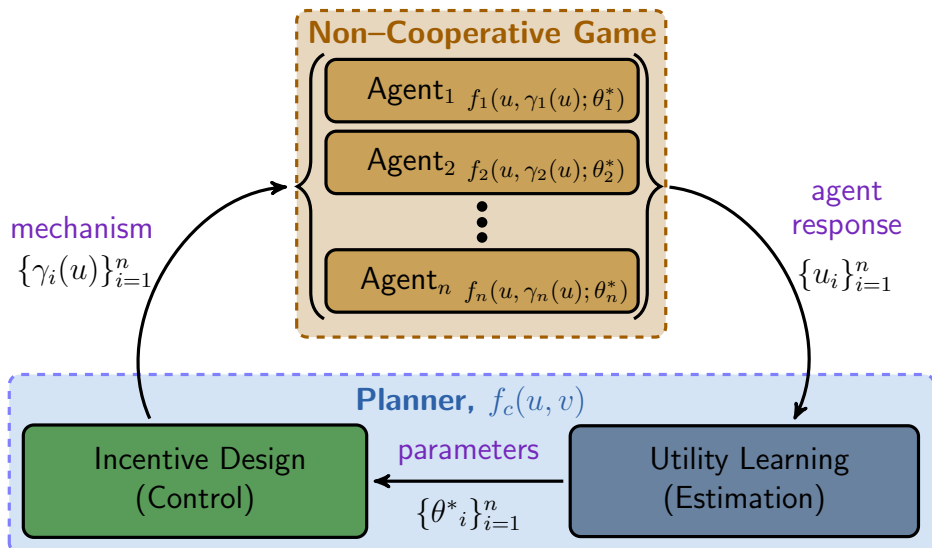


Figure 3.1: *Utility Learning and Incentive Design Algorithm Abstraction*—The agents play in a non-cooperative game induced by the incentive mechanisms $\{\gamma_i\}_{i=1}^n$. Their responses $\{u_i\}_{i=1}^n$ are observed by the planner who first formulates an estimate of the parameters $\{\theta_i^*\}_{i=1}^n$ of the agents’ cost functions. Given the estimated parameters, the planner designs the incentives $\{\gamma_i\}_{i=1}^n$ for the next round that induce the agents to play the desired u^d while satisfying $\gamma_i(u^d) = v_i^d$ for each $i \in \{1, \dots, n\}$. The goal is to derive an algorithm so that as the planner iterates through these steps, for each $i \in \{1, \dots, n\}$, the observed response u_i converges to the desired response u_i^d and the incentive mapping $\gamma_i(u)$ evaluates to the desired value v_i^d .

We remark that this framework encompasses the case where the basis functions of all the agents’ cost functions are identical ($\phi_{i,j} \equiv \phi_j, \forall i \in \{1, \dots, n\}$), the basis functions of the incentive mappings for all the agents are identical ($\psi_{i,j} \equiv \psi_j, \forall i \in \{1, \dots, n\}$), or even the case when the parameters of the incentive mappings are constrained to be identical ($\alpha_{i,j} = \alpha_j, \forall i \in \{1, \dots, n\}$).

The planner receives observations $u^{(k+1)}$ at each time k and uses past observations $\{u^{(t)}\}_{t=0}^{k+1}$ and past incentives $\{\gamma^{(t)}\}_{t=0}^k$ to formulate an estimate of $\theta_i^{(k+1)}$ for each $i \in \{1, \dots, n\}$. Given these estimates, the planner then finds the parameters $\alpha^{(k+1)}$ that induce the agents to play the desired response u^d . Figure 3.1 shows an abstraction of the multi-level game and depicts each step in the process. Our goal is to design an algorithm—in particular, the utility learning (estimation) and incentive design (control) steps—so that as the planner iterates through the process prescribed by the algorithm, each agent’s observed response u_i converge to the desired response u_i^d and the incentive mappings evaluate to the desired values v_i^d .

We will formulate a solution to the utility learning and incentive design problems first,

when the agents are rational and therefore play Nash at each iteration and second, when the agents are myopic and play according to some update rule whose structure is known to the planner. We refer to the former as the *Nash-play* case and the latter as the *myopic-play* case. The latter allows us to consider more realistic versions of the incentive design problem since the update rule may not require the agents to know the cost functions of all the other agents but instead only have information on their choices at each iteration. For both cases we consider the problem with and without noise.

3.2 Utility Learning Formulation

We begin by formulating the utility learning problem for both the Nash-play and myopic-play cases with the goal of placing both cases into a unified framework that we can then use to define an online algorithm for updating the estimate of $\theta^* = (\theta_1^*, \dots, \theta_n^*)$ at each iteration.

3.2.1 Utility Learning Under Nash-Play

Let us introduce the following compact notation: $f_i^{\gamma_i} \equiv f_i + \gamma_i$. We will use the notation $f_i^{\gamma_i}(u; \theta)$ when we need to make the dependence on the parameter θ explicit. At iteration k , given $\{u^{(t)}\}_{t=1}^{k+1}$ and $\{\alpha^{(t)}\}_{t=0}^k$, the planner forms an estimate $\theta_i^{(k)}$ for each $i \in \{1, \dots, n\}$. Our goal is to formulate an online algorithm to estimate θ_i^* for each $i \in \{1, \dots, n\}$ as each new observation $u^{(t)}$ is received. The observations $u^{(t)}$ are assumed to be Nash equilibria. To this end, we write the problem in a more generic form.

Recall that by Proposition 2.2.1, we know that for the incentivized game $(f_1^{\gamma_1}, \dots, f_n^{\gamma_n})$, $\omega(u) = 0$ and $D_{ii}^2 f_i^{\gamma_i}(u) \geq 0$ for each $i \in \{1, \dots, n\}$ are necessary conditions for a Nash equilibrium u . Hence, we expect that for each $t \in \{0, \dots, k\}$,

$$0_{p_i \times 1} = D_i f_i^{\gamma_i}(u^{(t+1)}) = D_i \Phi_i(u^{(t+1)})^T \theta_i^* + D_i \Psi_i(u^{(t+1)})^T \alpha_i^{(t)} \quad (3.2.1)$$

where

$$D_i \Phi_i(u^{(t+1)})^T = \begin{bmatrix} D_{i_1} \Phi_i(u^{(t+1)})^T \\ \vdots \\ D_{i_{p_i}} \Phi_i(u^{(t+1)})^T \end{bmatrix}.$$

In addition, for each $i \in \{1, \dots, n\}$, we have

$$0 \leq D_{ii}^2 f_i^{\gamma_i}(u^{(t+1)}) = D_{ii}^2 \Phi_i(u^{(t+1)}, \theta_i^*) + D_{ii}^2 \Psi_i(u^{(t+1)}, \alpha_i^{(t)}) \quad (3.2.2)$$

where

$$D_{ii}^2 \Phi_i(u^{(t+1)}, \theta_i^*) = \begin{bmatrix} D_{i_1 i_1}^2 \Phi_i(u^{(t+1)})^T \theta_i^* & \dots & D_{i_{p_i} i_1}^2 \Phi_i(u^{(t+1)})^T \theta_i^* \\ \vdots & \ddots & \vdots \\ D_{i_1 i_{p_i}}^2 \Phi_i(u^{(t+1)})^T \theta_i^* & \dots & D_{i_{p_i} i_{p_i}}^2 \Phi_i(u^{(t+1)})^T \theta_i^* \end{bmatrix} \quad (3.2.3)$$

and

$$D_{ii}^2 \Psi_i(u^{(t+1)}, \alpha_i^{(t)}) = \begin{bmatrix} D_{i_1 i_1}^2 \Psi_i(u^{(t+1)})^T \alpha_i^{(t)} & \cdots & D_{i_{p_i} i_1}^2 \Psi_i(u^{(t+1)})^T \alpha_i^{(t)} \\ \vdots & \ddots & \vdots \\ D_{i_1 i_{p_i}}^2 \Psi_i(u^{(t+1)})^T \alpha_i^{(t)} & \cdots & D_{i_{p_i} i_{p_i}}^2 \Psi_i(u^{(t+1)})^T \alpha_i^{(t)} \end{bmatrix}. \quad (3.2.4)$$

In an effort to unify notation across this subsection on agents playing Nash and the sequel on agents playing myopically, we define

$$y_i^{t+1} = -D_i \Psi_i(u^{(t+1)})^T \alpha_i^{(t)} \quad (3.2.5)$$

and

$$\Xi_i^t = D_i \Phi_i(u^{(t+1)}) \quad (3.2.6)$$

so that

$$y_i^{t+1} = (\Xi_i^t)^T \theta_i^*. \quad (3.2.7)$$

Let the admissible set of θ_i 's at iteration k be denoted by Θ_i^k . They are defined by adding the second-order conditions from the assumption that the observations at times $t \in \{1, \dots, k\}$ are Nash equilibria. These sets are nested, i.e.

$$\Theta_i^k \subseteq \Theta_i^{k-1} \subseteq \cdots \subseteq \Theta_i^0 \subseteq \Theta_i,$$

since at each iteration an additional constraint is added to the previous set. To make this more concrete, consider the case where each $p_i = 1$, $i \in \{1, \dots, n\}$. Then we can define

$$A_i(u^{(1,k)}) = \begin{bmatrix} D_{ii}^2 \Phi_i(u^{(1)})^T \\ \vdots \\ D_{ii}^2 \Phi_i(u^{(k)})^T \end{bmatrix} \quad \text{and} \quad b_i(u^{(1,k)}, \alpha_i^{(0,k-1)}) = \begin{bmatrix} D_{ii}^2 \Psi_i(u^{(1)})^T \alpha_i^{(0)} \\ \vdots \\ D_{ii}^2 \Psi_i(u^{(k)})^T \alpha_i^{(k-1)} \end{bmatrix}$$

where we use the fact that $D_{ii}^2 \Phi_i(u^{(t)}, \theta_i^*) = D_{i_1 i_1}^2 \Phi_i(u^{(t)})^T \theta_i^* = D_{ii}^2 \Phi_i(u^{(t)})^T \theta_i^*$ and we use the notation $u^{(1,k)} = (u^{(1)}, \dots, u^{(k)})$ and similarly for $\alpha_i^{(0,k-1)}$. Then using the above notation, the set of admissible θ_i 's are defined by

$$\Theta_i^k = \{\theta_i \in \Theta_i \mid A_i(u^{(1,k)})\theta_i + b_i(u^{(1,k)}, \alpha_i^{(0,k-1)}) \geq 0\} \subseteq \Theta_i. \quad (3.2.8)$$

Thus it is clear that the sets are nested. Similarly, for higher dimensional strategy spaces, i.e. $p_i > 1$, the nested sets are given by

$$\Theta_i^k = \{\theta_i \in \Theta_i \mid D_{ii}^2 \Phi_i(u^{(t)}, \theta_i) + D_{ii}^2 \Psi_i(u^{(t)}, \alpha_i^{(t-1)}) \geq 0, t \in \{1, \dots, k\}\} \subseteq \Theta_i. \quad (3.2.9)$$

Note, in this case, the sets are defined by semi-definite constraints which are still convex [BV04]. We remark that $\theta_i^* \in \Theta_i^k$ for all k since, by assumption, each observation $u^{(k)}$ is a local Nash equilibrium (again, see Proposition 2.2.1).

3.2.2 Utility Learning Under Myopic–Play

We now consider that the agents may not play exactly according to a Nash equilibrium strategy and instead, play according to some myopic update rule. For example, they may play according to *approximate myopic best response* (gradient play) [Rat+13; FL98; CBL06] where agents may update their strategies according to (2.6.2) or even a modified version in which they have their own *learning rates*:

$$u_i^{(k+1)} = u_i^{(k)} - h_i D_i f_i^{\gamma_i^{(k)}}(u^{(k)}) \quad (3.2.10)$$

where h_i represents player i 's learning rate. There are many other possible update rules including *approximate fictitious play* under which players play an approximate best response to the historical frequency of play and *partial approximate best response* under which a fixed portion of the population switches each period from their current action to an approximate best response to the aggregate statistic from the previous period (see, e.g., [FL98] for a more detailed description of both *fictitious play* and *partial approximate best response* as well as other learning dynamics).

We consider any update rule that can be written in the form

$$u_i^{(k+1)} = \tilde{\Phi}_i(u^{(k)})^T \theta_i^* + \tilde{\Psi}_i(u^{(k)})^T \alpha_i^{(k)} \quad (3.2.11)$$

for each $i \in \{1, \dots, n\}$ where $\tilde{\Phi}_i(u^{(k)})^T \in \mathbb{R}^{p_i \times m_i}$ and $\tilde{\Psi}_i(u^{(k)})^T \in \mathbb{R}^{p_i \times n_i}$. Note that $\tilde{\Phi}_i$ and $\tilde{\Psi}_i$ could on the entire past response which we denote by $u^{(0,k)}$ or some subset of the past responses $u^{(j,l)}$ for some $j, l \in \{0, \dots, k\}$ such that $j \leq l$. This would certainly be the case for an update rule determine by approximate fictitious play, for instance. In the remainder, we will simply show the dependence on $u^{(k)}$ and just remind the reader that the framework is flexible enough to encompass dependence on any subset of the past responses.

Let us consider an example. Suppose the agents are playing according to gradient play. Then, for a given player's incentivized cost $f_i^{\gamma_i}(u) = \Phi_i(u)^T \theta_i^* + \Psi_i(u)^T \alpha_i$, it is straightforward to write

$$u_i^{(k+1)} = u_i^{(k)} - h_i (D_i \Phi_i(u^{(k)})^T \theta_i^* + D_i \Psi_i(u^{(k)})^T \alpha_i^{(k)})$$

for each $i \in \{1, \dots, n\}$ in the form of (3.2.11) by either augmenting the parameter θ_i^* to contain the observation $u_i^{(k)}$ at each iteration and modifying the constraint set for θ appropriately or by allowing the left-hand side of (3.2.11) to be $\Delta u_i^{(k+1)} = (u_i^{(k+1)} - u_i^{(k)})$ and then defining $\tilde{\Phi}_i$ and $\tilde{\Psi}_i$ appropriately. We remark that the formulation of the update rule from the particular myopic–play details determines the relationship between $\gamma_i(u) = \Psi_i(u)^T \alpha_i$ and $\tilde{\Psi}_i(u)$. For that matter, it also determines the relationship between f_i and $\tilde{\Phi}_i$.

As before, we denote the set of admissible parameters for player i by Θ_i which we assume to be a compact subset of \mathbb{R}^{m_i} . Again, we assume that $\theta_i^* \in \Theta_i$ for each $i \in \{1, \dots, n\}$. In contrast to the Nash–play case, our admissible set of parameters is no long time varying; however, it could be time–varying in the event there are additional constraints that need to

be enforced as long as the true parameter θ_i^* remains in the admissible set at each iteration and the sets remain closed.

Keeping consistent with the notation of the previous sections, we let

$$y_i^{k+1} = u_i^{(k+1)} - \tilde{\Psi}_i(u^{(k)})^T \alpha_i^{(k)} \quad (3.2.12)$$

and

$$\Xi_i^k = \tilde{\Phi}_i(u^{(k)}) \quad (3.2.13)$$

so that (3.2.11) is rewritten as

$$y_i^{k+1} = (\Xi_i^k)^T \theta_i^*. \quad (3.2.14)$$

At iteration k , the planner uses data $\{u^{(t)}\}_{t=0}^{k+1}$ and $\{\alpha^{(t)}\}_{t=0}^k$ to estimate $\theta^{(k)}$ which is used to design $\alpha^{(k)}$. Notice that we assume the planner observes $u^{(0)}$ —this was not needed in the Nash–play case. The reason we need it in the myopic–play case is so we can determine Ξ_i^0 and y_i^1 .

Now we have massaged both the Nash–play and myopic–play cases in to the same basic form, i.e. for each case, at iteration k , the planner receives an observation

$$y_i^{k+1} = (\Xi_i^k)^T \theta_i^* \quad (3.2.15)$$

for each $i \in \{1, \dots, n\}$ and the planner’s goal is to estimate θ_i^* given the past observations and the past incentives.

3.3 Incentive Design Formulation

In this section, we formulate the incentive design problem for both the Nash–play and myopic–play cases. In both cases, at iteration k , the planner has past observations and incentives and has an estimate of each $\theta_i^{(k+1)}$ for $i \in \{1, \dots, n\}$. The data the planner has for the Nash–play case includes the choices of the players $\{u^{(t)}\}_{t=1}^{k+1}$ and the parameters of the incentives that have been issued $\{\alpha^{(t)}\}_{t=0}^k$. The data the planner has in the myopic–play case includes the choices of the players $\{u^{(t)}\}_{t=0}^{k+1}$ and the parameters of the incentives that have been issued $\{\alpha^{(t)}\}_{t=0}^k$. The planner will use the past data along with the parameter estimates to find an $\alpha^{(k+1)} = (\alpha_1^{(k+1)}, \dots, \alpha_n^{(k+1)})$ such that $\gamma_i(u^d) = v_i^d$ and for each $i \in \{1, \dots, n\}$, u_i^d is the response of player i at iteration $k + 1$. In the Nash–play case, this means u^d is the induced Nash equilibrium in the game

$$(f_1^{\gamma_1}(u; \theta_1^{(k+1)}), \dots, f_n^{\gamma_n}(u; \theta_n^{(k+1)}))$$

where $f_i^{\gamma_i}(u; \theta_i^{(k+1)})$ denotes the incentivized cost of player i parameterized by $\theta_i^{(k+1)}$. In the myopic–play case, this means for each $i \in \{1, \dots, n\}$,

$$u_i^d = \tilde{\Phi}_i(u^{(k)})^T \theta_i^{(k+1)} + \tilde{\Psi}_i(u^{(k)})^T \alpha_i^{(k+1)}.$$

Let us flesh out the details for each case.

3.3.1 Incentive Design Under Nash–Play

Given an estimate of $\theta_i^{(k+1)}$ for each $i \in \{1, \dots, n\}$, the planner seeks an incentive mapping $\gamma^{(k+1)} = (\gamma_1^{(k+1)}, \dots, \gamma_n^{(k+1)})$ such that it induces the desired Nash equilibrium u^d and $\gamma_i^{(k+1)}(u^d) = v_i^d$ for each $i \in \{1, \dots, n\}$. Given that γ_i has been parameterized, this amounts to finding $\alpha_i^{(k+1)}$ for each $i \in \{1, \dots, n\}$ such that u^d is a Nash equilibrium of the game

$$(\Phi_1(u)^T \theta_1^{(k+1)} + \Psi_1(u)^T \alpha_1^{(k+1)}, \dots, \Phi_n(u)^T \theta_n^{(k+1)} + \Psi_n(u)^T \alpha_n^{(k+1)})$$

and such that $\Psi_i(u^d)^T \alpha_i^{(k+1)} = v_i^d$ for each $i \in \{1, \dots, n\}$.

Recall the conditions defining a differential Nash equilibrium (see Definition 2.2.2): $\omega(u) = 0$ and $D_{ii}^2 f_i^{\gamma_i}(u) > 0$. By Theorem 2.2.1, these are also sufficient conditions for a local Nash equilibrium.

Assumption 3.3.1. *For every $\{\theta_i\}_{i=1}^n$ where $\theta_i \in \Theta_i$, there exist $\alpha_i \in \mathbb{R}^{s_i}$ for each $i \in \{1, \dots, n\}$ such that u^d is the induced differential Nash equilibrium in the game*

$$(f_1^{\gamma_1}(u; \theta_1), \dots, f_n^{\gamma_n}(u; \theta_n))$$

and $\gamma_i(u^d) = v_i^d$ where $\gamma_i(u) = \Psi_i(u)^T \alpha_i$.

We remark that the above assumption is not restrictive in the following sense. Finding α_i that induces the desired Nash equilibrium and results in γ_i evaluating to the desired incentive value amounts to finding $\alpha_i^{(k+1)}$ such that the first- and second-order sufficient conditions for a local Nash equilibrium are satisfied given our estimate of the player cost functions. Recall that for each $i \in \{1, \dots, n\}$ the first-order conditions are

$$\underbrace{\begin{bmatrix} D_i \Phi_i(u^d)^T \theta_i^{(k+1)} \\ -v_i^d \end{bmatrix}}_{\zeta_i^{k+1}} + \underbrace{\begin{bmatrix} D_i \Psi_i(u^d)^T \\ \Psi_i(u^d)^T \end{bmatrix}}_{\Lambda_i} \alpha_i^{(k+1)} = \underbrace{\begin{bmatrix} 0 \\ 0 \end{bmatrix}}_{0_{(p_i+1) \times 1}}. \quad (3.3.1)$$

and the second-order conditions are

$$0 < D_{ii}^2 (f_i^{\gamma_i}(u^d)) = D_{ii}^2 \Phi_i(u^d, \theta_i^{(k)}) + D_{ii}^2 \Psi_i(u^d, \alpha_i^{(k+1)}) \quad (3.3.2)$$

where we use the same notation as introduced in the previous section.

If Λ_i is full rank, i.e. has rank $p_i + 1$, then there exists a $\alpha_i^{(k+1)}$ that solves (3.3.1). If the number of basis function s_i satisfies $s_i > p_i + 1$, then the rank condition is not unreasonable and in fact, there are multiple solutions. In essence, by selecting s_i to be *large enough*, the planner is allowing for enough degrees of freedom to ensure there exists a set of parameters α that induce the desired result. Moreover, the problem of finding $\alpha_i^{(k+1)}$ reduces to a convex feasibility problem.

We remark on the case where $\gamma_i \equiv \gamma$ for each i —in particular, $\Psi_{i,j} \equiv \Psi_j$ and $s_i = s$ for each i . Let $p = \sum_{i=1}^n p_i$. In this case, $\alpha \in \mathbb{R}^s$ and the planner needs to find $\alpha^{(k+1)}$ such that $\gamma(u^d) = v^d$ which amounts to finding $\alpha^{(k+1)}$ such that

$$\underbrace{\begin{bmatrix} D_1 \Phi_1(u^d)^T \theta_1^{(k+1)} \\ \vdots \\ D_n \Phi_n(u^d)^T \theta_n^{(k+1)} \\ -v^d \end{bmatrix}}_{\zeta^{k+1}} + \underbrace{\begin{bmatrix} D_1 \Psi(u^d)^T \\ \vdots \\ D_n \Psi(u^d)^T \\ \Psi(u^d)^T \end{bmatrix}}_{\Lambda} \alpha^{(k+1)} = \underbrace{\begin{bmatrix} 0 \\ \vdots \\ 0 \\ 0 \end{bmatrix}}_{0_{(p+1) \times 1}} \quad (3.3.3)$$

and (3.3.2) both hold for each $i \in \{1, \dots, n\}$. Note that $\Lambda \in \mathbb{R}^{(p+1) \times s}$ and $\zeta^k \in \mathbb{R}^{(p+1) \times 1}$. If Λ is full rank, i.e. has rank $p+1$, then there exists a $\alpha^{(k+1)}$ that solves (3.3.3). Just as above, if the number of basis function s to satisfies $s > p+1$, then the rank condition is not unreasonable. Hence, even in this more constrained scenario, Assumption 3.3.1 is arguably not restrictive.

The convex feasibility problem defined by (3.3.1) and (3.3.2) (similarly, (3.3.3) and (3.3.2)) can be formulated as a least-squares type optimization problem:

$$\begin{aligned} \min_{\alpha_i^{k+1} \in \mathbb{R}^{s_i}} \quad & \|\zeta_i^{k+1} + \Lambda_i \alpha_i^{(k+1)}\|_2^2 \\ \text{s.t.} \quad & D_{ii}^2 \Phi_i(u^d, \theta_i^{(k+1)}) + D_{ii}^2 \Psi_i(u^d, \alpha_i^{(k+1)}) > 0, \forall i \in \{1, \dots, n\} \end{aligned} \quad (3.3.4)$$

A regularizer could be added to the problem in order to ensure a *sparse* α_i^{k+1} is found by adding the term $\lambda \|\alpha_i^{(k+1)}\|_1$ to the cost in (3.3.4) where λ is the regularization coefficient which is a well studied problem (see, e.g., [BV04, Chapter 4]).

If it is desirable that the induced Nash equilibrium is a *stable, non-degenerate differential Nash equilibrium* (i.e. a stable, isolated Nash), then the planner must add additional constraints to the feasibility problem defined by (3.3.1) and (3.3.2). In particular, second-order conditions on player cost functions must be satisfied, i.e. that the Hessian of the differential game form $d\omega$ is positive-definite. This reduces to ensuring

$$D^2 \Phi(u, \theta^{(k+1)}) + D^2 \Psi(u, \alpha^{(k+1)}) > 0 \quad (3.3.5)$$

where

$$D^2 \Phi(u, \theta^{(k+1)}) = \begin{bmatrix} D_{11}^2 \Phi_1(u, \theta_1^{(k+1)}) & D_{21}^2 \Phi_1(u, \theta_1^{(k+1)}) & \cdots & D_{n1}^2 \Phi_1(u, \theta_1^{(k+1)}) \\ D_{12}^2 \Phi_1(u, \theta_2^{(k+1)}) & D_{22}^2 \Phi_2(u, \theta_2^{(k+1)}) & \cdots & D_{n2}^2 \Phi_2(u, \theta_2^{(k+1)}) \\ \vdots & \cdots & \ddots & \vdots \\ D_{1n}^2 \Phi_n(u, \theta_n^{(k+1)}) & D_{2n}^2 \Phi_n(u, \theta_n^{(k+1)}) & \cdots & D_{nn}^2 \Phi_n(u, \theta_n^{(k+1)}) \end{bmatrix}, \quad (3.3.6)$$

$$D^2 \Psi(u, \alpha^{(k+1)}) = \begin{bmatrix} D_{11}^2 \Psi_1(u, \alpha_1^{(k+1)}) & D_{21}^2 \Psi_1(u, \alpha_1^{(k+1)}) & \cdots & D_{n1}^2 \Psi_1(u, \alpha_1^{(k+1)}) \\ D_{12}^2 \Psi_1(u, \alpha_2^{(k+1)}) & D_{22}^2 \Psi_2(u, \alpha_2^{(k+1)}) & \cdots & D_{n2}^2 \Psi_2(u, \alpha_2^{(k+1)}) \\ \vdots & \cdots & \ddots & \vdots \\ D_{1n}^2 \Psi_n(u, \alpha_n^{(k+1)}) & D_{2n}^2 \Psi_n(u, \alpha_n^{(k+1)}) & \cdots & D_{nn}^2 \Psi_n(u, \alpha_n^{(k+1)}) \end{bmatrix}. \quad (3.3.7)$$

The notation is consistent with (3.2.3) and (3.2.4). Notice that this constraint is a semi-definite constraint [BV04]. Just as above, finding $\alpha^{(k+1)}$ that induces u^d to be a stable, non-degenerate differential Nash equilibrium can be formulated as a constrained least-squares type optimization problem:

$$\begin{aligned} \min_{\alpha_i^{k+1} \in \mathbb{R}^{s_i}} \quad & \|\zeta_i^{k+1} + \Lambda_i \alpha_i^{(k+1)}\|_2^2 \\ \text{s.t.} \quad & D^2\Phi(u, \theta^{(k+1)}) + D^2\Psi(u, \alpha^{(k+1)}) > 0 \end{aligned} \quad (3.3.8)$$

By Assumption 3.3.1, for each $i \in \{1, \dots, n\}$, there is an $\alpha_i^{(k+1)}$ such that the cost is exactly minimized. The optimization problem (3.3.8) can be written as a semi-definite program and again a regularization term can be incorporated in order to find sparse parameters $\alpha_i^{(k+1)}$.

Ensuring the desired Nash equilibrium is a stable, non-degenerate differential Nash equilibrium means that, first and foremost, the desired Nash equilibrium is *isolated*. Thus, there is no nearby Nash equilibria to which the agents will converge. As we saw in Chapter 2, non-degenerate differential Nash equilibria are generic and structurally stable so that they are robust to small modeling errors and environmental noise. Further, stability ensures that if at each iteration players play according to a myopic approximate best response strategy (gradient play), then they will converge to the desired Nash equilibrium (see Proposition 2.6.1).

If a stable equilibrium is desired by the planner, we can consider a modified version of Assumption 3.3.1:

Assumption 3.3.1' (Modified—Stable Differential Nash). *For every $\{\theta_i\}_{i=1}^n$ where $\theta_i \in \Theta_i$, there exist $\alpha_i \in \mathbb{R}^{s_i}$ for each $i \in \{1, \dots, n\}$ such that u^d is the induced stable, non-degenerate differential Nash equilibrium in the game*

$$(f_1^{\gamma_1}(u; \theta_1), \dots, f_n^{\gamma_n}(u; \theta_n))$$

and $\gamma_i(u^d) = v_i^d$ where $\gamma_i(u) = \Psi_i(u)^T \alpha_i$.

3.3.2 Incentive Design Under Myopic Play

Given an estimate $\theta_i^{(k+1)}$ for each $i \in \{1, \dots, n\}$, the planner seeks an incentive mapping $\gamma^{(k+1)} = (\gamma_1^{(k+1)}, \dots, \gamma_n^{(k+1)})$ that induces the desired response u^d and such that $\gamma_i^{(k+1)}(u^d) = v_i^d$ for each $i \in \{1, \dots, n\}$. As before, given that γ_i has been parameterized, this amounts to finding $\alpha_i^{(k+1)}$ such that

$$u_i^d = \tilde{\Phi}_i(u^{(k)})^T \theta_i^{(k+1)} + \tilde{\Psi}_i(u^{(k)})^T \alpha_i^{(k+1)} \quad (3.3.9)$$

and such that $\Psi_i(u^d)^T \alpha_i^{(k+1)} = v_i^d$ for each $i \in \{1, \dots, n\}$.

Assumption 3.3.2. *For every $\{\theta_i\}_{i=1}^n$ where $\theta_i \in \Theta_i$, there exist $\alpha_i \in \mathbb{R}^{s_i}$ for each $i \in \{1, \dots, n\}$ such that u^d is the estimated collective response—that is, (3.3.9) is satisfied for each $i \in \{1, \dots, n\}$ —and such that $\Psi_i(u^d)^T \alpha_i^{(k+1)} = v_i^d$.*

Finding the parameters α at each iteration that induce the desired response amounts to solving a set of linear equations for each player. That is, for each $i \in \{1, \dots, n\}$, the planner must solve

$$\underbrace{\begin{bmatrix} u_i^d - \tilde{\Phi}_i(u^{(k)})^T \theta_i^{(k+1)} \\ v_i^d \end{bmatrix}}_{\tilde{\zeta}_i^{k+1}} = \underbrace{\begin{bmatrix} \tilde{\Psi}_i(u^{(k)})^T \\ \Psi_i(u^d)^T \end{bmatrix}}_{\tilde{\Lambda}_i^k} \alpha_i^{(k+1)} \quad (3.3.10)$$

for $\alpha_i^{(k+1)}$. The above set of equations will have a solution if the matrix $\tilde{\Lambda}_i^k$ has rank $p_i + 1$. Choosing the set of basis functions $\{\psi_{i,j}\}_{j=1}^{s_i}$ such that $s_i > p_i + 1$ makes this rank condition not unreasonable. One unfortunate difference between the Nash–play case and the present case of myopic–play is that in the former the planner could check the rank condition *a priori* given that it does not depend on the observations. On the other hand, $\tilde{\Lambda}_i^k$ depends on the observation at each iteration.

