Imperial College London

ELECTRONIC SOCIAL CAPITAL FOR SELF-ORGANISING MULTI-AGENT SYSTEMS

PATRICIO EMANUEL PETRUZZI

Thesis submitted for the degree of Doctor of Philosophy

Supervised by PROF. JEREMY PITT and DR. DÍDAC BUSQUETS

Intelligent Systems and Networks Group
Department of Electrical & Electronic Engineering
Imperial College London
August 2016

ABSTRACT

It is a recurring requirement in open systems, such as networks, distributed systems and socio-technical systems, that a group of agents must coordinate their behaviour for common good. In those systems – where agents are heterogeneous – unexpected behaviour can occur due to errors or malice. Agents whose practices free-ride the system can be accepted to a certain level; however, not only do they put the stability of the system at risk, but they also compromise the agents that behave according to the system's rules.

In social systems, it has been observed that *social capital* is an attribute of individuals that enhances their ability to solve collective action problems. Sociologists have studied collective action through human societies and observed that social capital plays an important role in maintaining communities though time as well as in simplifying the decision making in them. In this work, we explore the use of Electronic Social Capital for optimising self-organised collective action.

We developed a context-independent Electronic Social Capital framework to test this hypothesis. The framework comprises a set of handlers that capture events from the system and update three different forms of social capital: trustworthiness, networks and institutions. Later, a set of indicators are generated by the forms of social capital and used for decision-making. The framework was tested in different scenarios such as 2-player games, n-player games and public goods games. The experimental results show that social capital optimises the outcomes (in terms of long-term satisfaction and utility), reduces the complexity of decision-making and scales with the size of the population.

This work proposes an alternative solution using Electronic Social Capital to represent and reason with qualitative, instead of traditional quantitative, values. This solution could be embedded into socio-technical systems to incentivise collective action without commodifying the resources or actions in the system.

PUBLICATIONS

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

SELF-ORGANISING FLEXIBLE DEMAND FOR SMART GRID

Presents an experiment using *favours* as a form of social capital to facilitate the self-organisation of electricity demands. This paper was the first endeavour to model a straightforward form of social capital (Petruzzi et al. 2013).

SOCIAL CAPITAL AS A COMPLEXITY REDUCTION MECHANISM FOR DECISION MAKING IN LARGE SCALE OPEN SYSTEMS

Specifies a framework for electronic social capital and describes one of many possible implementation (Petruzzi et al. 2014c).

VISUALISATION OF SOCIAL CAPITAL

Describes a visualisation tool that allows to better understand the data generated from simulations in the PRESAGE2 platform and the electronic social capital framework. Some of the figures also appear in chapter 4 (Petruzzi et al. 2014b).

EXPERIMENTS WITH SOCIAL CAPITAL IN MULTI-AGENT SYSTEMS

Presents experimental results for the electronic social capital framework in cooperative and competitive scenarios (2-player games). Some of the figures and experimental results also appear in chapter 4 (Petruzzi et al. 2014a).

A GENERIC SOCIAL CAPITAL FRAMEWORK FOR OPTIMISING SELF-ORGANISED COLLECTIVE ACTION

Proposes an evolved *generic* electronic social capital framework and evaluates it in simultaneous n-player games. This paper covers much of the same grounds as chapter 3 and chapter 5 (Petruzzi et al. 2015).

Collective Intelligence and Algorithmic Governance of Socio-Technical Systems The idea of a decentralised community energy system was introduced in this paper. (Pitt et al. 2014a).

Inter-Institutional Social Capital for Self-Organising 'Nested Enterprises'

Presents an enhanced version of the electronic social capital framework implementing contextualised machine learning and evaluates it in a set of aggregated decentralised community energy systems. This paper is mainly included in chapter 6 (Petruzzi et al. 2016).

ACKNOWLEDGMENTS

The work presented in this thesis could not have been done without the help of many people. I would like to thank my supervisors for their continued guidance and support during my PhD.

Thanks to Prof. Jeremy Pitt for agreeing to have me as a his student and, for being a source of interesting ideas and discussions. Thank you for presenting to me the concepts of social capital and self-governing the commons, and giving me the unique opportunity to work on these ideas from very different disciplines. I appreciate his help, support and feedback on every paper we wrote; which made this thesis possible and, allowed me to travel to new places to present them.

Thanks to Dr. Dídac Busquets for his help and support while modelling these concepts into its 'electronic' equivalent. His help understanding sociologists and political scientists from a computer scientist perspective was invaluable.

I am very grateful to the project EPSRC Grand Challenge Project Autonomic Power System (EP/I031650/1), for their financial support and, feedback during the different presentations and project meetings.

To my colleagues from the Intelligent Systems and Networks group, with whom I have spent most of my time during the last few years and had many interesting conversations about research and other topics. Specials thanks to Sam Macbeth who provided all the support to start using the simulation platform.

Finally, thanks to all the people who made my time working on this thesis a lot more enjoyable – family and friends. I think they will appreciate more to hear it in person, which I will do, next time I see them.

DECLARATION

I, Patricio Emanuel Petruzzi, hereby declare that this work is my own, and where it is based on or derived from the work of others, I have acknowledged this and included a reference in the bibliography.

London, August 2016

CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF ALGORITHMS

ACRONYMS

CG	Cooperation Game
CPR	Common Pool Resource
CPRI	Common Pool Resource Institution
dCES	decentralised Community Energy System
DoF	Degrees of Freedom
EPGG	Electricity Public Goods Game
ESC	Electronic Social Capital
MAS	Multi-Agent System
PGG	Public Goods Game
PRESAGE2	Programming Environment for the Simulation of Agent Societies 2
SIC	Sociologically-Inspired Computing
SOEI	Self-Organising Electronic Institution
SVM	Support Vector Machine
UDDG	Unscrupulous Diner's Dilemma Game
UML	Unified Modelling Language

1 INTRODUCTION

1.1 MOTIVATION

Open systems, such as networks, distributed systems and socio-technical systems, often face the challenge that a group of participants must cooperate or coordinate their behaviour for the *common good*. Usually, these system participants are referred to as *agents*, and form opportunistic alliances to manage *resources* and perform *collective actions* that would otherwise be impossible or very costly to achieve individually. These systems also lack a central authority, which makes them dependant on self-organising mechanisms to succeed.

When agents of these systems share a resource, they must agree on how the appropriation of it is managed. Generally, the objective of allocating resources is to find a distribution of resources that is considered 'fair' and sustainable over time. The nature of the resource as *common* and the lack of a central authority to control it can encourage agents to free-ride the system by not contributing to it. This can be either the maintenance or the management of resources, depending on system's characteristics. A traditional approach is to use mechanism design (Hurwicz 1973) to model the situation as a strategic game and then use game theory, which allows the agents to analyse the game they are confronted with and decide on the best course of action or strategy to play. While game theory has been long applied to many problems, it imposes constraints and assumptions that might not be appropriate to model realistic scenarios (Binmore 2005). For example, it does not usually consider the possibility of refusing to engage in an interaction, even if the agent believes that it is more beneficial to do so.

Alternatively, Elinor Ostrom analysed how human societies are able to create institutions for the management and government of *resources* (Ostrom 1990). This allowed them to avoid the "tragedy of the commons" (Hardin 1968) –

depletion of the resource – predicted by a game-theoretic analysis. She defined institutions as "a set of working rules that are used to determine who is eligible to make decisions in some arena, what actions are allowed or constrained, ... [and] contain prescriptions that forbid, permit or require some action or outcome" (Ostrom 1990, p. 51). The "working rules" specify procedures for operational, collective and constitutional choices.

The formalisation of institutions in a computational form falls under the study of Self-Organising Electronic Institution (SOEI) (Pitt et al. 2012). A SOEI can be defined as a collection of agents, plus a specification of a dynamic, norm-governed system that encapsulates a set of operational, collective and constitutional choice rules, and the associated action situations, for realising self-organisation, self-regulation, and other self-* properties (i.e. the electronic 'equivalent' of Ostrom's self-governing institutions).

In later work (Ostrom and Ahn 2003), *social capital* has been analysed and defined as "an attribute of individuals that enhances their ability to solve collective action problems". Researchers also proposed that *social capital* created in one institution can enable other institutions to succeed, when both institutions are codependent (Ostrom 1990, pp.133–136). Ostrom's example of two codependent institutions is shown in Figure 1.1.

An irrigation system is actually composed of at least two different common resources: the water and the channel. Each of these resources represents a collective action problem to their users. Users must share the maintenance cost to keep the channel functional, which represents a provisioning problem.

On the other hand, the water use represents an appropriation problem that is associated with water allocation among users and the tools used to monitor the compliance with water rights. This example can also be extended to having two water appropriator institutions: one at the top and one at the bottom. In this example, the top enders could use all the water, leaving nothing for the bottom enders, but the social capital created in the channel institution can help manage the water institution and avoid the resource depletion. Despite the difficulties of managing different codependent institutions, Ostrom's analysis of communities across the world shows that some communities actually avoided the depletion of the resource ("tragedy of the commons").

Figure 1.1: Irrigation system example

In this thesis, we explore if the use of Self-Organising Electronic Institution (SOEI) with Electronic Social Capital (ESC) can enhance the ability to solve collective action problems in Multi-Agent System (MAS). We developed a context-independent ESC framework to test this hypothesis. Furthermore, we test whether ESC created in one SOEI can enable another SOEI to succeed, when both these institutions are codependent, as shown in the previous example.

1.2 METHODOLOGICAL BACKGROUND

In order to implement a theory from the social sciences in a computer context, we applied the methodology of Sociologically-Inspired Computing (SIC) (Jones et al. 2013). SIC endeavours to support systems engineering by developing formal and algorithmic models of social processes. On encountering an application problem, the general idea is to reflect on how people solve such problems, and use that as inspiration to reach a technical solution. We note that the paradigm of biologically-inspired computing operates in much the same vein (Andrews et al. 2010), only taking natural (biological) systems as its source of inspiration.

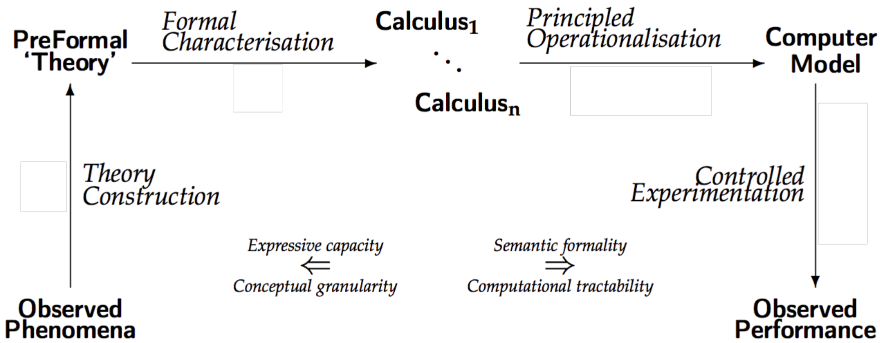

Figure 1.2: Methodology of SIC (Jones et al. 2013)

The methodology is illustrated in Figure 1.2. The steps involved are: when given a problem, identifying a theory from the social sciences regarding how people solve that (or an analogous) problem (*theory construction*); develop a formal model of that theory in appropriate calculus (*formal characterisation*), where by calculus we mean any formal language enabling symbolic representation and manipulation; implement that formal model (*principled operationalisation*); and then test the implementation to determine if it provides a solution to the original problem (*controlled experimentation*). Implicitly or explicitly, the methodology has been applied to Dennett's Intentional Stance (Dennett 1987) to produce the BDI agent architecture (Rao and Georgeff 1995): cognitive, psychological or physiological models provide decision-support systems based on trust (Neville and Pitt 2004), forgiveness (Vasalou et al. 2008) and emotions (Picard 1997); legal and organisational models provide a framework for agent societies (Artikis 2012) and learning by imitation provides for human-robot interaction (Demiris 2007).

1.3 THESIS OUTLINE

The body of this thesis is divided into five chapters and follows the Sociologically-Inspired Computing methodology. Figure 1.3 shows the order of chapters as it corresponds to the individual steps of the methodology.

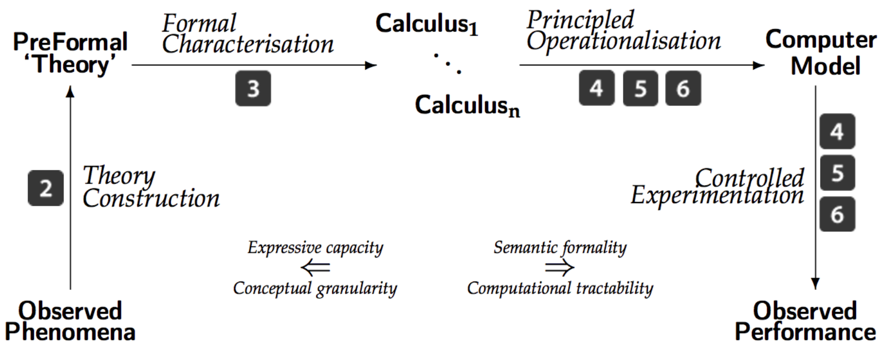

Figure 1.3: Chapters following the SIC methodology

CHAPTER 2 first discusses the different types of open MASs and *collective action* especially those subject to the problem of resource allocation. Moreover, a set of common characteristics shared by these systems is explained. Cooperation in MASs is introduced and contrasted with Ostrom's work on cooperation in human societies. Furthermore, Ostrom's work in self-governing institutions is introduced and how these institutions have been formalised electronically.

In this chapter, the notion of social capital is introduced and the concept is explored from sociological analyses of human societies to Ostrom's ideas on how social capital could be leveraged for solving collective action problems. The chapter ends with examples of basic forms of social capital used in computer systems.

CHAPTER 3 presents a formal specification of the ESC. Using an appropriate formalism was the first step towards simulating an open system. Moreover, the specification also includes a MAS to describe the environment in which further experiments will be conducted.

One of the many possible instantiations of the ESC framework is introduced. The internal components of the framework are described, including the events, data structures and the decision-making process. This chapter ends with an analysis on the complexity of the decision-making using the framework.

CHAPTER 4 addresses the first experiments using the ESC framework. In these experiments, the framework is used to decide whether to cooperate or defect in a strategic game. An experimental testbed is implemented, in which a population of agents repeatedly plays pairwise games, and uses the social capital framework as the basis for their action-selection.

CHAPTER 5 presents more experiments using the ESC framework. Firstly, a new test bed for *n*-player games is described and the different players for this scenario are specified. Four new experiments were carried out using this scenario. To conclude, some related work is presented and compared to the experiments performed.

CHAPTER 6 introduces further experiments analysing the role and nature of social capital using agents with learning capabilities. Both inter-agent and inter-institutional interactions are then situated in an extension of a Public Goods Game (PGG) for an electricity scenario. This type of game is commonly used to study the effects of free-riding in the system (free-riding reduces the positive benefits of cooperation in otherwise unregulated situations).

CHAPTER 7 summarises this work by drawing conclusions from the three different experiments. With these findings, we vindicate that Electronic Social Capital is an alternative solution to representing and reasoning with qualitative – instead of traditional quantitative – values. We also show that it incentivises *collective action* without commodifying the resources or actions in the system. Following this, we present several limitations of the different testbeds and the design decisions taken when defining the ESC framework. Afterwards, we suggest several lines of research that could continue this work.

1.4 CONTRIBUTION

This thesis has five main contributions:

- A generic framework to represent and reason using social capital. The framework decouples event processing and the updating of the social capital information from decision-making, thus providing a modular architecture to implement the framework.
- One of the many possible instantiations of the framework, which can be extended to include other forms of social capital.
- The design and implementation of a reusable experimental testbed, in which a set of experiments using the Electronic Social Capital framework in 2-player games were performed.
- A second experimental testbed facilitating experiments with large-scale agent populations with a correspondingly 'large' n-player game.
- A third experimental testbed with a new game Electricity Public Goods Game that models a collective action situation exhibited by aggregated decentralised Community Energy System and tests the role of social capital when the Common Pool Resource are codependent.

In conclusion, this thesis shows how theories of social capital can be formalised into their electronic 'equivalent' to represent and reason with qualitative, instead of traditional quantitative, values. It also shows that social capital facilitates cooperation, enhancing agents' ability to solve collective action problems in Multi-Agent System.

2 SOCIAL CAPITAL FOR COLLECTIVE ACTION

2.1 INTRODUCTION

In this chapter, we discuss the different types of open Multi-Agent System (MAS) that are subject to the problem of resource allocation, such as manufacturing and scheduling, public transport, grid computing or network routing. Moreover, a set of common characteristics shared between those systems is explained.

We will discuss how an allocation can be considered 'fair' by analysing the different approaches adopted. Furthermore, cooperation is briefly introduced and examples of how it has been addressed in MAS are included, Ostrom's work on cooperation in human societies follows. We introduce Ostrom's analysis on self-governing institutions for Common Pool Resource (CPR) management and examine how these institutions have been formalised electronically.

To conclude, we present the notion of social capital, which can play a key role in promoting cooperation among institutions. This chapter ends with some examples of basic forms of social capital used in computer systems.

2.2 OPEN MULTI-AGENT SYSTEMS

A wide diversity of systems, such as load balancing (Bernard et al. 2014), electronic spot markets (Anders et al. 2013) or job scheduling (Wu et al. 2011) can be viewed as MAS in which individual components act with each other and the environment to achieve individual and collective objectives. A sub-class of these systems is one that is *open* – i.e. a system in which agents can join or leave anytime and their individual objectives are unknown to others. In open MAS, it has been assumed (Huynh et al. 2006) that:

- Agents are likely to be self-interested and unreliable due to different ownership.

- The environment is not entirely known by the agents, mainly because that could be too costly or unfeasible.
- Agents' behaviours are not controlled by a central authority.

The success of these systems relies on the agents' abilities to cooperate and coordinate their behaviour, not only to fulfil their individual objectives but also to achieve the *common good*. The idea of the "common good" is a complex concept which can encompass a range of different qualities and values, such as the 'fair distribution of shared resources', 'those affected by rules participate in their selection', and 'sustainability'; moreover there are a variety of metrics for evaluating whether or not the value is realised, e.g. for "fairness" there is utilitarianism, proportionality, envy-free, Gini index, etc. Therefore, for this thesis, we *intuitively* consider the "common good" to be a beneficial outcome of active participation in a collective action situation with respect to some mutually agreed metrics; and indeed social capital is then an externality associated with structures, relations and behaviour that is both a product of and contributes towards achieving that outcome.

The attainment of this successful self-organisation makes the system tolerant to heterogeneity, conflicts and unexpected events (Pitt et al. 2012). However, without a central authority, heterogenous agents may exhibit selfish behaviour encouraged by conflicting goals, errors or malice. Agents whose practices free ride the system can be accepted to a certain level; however, they not only put the stability of the system at risk, but also compromise the agents that behave according to the system's rules. As a result, agents require incentives to participate, contribute or select an action that maximises the collective, rather than individual, utility.

Our main concern is how these systems self-organise when agents share a resource, and must agree on how the provision and/or appropriation of that resource is managed.

2.3 RESOURCE ALLOCATION IN OPEN MULTI-AGENT SYSTEMS

"Multi-Agent Resource Allocation is the process of distributing a number of items amongst a number of agents."

(Chevaleyre et al. 2006)

The allocation of resources is relevant to a wide range of domains, such as manufacturing (Brussel et al. 1998) and scheduling (Council 1998), public transport (Cantillon and Pesendorfer 2004), grid computing (Galstyan et al. 2005) and network routing (Feldmann et al. 2003). All of these examples share the following common characteristics (Chevaleyre et al. 2006):

2.3.1 *Resources*

The resources refer to *items* that are being distributed among the participating agents of the system. For instance, in a grid computing scenario, resources may refer to jobs or tasks that must be executed; in a network routing context, they could represent packets of information that must be delivered; and in a smart-grid context, they may relate to electricity.

Resources can be classified into *divisible* and *indivisible*, depending on their nature. When the resource is divisible, agents can receive fractions of it during the distribution. However, if the resource is indivisible, a distribution might assign the available resources to only some of the agents. In some cases, an indivisible resource could be also shared between different agents.

2.3.2 *Allocation*

Allocation refers to a particular distribution of resources between the agents. For example, for indivisible resources an allocation could be the subset of agents receiving the resource. When the resource is divisible, the allocation could be the fraction of the resources that each agent receives.

2.3.3 *Preferences*

Agents can have preferences regarding the allocations they receive. In some cases, agents can communicate their preferences by sending *demands* to other participants. The lack of a central authority and the heterogeneity of the system could question the *honesty* of such demands. For example, in an electricity scenario, agents could lie about their true necessities during peak times to receive better allocations.

2.3.4 *Allocation method*

The allocation method defines the procedures required to perform an allocation. Usually, the method can be divided into two categories: *centralised* or *distributed*. In the centralised method, one particular agent performs the role of allocator and decides how the resources will be distributed. Who occupies the role of allocator can be commonly agreed upon by participants, along with the allocation method and other externalities that may affect the process. In some cases, this role is also rotated among the participants. When the allocation method is distributed, the allocation is decided by negotiation and consensus among agents.

2.3.5 *Objectives*

The objective of *resource allocation* is to find an allocation that is either *feasible* or *optimal*. The feasibility refers to finding a suitable allocation that solves the problem – for example, when a set of tasks has to be performed. The optimality, on the other hand, relates to finding an allocation that maximises utility for the allocator (i.e. the auctioneer's revenue in combinatorial auctions) or that maximises agents' preferences (i.e. individual agent's utility).

In some scenarios, a combination of different objectives might be also possible. The main objective could include searching for an optimal solution amongst a set of feasible ones (i.e in auctioning, maximising the revenue and bidder satisfaction).

The objectives of these systems exhibit three further characteristics:

- They operate in time-slices, so the negotiations have to be conducted repeatedly and frequently, but they retain memory of past interactions.
- The *cost* of decision-making has to be taken into account, as it may have to be 'paid for', usually from the same resources that may be the subject of the allocation.
- They are time-constrained – i.e. the negotiations have to be concluded before the commencement of the next time slice.

It is important to highlight that the heterogeneity of the agents in the system and the semantics of the resources may make optimal allocation impossible to calculate. Therefore, any improvements towards optimal allocation can be considered a success.

2.4 EVALUATION OF RESOURCE ALLOCATION

> *"Equals should be treated equally, and unequals unequally,*
> *in proportion to the relevant similarities and differences"*
>
> (Aristotle 350BC)

Distributive justice is concerned with the *fair* allocation of goods to a set of actors in society. But what is *fair allocation*? Social scientists usually relate fairness with justice, in the sense that when resources are distributed fairly, societies are more just (Konow 2003). In this context, fairness can target keeping individuals 'happy' or 'satisfied', thus averting situations where individuals might misbehave and create a negative impact on society. Furthermore, fairness in computer systems is usually connected to maximising efficiency (Nagle 1987; Mo and Walrand 2000). This is especially the case for a network or grid computing cluster, where the 'happiness' of network devices or servers can not be considered, but their improved performance can be measured. There has been some interest in including social views of fairness in computer systems. For example, computational social choice connects aspects of the social sciences with Multi-Agent Systems, including fairness, among others (Chevaleyre et al. 2007).

In addition, three main families of distributive theories have been identified (Konow 2003):

1. *Equality and need:* This is concerned with the welfare of those least advantaged in the society and is based on the 'need' principle – i.e. equal satisfaction of basic needs.

2. *Utilitarianism and welfare economics:* This is based on maximising the global surplus of outcome, utility or satisfaction. It does not deal with individual outcomes, but with their aggregation.

3. *Equity and desert:* This links allocation with the actions of individuals and uses the equity principle – i.e. individuals should receive allocations that are proportional to their contribution to society.

However, there is still no formal agreement on a formal definition of 'fairness' and how it can be measured. In the field of resource allocation in MAS, the following properties of allocation methods and their outcomes have been summarised (Pitt et al. 2014b).

- *Proportional:* Each individual receives $(\frac{1}{n})^{th}$ of the resource allocation.

- *Envy-free:* No agent i prefers the allocation of agent j.

- *Equitable:* The utility of each agent's allocation is the same.

- *Efficient:* The greatest good for the majority is ensured.

- *Cost-effective:* This minimises the costs of calculating the 'best' allocation. It is especially used in systems where the allocation costs are paid with the resources.

- *Timely:* Computing the allocation ends quickly enough to avoid the loss of utility over time.

Furthermore, a computational model for self-organised 'fair' resource allocation was formalised (Pitt et al. 2014b) using the theory of distributive justice based on legitimate claims (Rescher 1966). The results of the experiments show sustainable 'fairness over time' and a better balance of utility and fairness compared to monistic or fixed approaches.

2.5 cooperation: a brief overview

An agent involved in an interaction with other agents faces the problem of deciding whether to cooperate (e.g. provision resources, consume the allocated amount, respect the system's norms) with the other agents or not. A traditional approach is to model the situation as a strategic game and use game theory, which allows the agents to analyse the game they are confronted with and decide on the best action or strategy to play. This usually involves computing the Nash equilibrium (Nash 1950), with which the agent is guaranteed that it cannot make itself better off by unilaterally changing its strategy. If all players have chosen their actions and no player can benefit by changing its action, then the current set of actions constitutes a Nash equilibrium.

The Logic of Collective Action (Olson 1965) analysed the premise of rational individuals. Although, when individuals within a group are considered *rational* and *self-interested*, and the group's and individuals' interests are aligned, it would be logical for them to act together to pursue a common objective, this is not always the case. Individuals will always have incentives to free-ride the system, but this could be counteracted by only providing the benefits to active participants.

Moreover, *The evolution of Cooperation* (Axelrod and Hamilton 1981) analysed cooperation in iterated 2-player games, in particular Prisoner's dilemma game (Rapoport and Chammah 1965). When the number of interactions among participants is not fixed in advance, other strategies [rather than not cooperate] might be stable as well. Furthermore, they emphasise that "the discrimination of others may be among the most important of abilities because it allows one to handle interactions with many individuals without having to treat them all the same" (Axelrod and Hamilton 1981).

In (Papadimitriou and Roughgarden 2005), it is suggested even less is known about computing equilibria for n-player games than for the special case of 2-player games. It is however shown in (Papadimitriou and Roughgarden 2005) that there is a polynomial-time algorithm for finding a Nash equilibrium in a certain type of n-player k-strategy game; but no such algorithm is known for general games, even when $k = 2$.

Elinor Ostrom analysed how human societies were able to create institutions for the management and government of resources (Ostrom 1990), which allowed them to avoid the "tragedy of the commons" (Hardin 1968) (depletion of the resources) predicted by a game-theoretical analysis. Furthermore, she added "What makes these models so dangerous – when they are used metaphorically as the foundations for policy – is that the constraints that are assumed to be fixed for the purpose of analysis are taken on faith as being fixed in empirical settings. . . [I'd] rather address the question of how to enhance the capabilities of those involved to change the [constraints] to lead to outcomes other than remorseless tragedies" (Ostrom 1990, pp.6–7).

Moreover, game theory imposes constraints and assumptions that might not be appropriate to model some scenarios (Binmore 2005). For example, game-theory only works when people play games *rationally*. So game theory neglects the problem that irrational players exist in the real world.

Inspired by economic sciences, Social Choice, Mechanism Design and Trust have been transposed to Multi-Agent System to address issues of cooperation.

Social Choice theory is concerned with the design and analysis of methods for aggregating the preferences of multiple agents (Chevaleyre et al. 2007). If we view the system as a society where autonomous agents have different objectives, different capabilities and hold different information, a clearly defined mechanism for aggregating their preferences is required to make collective decisions. An example of this situation includes voting procedures, which are used to aggregate the preferences of the voters to determine which candidate should win an election (Brandt et al. 2012).

Mechanism Design focus on implementing an optimal system-wide solution to a decentralised optimisation problem with self-interested agents with private information about their preferences for different outcomes (Larson and Sandholm 2004). Examples of applications in recent years include electronic market design and resource allocation problems.

Trust has long been acknowledged as an important notion in MAS, where agents may be self-interested, heterogeneous and dishonest (Sabater and Sierra 2005). Agents must rely in some agents and mistrust in other ones to achieve a goal. By relying on others, agents place their own interests at risk, which introduces the need for trust (Castelfranchi et al. 2006).

In this work, we will focus on cooperation in MAS to facilitate a good-enough system-wide solution to a decentralised satisfaction problem with self-interested agents with private information, but also a shared set of congruent values. As in the previous examples of cooperation theories transposed to MAS, we can ask ourselves: How do (groups of) people solve this sort of problem? One approach, used by communities, is to invent and self-organise sets of conventional rules to (voluntarily) regulate/organise their own behaviour, i.e., self-governing institutions (Ostrom 1990). In addition, communities also accredit value to complying with those rules, typically in the form of social capital (Ostrom and Ahn 2003). In the next section, Ostrom's work on self-governing the commons is introduced, followed by a consideration of her work on social capital in this context.

2.6 self-governing the commons

In Ostrom's work on self-governing institutions for CPR management, an institution is defined as "a set of working rules that are used to determine who is eligible to make decisions in some arena, what actions are allowed or constrained, ... [and] contain prescriptions that forbid, permit or require some action or outcome" (Ostrom 1990, p. 51). The "working rules" specify procedures for operational, collective and constitutional choice, and are respectively concerned with provision, appropriation and monitoring; determining the operational choice rules, rule enforcement and dispute resolution; and the eligibility for determining the collective choice rules.

These rules are role-based, mutually agreed upon, mutable and nested within each other in *decision arenas* or *action situations*. Distinguishing between nested situations requires a formal characterisation of institutionalised power (Jones and Sergot 1996), whereby an agent appointed to, or occupying a role in a particular action situation is empowered to bring about a fact of conventional or institutional significance by performing a designated action in that specific context. A role might be fixed in one action situation, but could be changed by the rules of another action situation within which it is nested.

Ostrom also observed that, on some occasions, the resource was not depleted and, on others, it was. Eight principles were identified as the necessary and sufficient conditions for a CPR managed by a self-governing institution to *endure* (Ostrom 1990, p. 90). These principles are summarised in Table 2.1.

Table 2.1: Design principles for managing CPR institutions

P1	Clearly defined boundaries
P2	Congruence between appropriation and provision rules and local conditions
P3	Collective choice arrangements
P4	Monitoring
P5	Graduated sanctions
P6	Conflict-resolution mechanisms
P7	Minimal recognition of rights to organise
P8	Nested enterprises

PRINCIPLE 1 defines the boundaries of CPR, specifying who has access to the resource and securing it from unauthorised access.

PRINCIPLE 2 defines the appropriation rules and puts restrictions on quantity, time and/or place. Provision is also regulated in terms of labour, material, quantities, etc. These rules must be congruent with the local environmental conditions.

PRINCIPLE 3 allows individuals who are affected by the operational rules the rights to modify them. This principle is very important, as the local participants usually have better knowledge about the resources and environmental conditions.

PRINCIPLE 4 applies to monitoring the status of the CPR and the participants' behaviour. The monitoring role can be given to the members themselves or to appointed agencies. There might be a *monitoring cost*, which is usually paid with provisions to the CPR.

PRINCIPLE 5 refers to appropriate sanctioning in the institution. A violation of a rule should be sanctioned according to the *severity* and *frequency* of the wrongdoing. Using graduated sanctions not only encourages individuals to comply with the institutional rules, but also to recover from unintentional misbehaviour.

PRINCIPLE 6 stipulates that participants need a mechanism for fast conflict resolution at a low *cost* – for example, when an operational rule is considered inappropriate in a rare and infrequent scenario.

PRINCIPLE 7 preserves the independence of the institution, allowing it to formulate its own rules. The institution should not be challenged by an external authority – i.e. a government.

PRINCIPLE 8 relates to CPRs that manage larger resources. In this case, nested structures of CPRs interact among themselves, with the smaller CPRs at the bottom.

To summarise, these principles relate to issues such as who belongs to the institution, congruence between rules for provision and appropriation and local conditions, whether those affected by the operational rules participate in the selection and modification of those rules (collective choice arrangements), graduated sanctions for violating rules, and layered or encapsulated systems.

2.7 SOCIAL CAPITAL

2.7.1 *Definition*

All types of capital involve the use of currently available resources to create other resources that will generate benefits in the future. The beneficiaries can be a small group of individuals, such as a small community, or larger groups, such as cities or countries. Capital is defined as "resources providing future benefits, in some measure, for a group of individuals" (Lachmann 1978).

Capital also comprises multiple forms, such as physical, human or social. Physical capital is entirely tangible and created by modifying materials to assist production (e.g. computers, railroads or factories). Human capital is less tangible, created by knowledge and skills obtained by an individual (Schultz

1961). The theory of social capital gained importance through the integration of classical sociological theory with the description of an intangible form of capital. Through the concept of social capital, researchers have tried to propose a synthesis between the value contained in communitarian approaches and the individualism professed by the *rational choice theory*. Social capital is therefore defined as a set of intangible collective resources that an individual or a group of individuals holds and "comes about through changes in the relations among them that facilitate action" (Coleman 1988).

Other relevant social capital definitions include:

- "The features of social organization, such as networks, norms and trust, that facilitate coordination and cooperation for mutual benefit" (Putnam 1993).
- "The aggregate of the actual or potential resources which are linked to possession of a durable network of more or less institutionalized relationships of mutual acquaintance and recognition..." (Bourdieu and Wacquant 1992).
- "An attribute of individuals that enhances their ability to solve collective action problems" (Ostrom and Ahn 2003).

Although it appears that there is still no agreement regarding on a single concrete and formal definition of 'social capital', researchers seem to agree on the qualitative value of connections and/or relations among individuals. These developed relations combined with institutionalised rules that are commonly agreed upon can facilitate the coordination (and cooperation) when performing collective actions.

Ostrom and Ahn observed that social capital has multiple forms, of which they identified three:

- *Trustworthiness*: This is distinct from trust and is related to reputation (i.e. a shared understanding of someone's willingness to honour agreements and commitments).
- *Networks*: These include strong and weak ties, and identify channels through which people communicate, as well as other social relations.

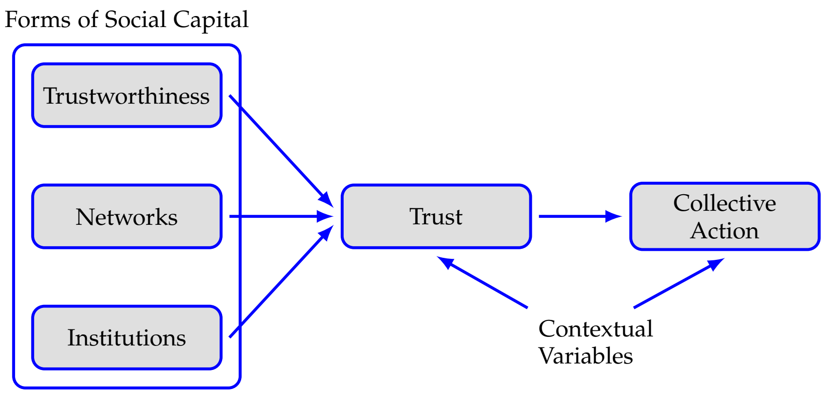

Figure 2.1: Ostrom and Ahn's social capital model

- *Institutions*: These are identified as sets of conventional rules by which people voluntarily, and mutually agree to, regulate their behaviour.