As before, the problem can be cast as a least–squares type optimization problem:

$$\max_{\alpha_i^{(k+1)} \in \mathbb{R}^{s_i}} \|\tilde{\zeta}_i^{k+1} - \tilde{\Lambda}_i^k \alpha_i^{(k+1)}\|_2^2 \quad (3.3.11)$$

Again, by Assumption 3.3.2, for each $i \in \{1, \dots, n\}$, there is an $\alpha_i^{(k+1)}$ such that the cost is exactly minimized. In addition, a regularization term $\lambda \|\alpha_i^{(k+1)}\|_1$ can be added to enforce a sparse solution if so desired.

3.4 Algorithm

Now that we have formulated the utility learning problem and the incentive design problem at each iteration, we formulate the algorithm for estimating the parameters and designing the incentive mechanism.

At each iteration, the algorithm will have two main steps: update the estimate of θ_i and choose parameters α_i . For each $i \in \{1, \dots, n\}$, consider the loss function

$$\ell(\theta_i^{(k)}) = \frac{1}{2} (y_i^{k+1} - (\Xi_i^k)^T \theta_i^{(k)})^2$$

that evaluates the error in the predicted observation and the true observation at time k for each player.

At iteration k , the parameter estimates of the θ_i 's are updated according to

$$\theta_i^{(k+1)} = P_{\theta_i^{(k)}} \left(\eta_k \nabla \ell(\theta_i^{(k)}) \right) \quad (3.4.1)$$

where $P_{\theta_i^{(k)}}$ is a prox–mapping associated with a distance generating function $\beta(\theta)$ (see Appendix 3.A). In the Nash–play case, the time–varying prox–mapping $P_{\theta_i^{(k)}}^{k+1}$ is used in

place of $P_{\theta_i^{(k)}}$. Note that if the prox-mapping that is used is associated with the distance generating function $\beta(\theta) = \frac{1}{2}\|\theta\|_2^2$, then the prox-mapping is the usual Euclidean projection, $P_\theta(\theta') = \Pi_{\Theta_i}(\theta - \theta')$, so that

$$\theta_i^{(k+1)} = \Pi_{\Theta_i^k} \left(\theta_i^{(k)} - \eta_k \nabla \ell(\theta_i^{(k)}) \right). \quad (3.4.2)$$

Given $\theta^{(k)} = (\theta_1^{(k)}, \dots, \theta_n^{(k)})$, define the set $\mathcal{A}(\theta^{(k)}, u, v)$ to be the set of $\alpha^{(k)} = (\alpha_1^{(k)}, \dots, \alpha_n^{(k)})$ such that u is a differential Nash equilibrium of $(f_1^{\gamma_1}(u), \dots, f_n^{\gamma_n}(u))$ and $\gamma_i(u) = v_i$ where $\gamma_i(u) = \Psi_i(u)^T \alpha_i^{(k)}$. Similarly, define $\mathcal{A}_s(\theta^{(k)}, u, v)$ to be the set of $\alpha^{(k)}$ that induce u to be a stable, non-degenerate differential Nash equilibrium where $\gamma_i(u) = v_i$. By Assumption 3.3.1 (respectively, Assumption 3.3.1'), $\mathcal{A}(\theta^{(k)}, u, v)$ (respectively, $\mathcal{A}_s(\theta^{(k)}, u, v)$) is non-empty. Further, it is straightforward to find an $\alpha^{(k)}$ belonging to $\mathcal{A}(\theta^{(k)}, u, v)$ (respectively $\mathcal{A}_s(\theta^{(k)}, u, v)$) by solving the convex problem stated in (3.3.4) (resp. (3.3.8)). Define $\mathcal{A}_m(\theta^{(k+1)}, u^{(k)}, u^d, v^d)$ to be the set of $\alpha_i^{(k+1)}$ that satisfy (3.3.10). Again, it is straightforward to find an $\alpha_i^{(k+1)}$ by solving the optimization problem posed in (3.3.11).

Suppose central planner has optimized its cost $f_c(u, v)$ to determine u^d and v^d . If the agents are playing according to Nash, then suppose that Assumption 3.3.1 (Assumption 3.3.1', respectively) holds. On the other hand, if the agents are playing myopically, then suppose that Assumption 3.3.2 holds. Then, following Algorithm 1, the planner has a procedure for updating the parameter estimates for each agent's cost function and for choosing the incentive mechanism given the estimated parameters. In the Nash-play case, the planner must use the time varying prox-mapping $P_{\theta_i^{(k)}}^{k+1}$ where in the myopic-play case they use $P_{\theta_i^{(k)}}$. Note that by replacing $\mathcal{A}(\theta^{(k+1)}, u^d, v^d)$ in line 17 of Algorithm 1 with $\mathcal{A}_s(\theta^{(k+1)}, u^d, v^d)$, allows the planner to choose an incentive such that the estimated response is a stable, non-degenerate differential Nash equilibrium.

3.5 Convergence Results

Let the set of basis functions for the agents' cost functions be denoted by \mathcal{F}_ϕ , that is $\phi_{i,j} \in \mathcal{F}_\phi$ where $\phi_{i,j} : U \rightarrow \mathbb{R}$ for each $j \in \{1, \dots, m_i\}$ for all $i \in \{1, \dots, n\}$.

Assumption 3.5.1. *All the functions in \mathcal{F}_ϕ are C^2 -Lipschitz continuous, i.e. for all $\phi \in \mathcal{F}_\phi$, $\phi \in C^2$ and there exists a real constant $K \geq 0$ such that, for all $x, y \in U$, $\|(\phi(x) - \phi(y))\|_{\mathbb{R}} \leq K\|x - y\|_U$ where $\|\cdot\|_{\mathbb{R}}$ and $\|\cdot\|_U$ are norms on \mathbb{R} and U respectively.*

Note that Assumption 3.5.1 implies that the derivative of any function in \mathcal{F}_ϕ is uniformly bounded.

Definition 3.5.1. *If, for each $i \in \{1, \dots, n\}$, there exists a constant $0 < c_{i,1} < \infty$ such that $\Xi_i^t(\Xi_i^t)^T \leq c_{i,1}I$ for all t , then we say the algorithm is stable.*

Algorithm 1: Online Utility Learning and Incentive Design

```

1: If myopic–play then
2:   Receive  $u^{(0)}$ 
3:   Choose arbitrary  $\theta_i^{(0)} \in \Theta_i$  for each  $i \in \{1, \dots, n\}$ 
4:   Choose  $\alpha_i^{(0)} \in \mathbb{R}^{m_i}$  for each  $i \in \{1, \dots, n\}$ 
5:   Issue incentive mapping  $\gamma_i(u) = \Psi_i(u)^T \alpha_i^{(0)}$ 
6:    $k \leftarrow 0$ 
7:   do
8:     for  $i = 1, \dots, n$  do
9:       Receive observation  $y_i^{k+1}$  and incur loss  $\ell(\theta_i^{(k)})$ 
10:      utility learning:
11:        If Nash–play then
12:           $\theta_i^{(k+1)} = P_{\theta_i^{(k)}}^{k+1} \left( \eta_k \nabla \ell(\theta_i^{(k)}) \right)$ 
13:        If myopic–play then
14:           $\theta_i^{(k+1)} = P_{\theta_i^{(k)}} \left( \eta_k \nabla \ell(\theta_i^{(k)}) \right)$ 
15:        incentive design:
16:          If Nash–play then
17:            Choose  $\alpha_i^{(k+1)} \in \mathcal{A}(\theta^{(k+1)}, u^d, v^d)$ 
18:          If myopic–play then
19:            Choose  $\alpha_i^{(k+1)} \in \mathcal{A}_m(\theta^{(k+1)}, u^{(k)}, u^d, v^d)$ 
20:          Issue incentive mapping  $\gamma_i(u) = \Psi_i(u)^T \alpha_i^{(k+1)}$ 
21:        end for
22:       $k \leftarrow k + 1$ 
23:    end do

```

In the Nash–play case, since $\Xi_i^t = D_i \Phi_i(u^{(t+1)})$, it is straightforward to see that such a constant exists for each player. In the myopic–play case, we will assume that the stability condition holds.

Definition 3.5.2. *If for each $i \in \{1, \dots, n\}$, there exists a constant $0 < c_{i,2} < \infty$ such that $c_{i,2} I \leq \Xi_i^t (\Xi_i^t)^T$ for all t , we will say the algorithm is persistently exciting. Define $c_1 = \max_{i \in \{1, \dots, n\}} c_{i,1}$ and $c_2 = \min_{i \in \{1, \dots, n\}} c_{i,1}$.*

Note that we borrow these concepts of persistence of excitation and stability from the adaptive control literature [GS84; KV86; SB89]. We make remarks after we present the results on the connections to these concepts.

From Lemma 3.A.2, we know that for a time–varying prox–mapping P_θ^{t+1} associated with a distance generating function $\beta(\theta)$ (modulus ν) and a function

$$V(\theta_1, \theta_2) = \beta(\theta_2) - \left(\beta(\theta_1) + \nabla \beta(\theta_1)^T (\theta_2 - \theta_1) \right),$$

for every $\theta_1 \in (\Theta^t)^\circ$, $\theta_2 \in \Theta^{t+1}$, and $g \in \mathbb{R}^m$,

$$V(P_{\theta_1}^{t+1}(g), \theta_2) \leq V(\theta_1, \theta_2) + g^T(\theta_2 - \theta_1) + \frac{1}{2\nu} \|g\|_*^2$$

where $\|\cdot\|_*$ denotes the dual norm to $\|\cdot\|$. For each $i \in \{1, \dots, n\}$, if we consider $\theta_1 = \theta_i^{(t)}$, $g = \eta_t \nabla \ell(\theta_i^{(t)})$, and $\theta_2 = \theta_i^*$, then we have

$$V(\theta_i^{(t+1)}, \theta_i^*) \leq V(\theta_i^{(t)}, \theta_i^*) + \eta_t (\theta_i^* - \theta_i^{(t)})^T \nabla \ell(\theta_i^{(t)}) + \frac{\eta_t^2}{2\nu} \|\nabla \ell(\theta_i^{(t)})\|_*^2. \quad (3.5.1)$$

Note that the time-varying prox-mapping is only needed for the Nash-play case; in the myopic-play case, the standard prox-mapping that projects onto Θ_i for each $i \in \{1, \dots, n\}$ at each iteration can be used. To make the notation a little more compact, we will define $V_t(\theta_i) \equiv V(\theta_i^{(t)}, \theta_i)$.

Theorem 3.5.1. *Suppose that a central planner follows Algorithm 1 for utility learning and incentive design with prox-mapping defined by the distance generating function $\beta(\theta_i) = \frac{1}{2} \|\theta_i\|_2^2$ (modulus $\nu = 1$). Further, suppose that the algorithm is persistently exciting, stable, and that the step-size η is chosen such that $\eta - \frac{\eta^2}{2} c_1 > \varepsilon$ for some $\varepsilon > 0$ such that $0 < \varepsilon < \frac{1}{2c_2}$ and where c_1 and c_2 are the stability and persistence of excitation bound respectively. Then for each $i \in \{1, \dots, n\}$, the $\theta_i^{(t)}$ converges exponentially fast to θ_i^* .*

Proof. Let $\beta(\theta_i) = \frac{1}{2} \|\theta_i\|_2^2$ so that $\nu = 1$ and $V_t(\theta_i^*) = V(\theta_i^{(t)}, \theta_i^*) = \frac{1}{2} \|\theta_i^* - \theta_i^{(t)}\|_2^2$. From Lemma 3.A.2, we have that

$$V_{t+1}(\theta_i^*) \leq V_t(\theta_i^*) - \eta (\theta_i^* - \theta_i^{(t)})^T \nabla \ell(\theta_i^{(t)}) + \frac{\eta^2}{2} \|\nabla \ell(\theta_i^{(t)})\|_2^2. \quad (3.5.2)$$

Hence,

$$V_{t+1}(\theta_i^*) \leq V_t(\theta_i^*) - \eta (\theta_i^* - \theta_i^{(t)})^T \Xi_i^t (\Xi_i^t)^T (\theta_i^* - \theta_i^{(t)}) + \frac{\eta^2}{2} (\theta_i^* - \theta_i^{(t)})^T \Xi_i^t (\Xi_i^t)^T \Xi_i^t (\Xi_i^t)^T (\theta_i^* - \theta_i^{(t)}) \quad (3.5.3)$$

$$\leq V_t(\theta_i^*) + (\theta_i^* - \theta_i^{(t)})^T \left(\frac{\eta^2}{2} \Xi_i^t (\Xi_i^t)^T - \eta I \right) \Xi_i^t (\Xi_i^t)^T (\theta_i^* - \theta_i^{(t)}) \quad (3.5.4)$$

$$\leq V_t(\theta_i^*) + (\theta_i^* - \theta_i^{(t)})^T \left(\frac{\eta^2}{2} c_1 - \eta \right) \Xi_i^t (\Xi_i^t)^T (\theta_i^* - \theta_i^{(t)}) \quad (3.5.5)$$

$$\leq V_t(\theta_i^*) - \varepsilon (\theta_i^* - \theta_i^{(t)})^T \Xi_i^t (\Xi_i^t)^T (\theta_i^* - \theta_i^{(t)}) \quad (3.5.6)$$

$$\leq V_t(\theta_i^*) - \varepsilon c_2 (\theta_i^* - \theta_i^{(t)})^T (\theta_i^* - \theta_i^{(t)}) \quad (3.5.7)$$

$$\leq V_t(\theta_i^*) (1 - 2c_2 \varepsilon) \quad (3.5.8)$$

where we used that $\Xi_i^t(\Xi_i^t)^T \leq c_1 I$ (consequence of Assumption 3.5.1), $\eta - \frac{\eta^2}{2}c_1 > \varepsilon$ by construction, and $c_2 I \geq \Xi_i^t(\Xi_i^t)^T$ (persistence of excitation). Since $0 < \varepsilon < \frac{1}{2c_2}$, we have that $1 - 2c_2\varepsilon < e^{-2c_2\varepsilon}$. Hence,

$$V_{t+1}(\theta_i^*) < e^{-2c_2\varepsilon} V_t(\theta_i^*) \quad (3.5.9)$$

so that

$$V_T(\theta_i^*) < e^{-2c_2T\varepsilon} V_0(\theta_i^*) \quad (3.5.10)$$

Therefore we have that $\theta_i^{(t)} \rightarrow \theta_i^*$ exponentially fast. The same argument holds for each $i \in \{1, \dots, n\}$. \square

We now relax the choice of $\beta(\theta) = \frac{1}{2}\|\theta\|_2^2$ and consider more general distance generating functions.

Remark 3.5.1. *In choosing other distance generating functions β —besides $\beta(\theta) = \frac{1}{2}\|\theta\|_2^2$ —that are informed by the geometry of Θ , we can greatly improve the dimension dependence of the algorithm's convergence rates. To provide some context, in [Nem+09], Nemirovski, et al. show precisely how the choice of distance generating function, informed by the geometry of the problem, can improve the convergence rate in a nice example. In particular, they consider the problem of estimating x^* which lives in $X = \{x \in \mathbb{R}^n \mid \sum_{i=1}^n x_i = 1, x \geq 0\}$, a standard simplex. They update their estimates of each x_i^* according to the standard Euclidean projection algorithm ($\|\cdot\| = \|\cdot\|_* = \|\cdot\|_2$ and $\beta(x) = \frac{1}{2}\|x\|_2^2$) and according to an ℓ_1 -norm prox-mapping update where $\beta(x) = \sum_{i=1}^n x_i \ln x_i$ is the entropy function, $\|\cdot\| = \|\cdot\|_1$ and $\|\cdot\|_* = \|\cdot\|_\infty$. In the standard Euclidean projection algorithm requires $O(n \ln n)$ operations to compute the prox-mapping whereas the ℓ_1 -norm requires only $O(n)$ operations. Furthermore, they show that better bounds on the expected loss can be achieved in the ℓ_1 -norm case.*

Theorem 3.5.2. *Suppose that a central planner follows Algorithm 1 for utility learning and incentive design using the prox-mapping associated with β (modulus ν). Further, suppose the algorithm it is persistently exciting and stable. Let the step-size η be chosen such that $\eta - \frac{\eta^2}{2\nu}\tilde{c}_1 > \varepsilon$ for some ε such that $0 < \varepsilon < \frac{1}{2c_2}$ with $c_2 = \min_{i \in \{1, \dots, n\}} c_{i,2}$ and where $0 < \tilde{c}_1 < \infty$ is such that $\|\Xi_i^t\|_*^2 \leq \tilde{c}_1$. Then, for each $i \in \{1, \dots, n\}$,*

$$\lim_{t \rightarrow \infty} \|(\Xi_i^t)^T(\theta_i^{(t)} - \theta_i^*)\|^2 = 0. \quad (3.5.11)$$

Moreover, $V_t(\theta_i^*)$ converges for each $i \in \{1, \dots, n\}$.

Proof. Since $\Xi_i^t(\Xi_i^t)^T \leq c_1 I$ (consequence of Assumption 3.5.1), there exists a $\tilde{c}_1 > 0$ such that for all t , $\|\Xi_i^t\|_*^2 \leq \tilde{c}_1$. From Lemma 3.A.2, we have that

$$V_{t+1}(\theta_i^*) \leq V_t(\theta_i^*) - \eta(\theta_i^* - \theta_i^{(t)})^T \nabla \ell(\theta_i^{(t)}) + \frac{\eta^2}{2\nu} \|\nabla \ell(\theta_i^{(t)})\|_*^2 \quad (3.5.12)$$

Hence,

$$V_{t+1}(\theta_i^*) \leq V_t(\theta_i^*) - \eta \|(\Xi_i^t)^T(\theta_i^* - \theta_i^{(t)})\|^2 + \frac{\eta^2}{2\nu} \|\Xi_i^t\|_*^2 \|(\Xi_i^t)^T(\theta_i^* - \theta_i^{(t)})\|^2 \quad (3.5.13)$$

$$\leq V_t(\theta_i^*) - \left(\eta - \frac{\eta^2}{2\nu} \|\Xi_i^t\|_*^2 \right) \|(\Xi_i^t)^T(\theta_i^* - \theta_i^{(t)})\|^2 \quad (3.5.14)$$

$$\leq V_t(\theta_i^*) - \left(\eta - \frac{\eta^2}{2\nu} \tilde{c}_1 \right) \|(\Xi_i^t)^T(\theta_i^* - \theta_i^{(t)})\|^2 \quad (3.5.15)$$

$$\leq V_t(\theta_i^*) - \varepsilon \|(\Xi_i^t)^T(\theta_i^* - \theta_i^{(t)})\|^2 \quad (3.5.16)$$

Thus,

$$\|(\Xi_i^t)^T(\theta_i^* - \theta_i^{(t)})\|^2 \leq \frac{V_t(\theta_i^*) - V_{t+1}(\theta_i^*)}{\varepsilon} \quad (3.5.17)$$

Summing from $t = 0$ to $t = T$, we get

$$\sum_{t=0}^T \|(\Xi_i^t)^T(\theta_i^* - \theta_i^{(t)})\|^2 \leq \frac{V_0(\theta_i^*) - V_{T+1}(\theta_i^*)}{\varepsilon} \leq \frac{V_0(\theta_i^*)}{\varepsilon} \quad (3.5.18)$$

so that

$$\lim_{T \rightarrow \infty} \sum_{t=0}^T \|(\Xi_i^t)^T(\theta_i^* - \theta_i^{(t)})\|^2 \leq \frac{V_0(\theta_i^*)}{\varepsilon} < \infty \quad (3.5.19)$$

which implies that

$$\lim_{t \rightarrow \infty} \|(\Xi_i^t)^T(\theta_i^* - \theta_i^{(t)})\|^2 = 0 \quad (3.5.20)$$

From (3.5.16) and the fact that $V_t(\theta_i^*)$ is always positive, we see that $V_t(\theta_i^*)$ is a decreasing sequence and hence, it converges. The analysis holds for each $i \in \{1, \dots, n\}$. \square

Remark 3.5.2. *Under the assumptions of Theorem 3.5.2, we can only show that observations converge to zero and that the prox-function $V_t(\theta_i)$ converges. Knowing the parameter values—a consequence of Theorem 3.5.1—allows for the opportunity to gain qualitative insights into how agents' preferences affect the outcome of their strategic interaction. On the other hand, as we have already remarked, the benefit in this case is that the distance generating function β can be chosen so that it is informed by the geometry Θ and this has the potential to improve convergence rates.*

In the myopic-play case, for Theorem 3.5.2 where we have used an arbitrary distance generating function β and $\Xi_i^t = \tilde{\Phi}_i(u^{(t)})$, it is automatic that the observed response converges to the desired response since the observed response is

$$u_i^{(t+1)} = \tilde{\Phi}_i(u^{(t)})^T \theta_i^* + \tilde{\Psi}_i(u^{(t)})^T \alpha_i^{(t)}$$

and the predicted induced response is

$$u_i^d = \tilde{\Phi}_i(u^{(t)})^T \theta_i^{(t)} + \tilde{\Psi}_i(u^{(t)})^T \alpha_i^{(t)}.$$

so that $\|u_i^{(t+1)} - u_i^d\|^2 = \|(\Xi_i^t)^T(\theta_i^* - \theta_i^{(t)})\|^2$.

On the other hand, the result of Theorem 3.5.1 is that $\theta_i^{(t)}$ converges to θ_i . This suggests that $u_i^{(t+1)}$ converges to u_i^d . Indeed,

$$\|u_i^{(t+1)} - u_i^d\|_2^2 \leq \|\tilde{\Phi}_i(u^{(t)})^T \theta_i^* - \tilde{\Phi}_i(u^{(t)}) \theta_i^{(t)}\|_2^2 \leq \|\tilde{\Phi}_i(u^{(t)})^T\|_{2,op}^2 \|\theta_i^* - \theta_i^{(t)}\|_2^2 \quad (3.5.21)$$

Then, since $\tilde{\Phi}_i(u^{(t)})$ is a bounded linear operator (by Assumption 3.5.1), we have that $\|u_i^{(t+1)} - u_i^d\|_2^2$ converges to zero. This implies that $\|u^{(t+1)} - u^d\|_2^2$ converges to zero where $u = (u_1, \dots, u_n)$.

Moreover, because of Assumption 3.3.2, we know that each $\alpha_i^{(t+1)}$ satisfies $\Psi_i(u^d)^T \alpha_i^{(t+1)} = v_i^d$. Let the set of basis functions for the incentive mapping be denoted by \mathcal{F}_ψ . That is, $\psi_{i,j} \in \mathcal{F}_\psi$ where $\psi_{i,j} : U \rightarrow \mathbb{R}$ for each $j \in \{1, \dots, s_i\}$ for all $i \in \{1, \dots, n\}$.

Assumption 3.5.2. *All the functions in \mathcal{F}_ψ are C^2 -Lipschitz continuous, i.e. for all $\psi \in \mathcal{F}_\psi$, $\psi \in C^2$ and there exists a real constant $K \geq 0$ such that, for all $x, y \in U$, $\|(\psi(x) - \psi(y))\|_{\mathbb{R}} \leq K \|x - y\|_U$ where $\|\cdot\|_{\mathbb{R}}$ and $\|\cdot\|_U$ are norms on \mathbb{R} and U respectively.*

Let $v_i^{(t+1)} = \Psi_i(u^{(t+1)})^T \alpha_i^{(t+1)}$. Then,

$$|v_i^{(t+1)} - v_i^d|^2 = |(\Psi_i(u^{(t+1)}) - \Psi_i(u^d))^T \alpha_i^{(t+1)}|^2 \leq \|\Psi_i(u^{(t+1)}) - \Psi_i(u^d)\|^2 \|\alpha_i^{(t+1)}\|_*^2$$

By Assumption 3.5.2, we know that

$$\|\Psi_i(u^{(t+1)}) - \Psi_i(u^d)\| \leq K_i \|u^{(t+1)} - u^d\|$$

for some constant K_i . Hence,

$$|v_i^{(t+1)} - v_i^d|^2 \leq K_i^2 \|\alpha_i^{(t+1)}\|_*^2 \|u^{(t+1)} - u^d\|^2.$$

Therefore, as a consequence of the fact that $\|u^{(t+1)} - u^d\|^2$ converging to zero, we know that $|v_i^{(t+1)} - v_i^d|^2$ converges to zero since $\|\alpha_i^{(t+1)}\|_*^2 < \infty$.

We summarize the above in the following result.

Corollary 3.5.1. *Suppose the agents play according to the myopic update rule (3.2.11). Under the assumptions of Theorem 3.5.1 (respectively, Theorem 3.5.2), for each $i \in \{1, \dots, n\}$, $\|u_i^{(t+1)} - u_i^d\|^2$ converges to zero. Furthermore, if Assumption 3.5.2 holds, then for each $i \in \{1, \dots, n\}$, $|v_i^{(t+1)} - v_i^d|^2$ converges to zero.*

In the Nash-play case, we can use the fact that non-degenerate differential Nash equilibria are structurally stable (see Theorem 2.3.1 in Chapter 2) to determine a bound on how close an equilibrium of the incentivized game

$$(f_1^{\gamma_1}(u; \theta_1^*), \dots, f_n^{\gamma_n}(u; \theta_n^*)) \quad (3.5.22)$$

with $\gamma_i(u) = \Psi_i(u)^T \alpha_i^{(t)}$ for each $i \in \{1, \dots, n\}$ is to the desired Nash equilibrium u^d given that we use the Algorithm 1 to estimate $\theta^{(t)} = (\theta_1^{(t)}, \dots, \theta_n^{(t)})$ and design $\alpha^{(t)} = (\alpha_1^{(t)}, \dots, \alpha_n^{(t)})$. Note that the observed Nash equilibrium $u^{(t+1)}$ is in the set of Nash equilibria of the game (3.5.22).

First, we define some notation. We write the differential game form ω for the incentivized game $(f_1^{\gamma_1}, \dots, f_n^{\gamma_n})$ (see Definition 2.2.1) now as a function of the parameter θ and the decision variable u , i.e. $\omega(\theta, u)$. By a slight abuse of notation, we will denote $D_1\omega(\theta, u)$ and $D_2\omega(\theta, u)$ as the local representation of the differential of ω with respect to θ and u respectively. If the parameters of the incentive mapping $\alpha^{(t)}$ at each iteration t are C^2 with respect to the players' strategies $u^{(t)} = (u_1^{(t)}, \dots, u_n^{(t)})$ and the cost function parameters $\theta^{(t)} = (\theta_1^{(t)}, \dots, \theta_n^{(t)})$, then differential of ω is well-defined. We remark that we formulated the optimization problem for finding the α 's as a constrained least-squares problem and there are well known results for determining when solutions to such problems are continuously dependent on parameter perturbations [Löt83; BS00]. There is still some work in applying such perturbation results to our setup. We leave this to future work and assume we have sufficiently smooth α 's.

We use the notation $p = \sum_{i=1}^n p_i$, $m = \sum_{i=1}^n m_i$, and $s = \sum_{i=1}^n s_i$.

Theorem 3.5.3. *Suppose that for each t , $\alpha^{(t)}(\theta, u) \in C^2(\mathbb{R}^m \times \mathbb{R}^p, \mathbb{R}^n)$ is chosen such that u^d is a non-degenerate differential Nash equilibrium. For $\|\theta^{(t)} - \theta^*\|$ sufficiently small, there is a Nash equilibrium u^* of the incentivized game $(f_1^{\gamma_1}(u; \theta_1^*), \dots, f_n^{\gamma_n}(u; \theta_n^*))$ with $\gamma_i(u) = \Psi_i(u)^T \alpha_i^{(t)}$ that is near the desired Nash equilibrium, i.e. there exists $\bar{\varepsilon} > 0$, such that for all $\theta^{(t)} \in B_{\bar{\varepsilon}}(\theta^*)$,*

$$\|u^* - u^d\| \leq \left(\sup_{0 \leq \lambda \leq 1} \|Dg((1 - \lambda)\theta^* + \lambda\theta^{(t)})\| \right) \|\theta^{(t)} - \theta^*\| \quad (3.5.23)$$

where

$$Dg(\theta) = -(D_2\omega)^{-1}(\theta, u^d) \circ D_1\omega(\theta, u^d). \quad (3.5.24)$$

Furthermore, if $\|Dg(\theta)\|$ is uniformly bounded by $M > 0$ on $B_{\bar{\varepsilon}}(\theta^*)$, then

$$\|u^* - u^d\| \leq M \|\theta^{(t)} - \theta^*\| \quad (3.5.25)$$

Proof. Consider the differential game form $\omega(\theta, u)$ which is given by

$$\omega(\theta, u) = \sum_{i=1}^n \sum_{l=1}^{p_i} \left(\sum_{j=1}^{m_i} D_{i_l} \phi_{i,j}(u) \theta_{i,j} + \sum_{k=1}^{n_i} D_{i_l} \psi_{i,k}(u) \alpha_{i,j} \right) du_i^l \quad (3.5.26)$$

where $D_{i_l}\phi$ denotes the derivative of the ϕ with respect to the l -th coordinate of player i 's strategy u_i —similarly, for $D_{i_l}\psi$ —and $\{du_i^l\}_{l=1}^{p_i}$ is a co-frame for U_i .

Since u^d is a non-degenerate differential Nash equilibrium, $D_2\omega(\theta^*, u^d)$ is an isomorphism. Thus, by the Implicit Function Theorem [Abr+88, Theorem 2.5.7], there exists a neighborhood W_0 of θ^* and a C^1 function $g : W_0 \rightarrow U$ such that for all $\theta \in W_0$,

$$\omega(\theta, g(\theta)) = 0.$$

Furthermore,

$$Dg(\theta) = -(D_2\omega)^{-1}(\theta, u^d) \circ D_1\omega(\theta, u^d)$$

Let $B_{\bar{\varepsilon}}(\theta^*)$ be the largest $\bar{\varepsilon}$ -ball inside of W_0 . Since $B_{\bar{\varepsilon}}(\theta^*)$ is convex, by Proposition [Abr+88, Proposition 2.4.7], we have that

$$g(\theta^{(t)}) - g(\theta^*) = \left(\int_0^1 Dg((1-\lambda)\theta^* + \lambda\theta^{(t)}) d\lambda \right) \cdot (\theta^{(t)} - \theta^*) \quad (3.5.27)$$

Hence,

$$\|u^* - u^d\| = \|g(\theta^{(t)}) - g(\theta^*)\| \leq \left(\sup_{0 \leq \lambda \leq 1} \|Dg((1-\lambda)\theta^* + \lambda\theta^{(t)})\| \right) \|\theta^{(t)} - \theta^*\| \quad (3.5.28)$$

Now, if $\|Dg(\theta^*)\|$ is uniformly bounded by $M > 0$ on $B_{\bar{\varepsilon}}(\theta^*)$, then it is straightforward to see from (3.5.28) that

$$\|u^* - u^d\| \leq M \|\theta^{(t)} - \theta^*\|. \quad (3.5.29)$$

□

As a consequence of Theorem 3.5.1 and Theorem 3.5.3, there exists a finite time t for which $\|\theta_i^* - \theta_i^{(t)}\|_2^2$ is sufficiently small for each $i \in \{1, \dots, n\}$ so that a Nash equilibrium of the incentivized game at time t is arbitrarily close to the desired Nash equilibrium u^d . There may be multiple Nash equilibria of the incentivized game; hence, if the agents converge to u^* so that $u^{(t+1)} = u^*$, then the observed Nash equilibrium is near the desired Nash equilibria. We know that for stable, non-degenerate differential Nash equilibria $u^{(t+1)}$, agents will converge *locally* if following the gradient flow determined by the differential game form ω .