They also suggested that *trust* itself was the 'glue' that enabled these various forms of social capital to be leveraged for solving collective action problems (see Figure 2.1) – for example, the sustainability of a CPR. Social capital generates 'reliance' trust; where reliance trust can be seen as a complexity-reducing decision-making shortcut that helps resolve collective action problems.

Inspired by an analysis of *trust* as comprising a belief component and an expectation component, a trust decision can be modelled reasoning about three components (Jones 2002):

- The belief that there is a rule.
- The expectation that someone else's behaviour will conform to that rule.
- The expectation that a third party will punish any behaviour that does not conform to that rule.

Although social capital has not been widely analysed in a computational context, some examples of individual forms of social capital in computer systems can be identified. In the next section, some of these systems are presented.

2.7.2 *Applications in computer systems*

Examples of computer systems which represent and reason with social capital in a computational form include forgiveness in e-commerce, legitimate claims for fair resource allocation in open networks, and demand-side self-organisation in SmartGrids (Pitt and Nowak 2014).

For instance, one feature of open systems is the expectation of error, but there are many dimensions of error, including a distinction between intentional and unintentional violations, levels of seriousness, and so on. Human society has evolved a standard mechanism for recovering from error in general: *forgiveness*. From the literature of psychology, four positive motivations for forgiveness can be identified, comprising twelve constituent signals. This has been formalised into a computational model of forgiveness (Vasalou and Pitt 2005), which uses fuzzy logic to compute a forgiveness decision from given weights associated with each of the twelve signals. The critical aspect of this forgiveness model is that some of the constituent signals – for example, a 'prior beneficial relationship' – are an indication of some form of social capital.

Similarly, in the model of distributive justice based on legitimate claims introduced in section 2.4, the representation of some of the claims – notably, the claims according to efforts and sacrifices, and according to socially useful services – provides a ranking based on the quantitative representation of a form of social capital.

In another resource allocation experiment, agents request time slots for accessing the resource and receive an allocation. Once all the allocations are made, agents can exchange these allocations among themselves to better satisfy their preferences for different time slots. During each exchange, the agents check to see if the received allocation is in their interest. If so, they count it as a "favour received" from the other agent. If it is not in their interest, they count it as a "favour done" for the other agent. Since the calculation of favours is internal for each agent, an exchange where two agents exchange one allocation for another that they prefer is perceived as a favour received by *both* of them. These favours are a form of social capital and it has been shown that a system with favours can outperform one without (in terms of the percentage of time-slot preferences that are satisfied) (Petruzzi et al. 2013).

While these examples demonstrate the potential benefit of some form of social capital in computer-assisted or computer-mediated decision-making, they are generally implicit, disjointed and concerned with informing individual – rather than community – based action.

2.8 self-organising electronic institutions

The formalisation of the principles of self-governing institutions in a computational form had been pursued in the context of Self-Organising Electronic Institution (SOEI) (Pitt et al. 2012).

A framework for dynamic norm-governed systems that allows agents to modify the rules or protocols of a norm-governed system at runtime has been described (Artikis 2012). This framework defines three components:

- A specification of a norm-governed system.
- A protocol stack for defining how to change the specification.
- A topological space for expressing the 'distance' between one specification instance and another.

This framework can be used to specify a wide variety of such systems. Accordingly, a SOEI can be defined as a collection of agents plus a specification of a dynamic norm-governed system that encapsulates a set of operational, collective and constitutional choice rules and the associated action situations for realising self-organisation, self-regulation, and other self-* properties (i.e. the electronic 'equivalent' of Ostrom's self-governing institutions).

The definition of Degrees of Freedom (DoF) is also included in the framework (Artikis 2012). A DoF can define:

- How participants are accepted into an institution.
- How they can be expelled.
- What action they must perform given a certain situation or what sanctions apply for misbehaviours.

In a SOEI the decision making process of an agent (e.g. choosing its actions, choosing what rules to use and vote for, etc.) greatly depends on the interaction with other agents. Thus, we believe that social capital can indeed play a key role in shaping an agent's behaviour, promoting cooperation, and helping them achieve any kind of collective action in such institutions. The essence of this thesis is to enhance the framework of SOEI with a complementary framework for electronic social capital.

2.9 summary

In this chapter, we overviewed different types of open Multi-Agent System (MAS) related to the problem of resource allocation. We explained the characteristics shared among these systems and discussed which properties of a system can be analysed when considering if allocation is 'fair'. We also briefly introduced cooperation and how it has been addressed in MAS.

Furthermore, we discussed how the problem of resource allocation is managed in human societies. We introduced Ostrom's analysis on self-governing institutions for CPR management and how these institutions have been formalised electronically. Moreover, the notion of social capital was introduced, and it was described how sociologists believe its features can facilitate collective action, reduce the complexity of decision-making and facilitate cooperation in human societies. Certain isolated examples of social capital in computer systems were identified and explained.

This leads to our main research question:

Can social capital support successful collective action
in self-organising Multi-Agent Systems?

We will examine the effects of using social capital with electronic institutions in MAS for CPR management. To test this premise, we developed a unified computational framework for representing and reasoning about social capital. The Electronic Social Capital (ESC) framework is introduced in the following chapter.

3 ELECTRONIC SOCIAL CAPITAL

3.1 INTRODUCTION

A specification of the Electronic Social Capital (ESC) using an appropriate formalisation was the first step towards simulating an open system. Moreover, the specification also includes a Multi-Agent System (MAS) to describe the environment where further experiments will be conducted. One of many possible instantiations follows.

The environment was implemented using PRESAGE2, which already provides the grounds for a multi-agent simulation. On top of this, reputation sources were implemented to provide additional information and the institutions were defined using Self-Organising Electronic Institution (SOEI).

An instantiation of the ESC framework is introduced. The internal components of the framework are described, including events, data structures and the decision-making process. This chapter ends with an analysis of the complexity of the decision-making process within the framework.

3.2 ELECTRONIC SOCIAL CAPITAL SPECIFICATION

3.2.1 *Multi-agent system*

Let M be a MAS consisting of a set of agents A. The system operates in a possibly infinite sequence of time slices t_0, t_1, $\ldots$,t_n,$\ldots$. At each time step agents have a 'game' to play involving some form of cooperation. Once all the agents have chosen their actions, the outcomes are calculated and distributed among them.

Formally, the state of the system M at a discrete time t is specified by:

$$M_t = \langle A, RS, I, G \rangle_t$$

where A is a set of agents, RS is a set of reputation sources, I is a set of institutions and G is a set of games.

Different kind of agents might participate in the system – i.e. agents which use social capital or not. All the agents can obtain information regarding reputation from the reputation sources denoted by RS. The system also contains a set of institutions I with different characteristics and a few associated behavioural rules. Finally, the system has a set of games G which, for the purpose of this work, will be 2-player, n-player and public goods games. Each of those components of the system will be explained in more detail in the following sections.

At each time step, the system operation is given by the following cycle:

1. Each agent manages its institutions' memberships (i.e. join, leave, vote, etc. . .).
2. Each agent performs an action, according to the game defined.
3. The outcome for each agent is computed.
4. Reputation sources are updated depending on the actions.
5. Each institution is updated regarding its members' actions (i.e. rewards, sanctions, etc. . .).

To summarise, during one time step agents are able to change their affiliation to institutions. They have to choose an action to perform depending on the rules of the game they are playing, and the whole system is then updated according to the actions performed. This includes calculating the outcomes of the actions, updating the reputation and updating the institutional status based in internal rules (i.e. penalising an agent for misbehaving).

3.2.2 *Agents with Electronic Social Capital*

Each agent participating in the system and implementing some form of social capital is then defined by:

$$a = \langle \gamma, \delta, \mu, \alpha, \beta \rangle$$

where γ is a set of event handlers, δ is a set of attributes, μ is a set of social capital indicators, α is a social decision function and β is an utility function.

Each form of social capital must define a set of $\langle \gamma, \delta, \mu \rangle$, which enables to add or remove forms of social capital without changing the electronic social capital structure. For clarity, we will start by explaining the set of attributes μ. Each attribute is formed by a data structure that holds social capital values $sc^{val} \in [0, 1]$. For example, the networks form of social capital can be represented as a graph where each vertex identifies an agent and the edges connecting them hold the sc^{val} between those agents. Institutions form of social capital can be modelled using a table with a unique sc^{val} for each institution. Other data structures might be more appropriate for different forms of social capital. Once explained how the social capital can be stored (δ), we can move on to how it can be updated (γ) and how it can be retrieved (μ).

The event handlers μ are responsible of updating the relevant social capital attributes according to a weight specified in each event w_{event}. Those events are triggered internally in the agent, depending the actions perceived by the agents. The following equation is used by the event handlers to update the sc^{val} in δ.

$$sc^{val,t+1} = \begin{cases} sc^{val,t} + w_{event} \cdot (1 - sc^{val,t}) & \textbf{if } \text{enhances cooperation} \\ sc^{val,t} \cdot (1 - w_{event}) & \textbf{otherwise} \end{cases}$$

In this equations, $sc^{val,t+1}$ is the social capital value and w_{event} is the weight assigned to the event used as a learning factor. The first expression is used to reinforce the social capital value for an agent or institution if its actions enhanced cooperation. Conversely, we apply the second expression for actions that diminish cooperation.

Furthermore, a set of indicators μ is produced to feed the a social decision function α using the social capital values sc^{val} stored. One attribute holding a data structure can produce multiple μ for a single agent. From the previous example, the graph data structure holding the values of the networks form of social capital could create one indicator for each edge of the vertex representing a single agent. Furthermore, the tuple holding the institutions social capital values can also produce multiple μ for a single agent, by using the sc^{val} of all the institutions that the specific agent participating in.

The social decision function α uses a set of social capital indicators $\mu \in [0, 1]$ to produce a boolean decision $\alpha_{output} \in [T, F]$. The method for transforming the set of social capital indicators into a boolean decision must be chosen during the implementation. The social decision could not only be a cooperation decision, but also a forgiveness decision, or any other social process that uses the social capital indicators to inform the outcome. Examples of implementations could be if a simple average of the indicators goes over a fixed threshold, a wighted average adding more importance to some of the indicators or more advanced methods.

At last, the utility function β produces a numeric output value $\beta_{output} \in \mathbb{R}$ to analyse the outcomes of the agents' actions. This function must be chosen accordingly to the game played. For example, if agents are playing a 2-player game such as Prisoners' Dilemma, the utility function can be the years in jail sentenced.

Figure 3.1 shows a schematic view of the electronic social capital included inside each social agent.

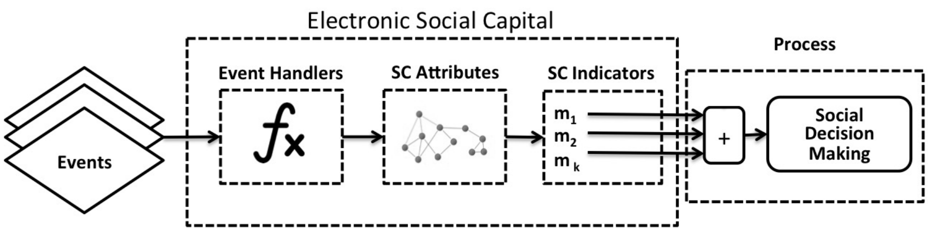

Figure 3.1: Schematic view of the Electronic Social Capital

3.2.3 *Reputation sources*

The modelling of reputation has engaged scientists from different fields, such as economics, psychology, sociology and computer science. Most of the computational models of reputation usually consider two main sources (Sabater and Sierra 2002):

1. Direct interactions between participant members of the system.
2. The information provided by other members about the interactions they had in the past.

We have used these reputation sources to model the reputation in our system as follows:

- *Direct Interaction Reputation (1)*: This is based on agents' individual experiences and are not shared with other agents. Within each social agent, the percentage of interactions where another agent cooperated is calculated using the following equation:

$$agent_k^{DIR} = \frac{1}{l} \cdot \sum_{i=1}^{l} A_i$$

  where l is the number of actions played against the other agent and $A \in [0,1]$ is the action played (1 if cooperated and 0 otherwise). This reputation source was defined to model reputation in 2-player games.

- *Agents' Reputation (2)*: This is also based on agents' individual experiences, but aggregated by a global entity. After each game is played, each agent reports if the game was successful or not for all the participants with which they have interacted. This feedback is aggregated into a value that shows the percentage of satisfaction achieved when interacting with an individual agent. It is calculated using the following equation:

$$agent_m^{AR} = \frac{1}{n} \cdot \sum_{j=1}^{n} \frac{1}{o} \cdot \sum_{i=1}^{o} S_{i,j}$$

  where n is the number of actions played by an agent m, o is the number of players in the game and $S \in [0,1]$ is the satisfaction reported by each player (1 if satisfied and 0 otherwise). This reputation source was defined to model reputation in n-player games.

- *Institutional Reputation (2)*: This is obtained from the interactions with institutions. Again, after each game, institutions submit data on whether the agents participating followed the rules. This value is also aggregated into a percentage of successful behaviour achieved when participating in an institution. It is calculated using the following equation:

$$agent_p^{IR} = \frac{1}{q} \cdot \sum_{j=1}^{q} \frac{1}{r} \cdot \sum_{i=1}^{r} R_{i,j}$$

  where q is the number of actions played by an agent p, r is the number of institutions which the agent is participating and $R \in [0,1]$ is factor that denotes if the agent followed the institution rules (1 if followed the institution rules and 0 otherwise). This reputation source was defined to model reputation in n-player games and linear public goods games from the institution point of view.

Agents' Reputation and *Institutional Reputation* models describe agents' behaviours from the other agents' and institutions' perspectives. However, it does not show how cooperative an agent is. For example, an agent who is not cooperating could report a successful interaction with another agent who also did not cooperate; since the other agent behaved reciprocally and as expected. The same situation could occur in an institution where no sanctions are applied and cooperation is not compulsory. The institution's feedback will be positive, as not cooperating would not be considered an institutional misbehaviour.

3.2.4 *Institutions*

The general definition of an institution is based on the SOEI (Artikis 2012) and an extension of the definition given previously in section 2.8:

$$I_t = \langle A, \mathcal{L}, G \rangle_t$$

where A is a set of agents, $\mathcal{L}$ is the specification instance and G is the game the institution regulates.

The Degrees of Freedom (DoF) of the rules in $\mathcal{L}$ are the following:

- *Allocation method*: Defines how resources must be distributed in the context of a resource allocation scenario.
- *Sanctions*: Sets whether rule violations are sanctioned or not.
- *Graduated sanctions*: Sets whether sanctions increase for recurring rule violators or not. Being expelled from the institution is the highest sanction.
- *Compulsory cooperation*: Sets whether cooperation is mandatory or not.
- *Democratised*: Allows its members to vote on certain managing aspects. Examples of voting situations are allowing a new member to join, expelling a current member, or choosing the institution's leader.

The corresponding ranges to the DoF listed above are shown in Figure 3.1. These characteristics are generalised and some might not apply in different games. This list provides an example of DoF.

Table 3.1: Rule Degrees of Freedom and range

DoF	*Range*
Allocation method	{ no-allocation, random, random-demand, average, contribution-ratio }
Sanctions	{ true, false }
Graduated sanctions	{ true, false }
Compulsory cooperation	{ true, false }
Democratised	{ true, false }

With this customisation of an institution, we included some of Ostrom's design principles regarding monitoring, sanctioning and membership (Ostrom 1990).

3.3 SYSTEM IMPLEMENTATION

3.3.1 *Simulation platform*

In this section, we explain the implementation of the MAS. We used Programming Environment for the Simulation of Agent Societies 2 (PRESAGE2) (Macbeth et al. 2014), a Java platform for rapid prototyping of Agent Societies. PRESAGE2 provides tools to simulate heterogeneous populations of software agents, including inter-agent communication, data logging, different types of networks and environment modelling.

PRESAGE2 was chosen as the base on which we built the different games. It is the evolution of the original Presage simulation platform (Neville and Pitt 2009). It was built based on modularity which allows flexibility and the reuse of components. A simulation in PRESAGE2 is discrete-time-driven where, at each loop, every agent perceives the environment and submits its actions. Besides the discrete-time-driven structure, the simulation platform does not specify how agents should be implemented, or what architecture they should use. PRESAGE2 simply queries for the result of an agent function at each time step. This enables operating at multiple levels of granularity and supports cases where agents must perform multiple actions at only one time step.

A simulation consists of four components: a set of agents, an environment state, a state transformer function, and a state observability function. In order to support modularity and component reuse, the platform is decomposed into monolithic functions that provide units of functionality. The state observability function is split into a set of functions named 'environment services' providing high-level access to a shared state while also defining observability. The 'state transformer' function consists of two sets of functions: firstly, a set of action handlers which define state changes for specific actions, and, secondly, a set of environment functions which define general state changes – for example for persistent effects on agents (momentum or current), or periodic (either random or triggered) events. Through specification of these sets of components one can define arbitrarily complex environments. Furthermore, components can be bundled together into 'modules' that provide specific functionality.