Corollary 3.5.2. *Suppose the assumptions of Theorem 3.5.3 hold and that u^* is stable. If for each $i \in \{1, \dots, n\}$, agent i follows the gradient of their cost $-D_i f_i$, then they will converge locally to u^* so that $u^{(t+1)} = u^*$. Moreover, there exists an $\bar{\varepsilon} > 0$ such that for all $\theta^{(t)} \in B_{\bar{\varepsilon}}(\theta^*)$,*

$$\|u^* - u^d\| \leq \left(\sup_{0 \leq \lambda \leq 1} \|Dg((1-\lambda)\theta^* + \lambda\theta^{(t)})\| \right) \|\theta^{(t)} - \theta^*\| \quad (3.5.30)$$

The above corollary is a direct application of the Theorem 3.5.3 and Proposition 2.6.1. The corollary essentially says that if the Nash equilibrium that the agents select is determined by the gradient flow $\dot{u} = -\omega(\theta^*, u)$ and they all initialize in a neighborhood of u^* , then $u^{(t+1)} = u^*$. The size of such a neighborhood can be approximated using techniques for computation

of region of attraction via a Lyapunov function [Sas99, Chapter 5]. This is in part due to the fact that in the case where $\alpha^{(t)}$ is chosen so that u^d is stable, i.e. $d\omega(\theta^{(t)}, u^d) > 0$, we have that $d\omega(\theta^*, u^*) > 0$ for $\theta^{(t)}$ near θ^* since the spectrum of $d\omega$ varies continuously.

Remark 3.5.3. *It is possible to explicitly construct the neighborhood W_0 obtained via the Implicit Function Theorem in Theorem 3.5.3 (see, e.g. [HH98, Theorem 2.9.10]). In particular, the Implicit Function Theorem tells us that*

$$L = \begin{bmatrix} I & 0 \\ D_1\omega(\theta, u) & D_2\omega(\theta, u) \end{bmatrix}$$

is invertible. Then, we can choose $R > 0$ such that on $B_{2R|L^{-1}|}(\theta^*, u^d)$, the derivative $D\omega$ (the local representation of the differential of ω with respect to (θ, u)) satisfies the Lipschitz condition

$$|D\omega(\theta', u) - D\omega(\theta, v)| \leq \frac{1}{2R|L^{-1}|^2}|(\theta', u) - (\theta, v)|.$$

Then we know that our implicit map $g : W_0 \rightarrow U$ is defined on W_0 where we have constructed $W_0 = B_R(\theta^*)$. This construction then allows us to explicitly determine the value of $\bar{\varepsilon}$ which in turn allows us to know how long to run Algorithm 1.

The result of Theorem 3.5.1 implies that the incentive value under u^* from Theorem 3.5.3 is arbitrarily close to the desired incentive value.

Corollary 3.5.3. *Under the assumptions of Theorem 3.5.1 and Theorem 3.5.3, there exists a finite T such that for all $t \geq T$,*

$$\|u^* - u^d\|_2^2 \leq MCe^{-2c_2t\varepsilon}, \quad \forall t \geq T \quad (3.5.31)$$

where $C = n \max_i \{2V_0(\theta_i^*)\}$ and u^* is the Nash equilibrium of the incentivized game (3.5.22) such that

$$\|u^* - u^d\| \leq M\|\theta^{(t)} - \theta^*\| \quad \forall \theta^{(t)} \in B_{\bar{\varepsilon}}(\theta^*). \quad (3.5.32)$$

Furthermore, if Assumption 3.5.2 holds, then for each $i \in \{1, \dots, n\}$,

$$|v_i^* - v_i^d|^2 \leq MK_i^2 \|\alpha_i^{(t)}\|_2^2 Ce^{-2c_2t\varepsilon}, \quad \forall t \geq T \quad (3.5.33)$$

where K_i is the Lipschitz bound on Ψ_i and $v_i^* = \Psi_i(u^*)^T \alpha_i^{(t)}$.

Proof. Choose T such that, for each $i \in \{1, \dots, n\}$, $2V_0(\theta_i^*)e^{-2\varepsilon c_2 T} < \bar{\varepsilon}$ so that we have

$$\|\theta_i^{(t)} - \theta_i\|_2^2 \leq 2V_0(\theta_i^*)e^{-2\varepsilon c_2 t}, \quad \forall t \geq T. \quad (3.5.34)$$

Thus, $\|\theta^{(t)} - \theta^*\|_2^2 \leq Ce^{-2\varepsilon c_2 t}$, for all $t \geq T$. By Theorem 3.5.3, we have that

$$\|u^* - u^d\|_2^2 \leq MCe^{-2\varepsilon c_2 t}, \quad \forall t \geq T \quad (3.5.35)$$

where M is the uniform bound on $\|Dg(\theta^*)\|$. We know that $v_i^d = \Psi_i(u^d)^T \alpha_i^{(t)}$ and $v_i^* = \Psi_i(u^*)^T \alpha_i^{(t)}$. Hence, by Assumption 3.5.2, we have that

$$|v_i^* - v_i^d|^2 = |(\Psi_i(u^*) - \Psi_i(u^d))^T \alpha_i^{(t)}|^2 \quad (3.5.36)$$

$$\leq \|\Psi_i(u^*) - \Psi_i(u^d)\| \|\alpha_i^{(t)}\|_2^2 \quad (3.5.37)$$

$$\leq K_i^2 \|\alpha_i^{(t)}\|_2^2 \|u^* - u^d\|_2^2 \quad (3.5.38)$$

$$\leq MK_i^2 \|\alpha_i^{(t)}\|_2^2 C e^{-2\epsilon c_2 t}, \quad \forall t \geq T. \quad (3.5.39)$$

□

We have argued that following Algorithm 1 with a particular choice of prox-mapping, the parameter estimates of each θ_i^* converge to the true values and as a consequence we can characterize the bound on how close the observed response and incentive value are to their desired values. Knowing the true parameter values for θ^* allows us to make qualitative insights into the rationale behind the observed responses. Relaxing the assumptions and choosing the prox-mapping to reflect the geometry of the feasible set, we can show more conservatively that the observations converge. The planner no longer gains access to the true parameter values, but has the opportunity to increase convergence rates of the utility learning and incentive design procedure. An interesting direction for future research would be to explore this tradeoff.

3.6 Uncertainty in Agent Play

In this section, we will use the unified framework that describes both the case where the agents play according to Nash and where the agents play myopically.

We will again consider the case where the agents have multiple inputs, $U_i \subset \mathbb{R}^{p_i}$, $p_i \geq 1$ for each $i \in \{1, \dots, n\}$; however, we will make the simplifying assumption that there is a different set of basis functions for each input. In particular, suppose each player's input is denoted by $u_i = (u_{i,1}, \dots, u_{i,p_i})$ and we assume that each input has its own set of basis functions $\{\phi_{i,l,j}\}_{j=1}^{m_{i,l}}$, $l \in \{1, \dots, p_i\}$ and $m_{i,l}$ denotes the dimension of $\theta_{i,l}^*$. We can state (3.2.7) as

$$\begin{bmatrix} y_{i,1}^{t+1} \\ \vdots \\ y_{i,p_i}^{t+1} \end{bmatrix} = \begin{bmatrix} (\xi_{i,1}^t)^T & 0 & 0 & \cdots & 0 \\ 0 & (\xi_{i,2}^t)^T & 0 & \cdots & 0 \\ \vdots & \ddots & \ddots & \cdots & \vdots \\ 0 & 0 & \cdots & \cdots & (\xi_{i,p_i}^t)^T \end{bmatrix} \begin{bmatrix} \theta_{i,1}^* \\ \vdots \\ \theta_{i,p_i}^* \end{bmatrix} \quad (3.6.1)$$

In this case, observations decouple and hence, the estimation for each set of parameters decouples. In this case, each $\theta_{i,l}^{(t)}$ is updated according to Algorithm 1 independently from the others. Since the notation is cumbersome and the details straightforward, we drop the additional index for each input. That is, instead of considering the update of our estimate of

$\theta_{i,l}^*$ in Algorithm 1, we simply consider the update of θ_i^* and note that due to the independence of the parameters for each input, that we could apply the update to each $\theta_{i,l}^*$.

In addition, we will consider *noisy* updates given by

$$y_i^{t+1} = (\xi_i^t)^T \theta_i^* + w_i^{t+1} \quad (3.6.2)$$

for each $i \in \{1, \dots, n\}$ where w_i^{t+1} is an independent, identically distributed (i.i.d) real stochastic process defined on a probability space (Ω, F, P) adapted to the sequence of increasing sub- σ -algebras $(F_t, t \in \mathbb{N})$, where F_t is the σ -algebra generated by the set $\{y_i^s, \alpha_i^{(s)}, w_i^s, s \leq t\}$ and such that

$$\mathbb{E}[w_i^{t+1} | F_t] = 0, \quad \forall t, \quad (3.6.3)$$

and

$$\mathbb{E}[(w_i^{t+1})^2 | F_t] = \sigma^2 > 0 \text{ a.s.}, \quad \forall t. \quad (3.6.4)$$

Note that F_t is also the σ -algebra generated by $\{y_i^s, \xi_i^s, s \leq t\}$ since w_i^t can be deduced from y_i^t and ξ_i^{t-1} through the relationship $w_i^t = y_i^t - (\xi_i^{t-1})^T \theta_i^*$ [KV86].

For simplicity, in this section we will assume that $U_i \subseteq \mathbb{R}^{p_i}$ with $p_i > 1$ and that there is a different set of basis functions for each input as described in (3.6.1). We will also drop the additional index corresponding to each input for each player since the same analysis can be applied to each of the inputs separately. That is, instead of considering p_i different observations

$$y_{i,l}^{t+1} = (\xi_{i,l}^t)^T \theta_{i,l}^* + w_{i,l}^{t+1} \quad \text{for } l \in \{1, \dots, p_i\}$$

we will perform the analysis for $y_i^{t+1} = (\xi_i^t)^T \theta_i^* + w_i^{t+1}$ and simply note that the same analysis holds for each input without carrying around the extra index l .

Theorem 3.6.1. *Suppose that for each $i \in \{1, \dots, n\}$, $\{w_i^t\}$ satisfies (3.6.3), (3.6.4), and (3.6.7). Furthermore, suppose that a central planner follows Algorithm 1 for utility learning and incentive design using the prox-mapping P_θ associated with β (modulus ν) and that the algorithm is persistently exciting and stable. Let the step-size η_t be selected such that $\sum_{t=1}^{\infty} \eta_t^2 < \infty$ and $\eta_t - \frac{\eta_t^2}{2\nu} \tilde{c}_1 > 0$ where $0 < \tilde{c}_1 < \infty$ is such that $\|\xi_i^t\|_*^2 \leq \tilde{c}_1$. Then, for each $i \in \{1, \dots, n\}$, $V_t(\theta_i^*)$ converges almost surely. Further, suppose that the sequence $\{r_t\}$ where $r_t = (\eta_t - \frac{\eta_t^2}{2\nu} \tilde{c}_1)^{-1}$ is a non-decreasing, non-negative sequence such that r_t is F_t measurable. If there exists constants $0 < K_1, K_2 < \infty$ and $0 < \bar{T} < \infty$ such that*

$$\frac{1}{T} r_{T-1} \leq K_1 + \frac{K_2}{T} \sum_{t=0}^{T-1} (y_i^{t+1} - (\xi_i^t)^T \theta_i^{(t)} - w_i^{t+1})^2, \quad \forall T \geq \bar{T}, \quad (3.6.5)$$

then

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[(y_i^{t+1} - (\xi_i^t)^T \theta_i^{(t)})^2 | F_t] = \sigma^2 \quad \text{a.s.} \quad (3.6.6)$$

If we make the additional assumption

$$\sup_t \mathbb{E}[(w_i^{t+1})^4 | F_t] < +\infty \text{ a.s.} \quad (3.6.7)$$

then

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} (y_i^{t+1} - (\xi_i^t)^T \theta_i^{(t)})^2 = \sigma^2 \text{ a.s.} \quad (3.6.8)$$

Proof. The proof follows a similar technique to that presented in [KV86, Chapter 13.4]. Starting from Lemma 3.A.2, we have

$$\begin{aligned} \mathbb{E}[V_{t+1}(\theta_i^*) | F_t] &\leq V_t(\theta_i^*) - \eta_t (\theta_i^* - \theta_i^{(t)})^T \xi_i^t (\mathbb{E}[y_i^{t+1} - (\xi_i^t)^T \theta_i^{(t)} | F_t]) \\ &\quad + \frac{\eta_t^2}{2\nu} \|\xi_i^t\|_*^2 \left((\mathbb{E}[y_i^{t+1} - (\xi_i^t)^T \theta_i^{(t)} | F_t])^2 + \sigma^2 \right) \end{aligned} \quad (3.6.9)$$

$$\begin{aligned} &\leq V_t(\theta_i^*) - \eta_t \left((\theta_i^* - \theta_i^{(t)})^T \xi_i^t - \frac{\eta_t}{2\nu} \|\xi_i^t\|_*^2 \mathbb{E}[y_i^{t+1} - (\xi_i^t)^T \theta_i^{(t)} | F_t] \right) \\ &\quad \cdot \mathbb{E}[y_i^{t+1} - (\xi_i^t)^T \theta_i^{(t)} | F_t] + \frac{\eta_t^2}{2\nu} \|\xi_i^t\|_*^2 \sigma^2 \end{aligned} \quad (3.6.10)$$

$$\leq V_t(\theta_i^*) - \left(\eta_t - \frac{\eta_t^2 \|\xi_i^t\|_*^2}{2\nu} \right) \left(\mathbb{E}[y_i^{t+1} - (\xi_i^t)^T \theta_i^{(t)} | F_t] \right)^2 + \frac{\eta_t^2}{2\nu} \|\xi_i^t\|_*^2 \sigma^2 \quad (3.6.11)$$

$$\leq V_t(\theta_i^*) - \left(\eta_t - \frac{\eta_t^2}{2\nu} \tilde{c}_1 \right) \left(\mathbb{E}[y_i^{t+1} - (\xi_i^t)^T \theta_i^{(t)} | F_t] \right)^2 + \frac{\eta_t^2}{2\nu} \tilde{c}_1 \sigma^2 \quad (3.6.12)$$

By the assumptions that $\eta_t - \frac{\eta_t^2}{2\nu} \tilde{c}_1 > 0$ and $\sum_{t=1}^{\infty} \eta_t^2 < \infty$, we can apply the almost supermartingale convergence theorem (Theorem 3.A.1) to get that

$$\sum_{t=1}^{\infty} \left(\eta_t - \frac{\eta_t^2}{2\nu} \tilde{c}_1 \right) \left(\mathbb{E}[y_i^{t+1} - (\xi_i^t)^T \theta_i^{(t)} | F_t] \right)^2 < \infty \text{ a.s.} \quad (3.6.13)$$

and that $V_t(\theta_i^*)$ converges almost surely.

Now, we argue (3.6.6) (the argument follows that which is presented in [GS84, Chapter 8]). To do this, we first show that

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} (y_i^{t+1} - (\xi_i^t)^T \theta_i^{(t)} - w_i^{t+1})^2 = 0 \text{ a.s.} \quad (3.6.14)$$

Note that (3.6.13) implies that

$$\lim_{T \rightarrow \infty} \sum_{t=0}^{T-1} \frac{(y_i^{t+1} - (\xi_i^t)^T \theta_i^{(t)} - w_i^{t+1})^2}{r_t} < \infty \text{ a.s.} \quad (3.6.15)$$

Where $r_t = (\eta_t - \frac{\eta_t^2}{2\nu} \tilde{c}_1)^{-1}$. Suppose that $r_t < K_3 < \infty$, i.e. that it is bounded. Hence, it is immediate from (3.6.15) that

$$\lim_{T \rightarrow \infty} \frac{1}{K_3} \sum_{t=0}^{T-1} (y_i^{t+1} - (\xi_i^t)^T \theta_i^{(t)} - w_i^{t+1})^2 < \infty \quad \text{a.s.} \quad (3.6.16)$$

so that (3.6.14) follows trivially. On the other hand, suppose r_t is unbounded. Then we can apply Kronecker's Lemma 3.A.3 to conclude that

$$\lim_{T \rightarrow \infty} \frac{1}{r_T} \sum_{t=0}^{T-1} (y_i^{t+1} - (\xi_i^t)^T \theta_i^{(t)} - w_i^{t+1})^2 = 0 \quad \text{a.s.}$$

Hence, from (3.6.5), we have that

$$\lim_{T \rightarrow \infty} \frac{\frac{1}{T} \sum_{t=0}^{T-1} (y_i^{t+1} - (\xi_i^t)^T \theta_i^{(t)} - w_i^{t+1})^2}{K_1 + \frac{K_2}{T} \sum_{t=0}^{T-1} (y_i^{t+1} - (\xi_i^t)^T \theta_i^{(t)} - w_i^{t+1})^2} = 0 \quad \text{a.s.} \quad (3.6.17)$$

so that (3.6.14) follows immediately. Note that

$$\mathbb{E}[(y_i^{t+1} - (\xi_i^t)^T \theta_i^{(t)})^2 | F_t] = \mathbb{E}[(y_i^{t+1} - (\xi_i^t)^T \theta_i^{(t)} - w_i^{t+1} + w_i^{t+1})^2 | F_t] \quad (3.6.18)$$

$$\begin{aligned} &= \mathbb{E}[(y_i^{t+1} - (\xi_i^t)^T \theta_i^{(t)} - w_i^{t+1})^2 + (w_i^{t+1})^2 \\ &\quad + 2(y_i^{t+1} - (\xi_i^t)^T \theta_i^{(t)} - w_i^{t+1})w_i^{t+1} | F_t] \end{aligned} \quad (3.6.19)$$

Since $y_i^{t+1} - w_i^{t+1}$ and $(\xi_i^t)^T \theta_i^{(t)}$ are F_t measurable and $\mathbb{E}[w_i^{t+1} | F_t] = 0$ almost surely, we have that

$$\mathbb{E}[(y_i^{t+1} - (\xi_i^t)^T \theta_i^{(t)})^2 | F_t] = (y_i^{t+1} - (\xi_i^t)^T \theta_i^{(t)} - w_i^{t+1})^2 + \mathbb{E}[(w_i^{t+1})^2 | F_t]. \quad (3.6.20)$$

Thus, (3.6.6) holds since $\mathbb{E}[(w_i^{t+1})^2 | F_t] = \sigma^2$ almost surely. Finally, if $\sup_t \mathbb{E}[(w_i^{t+1})^4 | F_t] < +\infty$ almost surely, then by Proposition 3.A.1, we have that

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} (y_i^{t+1} - (\xi_i^t)^T \theta_i^{(t)})^2 = \sigma^2 \quad \text{a.s.} \quad (3.6.21)$$

□

Let us state an alternative version of the above theorem after which we will make comments on the differences between the assumptions.

Theorem 3.6.2. *Suppose that for each $i \in \{1, \dots, n\}$, $\{w_i^t\}$ satisfies (3.6.3), (3.6.4), and (3.6.7). Furthermore, suppose that a central planner follows Algorithm 1 for utility learning and incentive design using the prox-mapping P_θ associated with β (modulus ν) and that*

the algorithm is persistently exciting and stable. Let the step-size η_t be selected such that $\sum_{t=1}^{\infty} \eta_t^2 < \infty$ and $\eta_t > 0$ where $0 < \tilde{c}_1 < \infty$ is such that $\|\xi_i^t\|_*^2 \leq \tilde{c}_1$. Then, for each $i \in \{1, \dots, n\}$, $V_t(\theta_i^*)$ converges almost surely. Further, suppose that each Θ_i^t is bounded and that the sequence $\{r_t\}$ where $r_t = (\eta_t)^{-1}$ is a non-decreasing, non-negative sequence such that r_t is F_t measurable. If there exists constants $0 < K_1, K_2 < \infty$ and $0 < \bar{T} < \infty$ such that

$$\frac{1}{\bar{T}} r_{T-1} \leq K_1 + \frac{K_2}{\bar{T}} \sum_{t=0}^{T-1} (y_i^{t+1} - (\xi_i^t)^T \theta_i^{(t)} - w_i^{t+1})^2, \quad \forall T \geq \bar{T}, \quad (3.6.22)$$

then

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[(y_i^{t+1} - (\xi_i^t)^T \theta_i^{(t)})^2 | F_t] = \sigma^2 \quad a.s. \quad (3.6.23)$$

If we make the additional assumption

$$\sup_t \mathbb{E}[(w_i^{t+1})^4 | F_t] < +\infty \quad a.s. \quad (3.6.24)$$

then

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} (y_i^{t+1} - (\xi_i^t)^T \theta_i^{(t)})^2 = \sigma^2 \quad a.s. \quad (3.6.25)$$

Note that added the assumption that each Θ_i^t is bounded and the modified assumption on the sequence r_t .

Proof. The proof is essentially the same as above with some minor modifications. Starting from Lemma 3.A.2, we have

$$\begin{aligned} \mathbb{E}[V_{t+1}(\theta_i^*) | F_t] &\leq V_t(\theta_i^*) - \eta_t (\theta_i^* - \theta_i^{(t)})^T \xi_i^t (\mathbb{E}[y_i^{t+1} - (\xi_i^t)^T \theta_i^{(t)} | F_t]) \\ &\quad + \frac{\eta_t^2}{2\nu} \|\xi_i^t\|_*^2 \left((\mathbb{E}[y_i^{t+1} - (\xi_i^t)^T \theta_i^{(t)} | F_t])^2 + \sigma^2 \right) \end{aligned} \quad (3.6.26)$$

$$\begin{aligned} &\leq V_t(\theta_i^*) - \eta_t \left(\mathbb{E}[y_i^{t+1} - (\xi_i^t)^T \theta_i^{(t)} | F_t] \right)^2 \\ &\quad + \frac{\eta_t^2}{2\nu} \|\xi_i^t\|_*^2 \left(\sigma^2 + (\mathbb{E}[y_i^{t+1} - (\xi_i^t)^T \theta_i^{(t)} | F_t])^2 \right) \end{aligned} \quad (3.6.27)$$

Hence, since Θ_i^t is bounded along with the stability assumption, we have that there exists some constant b such that $\mathbb{E}[y_i^{t+1} - (\xi_i^t)^T \theta_i^{(t)} | F_t] < b$. This gives us that

$$\mathbb{E}[V_{t+1}(\theta_i^*) | F_t] \leq V_t(\theta_i^*) - \eta_t \left(\mathbb{E}[y_i^{t+1} - (\xi_i^t)^T \theta_i^{(t)} | F_t] \right)^2 + \frac{\eta_t^2}{2\nu} \|\xi_i^t\|_*^2 (\sigma^2 + b^2) \quad (3.6.28)$$

By the assumptions that $\eta_t > 0$ and $\sum_{t=1}^{\infty} \eta_t^2 < \infty$, we can apply the almost supermartingale convergence theorem (Theorem 3.A.1) to get that

$$\sum_{t=1}^{\infty} \eta_t \left(\mathbb{E}[y_i^{t+1} - (\xi_i^t)^T \theta_i^{(t)} | F_t] \right)^2 < \infty \quad a.s. \quad (3.6.29)$$

and that $V_t(\theta_i^*)$ converges almost surely.

The remainder of the proof follows exactly what was shown for Theorem 3.6.1 only with $r_t = \eta_t^{-1}$; hence, we refer the reader to its proof. \square

Remark 3.6.1. *Let us remark on the difference between Theorem 3.6.1 and Theorem 3.6.2. The simplified assumptions on the step-size η_t allow for us to find choices for η_t for which it is much easier to check if condition (3.6.22) is satisfied. For instance, the step-size $\eta_t = t^{-1}$ trivially satisfies (3.6.22) and $\sum_t \eta_t^2 < \infty$. Further, the rate at which r_t increases is much slower since $r_t = \eta_t^{-1}$ instead of $r_t = (\eta_t - \frac{\eta_t^2}{2\nu} \tilde{c}_1)^{-1}$. This is desirable and is, in part, the purpose of (3.6.22), i.e. to ensure that the rate that r_t grows is proportionally upper bounded by the average error. The drawback to Theorem 3.6.2 is the additional assumption that Θ_i^t is bounded.*

Corollary 3.6.1. *Suppose the assumptions of Theorem 3.6.1 (or Theorem 3.6.2) hold with the exception that in Algorithm 1, we use the prox-mapping P_θ associated with $\beta(\theta) = \frac{1}{2}\|\theta\|_2^2$ (modulus $\nu = 1$). Then $V_t(\theta_i^*) = \frac{1}{2}\|\theta_i^* - \theta_i^{(t)}\|_2^2$ converges almost surely.*

In the myopic-play case—that is where $u_i^{(t+1)} = \tilde{\Phi}_i(u^{(t)})^T \theta_i^* + \tilde{\Psi}_i(u^{(t)})^T \alpha_i^{(t)}$ —as a consequence of (3.6.6), we expect that the average mean square error between the desired response u^d and the actual response $u^{(t)}$ converges to σ^2 almost surely. Indeed, (3.6.6) of Theorem 3.6.1—or (3.6.23) of Theorem 3.6.2—imply that

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[(u_i^{(t+1)} + w_i^{t+1} - u_i^d)^2 | F_t] = \sigma^2 \quad \text{a.s.}$$

Furthermore, note that we have designed $\alpha_i^{(t+1)}$ to satisfy $\gamma_i(u^d) = \Psi(u^{(t)})^T \alpha_i^{(t+1)}$ (or at the very least approximately if we solve the optimization problem posed in (3.3.11) as opposed to explicitly solving the linear system of equations).

On the other hand, in the Nash-play case, it is difficult to say much about the observed Nash equilibrium except in expectation. In particular, we can consider a modified version of Theorem 3.5.3 where we consider the differential game form in expectation, i.e. at iteration t , we have the local representation of the differential game form for the induced game

$$\tilde{\omega}(\theta, u) = \sum_{i=1}^n \mathbb{E} \left[D_i \Phi_i(u)^T \theta_i + D_i \Psi_i(u)^T \alpha_i^{(t)} + w_i^{t+1} | F_{t-1} \right] \quad (3.6.30)$$

Before, we knew that $D_2 \omega(\theta^*, u^d)$ was an isomorphism since u^d is a non-degenerate differential Nash equilibrium. Here, in order to apply the Implicit Function Theorem as in Theorem 3.5.3, we need that $D_2 \tilde{\omega}(\theta^*, u^d)$ is an isomorphism. Hence, we have the following result.

Proposition 3.6.1. *Suppose that $D_2\tilde{\omega}(\theta, u^d)$ is an isomorphism. There exists an $\epsilon > 0$, such that for all $\theta^{(t)} \in B_\epsilon(\theta^*)$,*

$$\|u^* - u^d\| \leq \left(\sup_{0 \leq \lambda \leq 1} \|Dg((1-\lambda)\theta^* + \lambda\theta^{(t)})\| \right) \|\theta^{(t)} - \theta^*\| \quad (3.6.31)$$

where

$$Dg(\theta^*) = -(D_2\tilde{\omega})^{-1}(\theta^*, u^d) \circ D_1\tilde{\omega}(\theta^*, u^d). \quad (3.6.32)$$

and u^* is a Nash equilibrium of the incentivized game $(f_1^{\gamma_1}(u; \theta_1^*), \dots, f_n^{\gamma_n}(u; \theta_n^*))$ with $\gamma_i(u) = \Psi_i(u)^T \alpha_i^{(t)}$ for each $i \in \{1, \dots, n\}$. Furthermore, if $\|Dg(\theta^*)\|$ is uniformly bounded by $M > 0$ on $B_\epsilon(\theta^*)$, then

$$\|u^* - u^d\| \leq M\|\theta^{(t)} - \theta^*\| \quad (3.6.33)$$

To apply Proposition 3.6.1 we would need a result ensuring that the parameter estimate $\theta^{(t)}$ converges to the true parameter value θ^* . One of the consequences of Theorem 3.6.1 (respectively Theorem 3.6.2) is that $V_t(\theta_i^*)$ converges almost surely and as a consequence of Corollary 3.6.1, $\|\theta_i^* - \theta_i^{(t)}\|_2^2$ converges almost surely. If it is the case that it converges almost surely to a value less than ϵ , then Proposition 3.6.1 would guarantee that a Nash equilibrium of the incentivized game is near the desired non-degenerate differential Nash equilibrium in expectation. We leave further exploration of the convergence of the parameter estimate $\theta^{(t)}$ as future work.

The results of Theorem 3.6.1 imply that the average mean square error between the observations and the predictions converges to σ^2 almost surely and if we recall, the observations are derived from noisy versions of the first-order conditions for Nash, i.e. we have the observation

$$y_i^{t+1} = (\xi_i^t)^T \theta_i^* + w_i^{t+1} = D_i \Phi_i(u^{(t+1)})^T \theta_i^* + w_i^{t+1}$$

and its predicted value

$$(\xi_i^t)^T \theta_i^{(t)} = D_i \Phi_i(u^{(t+1)})^T \theta_i^{(t)} = -D_i \Psi_i(u^{(t+1)})^T \alpha_i^{(t)}.$$

Thus, we have shown that

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[(D_i \Phi_i(u^{(t+1)})^T \theta_i^* + w_i^{t+1} + D_i \Psi_i(u^{(t+1)})^T \alpha_i^{(t)})^2 | F_t] = \sigma^2 \quad \text{a.s.}$$

or, equivalently,

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[(D_i \Phi_i(u^{(t+1)})^T (\theta_i^* - \theta_i^{(t)}) + w_i^{t+1})^2 | F_t] = \sigma^2 \quad \text{a.s.}$$

Remark 3.6.2 (Connections to Adaptive Control and Online Learning). *There are some not so subtle connections to the adaptive control literature. In particular, if in the myopic-play case the entries of $\tilde{\Phi}_i(u)$ and $\tilde{\Psi}(u)$ are linear in u_i , then the problem can easily be cast as a classical linear adaptive control problem [KV86, Chapter 10 & 13]. The Nash-play case has some distinctions in that the actual response is defined implicitly through the first- and second-order conditions for local Nash equilibria and is therefore new.*

Similarly, if we consider only the problem of estimating θ , then both the Nash-play case and the myopic-play case, even with nonlinearities in the observation updates, can be cast as an online estimation problem [GS84, Chapter 3 & 8] with the added difference that in the present work we did not select specific step-sizes. For example, if we make specific choices for the step-size η_t such as

$$\eta_t = \frac{a}{c + \|\xi_i^t\|_2^2},$$

the update scheme for $\theta_i^{(t)}$ reduces to the projection algorithm. Alternatively, if $\eta_t = \frac{\mu}{r_k}$ where $r_k = 1 + \sum_{t=0}^k (\xi_i^t)^T \xi_i^t$ and $\mu > 0$, then updates for θ follow the algorithm introduced in [KV86].

It is interesting to note that the persistence of excitation condition required to ensure convergence is informed by the choice of step-size. This is explored extensively in [GS84, Chapter 3.4]. However, because the conditions required for persistence of excitation depend on the data received at each time step (and this is true of Definition 3.5.2 as well), it is difficult to check when an algorithm is persistently exciting a priori. In [SB89] and [BS86], necessary and sufficient conditions are provided for persistence of excitation for model reference adaptive control. We are exploring extensions for the problems described in this chapter.

Our approach here is to try to remove the dependence of the results on the choice of step-size and to generalize to arbitrary choice of proximal mapping. This approach is similar to that which is done in the online learning literature [Rag+10; Nem+09]. The perceived gain is that this flexibility can allow for the proximal mapping to be informed by the geometry of the constraint set (feasible parameter space) thereby resulting in improved convergence rates.