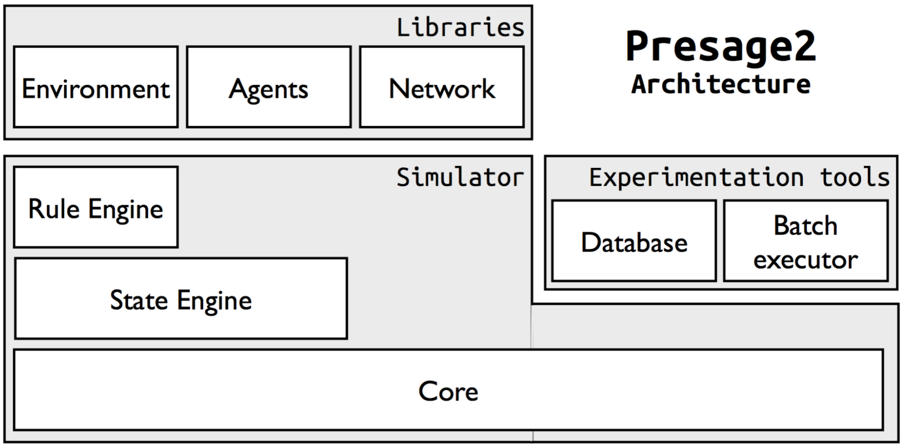

Figure 3.2: PRESAGE2 architectural block diagram

The simulation platform consists of several components, as shown in Figure 3.2. The state of the environment, and the state of agents is managed by a *state engine*. It provides a read-only interface to fetch data, as well as an interface to specify what changes should be made to the state at the end of the time step. The *rule engine* module is an extension to the PRESAGE2 state engine and incorporates a declarative rule-engine for the management and the modification of the shared state. The *network* module provides communication between agents via message actions. This module allows the configuration of constraints, which allows the simulation of different dynamic network topologies. The *environment* and *agent* modules provide the Java classes and tools to model different simulation scenarios and connect them to the whole PRESAGE2 environment. Lastly, two modules provide experimentation tools: the *database* and the *batch executor*. The first provides a standard connection to store simulation results into different storage providers as PostgreSQL, MySQL and MongoDB. The second provides a set of tools to launch simulations in batches with different simulation parameters.

PRESAGE2 EXTENSION The PRESAGE2 simulation platform had to be adapted to the requirements of this work. First, the *environment* had to be extended to implement the different games – 2-player, n-player and public goods games. This implementation is dependent on the experiments carried out and will be explained later with the associated experiments. The *agents* were extended to include the ESC framework, we called them *'social agents'*. Two more modules were developed: *reputation* and *institution* (the descriptions of both follow in the next section). Figure 3.3 shows the extended version of the platform components. The notation used at the formal specification of the system (subsection 3.2.1) is included at the top right of the modules.

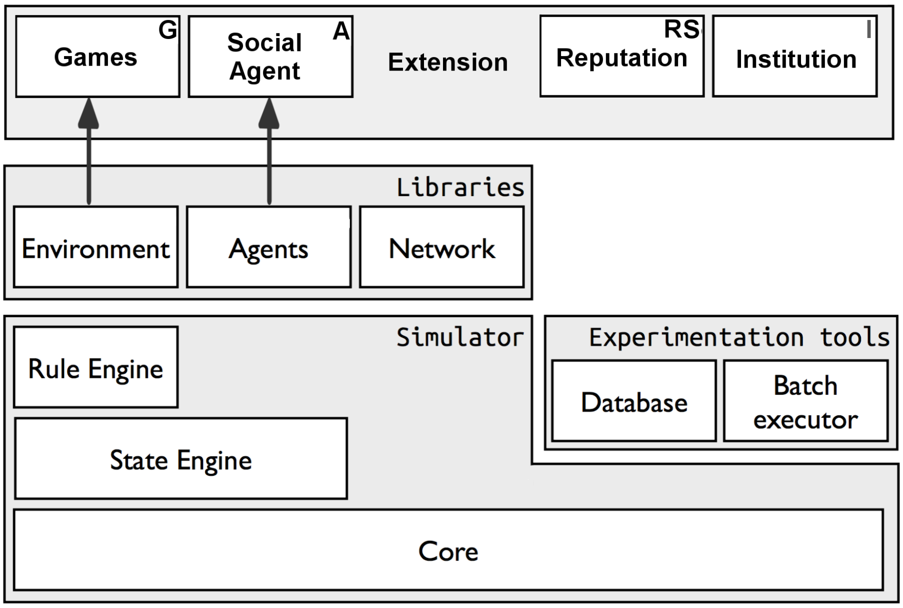

Figure 3.3: PRESAGE2 extension architectural block diagram

3.3.2 *Reputation sources*

Direct Interaction Reputation (1) was implemented within the *social agents* as the percentage of interactions where another agent acted as expected. Since it is not aggregated globally, there was no need to develop anything from the system's perspective.

Figure 3.4 shows the Unified Modelling Language (UML) diagram of the classes implemented to model the global *Agents' Reputation (2) and Institutional Reputation (2)*. The class *Reputation* will aggregate the individual outcomes of the agents' interactions between themselves or with the institutions.

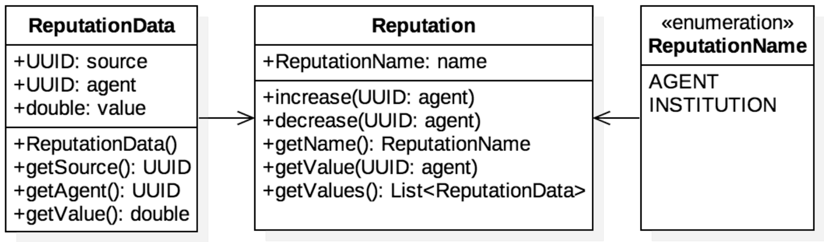

Figure 3.4: Reputation UML class diagram

3.3.3 *Institutions*

Figure 3.5 shows the UML diagram of the classes implemented to model the institutions. First, the class *InstitutionRules* defines all the DoF. These DoFs are implemented in the *Institution* class. This diagram only shows an example with the base characteristics; this must be extended to apply to different game scenarios.

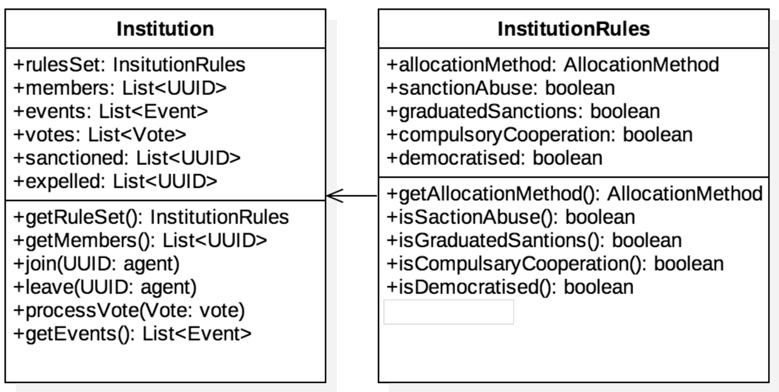

Figure 3.5: Institution UML class diagram

3.4 ELECTRONIC SOCIAL CAPITAL IMPLEMENTATION

We have instantiated three forms of social capital in this framework: trustworthiness, networks and institutions. Each form has an internal set of social capital handlers, attributes and indicators.

The trustworthiness form of social capital keeps the reputation information received from different sources. The network form stores information about all the actions that the other agents perform. This form of social capital can subjectively analyse the value of a relationship between two different agents among them. The third form of social capital, institutions, uses information on institutional actions performed in the system. A detailed description of each form of social capital is presented later in this chapter.

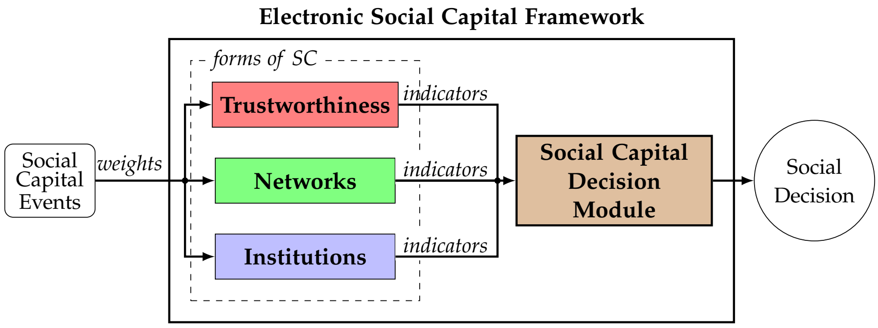

Figure 3.6: ESC framework inside an agent

Agents sense from the environment different events that occur in the system and translate them into social capital events. This information is the input of the ESC framework and includes data on about when an agent does or does not cooperate, what messages are sent or received, and all the institutional actions (such as joining, leaving, sanctioning, etc.). The three forms of social capital (trustworthiness, networks and institutions) will store the information received and aggregate it. When the agent needs information on another agent or an institution, it will query the social decision module. This module will combine all the indicators from the forms of social capital into a value that ranges zero to one, where zero is no cooperation and one is full cooperation, and transform into a boolean value using a threshold.

3.4.1 *Social capital events*

The framework comprises a set of context-independent events with an associated weight. When an agent perceives an action performed in the environment, the agent is responsible for translating the action into the appropriate social capital event. In the current version of the framework, the weights are fixed in the range of $(0, 1]$ and assigned by perceived significance (i.e. the weight of being expelled from an institution is higher than the weight for being sanctioned). Furthermore, the weights specify to which forms of social capital are bound. An event can be bound to more than one form of social capital. For example, a sanction performed by an institution is connected to the forms 'networks' and 'institutions'. This is, mainly because, at the networks form the social capital of the agent sanctioned is diminished, while at institutions form, the social capital of the sanctioning institution is increased.

The events implemented in the framework are:

- At trustworthiness: Regarding reputation data updates.
- At networks: Message sent and received; agent sanctioned, rewarded, expelled, voted; and some only relevant in resource allocation environments, such as resource allocated, demanded and appropriated.
- At institutions: Agent sanctioned, rewarded, expelled, joined and left (all shared with networks); and pure institutional (i.e. a vote or leader changed).

Implementation

Figure 3.7 shows the UML diagram of the main classes implemented to model the events. First, the abstract class *SocialEvent* must be extended to define a new event. Depending on which forms of social capital the event affects, it must also implement the appropriated interface. This solution simplifies the assignment of the events to different forms of social capital. This diagram shows only an example with three social events. More events can be added to the framework by extending and implementing these classes.

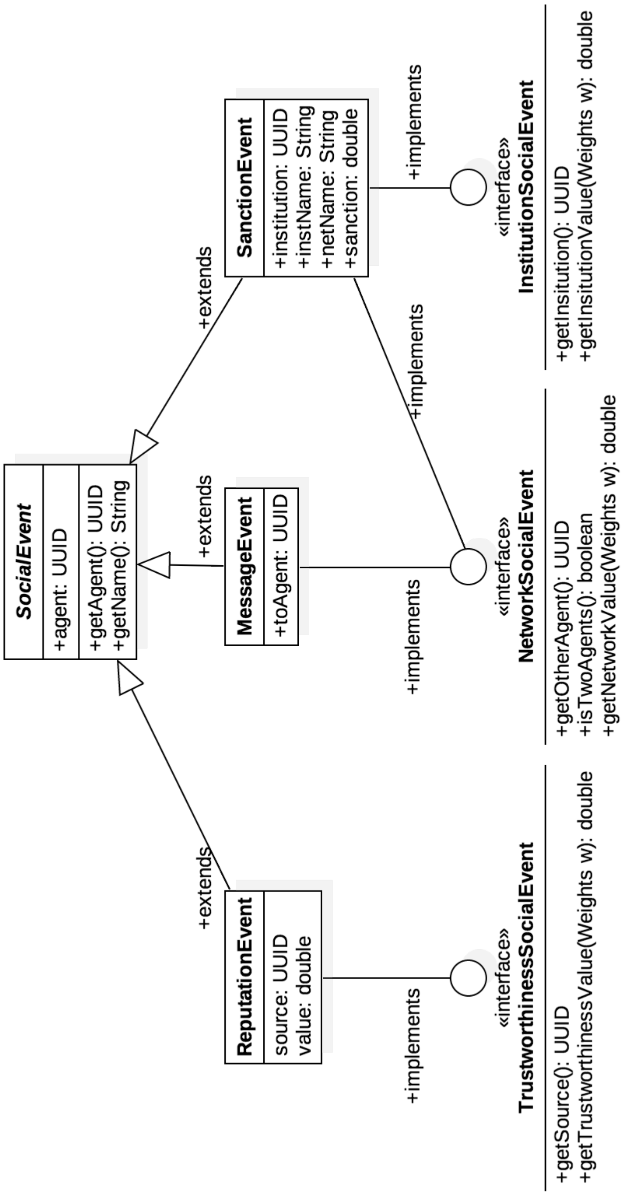

Figure 3.7: Social events UML class diagram

3.4.2 *Forms of social capital*

This section describes the forms of social capital implemented and the attributes used with their the data types.

Trustworthiness

This form of social capital collects reputation events. These events have a reputation source, an 'agent identification' and a reputation value. An update function normalises the reputation value to a value between zero and one and stores it. The data structure chosen to store this information is a *tuple*, which keeps the source, agent identification and normalised reputation value and updates them after each event is received. When a new reputation event is received, the previous value is overwritten, as the old reputation data is already included in the new one by the reputation sources. This form of social capital produces an indicator for each reputation source included in the system.

Networks

The networks form of social capital receives events regarding agents' interactions in the system. Examples of these events are whether an agent has performed the expected action or not, messages sent or received, and all the institutional actions performed (from an agent's perspective, not the institution's). Events have a specific weight based on their relevance. Sending a message, successfully cooperating or being banned from an institution do not have the same importance. Later, the event with its weight is used by the event handler to update the value at the social capital attributes. The data structure chosen to store the data is a graph, where each vertex represents an agent. The edges – with a value from zero to one assigned – represent the social capital between the agents. It is important to remark that this form of social capital updates information on the agents' actions even when the event is collected by an agent that does not participate in that action. For example, if two agents interact in some context where a third agent witnesses that action and its outcome, the *networks* form of social capital at the third agent will update the third agent's graph (the edge between the two agents) with the appropriate value for the outcome of that action.

Institutions

This form of social capital uses events related to the institutions, such as rules followed, sanctions being applied, agents joining or leaving, etc. A social capital value (between zero and one) is stored for each institution; this value represents the aggregated value of all the events involving a particular institution.

Implementation

Figure 3.8 shows the UML diagram of the main classes implemented to model the forms of social capital. The interface *SocialCapitalForm* defines all the methods that any form of social capital must implement, allowing the addition of forms in the future. The methods *accept* and *update* implement the update functions (γ) defined in section subsection 3.2.2. Furthermore, the method *getMetrics* provides the indicators (μ). One class for each form of social capital is also shown, where the attributes are defined (δ), and the appropriate methods to handle them. These methods are not shown in the diagram for clarity. Table 3.2 shows a simplification of the data types used in each form of social capital.

Table 3.2: Data types used in the different forms of social capital

Form of social capital	*Data types*
Trustworthiness	HashTable < UUID, HashTable < UUD, Double > >
Networks	Graph: - Vertex < UUID > - Edges < UUID, UUID, Double >
Institutions	HashTable < UUID, Double >

Event Handlers

When a social capital event is feed into the framework, the social capital attributes are updated according to a weight specified in each event. The networks and institutions forms of social capital use the equation introduced in section 3.2.2.

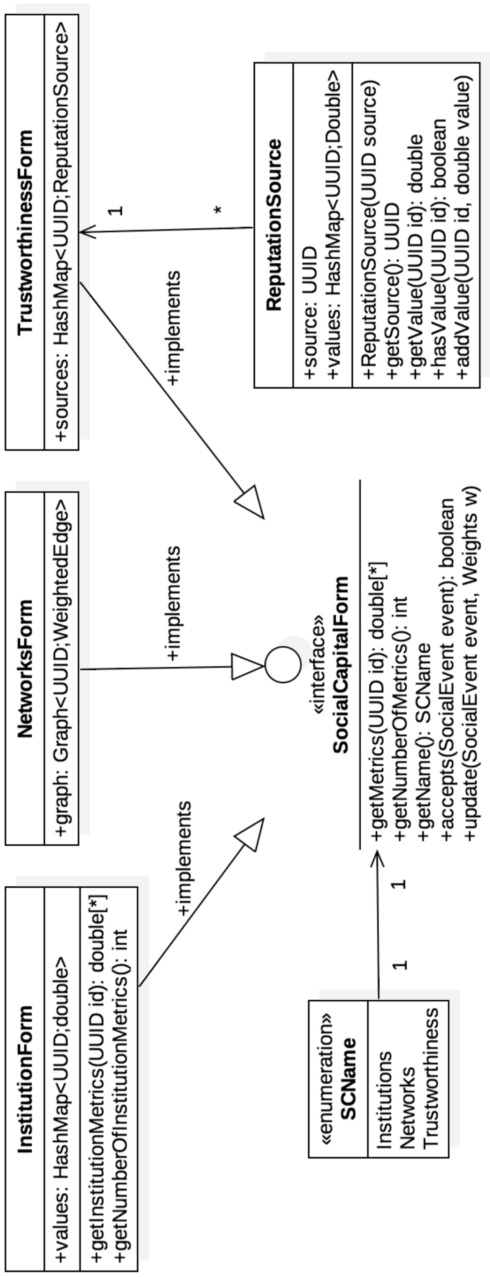

Figure 3.8: Forms of social capital UML class diagram

The trustworthiness form of social capital does not use this function because the new reputation event contains the updated value of the agents' reputations. In this case, the new reputation value replaces the old one stored within the appropriated source.

EXAMPLE When a sanction event with a weight $w = 0.00005$, is fed to the institutions form of social capital, the current $sc_i^{val,t}d$ is retrieved from the HashMap using the *id*. This value is set by default at 0.5. The new value $sc_i^{val,t+1}d$ is calculated applying the previous formula which results in $sc_i^{val,t+1}d = 0.50025$. This value is stored into the HashMap and will be used as one of the indicators provided by the institution *id* in the social decision module.

3.4.3 *Social decision module*

Given an agent, the generic model combines n social capital indicators from the forms of social capital to make a social decision, i.e. cooperate. In different contexts, a distinct set of indicators might be enabled. Formally, the cooperation decision uses the social capital $\mathcal{SC}(a)$ associated with an agent a as:

$$\mathcal{SC}(a) = \sum_{i=1}^{n} w_i v_i$$

Here, w_i is the weight of each social capital indicator (normalised such that they add up to 1) and v_i is the social capital value of each indicator. This value allows the agent to decide whether to cooperate or not with another agent.

Similarly, to compute the social capital $\mathcal{SC}(i)$ associated with a given institution i, the social decision module uses:

$$\mathcal{SC}(i) = \mathit{institution}(i)$$

where $\mathit{institution}(i)$ is the social capital indicator associated to i by the institutions form of social capital. The value of social capital of i then allows the agent to make decisions about the institution – for instance, whether to join it, leave it, recommend it to another agent, etc.

Implementation

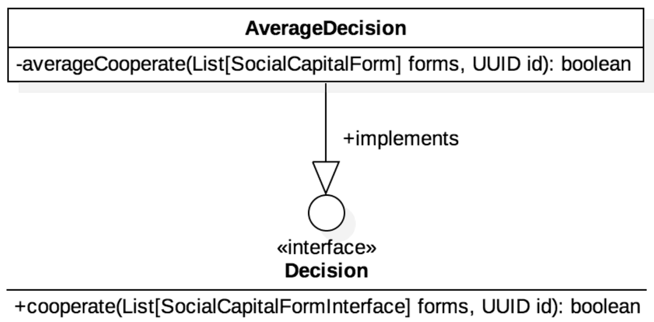

Figure 3.9: Decision-making module UML class diagram

Figure 3.9 shows the UML diagram of the main classes implemented for the initial decision module. The interface *Decision* defines the *cooperate* method that any decision-making class must implement, allowing to add new decision-making methods in the future. The class *AverageDecision* implements the methods and it calculates the average of all the indicators produced by the forms of social capital. It returns 'true' when the average is equal to or greater than 0.5.

EXAMPLE The decision module with the three forms of social capital described before might get three indicators: one from *Trustworthiness* – if only one reputation source is included, one from *Networks* – from the edge connecting the agent, and one from *Institutions* – if the agent participates in only one institution. Using the this implementation of the decision module, the three indicators received will be combined by a simple average, and the value used to decide to cooperate or not.

3.4.4 *Brief note on complexity*

The ESC framework implementation uses *online* learning through event handlers that update the social capital attributes. The weights used in this implementation are either constant or linear. Furthermore, the decision-making process combines the values of different indicators using a linear function. Therefore, the complexity of social decision-making in the social capital model (this implementation) is linear; and in general is linear in the number of attributes and independent of the number of players and their strategies.

3.5 SUMMARY

Based on the notion of social capital borrowed from social systems, a computational framework to represent and reason about its electronic equivalent was specified and implemented in this chapter.

To summarise, this chapter provides the following contributions:

- Presents a framework to represent and reason with social capital. The framework decouples event processing and the updating of the social capital information from the decision-making, thus providing a modular architecture to implement it.
- Provides one of many possible implementations of the framework. This indicates that the framework is generic and can be extended to include other forms of social capital.

The importance of the framework not only lies in the performance of the agents using social capital, but also in the fact that social capital itself can be a short-cut for reducing the complexity in the agent decision making for repetitive strategic games. It has been shown that computing the optimal strategy (i.e. Nash equilibrium) is computationally complex (Daskalakis et al. 2006). Thus, even if we assumed that agents have infinite computational power, it might not be feasible to compute it. On the other hand, the social capital framework presented in this chapter does not require complex computations, involving only straightforward updates in some data structures and aggregation using simple methods (e.g. averages and comparisons). Even in the case of the agents having pre-computed the Nash equilibrium, only needing to consult the best action, this framework would not be more complex (if at all) than such look up.

We contend that this approach to self-organised social arrangements would be equally beneficial for decision-making in the 2-player games, n-player games and public goods games. In the following chapters, experiments using those games will be presented. Different players trying to exploit the games by being not cooperative (Nash equilibrium, dominant or free-rinding strategy) will play against players using the electronic social capital framework in a cooperative strategy.

4

EXPERIMENTS IN 2-PLAYER GAMES

4.1 INTRODUCTION

A common requirement of distributed multi-agent systems is for the agents themselves to negotiate pairwise agreements on performing a joint action. In systems with endogenous resources, the *cost* of computing the decision-making has to be taken into account. If the computational resources expended in negotiating an optimal solution exceed the marginal benefits gained from that negotiation, then it would be more expedient and efficient to use the *memory* of past interactions to short-cut the complexity of decision-making in joint or collective actions of this kind. In social systems, it has been observed that social capital is an attribute of individuals that enhances their ability to solve collective action problems.

In this chapter, the Electronic Social Capital framework is used to decide whether to cooperate or defect in a strategic game. An experimental testbed was implemented, in which a population of agents repeatedly played pairwise games, and used the social capital framework as the basis for their action-selection.

4.2 TESTBED: COOPERATION GAME

The Cooperation Game (CG) is a strategic game where a population of agents is repeatedly randomly paired to play a game against each other. At each round, each player has a randomly designated opponent and a two-player strategic game to play. Table 4.1 shows the four pairwise games selected and their payoff matrix. Once paired, players must choose either to *cooperate*, *defect* or *refuse to play*. Then, the payoff matrices are applied and the agents receive or lose points depending on what they have played. If one of the players refuses to play, the game is cancelled and agents do not receive or lose any points. A global count of points is kept for all the players, and it is used to evaluate their performance over the time.

	C	D
C	2,2	-2,3
D	3,-2	1,1

Prisoner's dilemma

	C	D
C	1,1	-1,-1
D	-1,-1	1,1

Coordination

	C	D
C	1,1	-1,-1
D	-1,-1	-2,-2

Full convergence

	C	D
C	1,1	-1,-1
D	-1,2	-2,-2

Partial Convergence

Table 4.1: Payoff matrix for 2-player game (CG).

4.2.1 *Step-by-step algorithm*

Algorithm 4.1 shows the step-by-step procedure of the CG. At the beginning of each round, random pairs of agents are generated and a specific pairwise game is selected for each of the pairs. We called this a Match. Then, players select the action they want to play in that scenario (lines 9-11). At the end of the round, the players' actions are grouped with their matches and the outcome is calculated (line 13). Here, players will gain or lose points based on their actions and the game's payoff matrix (line 14). Institutions apply their sanctions based on the match results (line 15); agents that violate the rules of the institutions they are members of get sanctioned. Lastly, new institutions are created and the members of the institutions are updated, all based on agents' requests during that round. This process is repeated for every round until the end of the simulation.

4.2.2 *Institutions*

Another feature of the CG is institutions. They define how agents should play in a certain game and apply sanctions to theis members when they misbe-

```
A ← set of n agents
G ← set of m games
t ← 0
repeat
    p ← generate_random_pairs(A, G)
    for each pair (i, j, g) ∈ p do
        create match m_ij(g)
    end for
    for each agent i ∈ A do
        play action_i
    end for
    for each pair (i, j, g) ∈ p do
        mr_ij ← match_result(g, action_i, action_j)
        update_points(mr_ij)
        institutional_sanctions(mr_ij)
    end for
    create new institutions
    update institution membership
    t ← t + 1
until t == T_lim
```

Algorithm 4.1: Cooperation game

have. For example, an institution rule could be that agents playing prisoner's dilemma against other members of the institution should cooperate. If an agent defects, the institution will sanction it losing a stipulated amount of points. Players can create, join and leave institutions while playing the game. They can also invite others to join the ones they are currently members of.

Each institution has a ruleset (see section 2.8), which has the following Degrees of Freedom (DoF):

- *Vote to join*: Members have to vote to allow a new member to join an institution.
- Applies *only between members*: Institutional rules only apply when both players are members.
- *Graduated sanctions*: Sanctions increase for recurring rule violators. Being expelled from an institution is the highest sanction.
- *Points transferred*: When a sanction is applied, the player being defected receives the points taken from the sanctioned player.
- Number of *points sanctioned*: This refers to the number of points rule violators are sanctioned by an institution.

Some of Ostrom's design principles are included in this customisation of an institution, mainly those principles regarding monitoring, sanctioning and membership (Ostrom 1990).

4.2.3 *Reputation sources*

Reputation was modelled using "direct interactions between participant members of the system" (see subsection 3.2.3) and has been previously defined as 'Direct Interaction Reputation'. Agents keep a percentage of successful cooperation for all the other agents and it is updated after each game is played. By 'successful cooperation' we mean that the other agent has played *cooperate*. When we refer to 'no cooperation', we mean it has played *defect*. If the opponent agent chooses to play *refuse to play*, the reputation is not updated.

EXAMPLE If agent 1 and agent 2 have played 99 times against each other, where agent 1 played 33 times *cooperate*, 33 times *defect* and 33 times *refuse to play*, agent 2 reputation of agent 1 would be calculated as $\frac{1 \cdot 33 + 0 \cdot 33}{33+33} = 0.5$

4.2.4 *Social players*

Social players are agents who participate in the cooperation game and have some form of social capital included in their decision-making. Algorithm 4.2 describes their behaviour in a round of the simulation. When a round starts, players receive a random and limited number of results for the other players' matches (i.e. each agent will receive a different set of results). We use this to model agent observation, communication with other agents, or any form of publishing results that will deficiently spread the match results to other agents. Social agents update their social capital with this information (lines 5-7). Subsequently, they again update their social capital with all the information received by the institutions they are members of (lines 8-10). All institutions send information about who joined, left, was sanctioned, was rewarded or was expelled. In addition, if an institution called a vote to accept a new member, the vote is sent at this point. Next, the player has a probability q of creating an institution (line 11). The configuration of the institution is randomly generated based on the five customisable aspects explained before. Two institutions with the same characteristics are not allowed. With a probability r and with a probability s, social players will join or leave a random institution respectively (lines 12-13). In these cases, a Boltzmann distribution of the institutions based on their social capital value is used to choose one (when choosing which one to leave the value is inverted using $1 - value_{sc}$). The next step is to process the invitations to join institutions. In order to decide whether to accept the invitation or not, players check the social capital of the institution and of the player who sent the invitation. If these values are greater than a certain threshold, the invitation is accepted (lines 14-18). Following this, the agent checks the social capital of each institution it is a member of. If the value is lower than a threshold, the agent will leave the institution (lines 19-23). Afterwards, the current opponent is retrieved and the action to play is chosen according to this opponent's social capital. If the opponent's social capital value is lower than a threshold,

```
E ← set of n institutional events
R ← set of m other match results
I ← set of o institutions joined by the agent
J ← set of p invitations to join institutions
for each match result mr ∈ R do
    create social interaction si_mr
end for
for each event e ∈ E do
    create social interaction si_e
end for
create new institution, with probability q
join random institution, with probability r
leave random institution, with probability s
for each invitation j ∈ J do
    if accept(j) then
        join j
    end if
end for
for each institution i ∈ I do
    if not cooperate(i) then
        leave i
    end if
end for
opp_t ← current_match_opponent
if cooperate(opp_t) then
    play cooperate
else
    play refuse_to_play
end if
if opp_{t−1} played cooperate then
    send invitation to my institution
    create social interaction si_coop
end if
if opp_{t−1} played not_cooperate then
    create social interaction si_not_coop
end if
```

Algorithm 4.2: Social player in CG

social players refuse to play against this opponent (lines 25-29). Lastly, they receive the information about the last match's results and they update their social capital accordingly (lines 30-36).

4.2.5 *Probabilistic players*

Probabilistic players also use social capital but, even if the social capital advises them to cooperate, they play defect in a pre-defined percentage of rounds. Algorithm 4.2 also describes their behaviour in a round of the simulation before the percentage of defects is calculated. To decide if the agent will defect in the current round, a random number is generated and compared to the threshold of rounds for which this particular agents defects. If the number generated is lower than the threshold, the social capital is ignored and the action chosen is *defect*. However, if the value is greater than or equal to the threshold, the action chosen by the social capital is performed.

4.2.6 *Equilibrium players*

Equilibrium players choose their actions based on the Nash equilibrium at each game. They *defect* at prisoner's dilemma and coordination games, and *cooperate* at full and partial convergence.

4.2.7 *Random players*

Random players participate in the CG ignoring the payoff matrices of the game and without implementing any social capital. These players choose to *cooperate*, *defect* or *refuse to play* according to random selection.

4.3 EXPERIMENTS

In order to evaluate the social capital framework, the cooperation game defined above was tested with *social*, *probabilistic*, *random* and *equilibrium* players. As explained in the previous section, social players are agents that implement any

Player's name	Forms of social capital	% of defects
SocialPlayer-TNI	Trustworthiness, Networks, Institutions	0
SocialPlayer-NI	Networks, Institutions	0
SocialPlayer-TI	Trustworthiness, Institutions	0
SocialPlayer-TN	Trustworthiness, Networks	0
SocialPlayer-T	Trustworthiness	0
SocialPlayer-N	Networks	0
SocialPlayer-I	Institutions	0
Probabilistic-TNI-0.1	Trustworthiness, Networks, Institutions	10
Probabilistic-T-0.25	Trustworthiness	25
Probabilistic-N-0.25	Networks	25
Probabilistic-I-0.25	Institutions	25
Probabilistic-TNI-0.25	Trustworthiness, Networks, Institutions	25
Probabilistic-TNI-0.5	Trustworthiness, Networks, Institutions	50
Probabilistic-TNI-0.75	Trustworthiness, Networks, Institutions	75
RandomPlayer	NA	50
EquilibriumPlayer	NA	NA

Table 4.2: Different players participating at CG

form of social capital. In this set up, all the possible combinations of agents using one, two and the three forms of social capital were used. The social capital decision module used for *social* and *probabilistic* agents is the average of each active form of the agents' social capital.

Random players arbitrarily choose their actions. Equilibrium agents play the Nash equilibrium action in each game: they defect in prisoner's dilemma and coordination games, and cooperate in full and partial convergence. Equilibrium and random players do not participate in any institution and do not implement any form of social capital. Table 4.2 shows the characteristics of each of the players.

The simulation has been populated with 10 agents of each type, creating a total of 160 agents. The average values of ten simulations have been used for the results in this section.

4.3.1 *Experiment 1: Prisoner's dilemma*

In this setup, our first experiment was run the simulation where agents always play prisoner's dilemma. The results are shown in Figure 4.1. The plot shows the evolution of the average points achieved by different type of agents (for the sake of clarity, only the type obtaining the most points is shown, as well as the *equilibrium* and *random* agents). On the right-hand side of the figure, the average of the final score for each type of agent in the last round of the simulations is shown, as well as the percentage of points w.r.t the best score. It can be seen that, at the beginning *equilibrium* players outperform the rest. Due to the lack of interaction between agents, the forms of social capital do not have any information as yet, and *social* and *probabilistic* players get defected by *equilibrium* players. As the game evolves, each form of social capital starts collecting information and feeding their decision modules. After this, players with any form of social capital start refusing to play against the *equilibrium* players. At around round 8000, the *social* player with trustworthiness, networks and institutions forms of social capital outperforms the *equilibrium* player. *Random* players underperform all other players, but one. A huge difference can be seen between *social* players using any form of social capital and the rest of the players.

4.3.2 *Experiment 2: Full convergence*

Figure 4.2 shows the results when playing full convergence game at each round. *Social* and *equilibrium* players equally play cooperate in this scenario, and their scores are similar. The difference between them is that social capital prevents *social* players from being defected by *probabilistic* agents. Thus, when refusing to play against some *probabilistic* agents, *social* players minimise the amount of lost points. The damage control policy enforced by social capital helps agents obtain better results. *Random* players get negative results and are not shown in the figure.