There are a number of open questions in this area. For instance, in order to ensure we have estimated parameters that correspond to a game with stable, non-degenerate differential Nash equilibria requires having a constraint set defined by semi-definite constraints, i.e.

$$\Theta_i^k = \{\theta_i \in \Theta_i \mid D^2\Phi(u^{(t)}, \theta_i) + D^2\Psi(u^{(t)}, \alpha^{(t-1)}) \geq 0, t \in \{1, \dots, k\}\} \subset \Theta_i \quad (3.6.34)$$

where $D^2\Phi$ and $D^2\Psi$ are defined in (3.3.6) and (3.3.7) respectively. Given that proximal maps informed by the geometry—as the example in Remark 3.5.1 indicates—one interesting direction for future research is to investigate proximal maps that are informed by the geometry of a the positive semi-definite cone of matrices.

3.7 Discussion

By utilizing tools from online learning and adaptive control, we developed an iterative algorithm for learning the objective functions of competitive agents and designing incentives

to elicit from them a desired response. We consider both competitive agents who play according to a Nash equilibrium strategy and players that use a myopic update rule such as approximate best response. We are able to show in both cases, under reasonable assumptions, that the parameter estimates converge resulting in an incentive that induces a desired response. We provided convergence results for the algorithm in both the case with noise and without. We are ensuring that we have approximate *incentive compatibility* in the sense that the agent is acting rationally and the incentives we issue result in the players' choices converging asymptotically to the desired outcome.

There are a number of interesting directions for future research in this area. To highlight a few, we have left open the question of whether we can derive a solution that is budget balanced or individually rational—the latter meaning that players voluntarily participate. Such properties can be formulated as additional objective functions or constraints. For instance, if having a budget balanced solution is desired, the planner can add this as a constraint in the optimization problem to find (u^d, v^d) or include it in its objective function directly. We have explored some of these extensions in the context of designing incentives when there is perfect information [Coo+13; Rat+12]. Voluntary participation can be enforced by ensuring that each player's incentivized cost remains less than the outside option—the alternative to participating in the incentive program. Adding constraints enforcing voluntary participation could be done at each step in the algorithm where we choose the incentive design parameters $\alpha^{(k+1)}$ in such a way that in addition to all the proposed constraints, we also enforce that the estimated cost for each player under $\theta_i^{(k+1)}$ and $\alpha_i^{(k+1)}$ is less than the outside option.

These are classical questions that arise in the economic theory of incentives [LM02]. It is well known that mechanisms that achieve a socially optimal, budget balanced, individually rational, and incentive compatible solution generally do not exist [Arr50; MS83]. While we will touch on these concepts in greater detail in the sequel, in the context of the present chapter, an interesting direction for future research is not just in implementing each of these ideas in the algorithm for utility learning and incentive design, but rather to understand the classes of problems admitting solutions that satisfy a subset of these properties.

Furthermore, we have made preliminary efforts to test the proposed algorithm in practice. We constructed an experimental platform for inducing building occupants to consume shared resources more efficiently by involving them in a *social game*. Occupants vote on their desired use of shared resources such as lighting and heating, ventilation and air conditioning. Some function of their votes such as the average is implemented and occupants are rewarded points based on how efficient their votes are in comparison to the other occupants. We use a lottery mechanism to reward them. By implementing our utility learning scheme, we are able to get reasonably accurate predictions of the occupants' decision-making behavior and in simulation, our incentive design mechanism proves successful at inducing more efficient behavior [Rat+14e]. The number of occupants in the experiments is small (approximately 20); using tools from statistical learning, we created a scalable algorithm for estimating and predicting agents behavior as well as determining the stopping time required to obtain a prespecified accuracy bound [Jin+15]. However, there is much more to be done in terms of experimental validation as well as creating methods that are ready to transition to practice.

In particular, selecting the basis functions that accurately reflect how agents are making decisions is a very difficult task in practice. In applications we have explored [Rat+14c] and preliminary experiments [Rat+14e], we reasoned that agents were making decisions by trading off their comfort and their desire to win. In this simplistic setting, we get reasonably good predictions and, more importantly, due to the simple interpretation of each of the basis functions, we are able to make qualitative assessments about the reasoning behind certain decisions. Prediction and generalization could greatly improve with a more diverse set of basis functions, but at the loss of such qualitative insights.

We are investigating techniques for factoring in categorical data such as automated survey responses that perhaps occur prior to the incentive program or online. Such categorical data can help inform the decision-making model. Furthermore, we are currently exploring non-parametric methods for estimation as well as determining a set of basis functions that could be used in the parametric setting. This is a two pronged approach. On the one hand, we are using non-parametric methods on historical data to determine a set of basis functions that will be used in a online parametric setting. This approach has the advantage of pushing the computationally heavy work offline. On the other hand, we are developing a non-parametric version of the utility learning and incentive design algorithm that can be used directly online. We expect that employing non-parametric methods in either case will serve to improve predictions and generalization.

Generalization techniques in settings where there is a strong human-CPS coupling are difficult to obtain. In [Jin+15] we explored some methods in *transfer learning* [PY10]—or *learning to learn* [Bax97; Bax98], as it is sometimes called. In the transfer learning framework, learning can be done in one environment and be *transferred* to another. For instance, one goal of the social game, as described above, is to learn behavioral models and preferences of occupants in order to improve building automation. Learning the agents preferences in the competitive environment we created will not necessarily be applicable when the social game has ceased. Transfer learning is designed to address exactly this type of problem; tools from transfer learning can be used to build predictive models of behavior that are designed to have little generalization error.

Another interesting direction for future research is in balancing model-based approaches with data-driven approaches. The model-based approach—much like the parametric methods here—allows the planner to make qualitative insights which are important for shaping policy, regulations, and more broadly, system design. Data-driven approaches—much like non-parametric methods—are *evidence based* in the sense that observations are used to form an arguably more of an objective view of the system lacking some of the bias to which model-based approaches are prone. In addition, they have the ability to scale. Striking the right balance between model-based and data-driven methods is really key. We will return to this issue in Chapter 5.

Appendix 3.A Preliminaries

In this appendix, we introduce some of the mathematical preliminaries and results needed for this chapter. The following notation is taken from [Nem+09].

Let for any function $f : U \rightarrow \mathbb{R}$, for $u \in U$, let $\partial f(u)$ denote the set of *subgradients* of f , i.e. $g \in \partial f(u)$ at $u \in U$ if for all $u' \in U$,

$$f(u') \geq f(u) + g^T(u' - u).$$

Suppose that Θ is a compact subset of \mathbb{R}^m . We will say that a function $\beta : \Theta \rightarrow \mathbb{R}$ is a *distance generating* function modulus $\nu > 0$ with respect to $\|\cdot\|$, if β is convex and continuous on Θ_i , the set

$$\Theta^\circ = \{\theta \in \Theta \mid \partial\beta(\theta) \neq \emptyset\}$$

is convex (Θ° always contains the relative interior of Θ) and restricted to Θ° , β is continuously differentiable and *strongly convex* with parameter ν with respect to $\|\cdot\|$, i.e.,

$$(\theta' - \theta)^T(\nabla\beta(\theta') - \nabla\beta(\theta)) \geq \nu\|\theta' - \theta\|^2, \quad \forall \theta', \theta \in \Theta^\circ.$$

As an example, consider $\beta(\theta) = \frac{1}{2}\|\theta\|_2^2$ (modulus $\nu = 1$ with respect to $\|\cdot\|_2$, $\Theta^\circ = \Theta$).

We define a function $V : \Theta^\circ \times \Theta \rightarrow \mathbb{R}_+$ as follows:

$$V(\theta_1, \theta_2) = \beta(\theta_2) - (\beta(\theta_1) + \nabla\beta(\theta_1)^T(\theta_2 - \theta_1)) \quad (3.A.1)$$

We shall call functions V of the above form *prox-functions*—or Bregman divergence [Bre67]—associated with distance generating function $\beta(\theta)$. The function $V(\theta_1, \cdot)$ is nonnegative and is a strongly convex modulus ν with respect to $\|\cdot\|$.

We define a *prox-mapping* $P_\theta : \mathbb{R}^m \rightarrow \Theta^\circ$ associated with β and a point $\theta \in \Theta^\circ$, viewed as a parameter, by

$$P_\theta(g) = \arg \min_{\theta' \in \Theta^\circ} \{g^T(\theta' - \theta) + V(\theta, \theta')\} \quad (3.A.2)$$

We remark that the minimum in the right-hand side is obtained since β is continuous on Θ and Θ is compact, and all the minimizers belong to Θ° so that the minimizer is unique since $V(\theta, \cdot)$ is strongly convex on Θ° . Thus the prox-mapping is well-defined. Furthermore, it is a contraction [Mor65, Proposition 5.b].

We define the time varying prox-mapping $P_\theta^k : \mathbb{R}^m \rightarrow (\Theta^k)^\circ$ associated with β and a point $\theta \in (\Theta^k)^\circ$, viewed as a parameter, by

$$P_\theta^k(g) = \arg \min_{\theta' \in \Theta^k} \{g^T(\theta' - \theta) + V(\theta, \theta')\} \quad (3.A.3)$$

Let Π_{Θ_i} be the metric projection operator onto the set Θ , that is,

$$\Pi_{\Theta}(\theta) = \arg \min_{\theta' \in \Theta} \|\theta - \theta'\|_2. \quad (3.A.4)$$

Note that Π_{Θ_i} is a non-expanding operator, i.e.

$$\|\Pi_{\Theta_i}(\theta') - \Pi_{\Theta_i}(\theta)\|_2 \leq \|\theta' - \theta\|_2, \quad \forall \theta, \theta' \in \mathbb{R}^{m_i}$$

Then for $\beta(\theta) = \frac{1}{2}\|\theta\|_2^2$, we have that $P_\theta = \Pi_{\Theta}(\theta - \theta')$.

Lemma 3.A.1 ([Nem+09, Lemma 2.1]). *For every $\theta' \in \Theta, \theta \in \Theta^\circ$, and $g \in \mathbb{R}^m$, we have*

$$V(P_\theta(g), \theta') \leq V(\theta, \theta') + g^T(\theta' - \theta) + \frac{1}{2\nu} \|g\|_*^2 \quad (3.A.5)$$

where $\|\cdot\|_*$ is the dual norm to $\|\cdot\|$.

Note that for $\beta(\theta) = \frac{1}{2}\|\theta\|_2^2$, we have $V(\theta, \theta') = \frac{1}{2}\|\theta - \theta'\|_2^2$, $\nu = 1$, $\|\cdot\|_* = \|\cdot\|_2$.

We can extend the above lemma to the case of time-varying prox-mappings when the sets Θ^k contain the true value of the parameter θ^* we are trying to estimate. We will need Young's inequality which says that for any $v_1, v_2 \in \mathbb{R}^m$, we have that

$$v_1^T v_2 \leq \|v_1\|_* \|v_2\| \leq \frac{1}{2} \left(\frac{\|v_1\|_*^2}{\nu} + \nu \|v_2\|^2 \right). \quad (3.A.6)$$

Lemma 3.A.2. *For every $\theta^* \in \Theta^{k+1}, \theta^{(k)} \in (\Theta^k)^\circ$, and $g \in \mathbb{R}^m$, we have*

$$V(P_{\theta^{(k)}}^{k+1}(g), \theta^*) \leq V(\theta^{(k)}, \theta^*) + g^T(\theta^* - \theta^{(k)}) + \frac{1}{2\nu} \|g\|_*^2 \quad (3.A.7)$$

Proof. The proof is the essentially the same as the proof of Lemma 3.A.1 as presented in [Nem+09] with a few modifications.

Let $\theta^{(k)} \in (\Theta^k)^\circ$ and $\theta^{(k+1)} = P_{\theta^{(k)}}^{k+1}(g)$. Note that

$$\theta^{(k+1)} \in \arg \min_{\theta' \in \Theta^{k+1}} \{g^T(\theta' - \theta^{(k)}) + V(\theta^{(k)}, \theta')\} \quad (3.A.8)$$

or equivalently,

$$\theta^{(k+1)} \in \arg \min_{\theta' \in \Theta^{k+1}} \{\beta(\theta') - (\nabla\beta(\theta^{(k)}) - g)^T \theta'\} \quad (3.A.9)$$

where the latter form tells us that β is differentiable at $\theta^{(k+1)}$ and $\theta^{(k+1)} \in (\Theta^{k+1})^\circ$. Since $\nabla_2 V(\theta^{(k)}, \theta^{(k+1)}) = \nabla\beta(\theta^{(k+1)}) - \nabla\beta(\theta^{(k)})$, the optimality conditions for (3.A.8) imply that

$$(\nabla\beta(\theta^{(k+1)}) - \nabla\beta(\theta^{(k)}) + g)^T(\theta^{(k+1)} - \theta) \leq 0, \quad \forall \theta \in \Theta^{k+1} \quad (3.A.10)$$

Note that this is where the proof of Lemma 3.A.1 and the current proof are different. The above inequality holds here for all $\theta \in \Theta^{k+1}$ whereas in the proof of Lemma 3.A.1—using the notation of the current Lemma—the inequality would have held for all $\theta \in \Theta^k$. In particular, we need the inequality to hold for θ^* and it does since by assumption $\theta^* \in \Theta^{k+1}$ for each k .

Hence, for $\theta^* \in \Theta^{k+1}$, we have

$$\begin{aligned} V(\theta^{(k+1)}, \theta^*) - V(\theta^{(k)}, \theta^*) &= (\beta(\theta^*) - \nabla\beta(\theta^{(k+1)})^T(\theta^* - \theta^{(k+1)}) - \beta(\theta^{(k+1)})) \\ &\quad - (\beta(\theta^*) - \nabla\beta(\theta^{(k)})^T(\theta^* - \theta^{(k)}) - \beta(\theta^{(k)})) \\ &= (\nabla\beta(\theta^{(k+1)}) - \nabla\beta(\theta^{(k)}) + g)^T(\theta^{(k+1)} - \theta^*) + g^T(\theta^* - \theta^{(k+1)}) \\ &\leq g^T(\theta^* - \theta^{(k+1)}) - V(\theta^{(k)}, \theta^{(k+1)}) \end{aligned}$$

where the last inequality holds due to (3.A.10). By (3.A.6), we have that

$$g^T(\theta^{(k)} - \theta^{(k+1)}) \leq \frac{\|g\|_*^2}{2\nu} + \frac{\nu}{2}\|\theta^{(k)} - \theta^{(k+1)}\|^2. \quad (3.A.11)$$

Further, $\frac{\nu}{2}\|\theta^{(k)} - \theta^{(k+1)}\|^2 \leq V(\theta^{(k)}, \theta^{(k+1)})$ since $V(\theta^{(k)}, \cdot)$ is strongly convex. Thus,

$$\begin{aligned} V(\theta^{(k+1)}, \theta^*) - V(\theta^{(k)}, \theta^*) &\leq g^T(\theta^* - \theta^{(k+1)}) - V(\theta^{(k)}, \theta^{(k+1)}) \\ &= g^T(\theta^* - \theta^{(k)}) + g^T(\theta^{(k)} - \theta^{(k+1)}) - V(\theta^{(k)}, \theta^{(k+1)}) \\ &\leq g^T(\theta^* - \theta^{(k)}) + \frac{1}{2\nu}\|g\|_*^2 \end{aligned}$$

so that

$$V(P_{\theta^{(k)}}^{k+1}(g), \theta^*) \leq V(\theta^{(k)}, \theta^*) + g^T(\theta^* - \theta^{(k)}) + \frac{1}{2\nu}\|g\|_*^2. \quad (3.A.12)$$

□

The following classical results are needed for the proofs in Section 3.6.

Theorem 3.A.1 (Almost supermartingale convergence [RSS5]). *For each $n = 1, 2, \dots$, let z_n, μ_n, ζ_n , and τ_n be non-negative F_n -measurable random variables such that*

$$\mathbb{E}[z_{n+1}|F_n] \leq z_n(1 + \mu_n) + \zeta_n - \tau_n.$$

If $\sum_{n=1}^{\infty} \mu_n < \infty$ and $\sum_{n=1}^{\infty} \zeta_n < \infty$, then $\lim_{n \rightarrow \infty} z_n$ exists and is finite and $\sum_{n=1}^{\infty} \tau_n < \infty$ almost surely.

Lemma 3.A.3 (Kronecker's Lemma [KV86]). *Let $\{x_k\}$ and $\{r_k\}$ be two real valued sequences such that $r_k > 0$, $\lim_{k \rightarrow \infty} r_k = \infty$, and $\sum_{k=1}^{\infty} \frac{x_k}{r_k} < \infty$. Then, $\lim_{N \rightarrow \infty} \frac{1}{r_N} \sum_{k=1}^N x_k = 0$.*

Theorem 3.A.2 (Martingale Stability [KV86, Theorem 8.5.26]). *Suppose $\{x_k, F_k\}$ is a martingale difference sequence, (i.e. $F_k \subset F_{k+1}$ is an increasing sequence of σ -algebras, x_k is F_k -measurable, and $\mathbb{E}[x_{k+1}|F_k] = 0$ a.s.). In addition, suppose that for some $0 < p \leq 2$,*

$$\sum_{k=1}^{\infty} \frac{1}{k^p} \mathbb{E}[|x_k|^p | F_{k-1}] < \infty \quad a.s.$$

Then,

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{k=1}^T x_k = 0 \quad a.s.$$

Proposition 3.A.1 ([Nev75]). *Let $\{x_t\}$ be a zero conditional mean sequence of random variables adapted to $\{F_t\}$. If*

$$\sum_{t=0}^{\infty} \frac{1}{t^2} \mathbb{E}[x_t^2 | F_{t-1}] < \infty \quad a.s. \quad (3.A.13)$$

then

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{t=1}^N x_t = 0 \quad a.s. \quad (3.A.14)$$

Chapter 4

Privacy–Aware Incentive Design

We now consider the fact that in S-CPS, the planner often leans on the underlying CPS infrastructure to support its efforts to design mechanisms, both economic and physical, that aim to make the system more efficient. To reiterate the energy S-CPS vignette presented in Section 1.4.1, consider the set of non-cooperative, selfish agents that make up *society* to be consumers who are interested in consuming energy to maximize their own satisfaction. The planner we will consider will be the power company who has some objective such as implementing a direct load control (DLC) scheme to, for instance, correct for improper load forecasting. In order to do so, the power company takes advantage of access to high fidelity data from smart meters at the consumers’ homes thereby leading to increased exposure both to privacy and security risks.

Increasingly advanced metering infrastructure (AMI) is replacing older technology in the electricity grid. Smart meters measure detailed information about consumer electricity usage every half-hour, quarter-hour, or in some cases, every five minutes. This high-granularity data is needed to support energy efficiency efforts as well as demand-side management. In particular, high-granularity data is useful for customer segmentation [AR13], customizing offering to consumers in the form of pricing structure [Mot+12], detecting non-technical loss [Ami+15], and efficient operations management, e.g. peak load reduction, load shaping, direct load control [Don+14]. However, improper handling of this information could also lead to unprecedented invasions of consumer privacy [Sal+12; WT11; Har89; McK+12].

It has been shown that energy consumption data reveals a considerable amount of information about consumer activities. Furthermore, energy consumption data in combination with readily available sources of information can be used to discover even more about the consumer. Authors in [Lis+10] show that a privacy breach can be broadly implemented in two steps. First, energy usage data in combination with other sensors in the home—e.g. water and gas usage—can be used to track a person’s location, their appliance usage, and match individuals to observed events. In the second step, this learned information can be combined with demographic data—e.g. number, age, sex of individuals in the residence—to identify activities, behaviors, etc.

Given that smart grid operations inherently have privacy and security risks [Sal+12], it

would benefit the power company, to know the answer to the following questions: How do consumers in the population value privacy? How can we quantify privacy? How do privacy-aware policies impact smart grid operations? There have been a number of works making efforts to address these questions [Don+13b; Don+14; San+13; Raj+11]. In particular, it has been shown that there is a fundamental efficiency–privacy tradeoff in data collection policies in smart grid operations [Don+14; San+13].

Privacy is fundamentally subjective; it depends on the underlying preferences of the individual whose privacy is potentially being violated. To capture this fact, we propose an economic solution that allows for the power company to combine inference metrics with privacy-based service contracts to balance the efficiency–privacy tradeoff in the smart grid.

In general, contracts are essential for realizing the benefits of economic exchanges, such as the procurement of electric power from a strategic seller [TT14] and the design of demand response programs [FA00], among others [Ged94]. However, in the context of S-CPS operations and management, it is not enough to consider only the economic aspects of procuring power. Outside of contract design, there are a number of other methods to solicit hidden information for another party [CL14]. In addition, there are techniques that combine notions of privacy and mechanism design where privacy is considered a constraint or enforced through aggregation across players [Kea+12; PR13].

We consider privacy to be the good on which we design contracts, thereby allowing us to capture the fact that privacy is intrinsic to the consumer. We are seeing the rise of this phenomena in new consumer models for internet services provided by such companies as AT&T [Aue15; Luc15]. We formulate an optimization problem that allows for selection of a data collection policy and an amount to charge the consumer in order to maximize the utility of deploying smart grid technologies, while at the same time giving consumers—who have a variety of subjective preferences for their privacy—the option of selecting the privacy setting that fits their needs. This economic solution embeds in it the fact that the quality of service depends on how CPS technologies are being utilized by incorporating a detection theory framework that provides the means to quantify privacy, the good over which we propose to contract.

The optimal contracts—privacy setting and price—are incentive compatible and individually rational, meaning that consumers truthfully reveal their preferences and voluntarily participate. However, we show that the solution is not socially optimal in that consumers having a high valuation of privacy *free-ride* on the rest of society. Further, we assess loss risks due to privacy breaches given the optimally designed contracts. We design new contracts when these risks are explicitly considered by the power company. We show that there are inefficiencies when we consider privacy risk losses and thus, the power company has an incentive to offer compensation to the consumer, invest in security measures, and purchase insurance.

In this chapter we provide a theoretical framework in which economic tools can be combined with inference metrics in order to assess the potential impact of privacy breaches and, further, to provide qualitative insights within this framework. In Section 4.1, we design privacy-based service contracts when the privacy preferences of the consumer are unknown

to the power company and we introduce a DLC example to help concretize the formulation. In Section 4.2, we characterize the contract solution when the consumer is *risk-averse* and the risk of a privacy breach is explicitly modeled in the contract design. We return to the DLC example and show the affects of risk on performance. We argue that the power company has an incentive to invest in insurance or security and, in Section 4.3, we design insurance contracts offered by a third-party. Finally, in Section 4.4, we provide discussion. We remark that the results in this chapter extend those results appearing in our work [Rat+14f; Rat+14a].

4.1 Privacy Contracts

Our goal is to characterize the interactions between energy consumers and a power company that wants to obtain more data from consumers in order to improve the efficiency of their operations, while respecting the different privacy preferences of consumers. The status quo for most smart meter deployments does not provide fine-grained privacy options to consumers: consumers either have to accept the collection of data at predefined intervals (e.g., every 15 minutes) or opt-out (when available) from the installation of a smart meter on their premises.

In this section we propose a more fine-grained approach where power companies explore options for obtaining something they want (consumer data), while respecting the privacy preferences of individuals. For example, a consumer that does not feel comfortable allowing the power company to collect more than one sample per month of her electricity consumption might be charged the regular electricity cost in her monthly bill. On the other hand, another consumer may be willing to accept a significantly reduced electricity bill in exchange for allowing the power company greater data access. For instance, for a DLC program to work efficiently, high fidelity data is needed; perhaps the incentive to participate in the DLC program—which includes access to the needed data—is the reduction in energy bill the consumer experiences in practice.

We design privacy-based service contracts utilizing the theory of *screening* [Web11; BD05]. Specifically, a set of possible privacy settings are offered to consumers having privacy preferences that are unknown to the power company. The problem we consider is a classical principle-agent problem [BD05] with adverse selection in which a principle (the power company) desires that the agent (the consumer) perform a particular action but does not have access to the complete decision making process used by that agent to arrive at its chosen action. It is well known—by the Myerson-Satterthwaite Theorem [MS83] which is an extension in some sense of Arrow’s Impossibility Theorem [Arr50]—that there is no efficient way for two parties to trade a good when they have a secret and probabilistically varying valuations for it without forcing one party to trade at a loss. In particular, for a principle-agent problem, there is no mechanism that is socially optimal, budget balanced, incentive compatible and individually rational (voluntary participation). There are other related results including the Gibbard-Satterthwaite [All73; Sat75] and Green-Laffont [GL77]

results. As we mentioned above, we will design contracts that satisfy incentive compatibility and are individually rational.

We assume that a consumer's privacy preferences are characterized by the parameter θ —referred to as the consumer's *type*—that belongs to the set $\Theta = \{\theta_1, \dots, \theta_n\}$ where $\theta_i < \theta_{i+1}$ and $\theta_i \in \mathbb{R}$ for each $i \in \{1, \dots, n\}$. Note that the type is distinct from the private information that is subject to a privacy breach.

We model privacy as an economic good whose *quality* is the privacy-setting on the smart meter (e.g., sampling rate, noise injection). The privacy setting is a mapping $x : \Theta \rightarrow \mathbb{R}$. Since Θ is a finite set, we can denote the set of privacy settings by $\mathcal{X} = \{x_1, \dots, x_n\}$ where $x_i = x(\theta_i)$ and $x(\cdot)$ is selected by the power company.

Remark 4.1.1. *While we present the results for the case where the type space Θ is finite, there is an analogous framework for a continuum of types [BD05] and the same insights hold.*

We denote the consumer's utility by $\hat{f}(x, \theta)$ and make the following assumption:

Assumption 4.1.1. *The utility function $\hat{f} : \mathbb{R} \times \Theta \rightarrow \mathbb{R}$ is strictly increasing in $(x, \theta) \in \mathbb{R} \times \Theta$, concave and differentiable with respect to x .*

It is reasonable that the utility function of the consumer is increasing in (x, θ) since, for some privacy setting x , consumers with a high valuation of privacy experience greater satisfaction at a higher privacy setting than those with a lower valuation of privacy.

The consumer's type θ is unknown to the power company. However, we assume that the power company has a prior distribution over Θ . In particular, the power company faces type θ_i with probability p_i . Given this prior over Θ , the power company must design a *menu of contracts* $Y = \{y_i\}_{i=1}^n$ where $y_i = (x_i, t_i)$ with x_i the privacy setting being offered at the price t_i , i.e the price for preserving privacy. The menu of contracts is chosen to maximize the power company's profit which is given by

$$f_c(Y) = \sum_{i=1}^n p_i (t_i - g(x_i)) \quad (4.1.1)$$

where $g : x \mapsto g(x) \in \mathbb{R}$ is the unit cost of privacy setting x .

Assumption 4.1.2. *The cost function $g : \mathbb{R} \rightarrow \mathbb{R}$ is a strictly increasing, convex, and differentiable function.*

This assumption is reasonable since a low-privacy setting provides the power company with the high-granularity data necessary for efficient smart grid operations [Don+14; San+13; Raj+11]. For example, in [Don+14], the error in the DLC scheme as a function of sampling period increases approximately quadratically. We will use this fact in the DLC example later in this section. Furthermore, in [Raj+11] a framework for smart grid utility-privacy is presented and the measure of *utility* they propose is differentiable and convex.

Note that the contract mechanism will be successful only if consumers choose contracts designed for their types, i.e., consumers truthfully report their privacy preferences. In addition, consumers must choose to opt-in to a contract. Formally, voluntary participation is guaranteed by

$$\hat{f}(x_i, \theta_i) - t_i \geq 0, \quad (\text{IR-i})$$

for each $i \in \{1, \dots, n\}$ and (IR-i) is referred to as the *individual rationality* constraint. Roughly speaking, a consumer selects a contract only if she gets more profit by participating than otherwise. We assume that a consumer gets zero profit by not participating in the program (*outside option*). The framework is general enough to consider non-zero outside options. For the sake of clarity, we do not flesh out those details here.

In order to guarantee the consumer reports its type truthfully, the power company must enforce *incentive-compatibility* constraints:

$$\hat{f}(x_i, \theta_i) - t_i \geq \hat{f}(x_j, \theta_i) - t_j, \quad \forall j \neq i \quad (\text{IC-i,j})$$

for each $i \in \{1, \dots, n\}$. Simply put, these inequalities ensure that a consumer of type θ_i will prefer the contract $y_i = (x_i, t_i)$ over all others.

We now formulate an optimization problem to determine the optimal menu of contracts:

$$\left. \begin{array}{ll} \max_Y & f_c(Y) \\ \text{s.t.} & \hat{f}(x_i, \theta_i) - t_i \geq 0 \quad (\text{IR-i}) \\ & \hat{f}(x_i, \theta_i) - t_i \geq \hat{f}(x_j, \theta_i) - t_j \quad (\text{IC-i,j}) \\ & \forall j \neq i, \quad \forall i \in \{1, \dots, n\} \end{array} \right\} \quad (\text{P-1})$$

This maximization problem can be simplified by removing redundant constraints. First, all the individual rationality constraints, except (IR-1), are redundant. From constraints (IC-i,j) for all $i \in \{2, \dots, n\}$,

$$\hat{f}(x_i, \theta_i) - t_i \geq \hat{f}(x_{i-1}, \theta_i) - t_{i-1} \geq \hat{f}(x_{i-1}, \theta_{i-1}) - t_{i-1}. \quad (4.1.2)$$

Moreover, (IR-1) binds. Indeed, suppose not. Then

$$\hat{f}(x_2, \theta_2) - t_2 \geq \hat{f}(x_1, \theta_2) - t_1 \geq \hat{f}(x_1, \theta_1) - t_1 > 0 \quad (4.1.3)$$

where the middle inequality holds since $\hat{f}(x, \theta)$ is increasing in θ by Assumption 4.1.1. Hence, increasing t_1 and t_2 by a small $\varepsilon > 0$ would preserve (IR-1), not affect any of the incentive compatibility constraints, and raise profits thereby leading us to a contradiction.

We make the following assumption in order to reduce the $n(n-1)$ incentive compatibility constraints:

Assumption 4.1.3 (Spence-Mirrlees single-crossing condition). *For each $i \in \{1, \dots, n-1\}$, the marginal gain from increasing the privacy setting x is greater for type θ_{i+1} than type θ_i , i.e. $\hat{f}(x, \theta_{i+1}) - \hat{f}(x, \theta_i)$ is increasing in x .*

It is reasonable that the above assumption holds since consumers with a high-type value privacy much more than those with a low-type and hence, a small increase in their privacy setting value should provide them more utility as compared to a lower type consumer.

We define *local incentive compatibility constraints* as follows. For each $i \in \{2, \dots, n\}$, we define the *local downward incentive compatibility constraints*,

$$\hat{f}(x_i, \theta_i) - t_i \geq \hat{f}(x_{i-1}, \theta_i) - t_{i-1}, \quad (\text{LDIC-}i)$$

and for each $i \in \{1, \dots, n-1\}$, we define the *local upward incentive compatibility constraints*,

$$\hat{f}(x_i, \theta_i) - t_i \geq \hat{f}(x_{i+1}, \theta_i) - t_{i+1}. \quad (\text{LUIC-}i)$$

Proposition 4.1.1. *Given Assumption 4.1.3, monotonicity holds, i.e. $x_{i+1} \geq x_i$, and the local incentive compatibility constraints are necessary and sufficient for global incentive compatibility.*

Proof. Suppose the local constraints hold. Then

$$\hat{f}(x_i, \theta_i) - t_i - \hat{f}(x_{i-1}, \theta_i) + t_{i-1} \geq 0 \quad (4.1.4)$$

and

$$0 \geq \hat{f}(x_i, \theta_{i-1}) - t_i - \hat{f}(x_{i-1}, \theta_{i-1}) + t_{i-1} \quad (4.1.5)$$

so that

$$\hat{f}(x_i, \theta_i) - \hat{f}(x_i, \theta_{i-1}) \geq \hat{f}(x_{i-1}, \theta_i) - \hat{f}(x_{i-1}, \theta_{i-1}). \quad (4.1.6)$$

Thus, by Assumption 4.1.3, $x_i \geq x_{i-1}$ for each i (monotonicity).