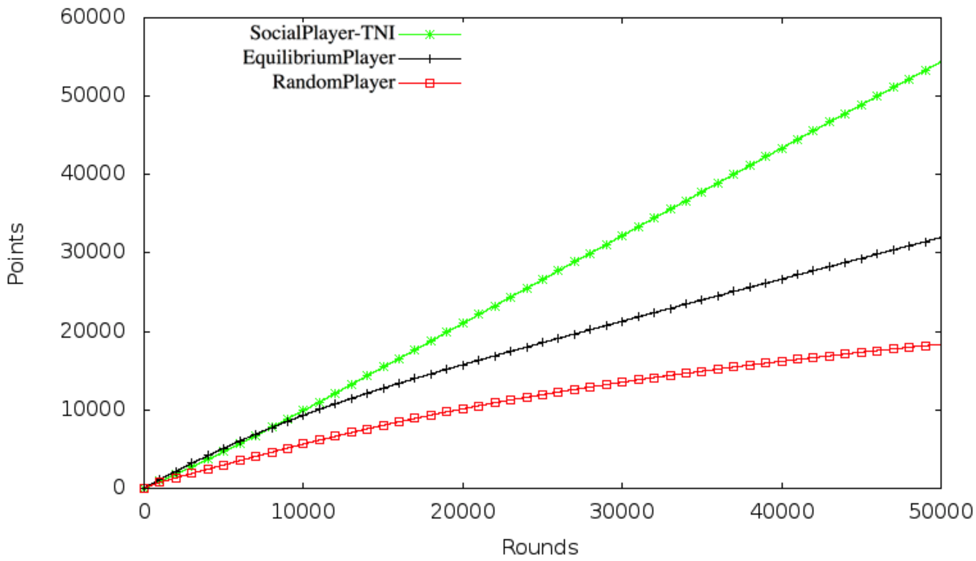

Agent Type	Score	%
SocialPlayer-TNI	**54404**	**100**
SocialPlayer-NI	53118	97
SocialPlayer-TI	52715	96
SocialPlayer-TN	47871	87
SocialPlayer-N	47754	87
SocialPlayer-T	46393	85
Probabilistic-TNI-0.1	45735	84
SocialPlayer-I	45216	83

Agent Type	Score	%
EquilibriumPlayer	**31991**	**58**
Probabilistic-T-0.25	30663	56
Probabilistic-N-0.25	30237	55
Probabilistic-TNI-0.25	28811	52
Probabilistic-TNI-0.75	26595	48
Probabilistic-TNI-0.5	25453	46
RandomPlayer	**18381**	**33**
Probabilistic-I-0.25	10927	20

Figure 4.1: Points obtained for players at prisoner's dilemma

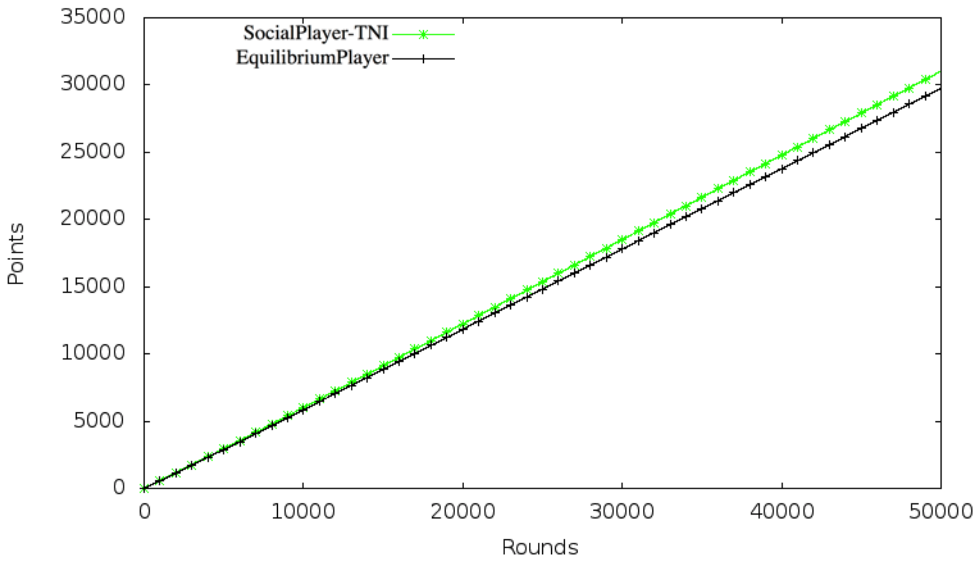

Agent Type	Score	%	Agent Type	Score	%
SocialPlayer-TNI	**31027**	**100**	Probabilistic-TNI-0.1	16238	52
SocialPlayer-TI	30791	99	Probabilistic-T-0.25	4044	12
SocialPlayer-NI	30736	99	Probabilistic-N-0.25	3794	12
EquilibriumPlayer	**29758**	**96**	Probabilistic-TNI-0.25	2083	7
SocialPlayer-I	28935	93	Probabilistic-I-0.25	-7608	-25
SocialPlayer-TN	27170	88	Probabilistic-TNI-0.5	-9478	-31
SocialPlayer-N	27109	87	**RandomPlayer**	**-10913**	**-35**
SocialPlayer-T	26366	85	Probabilistic-TNI-0.75	-17849	-38

Figure 4.2: Points obtained for players at full convergence

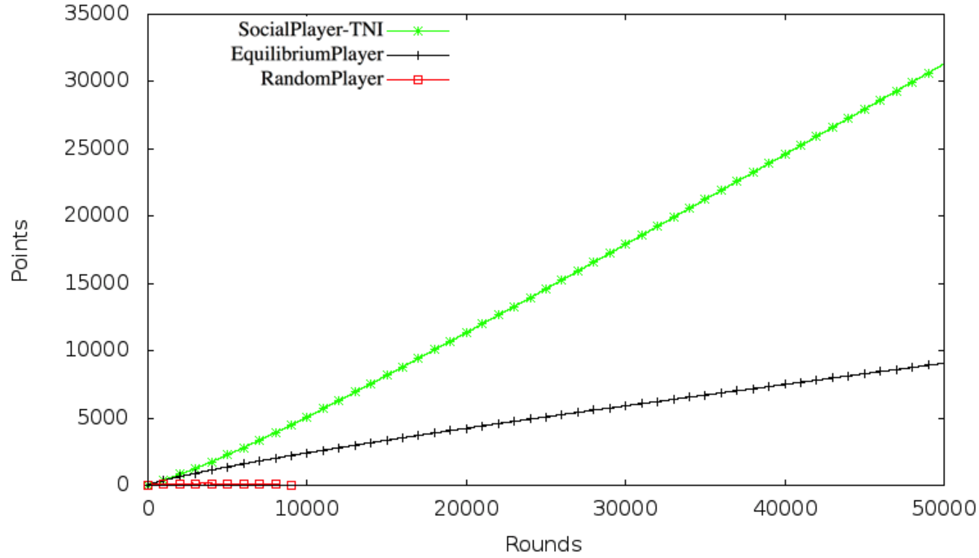

Agent Type	Score	%	Agent Type	Score	%
SocialPlayer-TI	**31318**	**100**	**EquilibriumPlayer**	**9080**	**29**
SocialPlayer-TNI	30242	97	Probabilistic-T-0.25	7097	23
SocialPlayer-NI	29866	95	Probabilistic-N-0.25	6719	21
SocialPlayer-TN	28571	91	Probabilistic-TNI-0.25	5579	18
SocialPlayer-N	28485	91	Probabilistic-TNI-0.5	1160	4
SocialPlayer-T	27350	87	Probabilistic-I-0.25	-285	-1
SocialPlayer-I	24042	77	Probabilistic-TNI-0.75	-589	-2
Probabilistic-TNI-0.1	20826	66	**RandomPlayer**	**-2506**	**-8**

Figure 4.3: Points obtained for players at all four games

4.3.3 *Experiment 3: All games*

In the last scenario (shown in Figure 4.3) the same type of agents as in the previous experiments now play all four pairwise games. At each round, when the opponent is randomly selected, one of the four games is also randomly chosen. When playing multiple games, actions in one game can affect the future interactions in other games. In this scenario, *equilibrium* players defecting in prisoner's dilemma and coordination games affect their ability to play in full and partial convergence. Because *social* and *probabilistic* players were defected by the *equilibrium* players, they refuse to play against them in future interactions, no matter what game they are playing. When the others refuse to play against *equilibrium* players, these agents lose the opportunity to win more points at those games. As with the prisoner's dilemma, in this case, the use of social capital by the *social* players allows them to achieve the best results.

4.3.4 *Experiment 4: Institutional choices*

Figure 4.4 shows the number of members participating in each institution type (as defined by the values of the different DoF) when playing prisoner's dilemma. In this simulation, the cooperation game was populated with 20 agents of each type that uses the institutions form of social capital (only the agents in the figure were at the simulation). The heat map shows the results of one simulation at round 25000. On the one hand, we can see how most *social* players participate in institutions where sanctions (points and graduated) are applied, as they always comply with the rules. On the other hand, *probabilistic* players participate in some institutions where sanction are applied but they all avoid the ones with graduated sanctions, as they repeatedly defect and the sanctions grow each time. Most of the players avoid institutions with no form of sanctioning for defectors.

4.4 VISUALISATION TOOL

Given the large amount of data generated by Presage2 using the Electronic Social Capital (ESC) framework simulations, and in order to make the task of

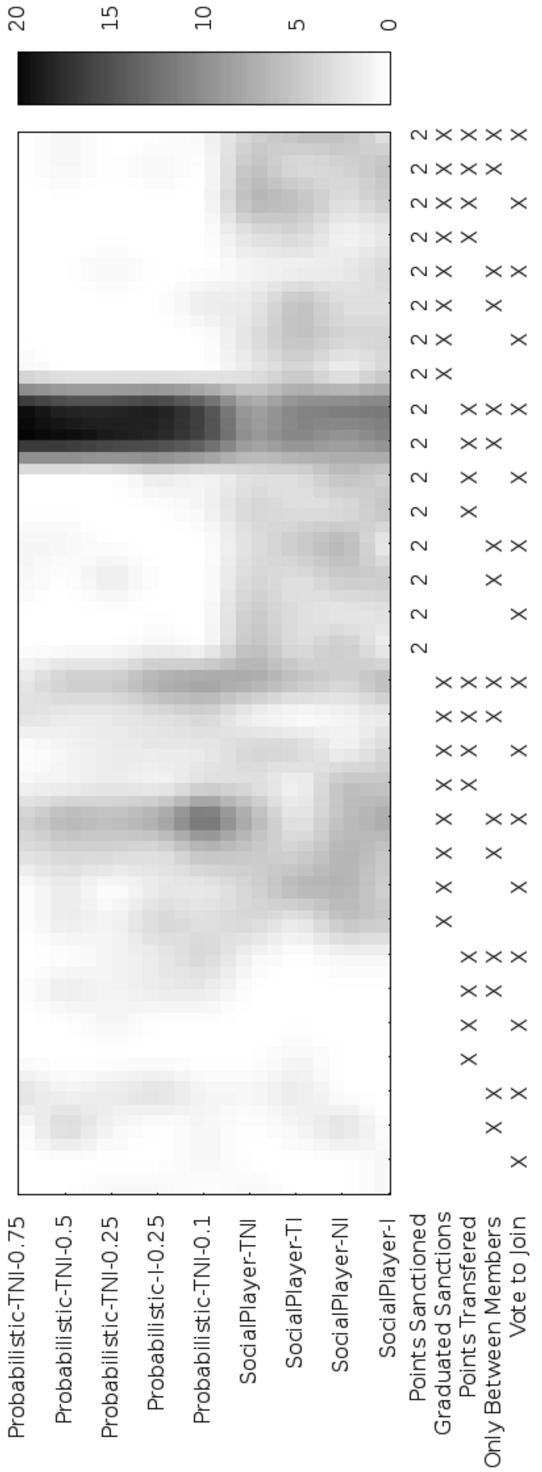

Figure 4.4: Number of members in each institution

analysing and understanding their results simpler, we developed a visualisation tool that allows one to graphically see what happens during each simulation. The tool has been developed so that it can sit on top of Presage2 for online visualisation (i.e. as the simulation is executed) or act as an offline visualiser that uses stored data from already executed simulations. A screenshot of the tool is shown in Figure 4.5.

The visualisation tool allows one to easily see the evolution of several metrics and aspects of the simulations. In particular, the tool we have developed for our ESC framework displays the following:

- Average utility of the different agents' types.
- Social capital accrued by the agents.
- Distribution of the agents among different institutions.
- Interaction and relationships among the agents.

At the top left of the figure, the plot shows the average number of points for the players, grouped according to their type, over the rounds played. The total points can be used to measure the performance of the different players in the simulation. We can observe that players using social capital in their decision-making do better than the others.

At the bottom left, the plot shows the accumulated social capital for the players using the social capital framework. Since the amount of social capital is subjective for each player, the results shown are the average for each player grouped by types. This figure shows how social capital increases among cooperators and decreases among defectors.

The heat map graph at the top right of the figure displays the number of players of each type that are members of each institution. This figure shows how the agents self-organise in different institutions (defined by the features shown in the $x-$axis), and helps us to understand which institutions are more popular among the players and if the different agent types have institutional preferences.

Finally, at the bottom right of the figure, we have a graph where the vertices are all the players in the simulation and the edges are weighted according to

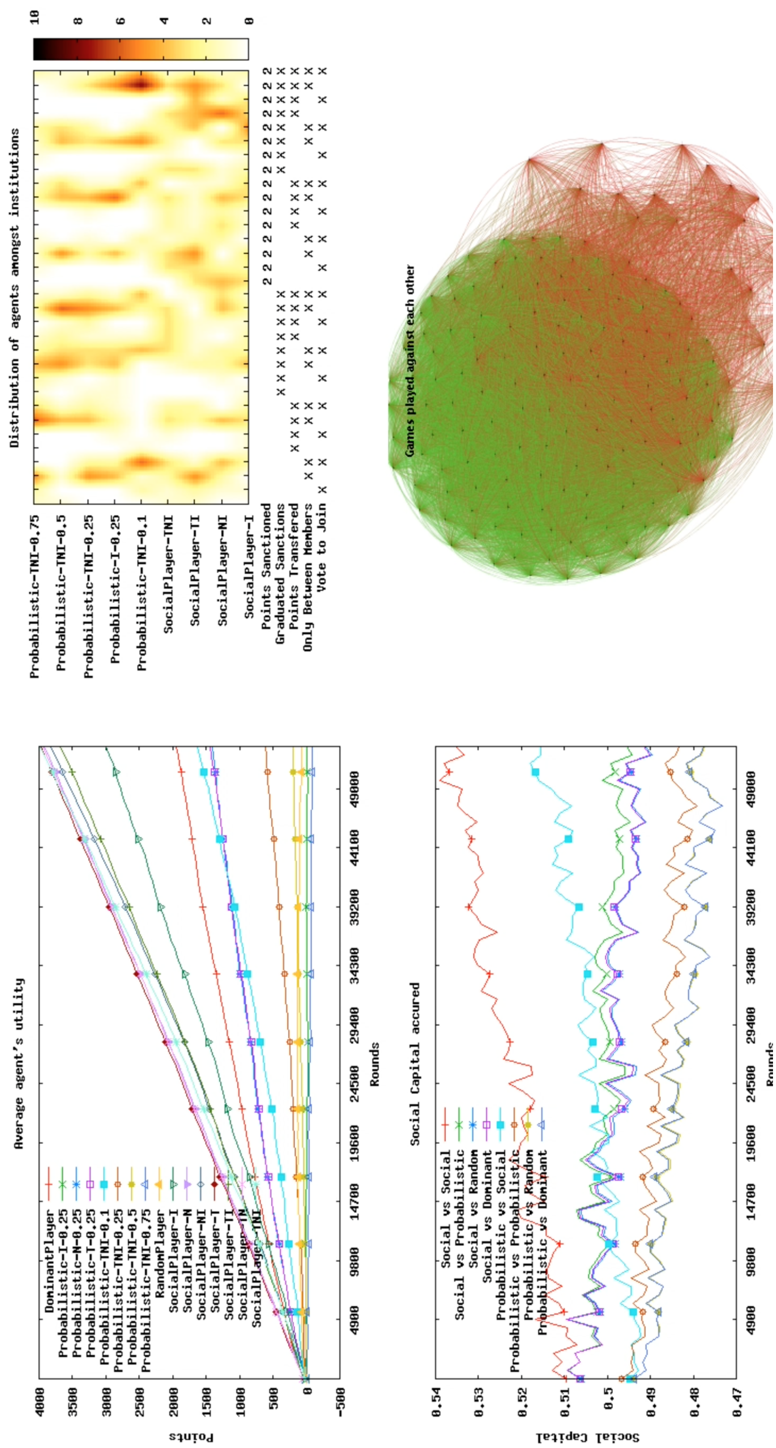

Figure 4.5: Visualising of social capital during a simulation

the number of games they have played against each other. This has been generated using Gephi (Bastian et al. 2009), and the communities were detected using the modularity tool (Blondel et al. 2008). We can observe that two distinctive communities arise, which show how cooperators (i.e. those using social capital, shown in green in the upper left "circle") tend to avoid playing against defectors (i.e. random and equilibrium players, shown in red in the lower left points cloud).

4.5 RELATED WORK

Action-selection for repeated games has been addressed using different approaches, from multi-agent learning (Bowling and Veloso 2002) to using the concept of trust (Mayer et al. 1995). The latter is the closest to our approach, as trustworthiness is actually one of the forms of social capital we have presented. Trust is built by the agents through repeated interactions with each other, and it is used to decide actions such as what agent to interact with, or whether to cooperate with a given agent (Pinyol and Sabater-Mir 2013). This idea has also been incorporated in some game theoretical models for repeated games (Mailath and Samuelson 2006).

While trust is indeed quite an appropriate metric to predict some agent's behaviour, it is based only on the interactions of agents (either experienced first-hand or communicated through reputation values). To include other sources of information relevant to social capital, we have presented a competitive scenario using the social capital framework. Our framework enhances the individual view of each agent by treating trustworthiness as one of the forms of social capital, and complementing it with information regarding the network of the agent, as well as the information coming from institutions.

Cooperative behaviour can also be achieved through repeated interactions and the agents' desire to avoid retaliation threats in the case of non-cooperative behaviour (Myerson 1997). In this case, however, cooperation does not arise from a willingness to do so, but rather from the objective of reducing the probability of being punished.

Our social capital framework also supports Axelrod's findings in his work *The evolution of Cooperation* (Axelrod and Hamilton 1981). Where 'success' in an iterated 2-player game can be connected to a strategy showing the following characteristics: never be the first to defect; reciprocate (return cooperation for cooperation and defect for defect); maximise its own score (not be envious); and clarity (don't be tricky). Our social agents implement those characteristics by cooperating at their first interaction, stop cooperating after being defected, having their internal utility function which does not include other players rewards, and not trying to maximise their utilities by ticking other players.

Norms can also lead to cooperative behaviour. Norms can be learnt from repeated social interactions using (e.g. reinforcement learning) (Sugawara 2011) or through aggregation techniques (e.g. ensemble methods) (Yu et al. 2013). After a sequence of interactions, agents might learn that cooperation is beneficial. This can then be explicitly stated as a norm, or implicitly internalised by the agents as a social convention. In the case of explicit norms, the compliance to them could be incorporated in the update of the social capital related to institutions.

4.6 SUMMARY

In this chapter, the first experiments using the ESC framework in 2-player games were reported.

The results show that the use of social capital benefits the agents playing against Nash equilibrium strategies. One of the main effects of using social capital is the facilitation to achieve win-win situations, where two agents involved in a pairwise interaction benefit from behaving cooperatively.

It has also been observed that social capital acts as a catalyst for self-organisation: agents decide with whom to interact according to their social capital, and they also use the social capital information to join or leave institutions. Thus, social capital will plays a key role in socio-technical systems where long-term collective goals are only be achievable through cooperation among participants.

In the next chapter, an extension of the current 2-player game to n-player game scenario will be presented, in order to substantiate the claim that the use of ESC does indeed support effective collective action.

5 EXPERIMENTS IN AN N-PLAYER GAME

5.1 INTRODUCTION

In this chapter, further experiments using the Electronic Social Capital (ESC) framework are presented. Firstly, a new testbed for n-player games is described, including the customisations regarding the institutions and reputation sources used. Moreover, the different players implemented for this scenario are specified; four new experiments were carried out with these players. To conclude, some related work is presented and compared to the experiments performed.

5.2 TESTBED: UNSCRUPULOUS DINER'S DILEMMA GAME

A social dilemma is a situation where one individual could benefit from acting selfishly unless everyone acts selfishly, in which case the whole group loses (Glance and Huberman 1994). In a scenario where a group of friends are dining in a restaurant with an unexpressed agreement to equally divide the bill, one person could order an expensive meal and enjoy a magnificent dinner at a bargain price. If everyone at the dinner does so, however, they will all end with an astronomical bill to pay. This situation is called the Unscrupulous Diner's Dilemma (Gneezy et al. 2004).

A detailed description of the Unscrupulous Diner's Dilemma is as follows. Several individuals go out to eat and are faced with the decision of what meal they would order: either expensive or inexpensive. Each meal has a price and the joy of eating associated with it. The price of each meal is denoted by p_e and p_i and the joy of eating is denoted by j_e and j_i respectively. Later, the bill is equally split between the participants at the meal. Let the total cost of the other people's meals be c and the number of participants at the meal be m. Then, the cost for ordering an expensive meal is $(c + p_e)/m$ and for an inexpensive meal is $(c + p_i)/m$. The utilities for each meal are $j_e - (c + p_e)/m$ for the expensive meal and $j_i - (c + p_i)/m$ for the inexpensive meal. Given the premise that $j_i - p_i > j_e - p_e$, everyone would be in a better situation ordering an inexpensive meal.

The problem is that individuals' strategies are arbitrary. One could order an expensive meal, obtaining the joy associated, and the individual price would be diminished by the sharing the bill with the other members of the group. The situation is symmetric for all the participants and it is only beneficial when choose the expensive meal. If all the participants order the expensive meal, each individual's utility is worse than if they ordered an inexpensive meal. With the lack of other information, ordering the expensive meal is the strictly dominant action – which means that ordering always the expensive meal is better than another strategy for an individual player, no matter how that player's opponents may play. The challenge, then, is for the participants to decide with whom they would prefer to eat together in order to maximise their individual utilities.

The Unscrupulous Diner's Dilemma Game (UDDG) is a strategic game where a population of agents is repeatedly faced with the decisions to choose with which group they would have a dinner and which meal they would order (either expensive or inexpensive). Table 5.1 shows the price, joy and satisfaction for each meal. These values were chosen respecting the constraint denoted before. Once the agents have chosen the group, the meal the payoff matrix is applied. A global count of money spent and joy acquired is kept for all the players and it is used to evaluate their performance over the time. Note that the satisfaction is depicted as the joy divided by the price.

	Price	Joy	Satisfaction
Inexpensive	1	1	1
Expensive	2	1.5	0.75

Table 5.1: Payoff matrix for n-player game (UDDG)

5.2.1 *Step-by-step algorithm*

```
1:  A ← set of n agents
2:  D ← set of d dinner_groups
3:  t ← 0
4:  repeat
5:      for each agent i ∈ A do
6:          choose dinner_group_i
7:      end for
8:      for each agent i ∈ A do
9:          order meal_i
10:     end for
11:     for each d ∈ D do
12:         bill_id ← create_bill(d)
13:         update_spent(bill_id)
14:         update_joy(bill_id)
15:         institutional_sanctions(bill_id)
16:     end for
17:     update institution membership
18:     t ← t + 1
19: until t == T_lim
```

Algorithm 5.1: Unscrupulous diner's dilemma game.

Algorithm 5.1 shows a step-by-step procedure of the game. At the beginning of each round, agents must choose in which group dinner they will participate. Then, agents select the meal they want to order in that group (lines 8-10). At the end of the round, each bill is calculated (line 12). Here, players will spend money and gain joy based on their actions and the game's payoff matrix (lines 13-14). Institutions apply their sanctions based on the bill (line 15); agents that violate the rules of the institutions they are members of get sanctioned. Finally, the members of the institutions are updated, all based on the agents' requests during that round. This process is repeated for every round until the end of the simulation.

5.2.2 *Institutions*

Another feature of the game are institutions. They define how agents should play and apply sanctions to their members when misbehaving – i.e. an institution's rule could be that agents playing against other members of the institution should order an inexpensive meal. If an agent orders an expensive meal, the institution will sanction it with a fine. Players can create, join and leave institutions while playing the game.

The system also includes a set of institutions with different characteristics. Each institution has a ruleset (see section 2.8) that has the following boolean Degrees of Freedom (DoF):

- *Vote*: This sets whether members have to vote for a new member to join the institution or to expel a member from the institution.
- *Compulsory cooperation*: This sets whether choosing an inexpensive meal is mandatory or not.
- *Sanction*: This sets whether the rule violations are sanctioned or not.
- *Graduated sanctions*: This sets whether sanctions increase for recurring rule violators or not. Being expelled from the institution is the highest sanction.

Some of the Ostrom's design principles are included in this customisation of an institution, mainly regarding monitoring, sanctioning and membership (Ostrom 1990).

5.2.3 *Reputation sources*

The reputation was modelled using "the information provided by other members about the interactions they had in the past." (see subsection 3.2.3). The two reputation sources used are:

- *Agents' Reputation* which is based on agents' individual experiences. After each meal, each agent reports if the meal was successful or not for all the participants at that meal. A meal is successful if the individual price paid by the agent (after the bill is split) is equal to or lower than the price of the meal it ordered. This feedback is aggregated into a value that shows the percentage of satisfaction when eating with an individual agent.

- *Institutional Reputation* and is obtained from the interaction with institutions. Again, after each meal, institutions submit if the member agents participating in the institution followed its rules. This value is also aggregated into a percentage of successful behaviour when participating in an institution.

The implementation of the reputation reveals an agent behaviour from the perspective of other agents and the institutions, it does not show how cooperative an agent is. For example, an agent ordering an expensive meal may report positive feedback about another agent ordering an expensive meal, because its individual satisfaction will be at worse-case scenario (all participants ordering expensive meals) as eating the expensive meal alone. The same situation can occur in an institution where no sanctions are applied and cooperation is not compulsory. The institution's feedback will be positive, because ordering an expensive meal is not considered an institutional misbehaviour.

5.2.4 *Social players*

Social players are agents who participate in the game and have some form of social capital included in their decision-making. Algorithm 5.2 describes their behaviour in a round of the simulation. When a round starts, players receive the bill for the last meal they participated in and they update their social capital with this information (line 5). This information is also partially reported to the system's reputation. Subsequently, they again update their social capital with all the information received by the institutions they are member of (lines 8-10). All institutions send information about who joined, left, was sanctioned or was expelled. In addition, if an institution called a vote to accept a new member or expel a current member, the vote is sent at this point. Next, the player has a probability r of joining a random institution (line 11). In this cases, a Boltzmann distribution of the institutions based on their social capital value is

```
E ← sequence of events
O ← set of o orders from last bill
I ← set of i institutions joined by the agent
for each order o ∈ O do
    create social capital event sc_o
    report reputation, with probability q
end for
for each event e ∈ E do
    create social capital event sc_e
end for
join random institution, with probability r
for each institution i ∈ I do
    if not cooperate(i) then
        leave i
    end if
end for
gd_t ← choose_group_dinner
if cooperate(gd_t) then
    order inexpensive meal
else
    order expensive meal
end if
```

Algorithm 5.2: Social player in UDDG used to choose one (when choosing which one to leave, the value is inverted using $1 - value_{sc}$). Following this, the agent checks the social capital of each institution of which it is a member. If the value is lower than a threshold, the agent will leave the institution (lines 12-26). After this, the current group dinner is chosen by selecting the one organised by the institution (it is member of) with the higher social capital. Finally, the agent combines the social capital of the participants at the dinner by calculating the average value. If the value is higher than a threshold, the agent orders the inexpensive meal. If not, the expensive meal is chosen.

5.2.5 *Social capital events*

The social capital events implemented are: regarding players' actions, *Cooperated* and *Not Cooperated* and several events related to institutions. These include an agent joining or leaving an institution (*Joined*, *Left*), an institution expelling or sanctioning an agent for misbehaving (*Expelled*, *Sanctioned*) and, an institution not sanctioning an agent where it should have (*Not Sanctioned*). Consequently, the 'update functions' are applied to the three forms of social capital.

To be able to analyse agents using one, two or the three forms of social capital (for example, only networks or networks and institutions), a flexible function that checks which forms of social capital are active in a specific agent was implemented. Once the values are retrieved, this function calculates the cubic average to combine the output metrics into a single output value.

5.2.6 *Dominant players*

Dominant players always choose to order an expensive meal. In the UDDG, the expensive meal is strictly dominant and the unique Nash equilibrium.

5.2.7 *Random players*

Random players participate in the UDDG by ignoring the payoff matrix of the game and without implementing any social capital. These players choose an *expensive* or *inexpensive* meal by random selection.

5.3 EXPERIMENTS

In order to evaluate the Electronic Social Capital framework in n-player games, the UDDG defined above was used with different types of agents. This included, in particular, the following agents: *social*, *random* and *dominant* players. As explained in the previous sections, *social* players are agents that implement any form of social capital. In this set up, all the possible combinations of agents using one, two and the three forms of social capital were used.

Four different experiments were undertaken with the following purposes:

- Experiment 1: Investigates the performance of social capital in a heterogeneous population of agents, with an equal number of each type.
- Experiment 2: Investigates the effect of scale, in terms of the size of the population, on the performance of social capital.
- Experiment 3: Examines the effect of having multiple forms of social capital.
- Experiment 4: Examines the effect of different institutional parameter values.

5.3.1 *Experiment 1: Comparative evaluation*

For the first experiment, the simulation was populated with 50 agents of each type, for a total of 150 agents. The average values of 50 simulations were used for the results in this section.

With this setup, the performance of social capital was investigated in a heterogeneous population of agents, with an equal number of each type. The results are shown in Figure 5.1.

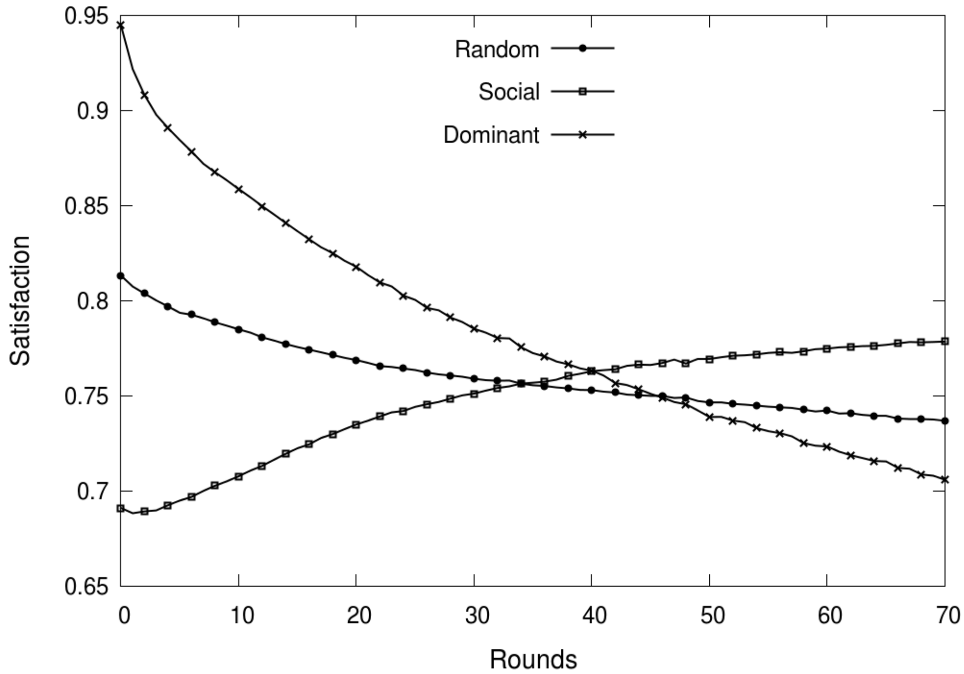

Figure 5.1: Players' satisfaction

The plot shows the evolution of the average satisfaction achieved by the different type of agents. At the beginning, the *dominant* players outperform the rest. Due to the lack of interaction between the agents, the forms of social capital do not have any information, and the *social* players share the cost of the expensive meals ordered by the *dominant* and *random* players. As the game evolves, each form of social capital starts collecting information. Then, *social* players start choosing to group with other *social* players in institutions where sanctions are applied and where repeated violators are expelled. At around round 42, the *social* players outperform the *dominant* players. An important difference can be seen between *social* players using social capital and the rest of the players.

5.3.2 *Experiment 2: Scaling with size of group*

For the next experiment, simulations were run with different number of players but with the same proportion of *social*, *dominant* and *random* players. With this setup, the effect of scale – in terms of size of the population – on the performance of social capital was investigated. Figure 5.2 shows the break points when *social* players outperform (in terms of satisfaction) *dominant* and *random* players. These results show that, as the system scales, the rounds to achieve the breakpoint increase less than linearly.

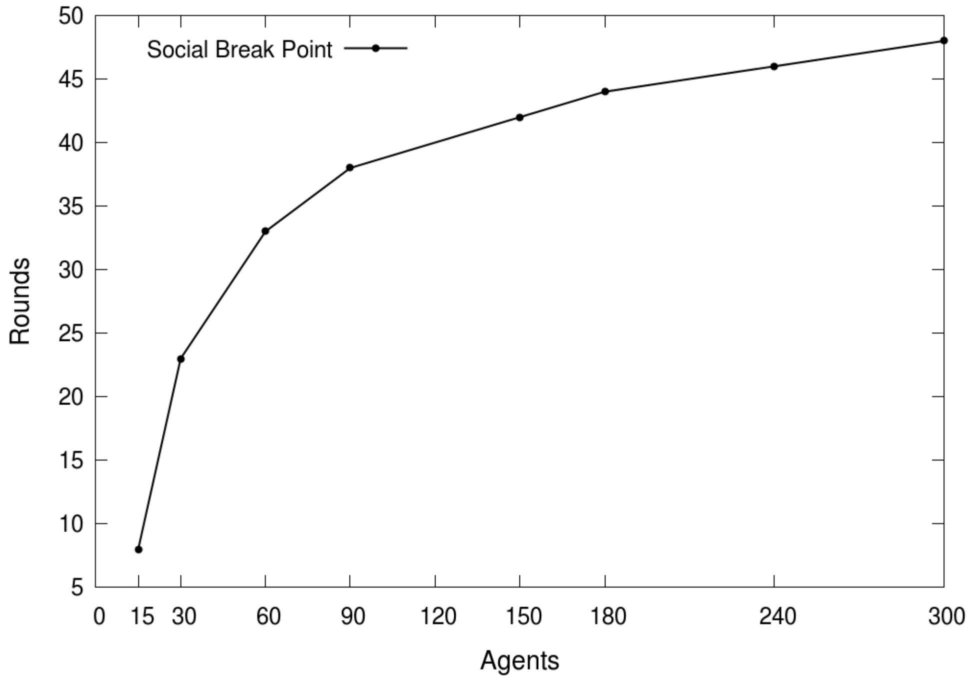

Figure 5.2: Social players' break points

5.3.3 *Experiment 3: Multiple forms of social capital*

In the next scenario, the effect of having multiple forms of social capital was examined. The experimental hypothesis is: it is expected that the more forms of social capital that an agent uses, the more satisfaction it will achieve. In fact, this is *not* what was observed. Figure 5.3 shows the results of *social* players using different combinations of the forms of social capital when playing against *dominant* and *random* players. The average values of 50 simulations were used for the results in this section. The 'T', 'N' and 'I' after the name of the players in the figure, denote the forms of social capital used by the agents (**T**rustworthiness, **N**etworks, **I**nstitutions).

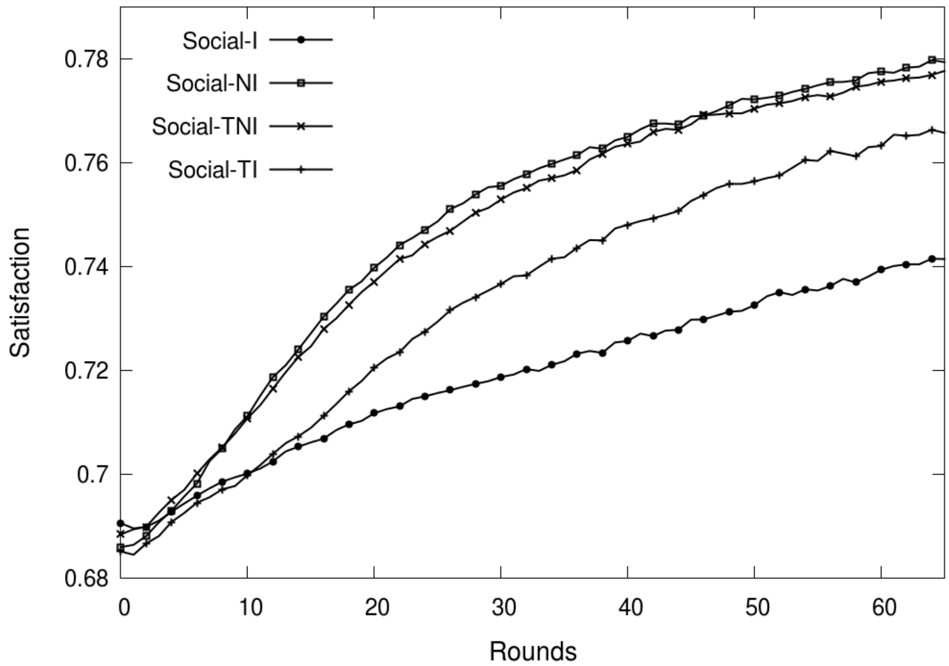

Figure 5.3: Satisfaction for different forms of social capital

In this case, *social* players using networks and institutions perform very similar to *social* players using all three forms of social capital. The lowest satisfaction outcome was achieved by players using only institutions form of social capital. Using the trustworthiness form combined with institutions improved the average satisfaction but not as much as with networks.