Now, consider (LDIC- i) and (LDIC-($i-1$)):

$$\hat{f}(x_i, \theta_i) - \hat{f}(x_{i-1}, \theta_i) \geq t_i - t_{i-1} \quad (4.1.7)$$

$$\hat{f}(x_{i-1}, \theta_{i-1}) - \hat{f}(x_{i-2}, \theta_{i-1}) \geq t_{i-1} - t_{i-2}. \quad (4.1.8)$$

Summing the above inequalities gives us

$$\hat{f}(x_i, \theta_i) - \hat{f}(x_{i-1}, \theta_i) + \hat{f}(x_{i-1}, \theta_{i-1}) - \hat{f}(x_{i-2}, \theta_{i-1}) \geq t_i - t_{i-2}. \quad (4.1.9)$$

This inequality along with monotonicity and Assumption 4.1.3, i.e.

$$\hat{f}(x_{i-1}, \theta_i) - \hat{f}(x_{i-2}, \theta_i) \geq \hat{f}(x_{i-1}, \theta_{i-1}) - \hat{f}(x_{i-2}, \theta_{i-1}), \quad (4.1.10)$$

imply that

$$\hat{f}(x_i, \theta_i) - \hat{f}(x_{i-2}, \theta_i) \geq t_i - t_{i-2}. \quad (4.1.11)$$

This is exactly (IC- $i, i-2$). We have shown that (LDIC- i) and (LDIC-($i-1$)) imply (IC- $i, i-2$). In addition, we can show that (IC- $i, i-1$) and (LDIC-($i-2$)) imply (IC- $i, i-3$) and so on. Thus inductively, the local downward incentive compatibility constraints imply the incentive compatibility constraints for all $i \geq j$. Similarly, we can argue that the local upward incentive compatibility constraints imply the incentive compatibility constraints for all $i \leq j$. Finally, necessity is straightforward. \square

We have reduced the constraint set of (P-1) to the local downward incentive compatibility constraints, local upward incentive compatibility constraints, and the individual rationality constraint for type θ_1 , (IR-1). We make the following claim:

Proposition 4.1.2. *The local downward incentive compatibility constraints bind at optimum.*

Proof. Suppose that the local downward incentive compatibility constraints is not binding at optimum for some type θ_i , i.e. for some $\varepsilon > 0$,

$$\hat{f}(x_i, \theta_i) - t_i - (\hat{f}(x_{i-1}, \theta_i) - t_{i-1}) > \varepsilon. \quad (4.1.12)$$

Then, for all $j \geq i$, we can increase t_j by ε without affecting any of the incentive compatibility constraints but increasing the profit by $(1 - \sum_{k=1}^{i-1} p_k)\varepsilon$. This leads to a contradiction. \square

Since $x_i \geq x_{i-1}$ and the local downward incentive compatibility constraints are binding at optimum, the local upward incentive compatibility constraints are automatically satisfied, i.e.

$$\hat{f}(x_i, \theta_{i-1}) - t_i \leq \hat{f}(x_{i-1}, \theta_{i-1}) - t_{i-1}. \quad (4.1.13)$$

Hence, we have reduced the problem to the following:

$$\left. \begin{array}{ll} \max_Y & f_c(Y) \\ \text{s.t.} & \hat{f}(x_1, \theta_1) - t_1 = 0 \\ & \hat{f}(x_i, \theta_i) - t_i = \hat{f}(x_{i-1}, \theta_i) - t_{i-1}, \quad \forall i \in \{2, \dots, n\} \\ & x_i \geq x_j \text{ whenever } \theta_i \geq \theta_j \end{array} \right\} \begin{array}{l} \text{(IR-1)} \\ \text{(LDIC-i)} \end{array} \quad \text{(P-2)}$$

We refer to the optimal solution of (P-2) as the *second-best* solution and we use the notation $Y^{\text{sb}} = \{(\hat{x}_i^{\text{sb}}, \hat{t}_i^{\text{sb}})\}_{i=1}^n$. Before we proceed, we will define the *first-best* solution. If the power company knows the type θ that it faces, then it solves

$$\max_{(x,t)} \{t - g(x) \mid \hat{f}(x, \theta) - t \geq 0\} \quad (4.1.14)$$

The solution to the above problem is denoted by $(\hat{x}^{\text{fb}}(\theta), \hat{t}^{\text{fb}}(\theta))$ where

$$\hat{x}^{\text{fb}}(\theta) = \arg \max_x \{\hat{f}(x, \theta) - g(x)\} \quad (4.1.15)$$

and $\hat{t}^{\text{fb}}(\theta) = \hat{f}(\hat{x}^{\text{fb}}(\theta), \theta)$. We denote the collection of solutions for the different types $Y^{\text{fb}} = \{(\hat{x}_i^{\text{fb}}, \hat{t}_i^{\text{fb}})\}_{i=1}^n$.

Writing down the Krush–Khun–Tucker (KKT) conditions to (P-2), we can formulate a system of equations to determine the optimal (second-best) solution Y^{sb} . There are a number of properties that arise from general contracting problems of this kind that we outline here and the details of which can be found in [BD05, Chapter 2].

First, the type that values privacy the highest, θ_n , gets the socially optimal contract (efficient allocation); this can be easily seen from (4.1.14) with $\theta = \theta_n$ and by examining the

first-order KKT conditions of (P-2). On the other hand, all other types get an inefficient allocation, i.e. $\hat{x}_i^{\text{sb}} \leq \hat{x}_i^{\text{fb}}$ for $i \in \{1, \dots, n-1\}$. The type which values privacy the lowest, θ_1 , gets *zero-surplus*; this is easily seen from (IR-1), i.e. $\hat{f}(\hat{x}_1^{\text{sb}}, \theta_1) = \hat{t}_1^{\text{sb}}$. All other types θ_i , $i \in \{2, \dots, n\}$ enjoy some *information rent* meaning they pay less than is socially optimal, and hence, *free-ride* on the rest of society; This can be easily seen by examining the local downward incentive compatibility constraints and (IR-1). Indeed, define the following function for the *information rent* of type θ_i as a function of $\mathbf{x} = (x_1, \dots, x_n)$:

$$\hat{f}_{ir}(\mathbf{x}, \theta_i) = \sum_{j=1}^{i-1} \left(\hat{f}(x_j, \theta_{j+1}) - \hat{f}(x_j, \theta_j) \right). \quad (4.1.16)$$

Since information rent is positive by Assumption 4.1.3 and $\hat{f}(\hat{x}_i^{\text{sb}}, \theta_i) \leq \hat{f}(\hat{x}_i^{\text{fb}}, \theta_i) = \hat{t}_i^{\text{fb}}$ where \hat{t}_i^{fb} is the socially optimal price, we have that

$$\hat{t}_i^{\text{sb}} = \hat{f}(\hat{x}_i^{\text{sb}}, \theta_i) - \hat{f}_{ir}(\hat{\mathbf{x}}^{\text{sb}}, \theta_i) \leq \hat{t}_i^{\text{fb}} - \hat{f}_{ir}(\hat{\mathbf{x}}^{\text{sb}}, \theta_i). \quad (4.1.17)$$

An interesting phenomena emerges when there are more than two types called *bunching*. Depending on the distribution of types, two or more adjacent types get the same second-best contract. Separating two adjacent types may give too much information rent to all higher types thereby resulting in a solution with greater inefficiencies from the point of view of society.

Remark 4.1.2. *It is fairly straightforward to extend to the case where the power company faces multiple consumers simultaneously. Suppose there are m consumers where consumer j 's type is θ^j taking its value in the finite set $\Theta^j = \{\theta_1^j, \dots, \theta_n^j\}$. Again, we say that the power company faces type θ_i^j with probability p_i^j . Thus, (P-2) becomes the following problem:*

$$\left. \begin{array}{l} \max_{\{Y^j\}_{j=1}^m} \quad \sum_{j=1}^m \sum_{i=1}^n p_i^j (t_i^j - g(x_i^j)) \\ \text{s.t.} \quad \hat{f}(x_1^j, \theta_1^j) - t_1^j = 0 \quad (IR-1, j) \\ \hat{f}(x_i^j, \theta_i^j) - t_i^j = \hat{f}(x_{i-1}^j, \theta_i^j) - t_{i-1}^j, \quad \forall i \in \{2, \dots, n\} \quad (LDIC-i, j) \\ x_i^j \geq x_k^j \text{ whenever } \theta_i^j \geq \theta_k^j \\ \forall j \in \{1, \dots, m\} \end{array} \right\} \quad (P-3)$$

where $Y^j = \{(x_i^j, t_i^j)\}_{i=1}^n$. The problem decomposes into a single problem per consumer since an individual consumer's privacy valuation θ and their utility are not dependent on any other consumer. This is a result of the fact that in this particular framework we assume that the privacy metric does not experience any of the network effects of information. We leave exploring these network effects to future work.

We now introduce an example to help the reader contextualize some of the concepts introduced. In previous work, we characterized the efficiency-privacy tradeoff for a DLC problem of thermostatically controlled loads (TCLs) [Don+14]. We showed that the ℓ_1 -norm of the error of the DLC (measured in terms of the ℓ_1 distance between the actual power

consumed by the TCLs and the desired power consumption) increases as a function of the sampling period, i.e. distance between samples, where a larger sampling period corresponds to a higher privacy-setting. Empirically the relationship between sampling period and ℓ_1 -norm error is approximately quadratic.

Example 4.1 (Direct Load Control). *Define the unit cost for implementing a privacy setting x to be given by*

$$g(x) = \frac{1}{2}\zeta x^2 \quad (4.1.18)$$

where $0 < \zeta < \infty$ so that it is proportional to the error of the DLC scheme.

Consider that there are two types of consumer: $\theta \in \{\theta_L, \theta_H\}$ where $0 < \theta_L < \theta_H$. Let the consumers' utility be given by

$$\hat{f}(x, \theta) = x\theta \quad (4.1.19)$$

where $x \in [0, 1]$, i.e. their utility is proportional to both the type and the privacy setting.

Using (4.1.14), we determine the first-best contracts are given by

$$(\hat{x}_H^{fb}, \hat{x}_L^{fb}) = \left(\frac{\theta_H}{\zeta}, \frac{\theta_L}{\zeta} \right), \quad (\hat{t}_L^{fb}, \hat{t}_H^{fb}) = \left(\frac{\theta_H^2}{\zeta}, \frac{\theta_L^2}{\zeta} \right). \quad (4.1.20)$$

Suppose the power company has the prior

$$P(\theta = \theta_H) = p, \quad P(\theta = \theta_L) = 1 - p. \quad (4.1.21)$$

For the two-type case, (P-2) reduces to the following two optimization problems:

$$\left. \begin{array}{l} \max_{x_L} \{ \hat{f}(x_L, \theta_L) - (1-p)g(x_L) - p\hat{f}(x_L, \theta_H) \} \\ \max_{x_H} \{ \hat{f}(x_H, \theta_H) - g(x_H) \} \end{array} \right\} \quad (\text{P-2-redux})$$

The solutions to the above problems are

$$(\hat{x}_H^{sb}, \hat{x}_L^{sb}) = \left(\frac{\theta_H}{\zeta}, \frac{1}{\zeta} \left[\theta_L - (\theta_H - \theta_L) \frac{p}{(1-p)} \right]_+ \right) \quad (4.1.22)$$

where $[\cdot]_+ = \max\{\cdot, 0\}$. □

We remark that for $p \geq \frac{\theta_L}{\theta_H} = \hat{p}^*$, the low-type receives zero allocation $\hat{x}_L^{sb} = 0$. This is referred to as the *shutdown* solution in which case the power company no longer offers a non-trivial contract to the low-type. The significance being that the probability of facing a high-type is above some critical threshold beyond which it is no longer beneficial to the power company to design non-trivial contracts for the low-type. This phenomenon is a direct result of the information asymmetry at the core of the problem.

In the above example, we can see the core properties of contracts under asymmetric information. First, the low-type gets *zero surplus*; indeed, the optimal price for the low-type is

$$\hat{t}_L^{sb} = \hat{f}(\hat{x}_L^{sb}, \theta_L) = \frac{\theta_L}{\zeta} \left[\theta_L - (\theta_H - \theta_L) \frac{p}{(1-p)} \right]_+. \quad (4.1.23)$$

In addition, the high-type gets some positive information rent due to the *positive externality* coming from the mere existence of the low-type:

$$\hat{t}_H^{\text{sb}} = \frac{\theta_H^2}{\zeta} - \frac{(\theta_H - \theta_L)}{\zeta} \left[\theta_L - (\theta_H - \theta_L) \frac{p}{(1-p)} \right]_+ \quad (4.1.24)$$

$$= \hat{t}_H^{\text{fb}} - \hat{f}_{ir}(\hat{\mathbf{x}}^{\text{sb}}, \theta_H) \quad (4.1.25)$$

where $\hat{f}_{ir}(\hat{\mathbf{x}}^{\text{sb}}, \theta_H) \geq 0$. This is desirable from the point of view of the high-type since it allows them to *free-ride* on society. However, it is undesirable from the point of view of society who bears the burden of this free-riding. Finally, the high-type gets the socially optimal allocation—see (4.1.20) and (4.1.22). We will return to this DLC example in the sequel, by introducing risk-averse consumers.

4.2 Effects of Privacy Loss Risk on Contracts

We are interested in analyzing the effect of loss risk (due to privacy breaches) in contracts.

Let us consider that an energy consumer of type θ suffers privacy breaches of cost $\ell(\theta)$, with probability $1 - \eta(x)$ where the probability of a privacy breach can be derived from a variety of inference metrics [Don+14; San+13].

4.2.1 Inferential Privacy Metrics

To give a concrete example, we construct a privacy metric—referred to as *inferential privacy*—by considering a detection theory framework [Don+14]. Suppose the consumer has some state ξ that they want to keep private. We assume that $\xi \in E$ where E is some finite set. This state ξ influences their device usage patterns x , which, in turn, determine their total energy consumption y . This is modeled in the following hierarchical Bayesian framework:

$$\xi \sim p_\xi \quad (4.2.1)$$

$$u|\xi \sim p_{u|\xi}(\cdot|\xi) \quad (4.2.2)$$

$$y|u, \xi \sim p_{y|u}(\cdot|x) \quad (4.2.3)$$

where p_ξ is a multinomial distribution, $p_{u|\xi}$ is the density of the distribution of device use patterns for consumers with private state ξ and $p_{y|u}$ is the density of a probability distribution that models the devices inside the household. Further, we let $p_{y|\xi}(y|\xi) = \int p_{y|u} p_{u|\xi}(u|\xi) du$.

Assumption 4.2.1. *Our adversary is able to observe the AMI signal y , and has knowledge of $p_\xi, p_{u|\xi}, p_{y|u}$. Additionally, the adversary has an arbitrary amount of computational power.*

This adversary has access to the measured data signal, and also holds priors on the consumer's private information ξ . He also knows how different consumer types use devices,

$p_{u|\xi}$, and also has access to models of the device's power consumption $p_{y|u}$. Although this adversary has quite a bit of knowledge about the consumers, he does not hold arbitrary side information.

Our privacy metric is the probability of error if an adversary tries to infer the private variable ξ .

Definition 4.2.1 (Inferential Privacy [Don+14]). *Under the hierarchical Bayesian model outlined in (4.2.1)-(4.2.3), an AMI protocol is α -inferentially private if, for any estimator $\hat{\xi} : \mathcal{Y} \rightarrow E$, we have $P(\hat{\xi}(y) \neq \xi) \geq \alpha$. This estimator is based on information in $p_\xi, p_{u|\xi}$, and $p_{y|u}$.*

We can define the maximum a posteriori (MAP) estimator $\hat{\xi}_{MAP}$, which maximizes $P(\hat{\xi}(y) = \xi)$.

Proposition 4.2.1 ([Don+14]). *Under the hierarchical Bayesian model outlined in (4.2.1)-(4.2.3), $P(\hat{\xi}(y) = \xi)$ is maximized by*

$$\hat{\xi}_{MAP}(y) = \arg \max_{i \in E} (p_\xi(i) \cdot p_{y|\xi}(y|i)) \quad (4.2.4)$$

The optimality of the MAP estimator with respect to the prior p_ξ immediately leads to a guarantee of privacy.

Proposition 4.2.2 ([Don+14]). *Under the hierarchical Bayesian model outlined in (4.2.1)-(4.2.3), the AMI protocol is α -inferentially private, where*

$$\alpha = P(\hat{\xi}_{MAP}(y) \neq \xi). \quad (4.2.5)$$

Furthermore, the AMI protocol is not α' -inferentially private for any $\alpha' > \alpha$.

In general, the optimal estimator is difficult to compute in practice. We can provide approximations using tools from statistics.

Definition 4.2.2 (Total Variation Distance). *The total variation distance between two densities p and q on a measure space (X, \mathcal{A}, μ) is given by*

$$\|p - q\|_{TV} = \sum_{A \in \mathcal{A}} \left| \int_A (p(x) - q(x)) \mu(dx) \right| = \frac{1}{2} \int_X |p(x) - q(x)| \mu(dx). \quad (4.2.6)$$

Definition 4.2.3 (Kullback-Leibler Divergence). *The Kullback-Leibler (KL) divergence between two densities p and q on a measure space (X, \mathcal{A}, μ) is given by*

$$D_{kl}(p||q) = \int p(x) \log \frac{p(x)}{q(x)} \mu(dx). \quad (4.2.7)$$

Similarly, we will define the KL-divergence between two random variable X and Y to be the KL-divergence between their densities, and it will be denoted $D_{kl}(X||Y)$.

We can use Le Cam’s method [LeC73] to provide the first approximation to the optimal privacy bound defined in Proposition 4.2.2. First, we state Le Cam’s lemma.

Lemma 4.2.1 ([LeC73][Tsy09, Lemma 2.3]). *Assume the hierarchical Bayesian model outlined in (4.2.1)-(4.2.3). Then for any estimator $\hat{\xi} : \mathcal{Y} \rightarrow E$ and any distinct $i, j \in E$, we have*

$$P(\hat{\xi}(y) \neq \xi | \xi = i) + P(\hat{\xi}(y) \neq \xi | \xi = j) \geq 1 - \|p_{y|\xi}(\cdot|i) - p_{y|\xi}(\cdot|j)\|_{TV}. \quad (4.2.8)$$

Our approximation directly follows from Le Cam’s lemma.

Proposition 4.2.3 (Le Cam Approximation [Don+14]). *Suppose the assumptions of Lemma 4.2.1 hold. Then $P(\hat{\xi}(y) \neq \xi)$ is bounded below:*

$$P(\hat{\xi}(y) \neq \xi) \geq \min(p_\xi(i), x_\xi(j))(1 - \|p_{y|\xi}(\cdot|i) - p_{y|\xi}(\cdot|j)\|_{TV}). \quad (4.2.9)$$

Thus the AMI protocol is α -inferentially private where

$$\alpha = \max_{i \neq j} [\min(p_\xi(i), p_\xi(j))(1 - \|p_{y|\xi}(\cdot|i) - p_{y|\xi}(\cdot|j)\|_{TV})]. \quad (4.2.10)$$

We remark that it will suffice to find an over approximation of the total variation distance, i.e. consider Pisker’s inequality [Tsy09, Lemma 2.5]

$$\|p - q\|_{TV} \leq \sqrt{\frac{1}{2} D_{kl}(p||q)}. \quad (4.2.11)$$

An alternative to Le Cam’s method is Fano’s inequality [Tsy09, Lemma 2.10].

Proposition 4.2.4 (Fano Approximation [Don+14]). *Provided the hierarchical Bayesian model outlined in (4.2.1)-(4.2.3), for any estimator $\hat{\xi} : \mathcal{Y} \rightarrow \Xi$, the probability of error $P(\hat{\xi}(y) \neq \xi)$ is bounded below:*

$$P(\hat{\xi}(y) \neq \xi) \geq \frac{1}{\log(r-1)} \left(\log r - \frac{1}{r^2} \sum_{i,j} D_{kl}(p_{y|\xi}(\cdot|i)||p_{y|\xi}(\cdot|j)) - \log 2 \right). \quad (4.2.12)$$

Hence, the AMI protocol is α -inferentially private where α is given by the right-hand side of (4.2.12).

We have shown in [Don+14] that income level can be inferred from energy consumption data. Further, we characterize in an example informed by real data how inferential privacy varies as a function of sample time.

In [San+13], an information theoretic approach is taken in order to create a metric for privacy and they investigate how the utility of consumption data decreases as privacy increases. On the other hand, differential privacy—originally introduced by Dwork in [Dwo11]—has been applied as a metric of privacy in the smart grid context where aggregation of consumption data across users is performed (see, e.g., [BM14; JB14]).

4.2.2 Contracts with Risk-Averse Consumers

Returning to the problem at hand, the consumer's expected utility is given by

$$f(x, \theta) = \hat{f}(x, \theta) - (1 - \eta(x))\ell(\theta). \quad (4.2.13)$$

The characteristics of the privacy breach are summarized in the following assumption.

Assumption 4.2.2. $\eta : \mathbb{R} \rightarrow [0, 1]$ (probability of avoiding a privacy breach) is strictly increasing with respect to the privacy setting x . The perceived loss $\ell : \Theta \rightarrow \mathbb{R}_{\geq 0}$ due to a privacy breach is increasing with respect to the type of each consumer.

Intuitively, the higher the privacy setting, the less likely a privacy breach will occur. Furthermore, a consumer with a high privacy valuation might experience a greater loss as compared to a consumer with a low privacy valuation.

The individual rationality constraint for the case where consumers are exposed to privacy risks is expressed as

$$\hat{f}(x, \theta) - t \geq (1 - \eta(x))\ell(\theta). \quad (4.2.14)$$

Recall that the optimal contract without risk for a consumer with the lowest valuation of privacy, $(\hat{x}_1^{\text{sb}}, \hat{t}_1^{\text{sb}})$, satisfies (IR-1) with strict equality, i.e. $\hat{f}(\hat{x}_1^{\text{sb}}, \theta_1) = \hat{t}_1^{\text{sb}}$. Thus, the optimal contract of the previous section violates (4.2.14) unless either $\ell(\theta_1) = 0$ or $\eta(\hat{x}_1^{\text{sb}}) = 1$. Consequently, consumers with the lowest privacy preferences θ_1 might get more profit by *opting out*.

On the other hand, when there are privacy risks, the local upward incentive compatibility constraint (LUIC-i) is expressed as

$$\hat{f}(x_i, \theta_i) - t_i \geq \hat{f}(x_{i+1}, \theta_i) - t_{i+1} + (\eta(x_{i+1}) - \eta(x_i))\ell(\theta_i). \quad (4.2.15)$$

Since $\hat{x}_i^{\text{sb}} \leq \hat{x}_{i+1}^{\text{sb}}$, $\eta(\hat{x}_{i+1}^{\text{sb}}) - \eta(\hat{x}_i^{\text{sb}}) \geq 0$. Hence, the inequality (4.2.15) might not be satisfied by the optimal contract that does not consider risk. Thus consumers with the privacy preferences θ_i for $i \in \{1, \dots, n-1\}$ may choose a contract designed for a higher type. Consequently, the power company will need to decrease the cost t_i or increase the privacy setting x_i in order to promote participation thereby decreasing the benefit and fees collected and thus, the social welfare. Hence, there is an incentive for the power company to purchase insurance or invest in security. Furthermore, security is inherently tied to privacy in this scenario since the security measures taken will impact the ability of an adversary to make inferences and hence the privacy metric.

4.2.3 Characterizing the Effects of Privacy Loss Risk

Suppose that f defined in (4.2.13) satisfies Assumption 4.1.1 and 4.1.3. Then the analysis in Section 4.1 holds when we replace \hat{f} with f .

From (4.2.13), we get the marginal utility with privacy loss risk:

$$\frac{\partial f}{\partial x}(x, \theta) = \frac{\partial \hat{f}}{\partial x}(x, \theta) + \frac{\partial \eta}{\partial x}(x) \ell(\theta). \quad (4.2.16)$$

Since f is strictly increasing by Assumption 4.1.1, we have

$$\frac{\partial f}{\partial x}(x, \theta) > 0. \quad (4.2.17)$$

Since the probability of a successful attack decreases with higher privacy settings, we have

$$\frac{\partial \eta}{\partial x}(x) > 0. \quad (4.2.18)$$

Hence, from (4.2.16) and the fact that $\ell(\theta) \geq 0$, we have that

$$\frac{\partial f}{\partial x}(x, \theta) \geq \frac{\partial \hat{f}}{\partial x}(x, \theta). \quad (4.2.19)$$

The following two propositions characterize the second-best contracts with and without privacy loss risk.

Proposition 4.2.5. *The privacy setting of all types $i \in \{1, \dots, n-1\}$ in contracts with and without privacy loss risk (x_i^{sb} and \hat{x}_i^{sb} resp.) satisfy the following. If $P_{i+1}\ell(\theta_{i+1}) < P_i\ell(\theta_i)$, then $x_i^{sb} \geq \hat{x}_i^{sb}$ where $P_i = \sum_{j=i}^n p_j$. Otherwise, $x_i^{sb} < \hat{x}_i^{sb}$.*

Proof. Substituting the constraints of (P-2) into the profit function $f_c(Y)$, we get that

$$f_c(Y) = p_1(\hat{f}(x_1, \theta_1) - g(x_1)) + \sum_{i=2}^n p_i \left(\hat{f}(x_i, \theta_i) - g(x_i) + \sum_{j=1}^{i-1} (\hat{f}(x_j, \theta_{j+1}) - \hat{f}(x_j, \theta_j)) \right). \quad (4.2.20)$$

For each $i \in \{1, \dots, n-1\}$, the following is the part of the profit function (4.2.20) that depends on x_i :

$$p_i(\hat{f}(x_i, \theta_i) - g(x_i)) - P_{i+1}(\hat{f}(x_i, \theta_{i+1}) - \hat{f}(x_i, \theta_i)). \quad (4.2.21)$$

Taking the derivative with respect to x_i , we get the first-order conditions in the case without risk,

$$p_i \left(\frac{\partial \hat{f}}{\partial x_i}(\hat{x}_i^{sb}, \theta_i) - \frac{\partial g}{\partial x_i}(\hat{x}_i^{sb}) \right) - P_{i+1} \left(\frac{\partial \hat{f}}{\partial x_i}(\hat{x}_i^{sb}, \theta_{i+1}) - \frac{\partial \hat{f}}{\partial x_i}(\hat{x}_i^{sb}, \theta_i) \right) = 0. \quad (4.2.22)$$

Replacing \hat{f} with f in (4.2.21) and taking the derivative with respect to x_i , we get the first-order conditions in the case with risk,

$$\begin{aligned} & p_i \left(\frac{\partial f}{\partial x_i}(x_i^{sb}, \theta_i) - \frac{\partial g}{\partial x_i}(x_i^{sb}) \right) - P_{i+1} \left(\frac{\partial f}{\partial x_i}(x_i^{sb}, \theta_{i+1}) - \frac{\partial f}{\partial x_i}(x_i^{sb}, \theta_i) \right) \\ & + \frac{\partial \eta}{\partial x_i}(x_i^{sb}) \left(p_i \ell(\theta_i) - P_{i+1} (\ell(\theta_{i+1}) - \ell(\theta_i)) \right) = 0. \end{aligned} \quad (4.2.23)$$

We now consider three cases. First, if $p_i \ell(\theta_i) - P_{i+1}(\ell(\theta_{i+1}) - \ell(\theta_i)) = 0$, then $\hat{x}_i^{\text{sb}} = x_i^{\text{sb}}$. Suppose now that $p_i \ell(\theta_i) - P_{i+1}(\ell(\theta_{i+1}) - \ell(\theta_i)) > 0$. This condition is equivalent to $P_{i+1} \ell(\theta_{i+1}) < P_i \ell(\theta_i)$. Then, we have that

$$\begin{aligned} & p_i \left(\frac{\partial f}{\partial x_i}(x_i, \theta_i) - \frac{\partial g}{\partial x_i}(x_i) \right) - P_{i+1} \left(\frac{\partial f}{\partial x_i}(x_i, \theta_{i+1}) - \frac{\partial f}{\partial x_i}(x_i, \theta_i) \right) > \\ & p_i \left(\frac{\partial \hat{f}}{\partial x_i}(x_i, \theta_i) - \frac{\partial g}{\partial x_i}(x_i) \right) - P_{i+1} \left(\frac{\partial \hat{f}}{\partial x_i}(x_i, \theta_{i+1}) - \frac{\partial \hat{f}}{\partial x_i}(x_i, \theta_i) \right) \end{aligned} \quad (4.2.24)$$

so that the first-order conditions with risk (4.2.23) imply

$$0 > p_i \left(\frac{\partial \hat{f}}{\partial x_i}(x_i^{\text{sb}}, \theta_i) - \frac{\partial g}{\partial x_i}(x_i^{\text{sb}}) \right) - P_{i+1} \left(\frac{\partial \hat{f}}{\partial x_i}(x_i^{\text{sb}}, \theta_{i+1}) - \frac{\partial \hat{f}}{\partial x_i}(x_i^{\text{sb}}, \theta_i) \right). \quad (4.2.25)$$

Hence, by (4.2.22) and the fact that the profit is concave, we have $x_i^{\text{sb}} > \hat{x}_i^{\text{sb}}$. On the other hand, if $p_i \ell(\theta_i) - P_{i+1}(\ell(\theta_{i+1}) - \ell(\theta_i)) < 0$, then the inequality in (4.2.24) is reversed. Thus, by (4.2.22) and (4.2.23), $x_i^{\text{sb}} < \hat{x}_i^{\text{sb}}$. \square

Proposition 4.2.6. *The privacy setting of an agent with type θ_n is higher with privacy loss risk, that is, $x_n^{\text{sb}} \geq \hat{x}_n^{\text{sb}}$.*

Proof. The piece of the profit that depends on x_n is $p_n(\hat{f}(x_n, \theta_n) - g(x_n))$ so that the first-order condition for x_n in the case without risk is given by

$$\frac{\partial \hat{f}}{\partial x_n}(\hat{x}_n^{\text{sb}}, \theta_n) - \frac{\partial g}{\partial x_n}(\hat{x}_n^{\text{sb}}) = 0. \quad (4.2.26)$$

Similarly, by replacing \hat{f} in (4.2.20) with f , the first-order condition for the case with risk is given by

$$\frac{\partial f}{\partial x_n}(x_n^{\text{sb}}, \theta_n) - \frac{\partial g}{\partial x_n}(x_n^{\text{sb}}) = 0. \quad (4.2.27)$$

Thus (4.2.19) implies that

$$0 \geq \frac{\partial \hat{f}}{\partial x_n}(x_n^{\text{sb}}, \theta_n) - \frac{\partial g}{\partial x_n}(x_n^{\text{sb}}). \quad (4.2.28)$$

Since $f(x, \theta_n) - g(x)$ is a concave function, its derivative with respect to x is a decreasing function of x . Hence, the optimal privacy setting without risk \hat{x}_n^{sb} , which satisfies (4.2.26), must be smaller than the privacy setting with risk, i.e. $x_n^{\text{sb}} \geq \hat{x}_n^{\text{sb}}$. This result is independent of the population distribution and the other types θ_i , $i \in \{1, \dots, n-1\}$. \square

Recall (4.1.16); in a similar way, we use the notation $f_{ir}(\mathbf{x}, \theta_i)$ to denote the information rent using the utility functions with risk f . The following result states that the optimal privacy setting for all consumer types excluding those with the highest valuation of privacy is decreasing with respect to the prior on types.

Proposition 4.2.7. *For each $i \in \{1, \dots, n-1\}$, suppose p_i is decreased by an amount $\varepsilon > 0$ and this probability mass is redistributed in any way to the types θ_j , $j \in \{i+1, \dots, n\}$, then x_i^{sb} and the information rent $f_{ir}(\mathbf{x}^{\text{sb}}, \theta_i)$ both decrease.*

Proof. Fix i in $\{1, \dots, n-1\}$. Consider two priors over the type space Θ : $\mathbf{p} = \{p_1, \dots, p_n\}$ and $\mathbf{p}' = \{p'_1, \dots, p'_n\}$ where $p'_i = p_i - \varepsilon$ and the probability mass $\varepsilon > 0$ is redistributed amongst p'_j for $j \in \{i+1, \dots, n\}$. We use the notation $x_i^{\text{sb}}(\mathbf{p})$ to denote the optimal privacy setting as a function of \mathbf{p} . The piece of the profit function dependent on x_i is given by

$$f_{p,i}(x_i, \mathbf{p}') = p'_i(f(x_i, \theta_i) - g(x_i)) - P'_{i+1}(f(x_i, \theta_{i+1}) - f(x_i, \theta_i))$$

so that the first-order conditions are

$$\begin{aligned} \frac{\partial f_{p,i}}{\partial x_i}(x_i, \mathbf{p}') &= p_i \left(\frac{\partial f}{\partial x_i}(x_i, \theta_i) - \frac{\partial g}{\partial x_i}(x_i) \right) - P_{i+1} \left(\frac{\partial f}{\partial x_i}(x_i, \theta_{i+1}) - \frac{\partial f}{\partial x_i}(x_i, \theta_i) \right) \\ &\quad - \varepsilon \left(\frac{\partial f}{\partial x_i}(x_i, \theta_i) - \frac{\partial g}{\partial x_i}(x_i) + \frac{\partial f}{\partial x_i}(x_i, \theta_{i+1}) - \frac{\partial f}{\partial x_i}(x_i, \theta_i) \right). \end{aligned} \quad (4.2.29)$$

Hence,

$$\frac{\partial f_{p,i}}{\partial x_i}(x_i^{\text{sb}}(\mathbf{p}), \mathbf{p}') = -\varepsilon \left(\frac{\partial f}{\partial x_i}(x_i^{\text{sb}}(\mathbf{p}), \theta_{i+1}) - \frac{\partial g}{\partial x_i}(x_i^{\text{sb}}(\mathbf{p})) \right). \quad (4.2.30)$$

Since $x_{i+1}^{\text{sb}}(\mathbf{p}) \geq x_i^{\text{sb}}(\mathbf{p})$ and $f(x, \theta_{i+1}) - g(x)$ is concave,

$$\frac{\partial f}{\partial x_i}(x_i^{\text{sb}}(\mathbf{p}), \theta_{i+1}) - \frac{\partial g}{\partial x_i}(x_i^{\text{sb}}(\mathbf{p})) > \frac{\partial f}{\partial x_i}(x_{i+1}^{\text{sb}}(\mathbf{p}), \theta_{i+1}) - \frac{\partial g}{\partial x_i}(x_{i+1}^{\text{sb}}(\mathbf{p})). \quad (4.2.31)$$

Further, by Assumption 4.1.3 and the first-order conditions for $x_{i+1}^{\text{sb}}(\mathbf{p})$, the right-hand side of the above inequality is greater than zero. Hence,

$$\frac{\partial f_{p,i}}{\partial x_i}(x_i^{\text{sb}}(\mathbf{p}), \mathbf{p}') < 0 \quad (4.2.32)$$

thereby indicating that $x_i^{\text{sb}}(\mathbf{p}) > x_i^{\text{sb}}(\mathbf{p}')$. Note that this result holds for the contracts without risk as well, $\hat{x}_i^{\text{sb}}(\mathbf{p}) > \hat{x}_i^{\text{sb}}(\mathbf{p}')$, by replacing f with \hat{f} in the above argument. The fact that the information rent decreases is straightforward. Indeed, Assumption 4.1.3 implies if $\hat{x}_i^{\text{sb}}(\mathbf{p}) > \hat{x}_i^{\text{sb}}(\mathbf{p}')$, then $f_{ir}((\hat{\mathbf{x}}^{\text{sb}})', \theta_i) \geq f_{ir}(\hat{\mathbf{x}}^{\text{sb}}, \theta_i)$. \square

As a consequence of the above proposition, t_i^{sb} is decreasing with respect to the prior. Furthermore, x_n^{sb} does not depend on the prior since the highest type always gets an efficient allocation. This applies regardless of the risk of privacy loss.

Results in Propositions 4.2.5 and 4.2.6 let us determine the impact of risk on the price t paid by the lowest and highest types.

Proposition 4.2.8. *The price for consumers with privacy valuation θ_n satisfies $t_n^{sb} > \hat{t}_n^{sb} - (1 - \eta(\hat{x}_{n-1}^{sb}))\ell(\theta_{n-1})$, if $p_n\ell(\theta_n) > (p_{n-1} + p_n)\ell(\theta_{n-1})$. The price for consumers with privacy valuation θ_1 satisfies*

$$\begin{cases} t_1^{sb} \geq \hat{t}_1^{sb} - (1 - \eta(\hat{x}_1^{sb}))\ell(\theta_1), & \text{if } P_2 \leq \frac{\ell(\theta_1)}{\ell(\theta_2)} \\ t_1^{sb} < \hat{t}_1^{sb} - (1 - \eta(\hat{x}_1^{sb}))\ell(\theta_1), & \text{otherwise} \end{cases} \quad (4.2.33)$$

where $P_2 = \sum_{j=2}^n p_j$.

Proof. Suppose $p_n\ell(\theta_n) > (p_{n-1} + p_n)\ell(\theta_{n-1})$. Then, from Proposition 4.2.5 and 4.2.6, $x_{n-1}^{sb} < \hat{x}_{n-1}^{sb}$ and $x_n^{sb} \geq \hat{x}_n^{sb}$. Hence, by Assumption 4.1.1 and Assumption 4.1.3, we have

$$\begin{aligned} t_n^{sb} &= f(x_n^{sb}, \theta_n) - (f(x_{n-1}^{sb}, \theta_n) - f(x_{n-1}^{sb}, \theta_{n-1})) \\ &\geq f(\hat{x}_n^{sb}, \theta_n) - (f(\hat{x}_{n-1}^{sb}, \theta_n) - f(\hat{x}_{n-1}^{sb}, \theta_{n-1})) \\ &\geq \hat{t}_n^{sb} + (\eta(\hat{x}_n^{sb}) - \eta(\hat{x}_{n-1}^{sb}))\ell(\theta_n) - (1 - \eta(\hat{x}_{n-1}^{sb}))\ell(\theta_{n-1}) \\ &\geq \hat{t}_n^{sb} - (1 - \eta(\hat{x}_{n-1}^{sb}))\ell(\theta_{n-1}) \end{aligned}$$

where the last inequality holds since the probability of a privacy breach $1 - \eta(x)$ is decreasing in x .

Now, suppose $\ell(\theta_2)P_2 \leq \ell(\theta_1)$ and note $P_1 = 1$. Then by Proposition 4.2.5, $\hat{x}_1^{sb} \leq x_1^{sb}$. Since $t_1^{sb} = f(x_1^{sb}, \theta_1)$, we have that

$$t_1^{sb} \geq f(\hat{x}_1^{sb}, \theta_1) \geq \hat{t}_1^{sb} - (1 - \eta(\hat{x}_1^{sb}))\ell(\theta_1). \quad (4.2.34)$$

On the other hand, if $\ell(\theta_2)P_2 > \ell(\theta_1)$, then $\hat{x}_1^{sb} > x_1^{sb}$ so that

$$t_1^{sb} \leq \hat{t}_1^{sb} - (1 - \eta(\hat{x}_1^{sb}))\ell(\theta_1). \quad (4.2.35)$$

□

Remark 4.2.1. *In general, for contract design problems, it is difficult to make any qualitative statements about price [BD05]. The above proposition gives conditions—as a function of the prior on types and the value of loss—such that the highest and lowest types pay more under the contracts with risk than they pay under the contract without risk less the expected loss. Further, it shows the alternate case for the lowest type; however, for type θ_n the reverse inequality cannot be achieved since $\hat{x}_n^{sb} \leq x_n^{sb}$ irrespective of the prior.*

Proposition 4.2.9. *For each $i \in \{2, \dots, n\}$, if $P_{j+1}\ell(\theta_{j+1}) \geq P_j\ell(\theta_j)$ for all $j \in \{1, \dots, i-1\}$, then information rent under the contract with privacy loss risk is less than without, i.e.*

$$f_{ir}(\mathbf{x}^{sb}, \theta_i) \leq \hat{f}_{ir}(\hat{\mathbf{x}}^{sb}, \theta_i) \quad (4.2.36)$$

where $\mathbf{x}^{sb} = (x_1^{sb}, \dots, x_n^{sb})$ and similarly for $\hat{\mathbf{x}}^{sb}$. Further, if $P_{i+1}\ell(\theta_{i+1}) < P_i\ell(\theta_i)$, then the price a consumer of type θ_i pays under the contract with privacy loss risk is higher than without, i.e. $t_i^{sb} \geq \hat{t}_i^{sb}$.

Proof. Fix $i \in \{2, \dots, n\}$. Suppose for all $j \in \{1, \dots, i-1\}$, $P_{j+1}\ell(\theta_{j+1}) \geq P_j\ell(\theta_j)$ so that Proposition 4.2.5 gives us $\hat{x}_j^{\text{sb}} \geq x_j^{\text{sb}}$. Thus, by Assumption 4.1.3, we have that $f_{ir}(\mathbf{x}^{\text{sb}}, \theta_i) \leq f_{ir}(\hat{\mathbf{x}}^{\text{sb}}, \theta_i) \leq \hat{f}_{ir}(\hat{\mathbf{x}}^{\text{sb}}, \theta_i)$ since

$$f_{ir}(\hat{\mathbf{x}}^{\text{sb}}, \theta_i) = \hat{f}_{ir}(\hat{\mathbf{x}}^{\text{sb}}, \theta_i) + \sum_{j=1}^{i-1} (1 - \eta(\hat{x}_j^{\text{sb}}))(\ell(\theta_j) - \ell(\theta_{j+1})).$$

Now, in addition, suppose that $P_{i+1}\ell(\theta_{i+1}) < P_i\ell(\theta_i)$ so that (by Proposition 4.2.5) $x_i^{\text{sb}} > \hat{x}_i^{\text{sb}}$. Then we have the following:

$$t_i^{\text{sb}} = f(x_i^{\text{sb}}, \theta_i) - f_{ir}(\mathbf{x}^{\text{sb}}, \theta_i) \geq f(\hat{x}_i^{\text{sb}}, \theta_i) - \hat{f}_{ir}(\hat{\mathbf{x}}^{\text{sb}}, \theta_i) \geq \hat{t}_i^{\text{sb}}. \quad (4.2.37)$$

□

Information rent characterizes how much the power company has to give up in profits in order to persuade consumers to report truthfully. The above shows that it decreases as we introduce risk into the model when the likelihood of the power company facing higher types is greater than the ratio of losses of lower types compared to higher types.

Example 4.1 (Direct Load Control—Continued). *Recall the DLC example that was introduced in Section 4.1. In addition to studying the effects of subsampling on DLC performance, we show that the probability of detecting a device is consuming power decreases with respect to sampling period by using the inferential privacy metric defined in (4.2.5) (see [Don+14]). Here, we take a linear approximation of the error. In particular, a higher privacy setting is less likely to be successfully attacked; hence, for the sake of the example, we take $1 - \eta(x) = m(1 - x)$ where $m > 0$ is a constant.*

As in the previous section, we model the consumer's risk aversion using the utility function in (4.2.13). When we consider risk-averse consumers, the first-best solution is given by

$$(t_H^{\text{fb}}, x_H^{\text{fb}}) = \left(\frac{\theta_H^2 - m\theta_H\ell(\theta_H)}{\zeta}, \frac{\theta_H - m\ell(\theta_H)}{\zeta} \right) \quad (4.2.38)$$

and

$$(t_L^{\text{fb}}, x_L^{\text{fb}}) = \left(\frac{\theta_L^2 - m\theta_L\ell(\theta_L)}{\zeta}, \frac{\theta_L - m\ell(\theta_L)}{\zeta} \right). \quad (4.2.39)$$

The second-best optimal contracts are given by

$$(t_L^{\text{sb}}, x_L^{\text{sb}}) = \left(\frac{\theta_L}{\zeta} \left[\frac{(m\ell(\theta_L) - pm\ell(\theta_H) - p\theta_H + \theta_L)}{(1-p)} \right]_+, \frac{1}{\zeta} \left[\frac{(m\ell(\theta_L) - pm\ell(\theta_H) - p\theta_H + \theta_L)}{(1-p)} \right]_+ \right) \quad (4.2.40)$$

and

$$(t_H^{\text{sb}}, x_H^{\text{sb}}) = \left(t_L^{\text{sb}} + x_H^{\text{sb}}\theta_H - x_L^{\text{sb}}\theta_H + m(x_H^{\text{sb}} - x_L^{\text{sb}})\ell(\theta_H), \frac{1}{\zeta}(\theta_H + m\ell(\theta_H)) \right). \quad (4.2.41)$$




Figure 4.1: (a) Comparison between the first-best and second-best solutions as a function of p for the case with (gray) and without (black) risk. The general shape of the curves is the same for different values of m ; changing m from 0 to 1 has the effect of shifting p^* closer to the origin as well as causing x_H^{sb} to decrease. (b) Optimal prices as a function of p for both the case with risk (gray) and without (black) as well as the information rent (* with risk, * without risk) as a function of p .

In Figure 4.1a, we show that as the probability of the high-type being drawn from the population increases, the privacy setting for the low-type decreases away from the first-best and socially optimal solution x_L^{fb} (resp. \hat{x}_L^{fb}). This occurs until the critical point

$$p^* = \frac{\theta_L + m\ell(\theta_L)}{\theta_H + m\ell(\theta_H)}.$$

As we discussed before, this critical point determines when the shut-down solution occurs.

In Figure 4.1b, we show the optimal prices for the first- and second-best solutions in both the case with risk and without. We see that for $p \leq p^*$ (resp. $p \leq \hat{p}^*$) we have positive information rent for the high-type. Essentially, when the probability of the existence of a low-type is large relative to that of a high-type, there is a positive externality—positive from the perspective of the high-type—on the high-type thereby allowing them to free-ride on society. People who value high privacy need to be compensated more to participate. Further, the low-type continues to get zero surplus since the individual rationality constraint of the low-type is always binding.

In Figure 4.2a and 4.2b respectively, we show the power company's expected profit and the social welfare

$$W(p, t_L, x_L, t_H, x_H) = f_c(t_L, x_L, t_H, x_H) + p(f(x_H, \theta_H) - t_H) + (1 - p)(f(x_L, \theta_L) - t_L), \quad (4.2.42)$$

which is the sum of expected profit of the power company and the consumer. Note we replace f with \hat{f} in the above equation for social welfare when considering a consumer's utility function without privacy loss risk. In this case we define the social welfare by \hat{W} .

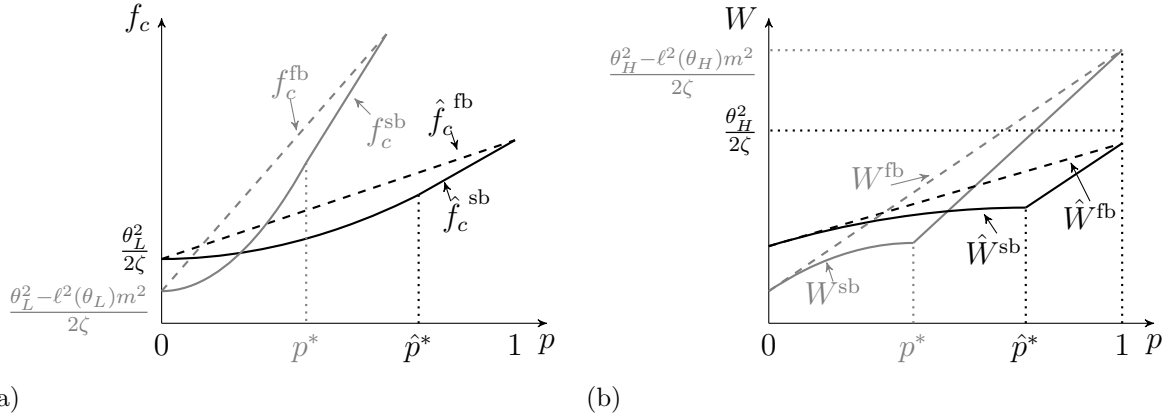


Figure 4.2: (a) Profit of the power company as a function of p for both the case with risk $f_c(p)$ (gray) and without $\hat{f}_c(p)$ (black). (b) Social Welfare as a function of p for both the case with risk W^{sb} (gray) and without risk \hat{W}^{sb} (black).

Notice the slope of the linear pieces of \hat{W}^{sb} and W^{sb} ; in particular, $\hat{W}^{sb}(p)$ for $p \geq \hat{p}^*$ is increasing at a slower rate than $W^{sb}(p)$ for $p \geq p^*$. Similarly, $f_c^{sb}(p)$ for $p \geq p^*$ increases at a faster rate than $\hat{f}_c^{sb}(p)$ for $p \geq \hat{p}^*$. This is in part due to the fact that $t_H^{sb}(p) - t_L^{sb}(p) > \hat{t}_H^{sb}(p) - \hat{t}_L^{sb}(p)$ as is shown in Figure 4.1b. Note that there are some values of p for which the power company's profit and the social welfare are lower with privacy loss risks thereby motivating compensation/insurance as a function of the population distribution or regulation in the form of subsidies or a privacy tax. \square

4.3 Insurance Contracts

In this section, we will design insurance contracts to be offered by a third-party insurance company to either the consumer or the power company which we will refer to more generally as the *agent*. In particular, we design insurance contracts that allow agents to purchase protection against attacks given they know the probability of a successful attack occurring. If the agent is the power company, then they may be purchasing insurance to account for compensation they offer to consumers in the event of a privacy breach or they may be purchasing insurance to account for privacy breaches the company itself experiences. On the other hand, if the agent is an electricity consumer, they may be purchasing insurance to protect against losses due to adversarial inferences made from their energy consumption data.

Using the theory of insurance contracts when there is asymmetric information and the probability that an adversary can gain access to the agent's private information, we analyze both the agent's choice on how much insurance to invest in as well as the insurer's decision about which contracts to offer to a population with both high-risk and low-risk agents. The

abstract analysis that follows is well known in the economics literature (see, e.g., [RS76; Jay78; MR78]). We borrow these tools to understand qualitatively how insurance investments will be made in a smart grid context and, as before, we take the novel view of privacy as a good.

4.3.1 Analysis of the Agent's Decision

We first analyze the decision the agent makes regarding selecting an amount of insurance given knowledge of the probability of a privacy breach $1 - \eta$. Let the agent's utility function be denoted by $f : \mathbb{R} \rightarrow \mathbb{R}$ and assume that f is increasing, twice differentiable and strictly concave. Let us suppose that the agent is *risk-averse* which means that the agent makes a decision under uncertainty and will try to minimize the impact of the uncertainty on its decision.

In addition, suppose the agent has initial wealth y , runs the risk of loss ℓ with probability $1 - \eta$. In the context of our problem, wealth represents *private information* that can be gained through analysis of energy consumption data or energy system data. A *loss* represents exposure of this private information.

The agent must decide how much insurance to buy. Let the cost of one unit of insurance be c and suppose that the insurer pays the agent β in the event that an adversary implements a successful attack resulting in an exposure of private information where β is the amount of insurance the agent agrees to buy. Then the agent wants to solve the following optimization problem:

$$\max_{\beta \geq 0} \{ \eta f(y - \beta c) + (1 - \eta) f(y + (1 - c)\beta - \ell) \}. \quad (4.3.1)$$

Let us consider the KKT necessary conditions for the optimization problem in (4.3.1). Suppose that β^* is a local optimum to the problem (4.3.1), then there exists a Lagrange multiplier λ such that

$$\left. \begin{aligned} 0 &= -\lambda - \eta c f'(y - \beta^* c) + (1 - \eta)(1 - c) f'(y + (1 - c)\beta^* - \ell) \\ 0 &= \lambda \beta^* \\ 0 &\geq \beta^* \\ 0 &\geq \lambda \end{aligned} \right\} \quad (4.3.2)$$

Combining the first and the last condition, we have

$$0 \geq -\eta c f'(y - \beta^* c) + (1 - \eta)(1 - c) f'(y + (1 - c)\beta^* - \ell) \quad (4.3.3)$$

We analyze the agent's decision by considering two cases and we present the results in the following propositions.

Proposition 4.3.1. *Suppose that the agent is offered privacy insurance at the rate $c = 1 - \eta$, i.e. at a rate equal to the probability of a successful attack. Then the agent will choose to purchase an amount of insurance equal to the loss, i.e. $\beta^* = \ell$.*

Proof. Since $c = 1 - \eta$, (4.3.3) reduces to

$$0 \geq \eta(1 - \eta)(f'(y + \eta\beta^* - \ell) - f'(y - \beta^*(1 - \eta))) \quad (4.3.4)$$

and since $(1 - \eta)\eta \geq 0$ this again reduces to

$$0 \geq f'(y + \eta\beta^* - \ell) - f'(y - \beta^*(1 - \eta)) \quad (4.3.5)$$

Recall that we assumed f to be a concave function and that a function is strictly concave if and only if its derivative f' is decreasing. Since $\ell > 0$,

$$f'(z) < f'(z - \ell). \quad (4.3.6)$$

The above inequality along with (4.3.5) implies that $\beta^* > 0$.

Now, we claim that $\beta^* = \ell$. Indeed, suppose that $0 < \beta^* < \ell$, then from (4.3.5) we have

$$f'(y + \eta\beta^* - \ell) \leq f'(y + \eta\beta^* - \beta^*). \quad (4.3.7)$$

This inequality violates (4.3.6). On the other hand, suppose that $0 \leq \ell \leq \beta^*$, then from (4.3.6) we have that

$$f'(y + \eta\beta^* - \beta^*) > f'(y + \eta\beta^* - \ell). \quad (4.3.8)$$

This violates the inequality (4.3.5). Hence, $\beta^* = \ell$ which is to say that the agent will purchase an amount of insurance equal to the loss of privacy she would endure under an attack. \square

Proposition 4.3.2. *Suppose that the agent is offered insurance at the rate $c > 1 - \eta$, i.e. at a rate higher than the probability of a successful attack. Then the agent will not purchase the full insurance, i.e. $\beta^* < \ell$.*

Proof. Suppose that the agent is offered privacy insurance at a rate $c > 1 - \eta$ and that the optimal choice for the agent is $\beta^* = \ell \geq 0$. Then, first-order optimality conditions imply that

$$-\eta f'(y - \ell c)c + (1 - \eta)f'(y - \ell c)(1 - c) = 0 \quad (4.3.9)$$

However, since $c > 1 - \eta$ and f is increasing, from (4.3.3) we have that

$$(-\eta c + (1 - \eta)(1 - c))f'(y - \ell c) < 0. \quad (4.3.10)$$

Thus the optimal amount of insurance β has to be less than the loss experienced, i.e. $\beta^* < \ell$. \square

4.3.2 Analysis of the Insurer's Decision

Let us now consider the design of privacy insurance contracts offered by a third-party insurance company. Insurance allows the consumer (when it is the agent in consideration) to *hedge their bet* against selecting a privacy-based service contract with a low-privacy setting which they may be incentivized by the power company to purchase. On the other hand, insurance allows the power company (when it is the agent in consideration) to hedge its bet against having a large number of consumers who have low valuations of privacy and therefore contribute to the level of riskiness of the power company. In addition, for the power company, there is a tradeoff between how much insurance they purchase versus how much security they invest in.

We assume the agent's utility function f is strictly concave, increasing and twice differentiable and for the sake of analysis we assume that $f(0) = 0$. We consider a scenario in which the insurer faces two types, i.e. either a high-risk agent θ_h or a low-risk agent θ_l . That is to say we are assuming that there is a portion of the population that is more likely to be attacked, i.e. the risky agents, possibly because they engage in high-risk behavior. For example, if the agent is the consumer, perhaps they selected a low-privacy setting on their smart meter, or if the agent is the power company, perhaps they have not invested in an appropriate level of security measures or they are not following the best practices recommendations, e.g. NIST-IR 7628 [Ell14]. In addition, there is a portion of the population that is less likely to be attacked, i.e. the low-risk agents. Just as in the previous sections, we can consider any finite number of types and make the same qualitative insights; however, we choose to present the results for only two types to make the presentation easier to follow.

The agent has an initial amount of wealth y and with probability $1 - \eta_j$ some of her private information is exposed resulting in a loss ℓ where $j = h, l$ indicates the agent's type. We assume that $1 - \eta_l < 1 - \eta_h$ which is to say that the likelihood a low-risk agent will experience a privacy loss is lower than the likelihood a high-risk agent will experience a loss, *ceteris paribus*. Furthermore, we will assume that the insurer has a prior over the distribution of types. In particular, we assume that the risky type θ_h occurs in the population with probability p and that $p > 0$.

Suppose we are given an insurance contract (α_a, α_n) where α_a is the compensation to the agent given that a successful attack occurred and α_n is the neutral case (no attack). Let X be a random variable representing the agent's wealth such that with probability $1 - \eta_i$ it takes value $y - \ell + \alpha_a$ and with probability η_i it takes value $y - \alpha_n$. Then, the agent's expected utility is given by

$$E[f(X)] = (1 - \eta_i)f(y - \ell + \alpha_a) + \eta_i f(y - \alpha_n). \quad (4.3.11)$$

Note that in the previous subsection we analyzed the agent's decision given a insurance contract of the form

$$(\alpha_a, \alpha_n) = ((1 - c)\beta, \beta c). \quad (4.3.12)$$

The insurer is a monopolist whose expected cost is given by

$$f_c(\alpha_a^h, \alpha_n^h, \alpha_a^l, \alpha_n^l) = p \left(-(1 - \eta_h)\alpha_a^h + \eta_h\alpha_n^h \right) + (1 - p) \left(-(1 - \eta_l)\alpha_a^l + \eta_l\alpha_n^l \right) \quad (4.3.13)$$

In the case of asymmetric information, i.e. the insurer does not know the agent's type, the optimization problem the insurer must solve is given by

$$\begin{aligned} \max_{\{\alpha_a^i, \alpha_n^i\}_{j=h,l}} \quad & f_c(\alpha_a^h, \alpha_n^h, \alpha_a^l, \alpha_n^l) & (\text{P-5}) \\ \text{s.t.} \quad & (1 - \eta_i)f(y - \ell + \alpha_a^i) + \eta_i f(y - \alpha_n^i) \geq (1 - \eta_i)f(y - \ell + \alpha_a^j) + \eta_i f(y - \alpha_n^j), & (\text{IC-i,j}) \\ & \text{for each } i, j \in \{h, l\}, i \neq j \\ & (1 - \eta_i)f(y - \ell + \alpha_a^i) + \eta_i f(y - \alpha_n^i) \geq (1 - \eta_i)f(y - \ell) + \eta_i f(y), & (\text{IR-i}) \\ & \text{for each } i \in \{h, l\} \end{aligned}$$

Constraints labeled (IC-i,j) are the incentive compatibility constraints and constraints (IR-i) are the individual rationality constraints. Both are similar to those presented in Section 4.1. Incentive compatibility ensures that the agent will report their type truthfully and the individual rationality constraint ensures that the agent will participate.

We can reduce the optimization problem (P-5) by reasoning about the constraint set defined by (IC-i,j) and (IR-i). Since $1 - \eta_l < 1 - \eta_h$, the incentive compatibility constraint for the risky type is active and the individual rationality constraint for the safe type is active, i.e. the constraint set for (P-5) reduces to the following two constraints:

$$(1 - \eta_h)f(y - \ell + \alpha_a^h) + \eta_h f(y - \alpha_n^h) = (1 - \eta_h)f(y - \ell + \alpha_a^l) + \eta_h f(y - \alpha_n^l) \quad (\text{IC-h})$$

and

$$(1 - \eta_l)f(y - \ell + \alpha_a^l) + \eta_l f(y - \alpha_n^l) = (1 - \eta_l)f(y - \ell) + \eta_l f(y). \quad (\text{IR-l})$$

Let us try to restate the problem in a way which allows us to characterize the solutions. Since we have assumed that f is strictly concave, increasing and twice differentiable, we can define F be its inverse, where $F' > 0$ and $F'' > 0$. Further, define

$$f_a^i = f(y - \ell + \alpha_a^i) \quad (4.3.14)$$

and

$$f_n^i = f(y - \alpha_n^i). \quad (4.3.15)$$

The transformed utility is given by

$$\begin{aligned} \tilde{f}_c(f_a^h, f_n^h, f_a^l, f_n^l) = & p \left(-\eta_h F(f_n^h) - (1 - \eta_h)F(f_a^h) + y - (1 - \eta_h)\ell \right) \\ & + (1 - p) \left(-\eta_l F(f_n^l) - (1 - \eta_l)F(f_a^l) + y - (1 - \eta_l)\ell \right). \end{aligned} \quad (4.3.16)$$

Then problem (P-5) transforms to the following optimization problem:

$$\begin{aligned} \max_{\{(f_a^i, f_n^i)\}_{i=h,l}} \quad & \tilde{f}_p(f_a^h, f_n^h, f_a^l, f_n^l) \\ \text{s.t.} \quad & (1 - \eta_h)f_a^h + \eta_h f_n^h = (1 - \eta_h)f_a^l + \eta_h f_n^l \\ & (1 - \eta_l)f_a^l + \eta_l f_n^l = (1 - \eta_l)f(y - \ell) + \eta_l f(y) \end{aligned} \quad (\text{P-6})$$

The Lagrangian of the optimization problem is

$$\begin{aligned} L(f_a^h, f_n^h, f_a^l, f_n^l, \lambda_1, \lambda_2) = & \tilde{f}_c(f_a^h, f_n^h, f_a^l, f_n^l) \\ & + \lambda_1((1 - \eta_h)f_a^h + \eta_h f_n^h - (1 - \eta_h)f_a^l - \eta_h f_n^l) \\ & + \lambda_2((1 - \eta_l)f_a^l + \eta_l f_n^l - (1 - \eta_l)f(y - \ell)). \end{aligned} \quad (4.3.17)$$

Proposition 4.3.3. *Given the probabilities $1 - \eta_j$, $j = h, l$ that an agent of type j will experience a privacy breach, if the insurer solves the optimization problem (P-6), then the high-risk agent will be fully insured and the low-risk agent will not be fully insured.*

Proof. We first show that the risky type will be fully insured. Taking the derivative of the Lagrangian with respect to f_a^h and f_n^h we have that

$$0 = -p(1 - \eta_h)F'(f_a^h) + \lambda_1(1 - \eta_h) \quad (4.3.18)$$

and

$$0 = -p\eta_h F'(f_n^h) + \lambda_1\eta_h. \quad (4.3.19)$$

Solving for λ_1 in the first equation and plugging it into the second, we get $f_a^h = f_n^h$ so that $\ell - \alpha_a^h = \alpha_n^h$, i.e. the amount the high-risk type pays for insurance is equal to the loss minus the compensation in the event of a privacy breach. Thus, the high-risk type will be fully insured.

Now, we show that the low-risk type will not be fully insured. Taking the derivative of the Lagrangian with respect to f_a^l and f_n^l , we get

$$0 = -(1 - \eta_l)(1 - p)F'(f_a^l) - \lambda_1(1 - \eta_h) + \lambda_2(1 - \eta_l) \quad (4.3.20)$$

and

$$0 = -(1 - p)\eta_l F'(f_n^l) - \lambda_1\eta_h + \lambda_2\eta_l. \quad (4.3.21)$$

From (4.3.18), we solved for $\lambda_1 = pF'(f_a^h)$. By plugging in λ_1 into (4.3.20), solving for λ_2 and plugging both λ_1 and λ_2 into (4.3.21), we have the following:

$$0 = F'(f_a^h)p \left(-\eta_h + \eta_l \frac{1 - \eta_h}{1 - \eta_l} \right) + \eta_l(1 - p)(F'(f_a^l) - F'(f_n^l)) \quad (4.3.22)$$

$$= F'(f_a^h)p \left(\frac{\eta_l - \eta_h}{1 - \eta_l} \right) + \eta_l(1 - p)(F'(f_a^l) - F'(f_n^l)). \quad (4.3.23)$$

Since $\eta_l > \eta_h$, we have that

$$0 > -\frac{F'(f_a^h)p}{\eta_l(1-p)} \left(\frac{\eta_l - \eta_h}{1 - \eta_l} \right) = F'(f_a^l) - F'(f_n^l). \quad (4.3.24)$$

Since F' is increasing, the above equation implies that

$$f_n^l - f_a^l > 0 \quad (4.3.25)$$

and hence, the low-risk type does not fully insure. \square

4.4 Discussion

Smart meters are increasingly becoming more and more capable of collecting data at high frequencies. As a result, we need to develop tools that allow consumers and power companies to benefit from these advances. Implementing privacy-aware data collection policies results in a reduction in the fidelity of the data and hence, a reduction in the efficiency of operations that depend on that data. This fundamental tradeoff provides an incentive for the power company to offer new service contracts.

We modeled electricity service as a product line differentiated according to privacy where consumers self-select the level of privacy that fits their needs and wallet. We derived privacy contracts with and without privacy loss risks, characterized the optimal contracts, and provided a comparative study. We showed that privacy loss risks can decrease the level service offered to each consumer type under certain conditions. Further, people who value a higher level of privacy *free-ride* on society and hence, need to be compensated accordingly to participate in the smart grid. We remark that the power company has an incentive to purchase insurance and invest in security in order to mitigate the effects of privacy loss risks. We leave questions regarding insurance versus security investment for future work.

Many open questions remain regarding the design of insurance contracts to be offered to the power company. We made some initial efforts to understand abstractly the level of insurance investment by power companies and consumers; however, there is much to be done in understanding how insurance investment varies as a function of the distribution of types and the selected privacy metric. Moreover, it is worthwhile to investigate the tradeoff experienced by the power company between investing in insurance and privacy-based service contracts.

Other researchers have used contract theory for demand-side management such as DLC and demand response programs. It would be interesting to consider the design of contracts with multiple goods—e.g. privacy setting, DLC options—in a multidimensional screening problem. In such a setting, we may also model the consumer's private information (type) as multidimensional vector thereby increasing the practical relevance of the model. Noting that Assumption 4.1.3 is often referred to as the *sorting condition*, we remark that one of the major difficulties in extending to the multidimensional case is the lack of being able to

sort or compare across the different goods and their qualities [Bas05]; however, a potential solution is to create a partial order of the multiple goods (benefits and privacy) available to consumers.

Another interesting direction for future research is in examining the network effects of information on privacy. In particular, we have implicitly assumed that the adversary uses only data from the individual that is being attacked. Perhaps information about other like consumers—for instance, based on their type—can be used to infer information about a particular consumer. These network effects can greatly impact the ability of an adversary to infer private information. Differential privacy [Dwo11] is a tool for providing privacy guarantees on the aggregate and perhaps could be used along side of inferential privacy to address the network effects of information. This is an idea worth exploring.

We have begun the conversation on how to combine inference bounds from statistics and economic mechanisms for managing the efficiency–privacy. A similar discussion could be started regarding security in the smart grid. One starting point might be in addressing financial attacks such as electricity theft. Electricity theft is a nice example of how regulations such as fines imposed by the regulating body (e.g. the utility commission) directly effect the consumer–power company interaction. Electricity theft can be modeled using *moral hazard* where we consider the hidden action to be the amount of electricity stolen. Incentives can be designed to thwart such theft along with using detection theory to classify *fraudulent* users and *normal* users. In this scenario, it would be interesting to consider the tradeoff between investing in security tools for theft detection and incentive mechanisms used to dissuade consumers from stealing.

In conclusion, there are multiple future research directions to be explored. Our model provides a general a mathematical framework for considering privacy as part of a service contract between a power company and its consumers. This line of research informs data collection schemes and privacy policy in the smart grid.

Chapter 5

Conclusion and Future Directions

This dissertation contributes foundational tools and techniques for characterizing the outcome of strategic interactions between non-cooperative agents engaging with a larger sociotechnical system which we refer to as a S-CPS (Chapter 2), designing mechanisms to shape the outcome to a more desirable or socially optimal one in an adaptive, online manner (Chapter 3), and designing of vulnerability-aware mechanisms (Chapter 4).

5.1 Emerging Tools for S-CPS

The running theme throughout the text is the problem of adverse selection as it arises in S-CPS, i.e. designing economic mechanisms when the underlying dynamics (generated by the decision making process of humans coupling to a CPS) of the system are unknown to the designer. In Chapter 2, we derived a characterization of local Nash equilibria that has a differentiable structure—therefore, amenable to computation—and is structurally stable and generic. This is done in support of creating a computationally tractable way of learning agents’ objective functions and designing incentives to change their behavior. In Chapter 3, we do just this; we create an online algorithm for utility learning and incentive design and provide convergence results by adapting tools from adaptive control and online convex optimization and learning.

Given the motivating problem of a planner taking advantage of new CPS technologies to integrate the consumer into the closed-loop behavior of the broader S-CPS, we consider the inherent efficiency–vulnerability tradeoff that arises in such systems. In particular, in Chapter 4, we focus on the demand-side of the smart grid by considering privacy breaches that arise due to streaming data from smart meters being used for operations such as DLC. Since privacy is inherently subjective and intrinsic to the individual, this problem of adverse selection arises naturally. In particular, how a person values her data and the information that can be interpreted from it, is unknown to the power provider; however, this data is also valuable to the provider. We present a solution that combines game theory (economics) and statistical learning (detection theory) for addressing fundamental problems that arise

in S-CPS. The tools we derived are not restricted to the demand-side of the smart grid or even the energy systems domain.

Owing to the greater use and integration of CPS technologies into management and operations, we are seeing novel vulnerabilities emerge and these vulnerabilities are tightly coupled with the socioeconomics of the system. Hence, there is a need for new tools and techniques that combine knowledge from game theory and statistical learning that can account for the socioeconomics and the variability and uncertainty these CPS technologies introduce.

In this regard, we believe that the contributions of this dissertation are the first steps towards an emerging systems theory for S-CPS. They present techniques that can be used to characterize and analyze the outcomes of strategic interactions and can be leveraged in the design of mechanisms for closing-the-loop around decision making agents—in particular, humans. These methods can be used to not only consider the design of closed-loop systems in isolation, but also understand, model, and design the interaction between these tightly coupled decision-making loops. By providing classical tools from economics and emerging tools in statistics, the developed techniques offer opportunities to not only obtain qualitative insights that can support policy development but quantitative metrics for evaluating quality of service and level of vulnerability protection a particular S-CPS offers. Furthermore, they expose a number of new directions for future research, some that can be tackled in the short term and many more for the horizon.

5.2 Future Plans and Frontiers

Critical infrastructures are quickly becoming large-scale, networked S-CPS with many sensors, actuators and decision makers. The complexity of these systems along with socioeconomic considerations has led to the need for new scalable and tractable methods that can address societal-level problems. While in this thesis we provided solutions to some fundamental problems that arise in S-CPS, there is much to be done in terms of extending these concepts to the general S-CPS framework as presented in Chapter 1.

5.2.1 Merging Game-Theory and Statistical Learning

There is a need for game-theoretic tools that incorporate the use of data-driven analytics for sustainable operation of S-CPS. The development of scalable, tractable algorithms that apply to problems with many competitive agents and streaming temporal data sources is necessary for integrating ubiquitous data sources into the sustainable operation of S-CPS.

There is a dichotomy between the data-driven approaches in which scalable tools exist but typically only correlation can be inferred and the model-based paradigm in which causal relationships are observed yet scale seems an insurmountable problem. This issue is really at the heart of the need for new tools at the intersection of game theory and statistical learning. In particular, game theory captures complex socioeconomic interactions while

statistical learning aids in reducing dimensionality and provides tools for analysis of large-scale, time-series data. Through a synergistic joining of the two, data-driven game-theoretic models can be leveraged in the design of economic/physical controls for sustainability and resilience.

Competitive environments in critical infrastructures emerge in many forms such as leader-follower (e.g. physical therapist incentivizing consistent exercise routines), simultaneous play (e.g. populations of drivers selecting routes), as well as dynamic (e.g. bidding in electricity markets) games. Further research is needed to advance the frontiers of game theory by leveraging tools from dynamical systems theory to find generic and structurally stable representations of equilibria in competitive environments arising naturally in S-CPS.

Due to the importance of agent preferences in the design of incentives and contracts which are efficient, stable, and resilient, there is a need to understand the type space of agents in and across various S-CPS. To this end, statistical learning provides tools to derive low-dimensional models that integrate information in high-dimensional time-series data with the effect of reducing complexity. On the other hand, dynamical systems theory provides tools for determining key parameters via bifurcation analysis as well as assessing stability and efficiency in this low-dimensional space.

Another interesting direction for future research is the construction a modular tool set that allows subsystems (local interactions perhaps) to be modeled and then integrated into the analysis of the overall S-CPS. Along this line of thinking, recall that in Chapter 1 microscopic versus macroscopic decision-making was discussed. As the number of interacting parties starts to increase, new behaviors arise that were not readily observable at the microscopic—or individual human-CPS coupling—level. Techniques for micro-to-macro and macro-to-micro decision-making need to be developed along with improved methods of managing large amounts of data in this context. Further, empirical validation of models of complex interactions that are fundamental to large-scale infrastructure is necessary.

5.2.2 Vulnerability-Aware Incentives and the Emerging Data Market

The modernization of critical infrastructures into S-CPS leads to new opportunities and new vulnerabilities. A sustainable society requires secure and privacy-aware data collection. There have been a number of reports and best practice recommendations released in the last few years including a report from the National Institute of Standards and Technology providing qualitative guidelines for data collection policies [Eli14] as well as the reports from the President's Council of Advisors on Science and Technology on cyber-security [HL13] and big data and privacy [HL14]. There is a need to complement such policies via quantifying both privacy and security in S-CPS. Moreover, understanding the efficiency-vulnerability and the development metrics for privacy and security risk are imperative for the development of new service models, e.g. controls and economic mechanisms, and policy.

We are already seeing the emergence of new service models, e.g. contracts and other

economic mechanisms, and controls that support operation and management of S-CPS. An example in point is the new pay-as-you-go auto insurance model that requires the installation of a tracking device that broadcasts information—here we see that technology allows for an adaptive service model yet increases the opportunity for a privacy breach. Expanding this new set of service models by creating service models that factor in awareness of vulnerabilities will not only enhance efficiency and sustainability, but also improve resilience.

A fundamental aspect of emerging S-CPS is the web of networked control systems, consumers and providers, third-party solutions providers such as data aggregators, and the information flows between them. This exchange and trade of data between these actors and its value in the context of S-CPS is the *data market*. To understand this data market, traditional methods of modeling simply the interaction between the user population, e.g. energy consumers, and the provider of resources, e.g. power company, are inadequate; the motivations and information sets of third-party solution providers, regulators, and policy makers need to be modeled as well.

Furthermore, the incentive structure of S-CPS is collectively the motivations of actors, information asymmetries, and economic mechanisms. For example, in the current energy system, neither the utility companies nor third-party solution providers are adequately incentivized to invest in cyber-security of the composite system (referred to as underinvestment in the common good). There is a dire need for tools to design contracts and other mechanisms to help shape the market such that participating parties are motivated to improve the overall operation and resilience of S-CPS.

The contributions of this thesis are a first step towards understanding the data market and the incentive structure of S-CPS as a set of interactions between multiple stakeholders along with feedback to regulatory and other entities as well as new business opportunities. There are a number of potential game changers that could arise in this line of research. Due to the emerging data market, we see the increased awareness of the value of S-CPS data—hence, the emergence of companies that are capitalizing on access to this data—but also an increase in risks associated with access to streaming data. On the consumer side, creating opt-in security and privacy aware mechanisms that incorporate individual preferences will allow consumers to be adequately compensated for the use of their data. More broadly, realignment of incentives will lead to increased awareness of and investment in security and privacy mechanisms thereby increasing the overall resilience of S-CPS. As more actors come into play in this data market, the influence of third-party companies, such as data aggregators, will become much more substantial. Such an evolution has already begun in the transportation and energy CPS infrastructures and has tremendous potential to revolutionize the both the energy and transportation S-CPS data market and industry landscape.

5.2.3 Smart Urban Spaces

We are going through a period of greater urbanization; city centers are sprawling into mega connected cities. As urban centers rapidly grow, we are seeing that many services and

resources are poorly distributed and require significant effort to manage.

The interplay between these greater infrastructure systems is often—an unavoidable issue in smart urban spaces—ignored. It is not enough to only consider the data market and incentive structure of S-CPS such as transportation systems and the smart grid in isolation; we must consider their interconnections and interdependencies. New service models are emerging that capitalize on existence of CPS technologies and integrate previously isolated infrastructure systems. For example, the fact that there is a push towards a higher penetration of electric vehicles means that transportation infrastructures and the power grid no longer operate in isolation. New technologies including novel smart sensing and actuation modalities that can enable urban CPS networks to dynamically respond to human mobility and behavior patterns are being integrated into S-CPS. Metering at electric vehicle stations along with cellphone GPS data can provide more granular information about human mobility. However, this also comes at the risk of greater privacy loss. Due to the emerging interdependencies between S-CPS, there is an increase in exposure to vulnerabilities and opportunities for attack. As a result, there is a need for tools and techniques that are modular not only within a single S-CPS but across various types of S-CPS.

One interesting direction for future research is in the development of new interfaces that adapt in real-time and conform to individual preferences, needs, and wallets. Such interfaces will replace traditional static service models with options that are flexible and reflect the heterogeneous population of users as well as factors that influence their day-to-day decision making such as mobility, comfort, access, privacy and security.

Bibliography

- [AE05] A Akbas and M Ergun. “Dynamic traffic signal control using a nonlinear coupled oscillators approach”. In: *Canadian Journal of Civil Engineering* 32.2 (2005), pp. 430–441. DOI: [10.1139/104-121](https://doi.org/10.1139/104-121) (cit. on p. 14).
- [AR13] A. Albert and R. Rajagopal. “Smart Meter Driven Segmentation: What Your Consumption Says About You”. In: *IEEE Transactions on Power Systems* 28.4 (2013), pp. 4019–4030. DOI: [10.1109/TPWRS.2013.2266122](https://doi.org/10.1109/TPWRS.2013.2266122) (cit. on p. 82).
- [Abr+88] R. Abraham, J. E. Marsden, and T. Ratiu. *Manifolds, Tensor Analysis, and Applications*. 2nd. Springer, 1988. DOI: [10.1007/978-1-4612-1029-0](https://doi.org/10.1007/978-1-4612-1029-0) (cit. on pp. 19, 21, 22, 34, 41–44, 66).
- [Ake70] G. A. Akerlof. “The Market for ‘Lemons’: Quality Uncertainty and the Market Mechanism”. In: *The Quarterly Journal of Economics* 84.3 (1970), pp. 488–500. DOI: [10.2307/1879431](https://doi.org/10.2307/1879431) (cit. on p. 47).
- [All73] G. Allan. “Manipulation of Voting Schemes: A General Result”. In: *Econometrica* 41.4 (1973), pp. 587–601. DOI: [10.2307/1914083](https://doi.org/10.2307/1914083) (cit. on p. 84).
- [Ami+13] S. Amin, G. A. Schwartz, and S. S. Sastry. “Security of interdependent and identical networked control systems”. In: *Automatica* 49.1 (2013), pp. 186–192. DOI: [10.1016/j.automatica.2012.09.007](https://doi.org/10.1016/j.automatica.2012.09.007) (cit. on p. 5).
- [Ami+15] S. Amin, G. Schwartz, A. Cardenas, and S. Sastry. “Game-Theoretic Models of Electricity Theft Detection in Smart Utility Networks: Providing New Capabilities with Advanced Metering Infrastructure”. In: *IEEE Control Systems Magazine* 35.1 (2015), pp. 66–81. DOI: [10.1109/MCS.2014.2364711](https://doi.org/10.1109/MCS.2014.2364711) (cit. on p. 82).
- [Arr50] K. J. Arrow. “A Difficulty in the Concept of Social Welfare”. In: *Journal of Political Economy* 58.4 (1950), pp. 328–346 (cit. on pp. 76, 84).
- [Aue15] D. Auerbach. *Privacy Is Becoming a Premium Service: AT&T wants customers to pay the company not to spy on them. And it’s not an outlier.* 2015. URL: http://www.slate.com/articles/technology/bitwise/2015/03/at_t_gigapower_the_company_wants_you_to_pay_it_not_to_sell_your_data.html (cit. on p. 83).

- [BD05] P. Bolton and M. Dewatripont. *Contract theory*. MIT press, 2005. DOI: [10.1002/mde.1332](https://doi.org/10.1002/mde.1332) (cit. on pp. 84, 85, 88, 98).
- [BH12] A. Bressan and K. Han. “Nash equilibria for a model of traffic flow with several groups of drivers”. In: *ESAIM: Control, Optimisation and Calculus of Variations* 18.04 (2012), pp. 969–986. DOI: [10.1051/cocv/2011198](https://doi.org/10.1051/cocv/2011198) (cit. on p. 13).
- [BM14] M. Backes and S. Meiser. “Differentially Private Smart Metering with Battery Recharging”. In: *Lecture Notes in Computer Science* (2014), pp. 194–212. DOI: [10.1007/978-3-642-54568-9_13](https://doi.org/10.1007/978-3-642-54568-9_13) (cit. on p. 93).
- [BS00] J. F. Bonnans and A. Shapiro. *Perturbation analysis of optimization problems*. Springer Science & Business Media, 2000. DOI: [10.1007/978-1-4612-1394-9](https://doi.org/10.1007/978-1-4612-1394-9) (cit. on p. 65).
- [BS86] S. Boyd and S. Sastry. “Necessary and sufficient conditions for parameter convergence in adaptive control”. In: *Automatica* 22.6 (1986), pp. 629–639. DOI: [10.1016/0005-1098\(86\)90002-6](https://doi.org/10.1016/0005-1098(86)90002-6) (cit. on p. 75).
- [BT10] H. Broer and F. Takens. “Chapter 1 - Preliminaries of Dynamical Systems Theory”. In: *Handbook of Dynamical Systems*. Ed. by F. T. Henk Broer and B. Hasselblatt. Vol. 3. Handbook of Dynamical Systems. Elsevier Science, 2010, pp. 1–42. DOI: [10.1016/S1874-575X\(10\)00309-7](https://doi.org/10.1016/S1874-575X(10)00309-7) (cit. on pp. 23, 46).
- [BV04] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge university press, 2004 (cit. on pp. 52, 56, 57).
- [BW14] J. Bhuiyan and C. Warzel. “God View”: Uber Investigates Its Top New York Executive For Privacy Violations. 2014. URL: <http://www.buzzfeed.com/johanabhuiyan/uber-is-investigating-its-top-new-york-executive-for-privacy#.hqqn30e> (cit. on p. 4).
- [Ba95] T. Başar, G. J. Olsder, G. Clsder, T. Basar, T. Baser, and G. J. Olsder. *Dynamic noncooperative game theory*. Vol. 200. SIAM, 1995. DOI: [10.1137/1.9781611971132](https://doi.org/10.1137/1.9781611971132) (cit. on p. 13).
- [Bas05] S. Basov. “Multidimensional Screening”. In: *Studies in Economic Theory*. Vol. 22. Springer, 2005. DOI: [10.1007/b138910](https://doi.org/10.1007/b138910) (cit. on p. 108).
- [Bax97] J. Baxter. “A Bayesian/Information Theoretic Model of Learning to Learn via Multiple Task Sampling”. In: *Machine Learning* 28.1 (1997), pp. 7–39. DOI: [10.1023/a:1007327622663](https://doi.org/10.1023/a:1007327622663) (cit. on p. 77).
- [Bax98] J. Baxter. “Theoretical Models of Learning to Learn”. In: *Learning to Learn* (1998), pp. 71–94. DOI: [10.1007/978-1-4615-5529-2_4](https://doi.org/10.1007/978-1-4615-5529-2_4) (cit. on p. 77).
- [Baş87] T. Başar. “Relaxation techniques and asynchronous algorithms for on-line computation of non-cooperative equilibria”. In: *Journal of Economic Dynamics and Control* 11.4 (1987), pp. 531–549. DOI: [10.1109/CDC.1987.272779](https://doi.org/10.1109/CDC.1987.272779) (cit. on pp. 10, 35).

- [Ber99] D. P. Bertsekas. *Nonlinear programming*. Athena Scientific, 1999 (cit. on pp. [11](#), [17](#), [18](#), [33](#), [35](#), [41](#)).
- [Bre67] L. M. Bregman. “The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming”. In: *USSR computational mathematics and mathematical physics* 7.3 (1967), pp. 200–217. DOI: [10.1016/0041-5553\(67\)90040-7](#) (cit. on p. [78](#)).
- [CB13] G. Cavraro and L. Badia. “A game theory framework for active power injection management with voltage boundary in smart grids”. In: *European Control Conference*. 2013, pp. 2032–2037 (cit. on p. [15](#)).
- [CBL06] N. Cesa-Bianchi and G. Lugosi. *Prediction, learning, and games*. Cambridge University Press, 2006. DOI: [10.1017/CB09780511546921](#) (cit. on pp. [48](#), [53](#)).
- [CL14] C. P. Chambers and N. S. Lambert. “Dynamically eliciting unobservable information”. In: *Proceedings of the 15th ACM Conference on Economics and computation*. ACM. 2014, pp. 987–988. DOI: [10.1145/2600057.2602859](#) (cit. on p. [83](#)).
- [Can+10] U. O. Candogan, I. Menache, A. Ozdaglar, and P. A. Parrilo. “Near-optimal power control in wireless networks: a potential game approach”. In: *Proceedings of the 29th IEEE Conference on Information Communications*. 2010, pp. 1–9. DOI: [10.1.1.208.794](#) (cit. on p. [12](#)).
- [Can+11] O. Candogan, I. Menache, A. Ozdaglar, and P. A. Parrilo. “Flows and Decompositions of Games: Harmonic and Potential Games”. In: *Mathematics of Operations Research* 36.3 (2011), pp. 474–503. DOI: [10.1287/moor.1110.0500](#) (cit. on p. [29](#)).
- [Can+13] O. Candogan, A. Ozdaglar, and P. A. Parrilo. “Dynamics in near-potential games”. In: *Games and Economic Behavior* 82 (2013), pp. 66–90. DOI: [10.1016/j.geb.2013.07.001](#) (cit. on p. [29](#)).
- [Cha+11] N. Chaturvedi, A. Sanyal, and N. McClamroch. “Rigid-Body Attitude Control”. In: *IEEE Control Systems Magazine* 31.3 (2011), pp. 30–51. DOI: [10.1109/MCS.2011.940459](#) (cit. on p. [14](#)).
- [Con+04] J. Contreras, M. Klusch, and J. Krawczyk. “Numerical Solutions to Nash–Cournot Equilibria in Coupled Constraint Electricity Markets”. In: *IEEE Transactions on Power Systems* 19.1 (2004), pp. 195–206. DOI: [10.1109/TPWRS.2003.820692](#) (cit. on pp. [10](#), [35](#)).
- [Coo+13] S. Coogan, L. J. Ratliff, D. Calderone, C. Tomlin, and S. S. Sastry. “Energy management via pricing in LQ dynamic games”. In: *Proceedings of the 2013 American Control Conference*. 2013, pp. 443–448. DOI: [10.1109/ACC.2013.6579877](#) (cit. on pp. [13](#), [38](#), [76](#)).

- [Coo+15] S. Coogan, G. Gomes, E. Kim, M. Arcak, and P. Varaiya. “Offset optimization for a network of signalized intersections via semidefinite relaxation”. In: *Proceedings of the 54th IEEE Conference on Decision and Control*. 2015 (cit. on p. 14).
- [Cár+09] A. Cárdenas, S. Amin, B. Sinopoli, A. Giani, A. Perrig, and S. Sastry. “Challenges for securing cyber physical systems”. In: *Workshop on future directions in cyber-physical systems security*. 2009. DOI: [10.1.1.152.5198](https://doi.org/10.1.1.152.5198) (cit. on p. 5).
- [DB11] F. Dörfler and F. Bullo. “On the Critical Coupling for Kuramoto Oscillators”. In: *SIAM Journal on Applied Dynamical Systems* 10.3 (2011), pp. 1070–1099. DOI: [10.1137/10081530x](https://doi.org/10.1137/10081530x) (cit. on p. 14).
- [DB12] F. Dörfler and F. Bullo. “Synchronization and Transient Stability in Power Networks and Nonuniform Kuramoto Oscillators”. In: *SIAM Journal on Control and Optimization* 50.3 (2012), pp. 1616–1642. DOI: [10.1137/110851584](https://doi.org/10.1137/110851584) (cit. on pp. 14, 33).
- [DB78] S. Daan and C. Berde. “Two coupled oscillators: Simulations of the circadian pacemaker in mammalian activity rhythms”. In: *Journal of Theoretical Biology* 70.3 (1978), pp. 297–313. DOI: [10.1016/0022-5193\(78\)90378-8](https://doi.org/10.1016/0022-5193(78)90378-8) (cit. on p. 14).
- [DS67] N. Dunford and J. T. Schwartz. *Linear operators*. Ed. by 4. New York: Interscience Publishers, Inc., 1967 (cit. on p. 34).
- [Don+13a] R. Dong, L. J. Ratliff, H. Ohlsson, and S. Sastry. “Energy disaggregation via adaptive filtering”. In: *Proceedings of the 51st Annual Allerton Conference on Communication, Control, and Computing*. 2013, pp. 173–180. DOI: [10.1109/Allerton.2013.6736521](https://doi.org/10.1109/Allerton.2013.6736521) (cit. on p. 9).
- [Don+13b] R. Dong, L. Ratliff, H. Ohlsson, and S. S. Sastry. “Fundamental Limits of Nonintrusive Load Monitoring”. In: *Proceedings of the 3rd ACM International Conference on High Confidence Networked Systems* (2013), pp. 11–18. DOI: [10.1145/2566468.2566471](https://doi.org/10.1145/2566468.2566471) (cit. on pp. 9, 83).
- [Don+14] R. Dong, A. A. Cárdenas, L. J. Ratliff, H. Ohlsson, and S. S. Sastry. “Quantifying the Utility-Privacy Tradeoff in the Smart Grid”. In: *arXiv* 1406.2568 (2014) (cit. on pp. 9, 82, 83, 85, 89, 91–93, 99).
- [Dor+13] D. Dorsch, H. Jongen, and V. Shikhman. “On Structure and Computation of Generalized Nash Equilibria”. In: *SIAM Journal on Optimization* 23.1 (2013), pp. 452–474. DOI: [10.1137/110822670](https://doi.org/10.1137/110822670) (cit. on p. 10).
- [Dwo11] C. Dwork. “Differential privacy”. In: *Encyclopedia of Cryptography and Security*. Springer, 2011, pp. 338–340. DOI: [10.1007/11787006_1](https://doi.org/10.1007/11787006_1) (cit. on pp. 93, 108).

- [Dör+13] F. Dörfler, M. Chertkov, and F. Bullo. “Synchronization in complex oscillator networks and smart grids”. In: *Proceedings of the National Academy of Sciences* 110.6 (2013), pp. 2005–2010. DOI: [10.1073/pnas.1212134110](https://doi.org/10.1073/pnas.1212134110) (cit. on pp. 14, 33).
- [Eis82] T. Eisele. “Nonexistence and nonuniqueness of open-loop equilibria in linear-quadratic differential games”. In: *Journal of Optimization Theory and Applications* 37.4 (1982), pp. 443–468. DOI: [10.1007/BF00934951](https://doi.org/10.1007/BF00934951) (cit. on p. 37).
- [Eke74] I. Ekeland. “Topologie différentielle et théorie des jeux”. In: *Topology* 13.4 (1974), pp. 375–388 (cit. on pp. 10, 20).
- [Ell14] M. Ellison. *NISTIR 7628 User’s Guide: A White Paper developed by the Smart Grid Interoperability Panel*. Smart Grid Cybersecurity Committee. 2014 (cit. on pp. 4, 104, 111).
- [Ene] *Data Privacy and the Smart Grid: A Voluntary Code of Conduct*. Tech. rep. United States Department of Energy, 2015. URL: <http://energy.gov/oe/downloads/data-privacy-and-smart-grid-voluntary-code-conduct> (cit. on p. 4).
- [FA00] M. Fahrioglu and F. Alvarado. “Designing incentive compatible contracts for effective demand management”. In: *IEEE Transactions on Power Systems* 15.4 (2000), pp. 1255–1260. DOI: [10.1109/59.898098](https://doi.org/10.1109/59.898098) (cit. on p. 83).
- [FL98] D. Fudenberg and D. K. Levine. *The theory of learning in games*. Vol. 2. MIT press, 1998 (cit. on pp. 49, 53).
- [Fac+07] F. Facchinei, A. Fischer, and V. Piccialli. “On generalized Nash games and variational inequalities”. In: *Operations Research Letters* 35.2 (2007), pp. 159–164. DOI: [10.1016/j.orl.2006.03.004](https://doi.org/10.1016/j.orl.2006.03.004) (cit. on p. 10).
- [Faw+14] H. Fawzi, P. Tabuada, and S. Diggavi. “Secure estimation and control for cyber-physical systems under adversarial attacks”. In: *IEEE Transactions on Automatic Control* 59.6 (2014), pp. 1454–1467. DOI: [10.1109/TAC.2014.2303233](https://doi.org/10.1109/TAC.2014.2303233) (cit. on p. 5).
- [Flå02] S. D. Flåm. “Equilibrium, evolutionary stability and gradient dynamics”. In: *International Game Theory Review* 4.04 (2002), pp. 357–370. DOI: [10.1142/S0219198902000756](https://doi.org/10.1142/S0219198902000756) (cit. on p. 10).
- [Flå98] S. D. Flåm. “Restricted attention, myopic play, and the learning of equilibrium”. In: *Annals of Operations Research* 82 (1998), pp. 473–482. DOI: [10.1023/A:1018987409039](https://doi.org/10.1023/A:1018987409039) (cit. on p. 10).
- [Flå99] S. D. Flåm. “Learning Equilibrium Play: A Myopic Approach”. In: *Computational Optimization and Applications* 14.1 (1999), pp. 87–102. DOI: [10.1023/A:1008709129421](https://doi.org/10.1023/A:1008709129421) (cit. on p. 10).

- [Fri+12] P. Frihauf, M. Krstic, and T. Başar. “Nash Equilibrium Seeking in Noncooperative Games”. In: *IEEE Transactions on Automatic Control* 57.5 (2012), pp. 1192–1207. DOI: [10.1109/TAC.2011.2173412](https://doi.org/10.1109/TAC.2011.2173412) (cit. on p. 10).
- [GG73] M. Golubitsky and V. Guillemin. *Stable Mappings and Their Singularities*. Springer-Verlag, 1973. DOI: [10.1007/978-1-4615-7904-5](https://doi.org/10.1007/978-1-4615-7904-5) (cit. on p. 46).
- [GL77] J. Green and J.-J. Laffont. “Characterization of satisfactory mechanisms for the revelation of preferences for public goods”. In: *Econometrica* (1977), pp. 427–438. DOI: [10.2307/1911219](https://doi.org/10.2307/1911219) (cit. on p. 84).
- [GR01] C. Godsil and G. Royle. “Algebraic Graph Theory”. In: *Graduate Texts in Mathematics* (2001). DOI: [10.1007/978-1-4613-0163-9](https://doi.org/10.1007/978-1-4613-0163-9) (cit. on p. 30).
- [GS84] G. C. Goodwin and K. S. Sin. *Adaptive filtering prediction and control*. Englewood Cliffs, NJ: Prentice-Hall, 1984 (cit. on pp. 48, 60, 70, 75).
- [Ge+08] X. Ge, M. Arcaç, and K. Salama. “Nonlinear analysis of cross-coupled oscillator circuits”. In: *Proceedings of the 47th IEEE Conference on Decision and Control*. 2008, pp. 13–18. DOI: [10.1109/CDC.2008.4738885](https://doi.org/10.1109/CDC.2008.4738885) (cit. on p. 14).
- [Ged94] T. Gedra. “Optional forward contracts for electric power markets”. In: *IEEE Transactions on Power Systems* 9.4 (1994), pp. 1766–1773. DOI: [10.1109/59.331429](https://doi.org/10.1109/59.331429) (cit. on p. 83).
- [Gib+76] C. Gibson, K. Wirthmüller, A. du Plessis, and E. Looijenga. *Topological Stability of Smooth Mappings*. Vol. 552. Springer-Verlag, 1976. DOI: [10.1007/BFb0095244](https://doi.org/10.1007/BFb0095244) (cit. on p. 24).
- [Gli52] I. L. Glicksberg. “A further generalization of the Kakutani fixed point theorem, with application to Nash equilibrium points”. In: *Proceedings of the American Mathematics Society* 3.1 (1952), pp. 170–174. DOI: [10.2307/2032478](https://doi.org/10.2307/2032478) (cit. on p. 12).
- [Got+10] T. Goto, T. Hatanaka, and M. Fujita. “Potential game theoretic attitude coordination on the circle: Synchronization and balanced circular formation”. In: *IEEE International Symposium on Intelligent Control*. 2010, pp. 2314–2319. DOI: [10.1109/ISIC.2010.5612896](https://doi.org/10.1109/ISIC.2010.5612896) (cit. on pp. 15, 30).
- [HH98] J. H. Hubbard and B. B. Hubbard. *Vector calculus, linear algebra, and differential forms: a unified approach*. Prentice Hall, 1998 (cit. on p. 67).
- [HL13] J. P. Holdren and E. S. Lander. *Immediate opportunities for strengthening the nation’s cybersecurity*. President’s Council of Advisors on Science and Technology. 2013 (cit. on p. 111).
- [HL14] J. P. Holdren and E. S. Lander. *Big data and privacy: a technological perspective*. President’s Council of Advisors on Science and Technology. 2014 (cit. on p. 111).

- [HMC03] S. Hart and A. Mas-Colell. “Uncoupled Dynamics Do Not Lead to Nash Equilibrium”. In: *American Economic Review* 93.5 (2003), pp. 1830–1836. DOI: [10.1257/000282803322655581](https://doi.org/10.1257/000282803322655581) (cit. on p. 34).
- [Har89] G. Hart. “Residential energy monitoring and computerized surveillance via utility power flows”. In: *IEEE Technology Society Magazine* 8.2 (1989), pp. 12–16. DOI: [10.1109/44.31557](https://doi.org/10.1109/44.31557) (cit. on p. 82).
- [Hat02] A. Hatcher. *Algebraic Topology*. Cambridge University Press, 2002 (cit. on p. 28).
- [Hir76] M. W. Hirsch. *Differential topology*. Springer New York, 1976. DOI: [10.1007/978-1-4684-9449-5](https://doi.org/10.1007/978-1-4684-9449-5) (cit. on pp. 24, 45, 46).
- [JB14] M. Jelasity and K. P. Birman. “Distributional differential privacy for large-scale smart metering”. In: *Proceedings of the 2nd ACM workshop on Information hiding and multimedia security* (2014). DOI: [10.1145/2600918.2600919](https://doi.org/10.1145/2600918.2600919) (cit. on p. 93).
- [Jad+04] A. Jadbabaie, N. Motee, and M. Barahona. “On the stability of the Kuramoto model of coupled nonlinear oscillators”. In: *Proceedings of the 2004 American Control Conference*. Vol. 5. 2004, pp. 4296–4301. DOI: [10.1.1.239.7971](https://doi.org/10.1.1.239.7971) (cit. on p. 14).
- [Jay78] G. D. Jaynes. “Equilibria in monopolistically competitive insurance markets”. In: *Journal of Economic Theory* 19.2 (1978), pp. 394–422. DOI: [10.1.1.184.1553](https://doi.org/10.1.1.184.1553) (cit. on p. 102).
- [Jin+15] M. Jin, L. J. Ratliff, I. Konstantakopoulos, C. Spanos, and S. Sastry. “REST: a reliable estimation of stopping time algorithm for social game experiments”. In: *Proceedings of the ACM/IEEE Sixth International Conference on Cyber-Physical Systems*. ACM. 2015, pp. 90–99. DOI: [10.1145/2735960.2735974](https://doi.org/10.1145/2735960.2735974) (cit. on pp. 76, 77).
- [Jon+01] H. T. Jongen, P. Jonker, and F. Twilt. “Nonlinear Optimization in Finite Dimensions”. In: *Nonconvex Optimization and Its Applications* (2001). DOI: [10.1007/978-1-4615-0017-9](https://doi.org/10.1007/978-1-4615-0017-9) (cit. on p. 41).
- [KK02] E. Klavins and D. E. Koditschek. “Phase regulation of decentralized cyclic robotic systems”. In: *The International Journal of Robotics Research* 21.3 (2002), pp. 257–275. DOI: [10.1.1.6.4024](https://doi.org/10.1.1.6.4024) (cit. on p. 10).
- [KV86] P. R. Kumar and P. Varaiya. *Stochastic systems: estimation, identification and adaptive control*. Englewood Cliffs, NJ: Prentice-Hall, 1986. DOI: [10.1002/acs.4480030109](https://doi.org/10.1002/acs.4480030109) (cit. on pp. 48, 60, 69, 70, 75, 80).
- [Kea+12] M. Kearns, M. M. Pai, A. Roth, and J. Ullman. “Mechanism Design in Large Games: Incentives and Privacy”. In: *arXiv* 1207.4084 (2012) (cit. on p. 83).

- [Kri+14] W. Krichene, J. Reilly, S. Amin, and A. Bayen. “Stackelberg Routing on Parallel Networks With Horizontal Queues”. In: *IEEE Transactions on Automatic Control* 59.3 (2014), pp. 714–727. DOI: [10.1109/TAC.2013.2289709](https://doi.org/10.1109/TAC.2013.2289709) (cit. on p. 12).
- [Kri+15] W. Krichene, S. Krichene, and A. Bayen. “Convergence of Mirror Descent Dynamics in the Routing Game”. In: *European Control Conference*. 2015 (cit. on p. 27).
- [Kur75] Y. Kuramoto. “Self-entrainment of a population of coupled non-linear oscillators”. In: *Lecture Notes in Physics* (1975), pp. 420–422. DOI: [10.1007/bfb0013365](https://doi.org/10.1007/bfb0013365) (cit. on p. 14).
- [LB87] S. Li and T. Başar. “Distributed algorithms for the computation of noncooperative equilibria”. In: *Automatica* 23.4 (1987), pp. 523–533. DOI: [10.1016/0005-1098\(87\)90081-1](https://doi.org/10.1016/0005-1098(87)90081-1) (cit. on p. 10).
- [LM02] J.-J. Laffont and D. Martimort. *The Theory of Incentives: The Principal-Agent Model*. Princeton university press, 2002 (cit. on pp. 2, 76).
- [Law10a] N. Lawson. *Reverse-engineering a smart meter*. 2010. URL: <http://rdist.root.org/2010/02/15/reverse-engineering-a-smart-meter/> (cit. on p. 4).
- [Law10b] N. Lawson. *Smart meter crypto flaw worse than thought*. 2010. URL: <http://rdist.root.org/2010/01/11/smart-meter-crypto-flaw-worse-than-thought/> (cit. on p. 4).
- [LeC73] L. LeCam. “Convergence of Estimates Under Dimensionality Restrictions”. In: *The Annals of Statistics* 1.1 (1973), pp. 38–53 (cit. on p. 93).
- [Lee+08] J.-H. Lee, R. Diersing, and C.-H. Won. “Satellite attitude control using statistical game theory”. In: *Proceedings of the 2008 American Control Conference*. 2008, pp. 4856–4861. DOI: [10.1109/ACC.2008.4587263](https://doi.org/10.1109/ACC.2008.4587263) (cit. on p. 15).
- [Lee12] J. M. Lee. *Introduction to smooth manifolds*. 2nd. Springer, 2012. DOI: [10.1007/978-1-4419-9982-5](https://doi.org/10.1007/978-1-4419-9982-5) (cit. on pp. 26, 28, 30, 41, 44, 45).
- [Lis+10] M. Lisovich, D. Mulligan, and S. Wicker. “Inferring Personal Information from Demand-Response Systems”. In: *IEEE Security & Privacy* 8.1 (2010), pp. 11–20. DOI: [10.1109/MSP.2010.40](https://doi.org/10.1109/MSP.2010.40) (cit. on p. 82).
- [Luc15] V. Luckerson. *How AT&T Wants You to Pay For Your Privacy*. 2015. URL: <http://time.com/3713931/att-privacy-charge/> (cit. on p. 83).
- [Löt83] P. Lötstedt. “Perturbation bounds for the linear least squares problem subject to linear inequality constraints”. In: *BIT Numerical Mathematics* 23.4 (1983), pp. 500–519. DOI: [10.1007/BF01933623](https://doi.org/10.1007/BF01933623) (cit. on p. 65).

- [ME05] A. Muhammad and M. Egerstedt. “Decentralized coordination with local interactions: Some new directions”. In: *Cooperative Control*. Springer, 2005, pp. 153–170. DOI: [10.1.1.488.1935](https://doi.org/10.1.1.488.1935) (cit. on p. 10).
- [MR78] M. Mussa and S. Rosen. “Monopoly and product quality”. In: *Journal of Economic Theory* 18.2 (1978), pp. 301–317. DOI: [10.1016/0022-0531\(78\)90085-6](https://doi.org/10.1016/0022-0531(78)90085-6) (cit. on p. 102).
- [MS83] R. B. Myerson and M. A. Satterthwaite. “Efficient mechanisms for bilateral trading”. In: *Journal of Economic Theory* 29.2 (1983), pp. 265–281. DOI: [10.1016/0022-0531\(83\)90048-0](https://doi.org/10.1016/0022-0531(83)90048-0) (cit. on pp. 76, 84).
- [MS96] D. Monderer and L. S. Shapley. “Potential games”. In: *Games and economic behavior* 14.1 (1996), pp. 124–143. DOI: [10.1006/game.1996.0044](https://doi.org/10.1006/game.1996.0044) (cit. on pp. 26, 27).
- [McK+12] E. McKenna, I. Richardson, and M. Thomson. “Smart meter data: Balancing consumer privacy concerns with legitimate applications”. In: *Energy Policy* 41 (2012), pp. 807–814. DOI: [10.1016/j.enpol.2011.11.049](https://doi.org/10.1016/j.enpol.2011.11.049) (cit. on p. 82).
- [Mil63] J. W. Milnor. *Morse theory*. Princeton university press, 1963 (cit. on p. 18).
- [Moi+10] R. Molioli, P. Vargas, and P. Husbands. “Exploring the Kuramoto model of coupled oscillators in minimally cognitive evolutionary robotics tasks”. In: *IEEE Congress on Evolutionary Computation*. 2010, pp. 1–8. DOI: [10.1109/CEC.2010.5586486](https://doi.org/10.1109/CEC.2010.5586486) (cit. on p. 14).
- [Mor65] J.-J. Moreau. “Proximité et dualité dans un espace hilbertien”. In: *Bulletin de la Société mathématique de France* 93 (1965), pp. 273–299 (cit. on p. 78).
- [Mot+12] A. Motamedi, H. Zareipour, and W. Rosehart. “Electricity Price and Demand Forecasting in Smart Grids”. In: *IEEE Transactions on Smart Grid* 3.2 (2012), pp. 664–674. DOI: [10.1109/TSG.2011.2171046](https://doi.org/10.1109/TSG.2011.2171046) (cit. on p. 82).
- [Nem+09] A. Nemirovski, A. Juditsky, G. Lan, and A. Shapiro. “Robust Stochastic Approximation Approach to Stochastic Programming”. In: *SIAM Journal on Optimization* 19.4 (2009), pp. 1574–1609. DOI: [10.1137/070704277](https://doi.org/10.1137/070704277) (cit. on pp. 48, 62, 75, 78, 79).
- [Nev75] J. Neveu. *Discrete-Parameter Martingales*. North-Holland Publishing Company, 1975 (cit. on p. 81).
- [PR13] M. Pai and A. Roth. “Privacy and Mechanism Design”. In: *arXiv* 1306.2083 (2013) (cit. on p. 83).
- [PY10] S. J. Pan and Q. Yang. “A Survey on Transfer Learning”. In: *IEEE Transactions on Knowledge and Data Engineering* 22.10 (2010), pp. 1345–1359. DOI: [10.1109/tkde.2009.191](https://doi.org/10.1109/tkde.2009.191) (cit. on p. 77).

- [Pal+07] D. Paley, N. E. Leonard, R. Sepulchre, D. Grünbaum, J. K. Parrish, et al. “Oscillator models and collective motion”. In: *IEEE Control Systems Magazine* 27.4 (2007), pp. 89–105. DOI: [10.1109/CDC.2005.1582776](https://doi.org/10.1109/CDC.2005.1582776) (cit. on pp. 14, 30, 31).
- [Pal63] R. S. Palais. “Morse theory on Hilbert manifolds”. In: *Topology* 2.4 (1963), pp. 299–340. DOI: [10.1016/0040-9383\(63\)90013-2](https://doi.org/10.1016/0040-9383(63)90013-2) (cit. on p. 43).
- [Par+01] J.-B. Park, B. H. Kim, J.-H. Kim, M.-H. Jung, and J.-K. Park. “A continuous strategy game for power transactions analysis in competitive electricity markets”. In: *IEEE Transactions on Power Systems* 16.4 (2001), pp. 847–855. DOI: [10.1109/59.962436](https://doi.org/10.1109/59.962436) (cit. on p. 12).
- [Pas+12a] F. Pasqualetti, F. Dörfler, and F. Bullo. “Attack Detection and Identification in Cyber-Physical Systems—Part I: Models and Fundamental Limitations”. In: *arXiv* 1202.6144 (2012) (cit. on p. 5).
- [Pas+12b] F. Pasqualetti, F. Dörfler, and F. Bullo. “Attack Detection and Identification in Cyber-Physical Systems—Part II: Centralized and Distributed Monitor Design”. In: *arXiv* 1202.6049 (2012) (cit. on p. 5).
- [Pes75] C. S. Peskin. *Mathematical aspects of heart physiology*. Courant Institute of Mathematical Sciences, New York University, 1975 (cit. on p. 14).
- [Pol97] E. Polak. *Optimization: algorithms and consistent approximations*. Springer New York, 1997. DOI: [10.1007/978-1-4612-0663-7](https://doi.org/10.1007/978-1-4612-0663-7) (cit. on pp. 11, 13, 17, 18, 33, 35, 37).
- [RS76] M. Rothschild and J. Stiglitz. “Equilibrium in competitive insurance markets: An essay on the economics of imperfect information”. In: *The Quarterly Journal of Economics* 90.4 (1976), pp. 629–649. DOI: [10.1.1.323.5371](https://doi.org/10.1.1.323.5371) (cit. on p. 102).
- [RS85] H. Robbins and D. Siegmund. “A Convergence Theorem for Non Negative Almost Supermartingales and Some Applications”. English. In: *Herbert Robbins Selected Papers*. Ed. by T. Lai and D. Siegmund. Springer New York, 1985, pp. 111–135. DOI: [10.1007/978-1-4612-5110-1_10](https://doi.org/10.1007/978-1-4612-5110-1_10) (cit. on p. 80).
- [Rag+10] M. Raginsky, A. Rakhlin, and S. Yüksel. “Online convex programming and regularization in adaptive control”. In: *Proceedings of the 49th IEEE Conference on Decision and Control*. 2010, pp. 1957–1962. DOI: [10.1109/CDC.2010.5717262](https://doi.org/10.1109/CDC.2010.5717262) (cit. on pp. 48, 75).
- [Raj+11] S. Rajagopalan, L. Sankar, S. Mohajer, and H. Poor. “Smart meter privacy: A utility-privacy framework”. In: *Proceedings of the 1st IEEE International Conference on Smart Grid Communications*. 2011, pp. 190–195. DOI: [10.1109/SmartGridComm.2011.6102315](https://doi.org/10.1109/SmartGridComm.2011.6102315) (cit. on pp. 83, 85).

- [Rat+12] L. Ratliff, S. Coogan, D. Calderone, and S. Sastry. “Pricing in linear-quadratic dynamic games”. In: *Proceedings of the 50th Annual Allerton Conference on Communication, Control, and Computing*. 2012, pp. 1798–1805. DOI: [10.1109/Allerton.2012.6483440](https://doi.org/10.1109/Allerton.2012.6483440) (cit. on pp. 38, 76).
- [Rat+13] L. J. Ratliff, S. A. Burden, and S. S. Sastry. “Characterization and Computation of Local Nash Equilibria in Continuous Games”. In: *Proceedings of the 51st Annual Allerton Conference on Communication, Control, and Computing*. 2013. DOI: [10.1109/Allerton.2013.6736623](https://doi.org/10.1109/Allerton.2013.6736623) (cit. on pp. 17, 19, 34, 53).
- [Rat+14a] L. Ratliff, R. Dong, H. Ohlsson, A. Cardenas, and S. Sastry. “Privacy and customer segmentation in the smart grid”. In: *Proceedings of the IEEE 53rd Annual Conference on Decision and Control*. 2014, pp. 2136–2141. DOI: [10.1109/CDC.2014.7039714](https://doi.org/10.1109/CDC.2014.7039714) (cit. on p. 84).
- [Rat+14b] L. J. Ratliff, S. A. Burden, and S. S. Sastry. “Genericity and Structural Stability of Non-Degenerate Differential Nash Equilibria”. In: *Proceedings of the 2014 American Controls Conference*. 2014, pp. 3990–3995. DOI: [10.1109/ACC.2014.6858848](https://doi.org/10.1109/ACC.2014.6858848) (cit. on pp. 21, 23).
- [Rat+14c] L. J. Ratliff, R. Dong, H. Ohlsson, and S. S. Sastry. “Incentive Design and Utility Learning via Energy Disaggregation”. In: *Proceedings of the 19th World Congress of the International Federation of Automatic Control*. 2014, pp. 3158–3163 (cit. on p. 77).
- [Rat+14d] L. J. Ratliff, S. A. Burden, and S. S. Sastry. “On the Characterization of Local Nash Equilibria in Continuous Games”. In: *IEEE Transactions on Automatic Control* (2014) (cit. on pp. 17, 19, 21, 34).
- [Rat+14e] L. J. Ratliff, M. Jin, I. C. Konstantakopoulos, C. Spanos, and S. S. Sastry. “Social game for building energy efficiency: Incentive design”. In: *Proceedings of the 52nd Annual Allerton Conference on Communication, Control, and Computing*. IEEE. 2014, pp. 1011–1018. DOI: [10.1109/ALLERTON.2014.7028565](https://doi.org/10.1109/ALLERTON.2014.7028565) (cit. on pp. 76, 77).
- [Rat+14f] L. Ratliff, C. Barreto, R. Dong, H. Ohlsson, A. A. Cárdenas, and S. S. Sastry. “Effects of Risk on Privacy Contracts for Demand-Side Management”. In: *arxiv* 1409.7926v3 (2014) (cit. on p. 84).
- [Ros65] J. B. Rosen. “Existence and Uniqueness of Equilibrium Points for Concave N-Person Games”. In: *Econometrica* 33.3 (1965), p. 520. DOI: [10.2307/1911749](https://doi.org/10.2307/1911749) (cit. on pp. 10, 18, 41).
- [SA05] J. Shamma and G. Arslan. “Dynamic fictitious play, dynamic gradient play, and distributed convergence to Nash equilibria”. In: *IEEE Transactions on Automatic Control* 50.3 (2005), pp. 312–327. DOI: [10.1109/TAC.2005.843878](https://doi.org/10.1109/TAC.2005.843878) (cit. on pp. 10–12, 34).

- [SB89] S. S. Sastry and M. Bodson. *Adaptive Control*. Englewood Cliffs, NJ: Prentice–Hall, 1989 (cit. on pp. 48, 60, 75).
- [Sal+12] M. Salehie, L. Pasquale, I. Omoronyia, and B. Nuseibeh. “Adaptive security and privacy in smart grids: A software engineering vision”. In: *International Workshop on Software Engineering for the Smart Grid*. 2012, pp. 46–49. DOI: [10.1109/SE4SG.2012.6225718](https://doi.org/10.1109/SE4SG.2012.6225718) (cit. on p. 82).
- [San+13] L. Sankar, S. R. Rajagopalan, S. Mohajer, and H. V. Poor. “Smart meter privacy: A theoretical framework”. In: *IEEE Transactions on Smart Grid* 4.2 (2013), pp. 837–846. DOI: [10.1109/TSG.2012.2211046](https://doi.org/10.1109/TSG.2012.2211046) (cit. on pp. 83, 85, 91, 93).
- [Sas99] S. Sastry. *Nonlinear Systems*. Interdisciplinary Applied Mathematics. Springer New York, 1999. DOI: [10.1007/978-1-4757-3108-8](https://doi.org/10.1007/978-1-4757-3108-8) (cit. on p. 67).
- [Sat75] M. A. Satterthwaite. “Strategy-proofness and Arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions”. In: *Journal of Economic Theory* 10.2 (1975), pp. 187–217. DOI: [10.1016/0022-0531\(75\)90050-2](https://doi.org/10.1016/0022-0531(75)90050-2) (cit. on p. 84).
- [Sep+05a] R. Sepulchre, D. Paley, and N. Leonard. “Collective Motion and Oscillator Synchronization”. English. In: *Cooperative Control*. Ed. by V. Kumar, N. Leonard, and A. Morse. Vol. 309. Lecture Notes in Control and Information Science. Springer Berlin Heidelberg, 2005, pp. 189–205. DOI: [10.1007/978-3-540-31595-7_11](https://doi.org/10.1007/978-3-540-31595-7_11) (cit. on p. 14).
- [Sep+05b] R. Sepulchre, D. Paley, and N. E. Leonard. “Graph Laplacian and Lyapunov design of collective planar motions”. In: *Proceedings of the International Symposium on Nonlinear Theory and Its Application*. 2005. DOI: [10.1.1.125.2707](https://doi.org/10.1.1.125.2707) (cit. on p. 31).
- [Sep+07] R. Sepulchre, D. Paley, and N. Leonard. “Stabilization of Planar Collective Motion: All-to-All Communication”. In: *IEEE Transactions on Automatic Control* 52.5 (2007), pp. 811–824. DOI: [10.1109/TAC.2007.898077](https://doi.org/10.1109/TAC.2007.898077) (cit. on pp. 14, 30).
- [Sma75] S. Smale. “Global analysis and economics”. In: *Synthese* 31.2 (1975), pp. 345–358. DOI: [10.1007/BF00485983](https://doi.org/10.1007/BF00485983) (cit. on p. 10).
- [Spa81] E. H. Spanier. *Algebraic Topology*. Springer New York, 1981. DOI: [10.1007/978-1-4684-9322-1](https://doi.org/10.1007/978-1-4684-9322-1) (cit. on p. 28).
- [Ste10] N. Stein. “Games on Manifolds”. Unpublished Notes (Personal Correspondence). 2010 (cit. on pp. 16, 27, 28).
- [TK14] Z. Tufekci and B. King. *We Can’t Trust Uber*. 2014. URL: http://www.nytimes.com/2014/12/08/opinion/we-cant-trust-uber.html?_r=0 (cit. on p. 4).

- [TT14] H. Tavafoghi and D. Teneketzis. “Optimal Energy Procurement from a Strategic Seller with Private Renewable and Conventional Generation”. In: *arXiv* 1401.5759v1 (2014) (cit. on p. 83).
- [Tho74] R. Thom. “L’optimisation Simultanée et la Théorie des Jeux en Topologie Différentielle”. In: *Comptes rendus des Journées Mathématiques de la Société Mathématique de France* 3 (1974), pp. 63–70 (cit. on p. 10).
- [Tsy09] A. B. Tsybakov. “Introduction to Nonparametric Estimation”. In: *Springer Series in Statistics* (2009). DOI: [10.1007/b13794](https://doi.org/10.1007/b13794) (cit. on p. 93).
- [UR94] S. Uryasev and R. Rubinstein. “On relaxation algorithms in computation of noncooperative equilibria”. In: *IEEE Transactions on Automatic Control* 39.6 (1994), pp. 1263–1267. DOI: [10.1109/9.293193](https://doi.org/10.1109/9.293193) (cit. on p. 35).
- [Var04] H. Varian. “System Reliability and Free Riding”. English. In: *Economics of Information Security*. Ed. by L. Camp and S. Lewis. Vol. 12. Advances in Information Security. Springer US, 2004, pp. 1–15. DOI: [10.1007/1-4020-8090-5_1](https://doi.org/10.1007/1-4020-8090-5_1) (cit. on p. 7).
- [WG14] E. Weise and J. Guynn. *Uber tracking raises privacy concerns*. 2014. URL: <http://www.usatoday.com/story/tech/2014/11/19/uber-privacy-tracking/19285481/> (cit. on p. 4).
- [WKD91] J.-Y. Wen and K. Kreutz-Delgado. “The attitude control problem”. In: *IEEE Transactions on Automatic Control* 36.10 (1991), pp. 1148–1162. DOI: [10.1109/9/9.90228](https://doi.org/10.1109/9/9.90228) (cit. on p. 14).
- [WT11] S. Wicker and R. Thomas. “A Privacy-Aware Architecture for Demand Response Systems”. In: *Proceedings of the 44th International Conference on System Sciences*. 2011, pp. 1–9. DOI: [10.1109/HICSS.2011.24](https://doi.org/10.1109/HICSS.2011.24) (cit. on p. 82).
- [Wal26] L. Walras. *Eléments D’économie Politique Pure ou Théorie de la Richesse Sociale*. Pichon and Durand-Auzias, 1926 (cit. on p. 11).
- [Wan+08] R. Wang, C. Li, L. Chen, and K. Aihara. “Modeling and Analyzing Biological Oscillations in Molecular Networks”. In: *Proceedings of the IEEE* 96.8 (2008), pp. 1361–1385. DOI: [10.1109/JPROC.2008.925448](https://doi.org/10.1109/JPROC.2008.925448) (cit. on p. 14).
- [Wan+15] Y. Wang, Y. Hori, S. Hara, and F. Doyle. “Collective Oscillation Period of Inter-Coupled Biological Negative Cyclic Feedback Oscillators”. In: *IEEE Transactions on Automatic Control* 60.5 (2015), pp. 1392–1397. DOI: [10.1109/TAC.2014.2342072](https://doi.org/10.1109/TAC.2014.2342072) (cit. on p. 14).
- [Web11] T. A. Weber. “Optimal control theory with applications in economics”. In: *MIT Press Books* 1 (2011). DOI: [10.7551/mitpress/9780262015738.001.0001](https://doi.org/10.7551/mitpress/9780262015738.001.0001) (cit. on p. 84).

- [Yin+10] H. Yin, P. Mehta, S. Meyn, and U. Shanbhag. “Learning in mean-field oscillator games”. In: *Proceedings of the 49th IEEE Conference on Decision and Control*. 2010, pp. 3125–3132. DOI: [10.1109/CDC.2010.5717142](https://doi.org/10.1109/CDC.2010.5717142) (cit. on p. 15).
- [Yin+12] H. Yin, P. Mehta, S. Meyn, and U. Shanbhag. “Synchronization of Coupled Oscillators is a Game”. In: *IEEE Transactions on Automatic Control* 57.4 (2012), pp. 920–935. DOI: [10.1109/TAC.2011.2168082](https://doi.org/10.1109/TAC.2011.2168082) (cit. on p. 15).
- [ZS01] J. Zhang and S. Sastry. “Aircraft conflict resolution: Lie-Poisson reduction for game on SE(2)”. In: *Proceedings of the 40th IEEE Conference on Decision and Control*. Vol. 2. 2001, pp. 1663–1668. DOI: [10.1109/.2001.981140](https://doi.org/10.1109/.2001.981140) (cit. on p. 15).
- [Zet14] K. Zetter. *Hackers Can Mess With Traffic Lights to Jam Roads and Reroute Cars*. 2014. URL: <http://www.wired.com/2014/04/traffic-lights-hacking/> (cit. on p. 4).
- [Zhu+12a] Q. Zhu, J. Zhang, P. Sauer, A. Dominguez-Garcia, and T. Başar. “A game-theoretic framework for control of distributed renewable-based energy resources in smart grids”. In: *Proceedings of the 2012 American Control Conference*. 2012, pp. 3623–3628. DOI: [10.1109/ACC.2012.6315275](https://doi.org/10.1109/ACC.2012.6315275) (cit. on p. 13).
- [Zhu+12b] Q. Zhu, C. J. Fung, R. Boutaba, and T. Başar. “GUIDEX: A Game-Theoretic Incentive-Based Mechanism for Intrusion Detection Networks”. In: *IEEE Journal of Selected Areas in Communications* (2012), pp. 2220–2230. DOI: [10.1109/JSAC.2012.121214](https://doi.org/10.1109/JSAC.2012.121214) (cit. on p. 38).
- [Zou+12] A.-M. Zou, K. Kumar, and Z.-G. Hou. “Attitude Coordination Control for a Group of Spacecraft Without Velocity Measurements”. In: *IEEE Transactions on Control Systems Technology* 20.5 (2012), pp. 1160–1174. DOI: [10.1109/TCST.2011.2163312](https://doi.org/10.1109/TCST.2011.2163312) (cit. on p. 14).