The question is, then, if having *many* forms of social capital performs similar to only having *some* forms of social capital, why do we have many? The answer could be that it depends on context, and various combinations of social capital are likely to be better suited, or not, to different social decision-making processes. The issue, then, is to tune each process so that the 'best' social capital indicators are enabled and/or have the appropriate weight.

In some ways, this is not dissimilar to the work on the formal model of legitimate claim for distributive justice (Pitt et al. 2014b). Thus, each process needs to work out which indicators of social capital are relevant (in context), how to accommodate them if there is plurality, and how to reconcile them if there is conflict.

5.3.4 *Experiment 4: Exploring the parameter space*

In the last experiment, the influence of a parameter on the framework was investigated. The parameter chosen was the *weight* of the sanction event. As shown in Figure 5.4, *social* players using only the institutions form of social capital play against *dominant* and *random* players.

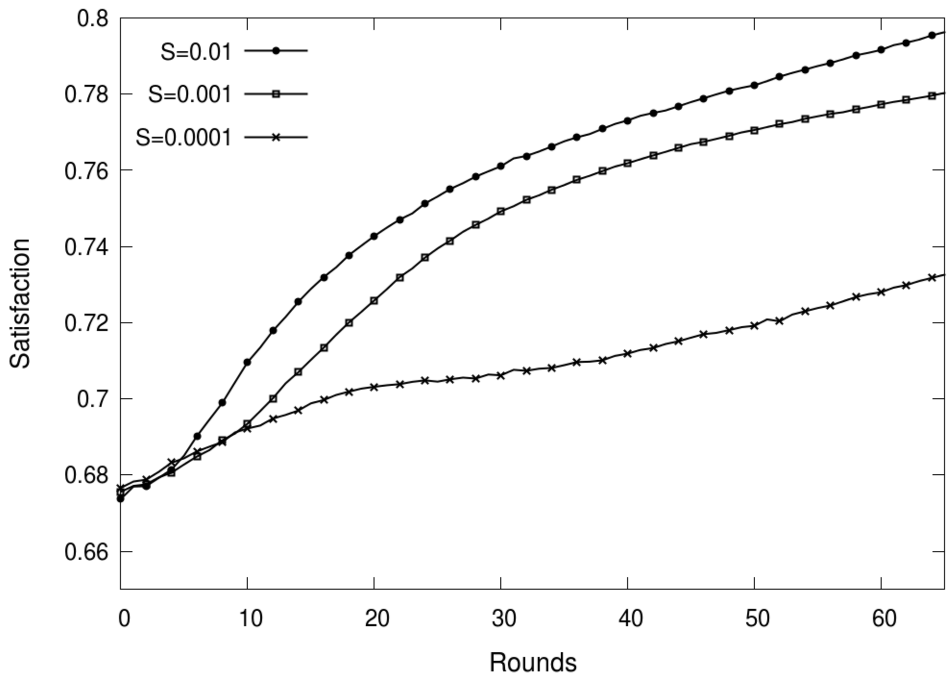

Figure 5.4: Satisfaction with different social capital values

The plot shows the satisfaction of social players when the value of the institution event *sanction* was increased by a 10 or a 100 times. As the value increases, agents need less rounds to spot the 'good' institutions. Having a value too low for this can make other events, such as joining or leaving the institution, too 'shallow' – in this example, sanctioning (when value is 0.0001). If the values are proportional to other events, the results are similar.

5.4 RELATED WORK

The notion of trust is widely studied in the field of multi-agent systems (Pinyol and Sabater-Mir 2013). However, most of those works define trust as a function or a value, which typically computes the probability of a beneficial outcome. The use of trust in this context refers to it as the 'glue' that allows the different forms of social capital to be combined together. Therefore, the notion of trust that is most compatible with that presented here is which presents different situations in which it could (objectively) said that *A trusts B* (Jones 2002). It then identifies two common features of these situations; firstly, *A* has a belief that there is a rule and, secondly, that *A* has an expectation that *B*'s behaviour will conform to that rule. This is what is seen in this framework and the notion that 'trust' is the glue between social capital and (successful) collective action. The belief that there is a rule (or there are rules) is captured by the social capital attribute of institutions, and the expectation of conformance is captured, in effect, by the output of the corresponding decision-making process. This further underlines the point about what social capital *does*: since the UDDG is based on *mutual* trust, then as *B* has the same belief and expectations (except reciprocally concerning *A*), social capital effectively coordinates their expectations.

Furthermore, this framework vindicates for various arguments about forgiveness. Given a definition of trust as a willingness to expose oneself to risk (and coordinating expectations is a de-risking exercise), then the question is: what is to be done when the trust decision is wrong? The common answer in the trust literature is to trash the reputation, which is why the two concepts are found conjoined (as in "trust and reputation"). But what it is seen here is that the reputation (or trustworthiness) is just one element of the social capital attributes which is an input to the trust decision itself, and a wrong decision can update all the different forms of social capital.

In addition, the proper counterpoint to trust is *forgiveness* (Vasalou et al. 2008). In fact, forgiveness – as defined there – could be construed (in the current framework) as another social decision-making process using the social capital indicators as its input (e.g. beneficial historical relationship (Vasalou et al. 2008) – i.e., social capital).

Social capital is defined as in terms of an individual's ability to access and mobilise resources in online communities and social networks (Smith 2011). They use NetLogo to implement an experimental testbed to generate networks, distribute resources, simulate exchange, and compute 'social capital' over time. However, their version of social capital is a measure of an individual's access to social resources at a certain time, and by computation of social capital they effectively mean only a formal metric for a social network which gives a function for computing it. In this work, computational social capital involves all four of: data structures for representing social capital attributes, operations for updating those attributes, metrics for computing social capital indicators from the attributes, and algorithms for decision-making processes which compute decision outputs from those indicators.

It is observed that social capital is a multivariate concept, and, as these experiments have also demonstrated, some variables and metrics may be more significant than others (Daniel et al. 2003). One way to model the interactions and dependencies between the variables constituting social capital is to use a Bayesian Belief Network (Daniel et al. 2003). They use qualitative expert knowledge to generate conditional probability tables which are refined used real world data, in order to measure social capital in virtual communities, rather than use it as a basis for collective action (as in this framework).

5.5 SUMMARY

In the previous chapters, a framework for ESC was specified; based on its definition as "attributes of individuals that help them resolve collective action problems". Experimental results were presented, which showed that this can be indeed the case for 2-player games.

The current chapter builds on these previous chapters but makes the following substantive stand-alone contributions:

- It describes a new experimental testbed, facilitating experiments with large-scale agent populations with a correspondingly 'large' n-player collective action situation.

- It reports three new experimental results, notably that social capital:

– Optimises the outcomes (in terms of long-term satisfaction and utility) of collective decision-making in competitive environments, compared to alternative simplistic strategies.

– Reduces the complexity of that decision-making (the social capital processes are all lower complexity than computing the optimum strategy for n-player and k-strategy games with multiple criteria (Papadimitriou and Roughgarden 2005)).

– Scales with the size of the population (contra (Olson 1965)), because its complexity is independent of population size.

Throughout the experiments, it was observed that social capital facilitates the self-organisation of 'like-minded' individuals into groups, and incentivises those individuals, who may initially be disinclined to conform to the group norms.

In conclusion, it is this self-organisation into groups that may be the most significant outcome of these experiments, especially if we understand the organisation of these groups in terms of *communities*.

It has been argued that the value of communities is that they can resolve certain types of collective action problems, or reap the benefits from other forms of collective action, that are otherwise resistant to purely market-based or (top-down) policy-based solutions (Bowles and Gintis 2002). This might be because communities can leverage relational information, as provided by social capital, in ways that market-based economies, if they are solely reliant on transactional information, cannot.

It is hence predicted that self-organised community systems will be of increasing importance as a mechanism for solving collective action problems in the digital society and achieving satisfactory outcomes for citizens. Electronic forms of social capital will be an essential feature of such systems, along with self-governance and collective attention (Pitt and Diaconescu 2015).

6 EXPERIMENTS IN MULTIPLE CONTEXTS

6.1 INTRODUCTION

For many years, energy systems have implemented a *centralised* structure to deliver electricity from power stations to their customers. It should be noted, however, that the outbreak of renewable energy is changing the paradigm in electricity networks by including new participants, and allowing each of them to become prosumers (producers + consumers). This new paradigm will enable the *de-centralisation* of the grid into smaller *community energy systems* that will interact and will be codependent on each other.

In the absence of central authority, systems rely on *self-organisation* to coordinate their behaviour and agree on the expectations of others. There are many scenarios where humans have self-organised for the common good. Based on extensive field work in fisheries, forests and grazing land, Elinor Ostrom identified eight design principles for successful common pool management (Ostrom 1990). In all these occurrences, participants united in what she called an *institution* and defined a set of rules which regulated and constrained the provision and appropriation of the resources.

When the common pool resource is too large to be managed locally, Ostrom proposed a layered structure of nested enterprises, with small common pool resource at the base level. In further work (Ostrom and Ahn 2003), they proposed that *social capital* created in one institution could enable other institutions to succeed, when these institutions are codependent.

In the previous chapters, it has been shown that formally representing and reasoning with Electronic Social Capital can reduce the complexity of decision-making, enhance participation and improve outcomes of collective action. However, it was assumed that these institutions were essentially independent, and so the members were limited in their ability to fully exploit their social capital across different institutions.

In this chapter, the role and nature of social capital using agents in the context of Ostrom's institutional principles is investigated. In particular, the eighth principle: multiple layers of nested enterprises. Both inter-agent and inter-institutional interactions are then situated in an extension of a Public Goods Game (PGG) for an electricity scenario. These type of games are commonly used to study the effects of free-riding in the system, which reduces the positive benefits of cooperation in otherwise unregulated situations.

Therefore, this chapter is structured as follows. The proposed Electricity Public Goods Game (EPGG), which models decentralised Community Energy System (dCES), is described in section 6.2. Enhancements in the Electronic Social Capital framework are explained in section 6.3. Experimental results are presented and analysed in section 6.4. A discussion of related work follows in section 6.5.

6.2 testbed: electricity public goods game

A traditional approach in energy systems is to have a central generator providing for a set of consumers (e.g. households or factories). This structure has been very efficient for many years when electricity was generated in big generator plants (e.g. nuclear or hydroelectric). In the last few years, the cost reduction of small size renewable generators has enabled the introduction of new participants into the grid. Therefore, a dCES is characterised by heterogeneous members that not only consume electricity but also participate in the generation and distribution of energy. In this chapter, the concept of PGG is introduced and a variation is proposed to simulate a dCES.

In a PGG, a set of players must choose how many resources they will provide for a common pool. Once provisioned, the total resources are multiplied by an incremental factor and distributed among the participants. When all the players collaborate and offer resource provisions to the system, they all benefit by receiving more resources. However, a selfish player is tempted to free-ride the system by collecting the distributed resources without provisioning, which reduces the benefits and promotes free-riding among the other players. If this situation is not managed, it will provoke a vicious cycle that leads to the failure of the system. A more detailed analysis of conditional cooperation on PGG can be found in (Gächter 2007).

In this chapter, a variation of a PGG for an EPGG is proposed. The game consists of a set of n players that possesses a quantity of electricity $g_i \in [0, 1]$, and requires another quantity of electricity $r_i \in [0, 1]$ (which are individually and randomly assigned). Each player decides to contribute a part c_i to the common pool or to store it s_i for the future. The amount of energy contributed and stored must not be greater than the amount available $(c_i + s_i \leq g_i)$. When the energy is stored, a discount factor α is applied and the available energy at the next time step will be increased with the stored energy $(g_{i+1} + s_i * \alpha)$, simulating the costs of energy storage infrastructures such as batteries or capacitors. Alternatively, if players contribute to the common pool, another discount factor β is applied, which represents the network infrastructure costs.

The contributions from the common pool are summed up and the total available energy e_i for each player i is given by:

$$e_i = f(\beta \cdot \sum_{j=1}^{n} c_j) + (g_i - c_i - s_i) + (S_i \cdot \alpha)$$

where the function f is the allocation method chosen at the common pool, $(g_i - c_i - s_i)$ is the amount of energy not contributed or stored and S_i is the accumulated energy stored from the past rounds. It is assumed that the cost of storage is higher than the cost of the network $(\alpha > \beta)$ in the game. If the available energy is lower than the required $(e_i < r_i)$, the difference is consumed from the grid.

The ideal scenario would be that all the players collaborate on provisioning to the common pool, which has lower costs than storage, to balance the differences between the generation and the consumption of electricity. Nevertheless, this behaviour can not be enforced due to the lack of central authority. Players must self-organise by creating a *Common Pool Resource Institution (CPRI)* to manage the common pool. These CPRI have a set of *conventional rules* to regulate how the resource is managed and distributed (i.e. how it is allocated based on individual demands).

Furthermore, players are physically distributed in different locations. Members of different locations cannot interact directly; they must form local CPRI

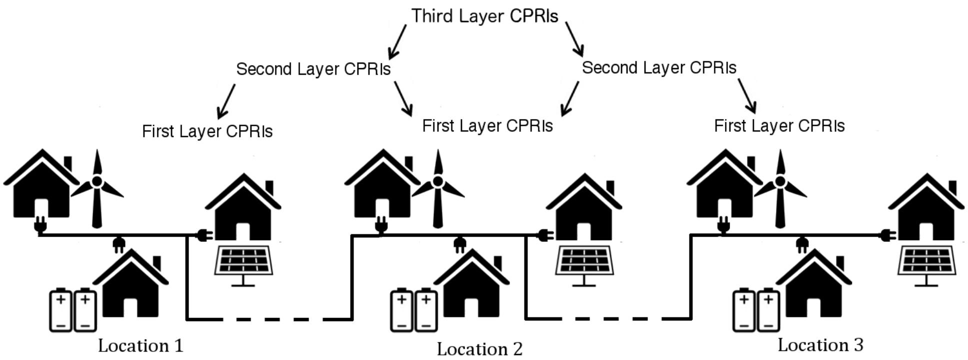

Figure 6.1: CPRIs structure in EPGG

that will interact among them. A layered structure of nested CPRI is defined where each new layer connects two adjacent locations from below. Figure 6.1 illustrates the structure of the game.

The EPGG game shares some similarities with the example of the irrigation system described in section 1.1. In irrigation systems, top-enders must agree with bottom-enders how water is appropriated without depleting it at the top. Similarly, in the EPGG there is a resource surplus at some locations, emulating the top-enders, but this situation is only temporary. Depending on the energy generation conditions, the surplus will indistinctly change between locations; making each location's long-term efficiency codependent on the others.

6.2.1 *Step-by-step algorithm*

The EPGG game was time-driven implemented. At the beginning of each round, players have a quantity of electricity available and another quantity needed. CPRI must be created to start accepting contributions, and the institutional rules of the CPRI must be defined. Players can join and leave any CPRI at the beginning of each round. This allows players to leave when they believe that the rules of the current CPRI are not favourable and, to create a new one when all CPRIs available are disadvantageous. Each player must inform how much will provision and/or it will demand to all the CPRIs it is associated with. With this information, CPRIs can calculate their surplus or need for electricity and perform the same provision and/or demand to the upper level of CPRIs.

```
A ← set of n agents
C ← set of c CPRIs
max_lvl ←maximum nested levels
t ← 0
repeat
   for each agent i ∈ A do
      calculate g_i {resource available}
      calculate r_i {resource need}
   end for
   for each agent i ∈ A do
      manage CPRI membership
   end for
   for each agent i ∈ A do
      request_provision(i)
      request_demand(i)
   end for
   lvl ← 0
   repeat
      for each CPRI c ∈ C where level = lvl do
         request_provision(c)
         request_demand(c)
      end for
      lvl ← lvl + 1
   until lvl == max_lvl
   repeat
      for each CPRI c ∈ C where level = lvl do
         a_c ← allocate(method)
      end for
      lvl ← lvl − 1
   until lvl == 0
   for each c ∈ C do
      update reputation
      update events
      institutional sanctions
      process votes
   end for
   for each agent i ∈ A do
      request_consumption(i)
   end for
   t ← t + 1
until t == T_lim
```

Algorithm 6.1: Electricity public goods game

The decision on how much to provision and demand between nested CPRIs is done by one of the players acting as a CPRI leader. Following this, by a top-down hierarchy, each CPRI allocates the player's demands using the allocation method chosen when the CPRI was created. Allocations from top-level CPRIs are included as available resources in the lower levels and reallocated recursively until the bottom level. With the allocation results, CPRIs update their internal state processing events and sanctions, updating reputations and calling for votes if needed. A description of the CPRIs's features follows in the next section. Finally, each player will consume from the grid the required electricity if the electricity allocated from the CPRIs and saved is not enough to cover its needs. Algorithm 6.1 shows a detailed step-by-step procedure of the game.

6.2.2 *Institutions*

An institutional ruleset must be defined when a CPRI is constituted. The features defined in this set are the following:

- Allocation method: This defines how the resources are allocated, mainly affecting situations of *scarcity* when all the demands cannot be allocated. We implemented the following methods: *no allocation*, which is self-explanatory; *random*, which ignores the demands and allocates a random amount to each player; *random demand*, which allocates according to the players' demands in random order while resources are available; *average*, which allocates equally to each player; and *contribution ratio* which allocates what players demanded by ordering their priorities based on their $\frac{provisioned}{allocated}$ ratio from previous rounds. The similarities of our allocation methods and the properties of allocation methods proposed by (Pitt et al. 2014b) (described in section 2.4) follow:
  - *Random* is *cost-effective*.
  - *Random demand* is also *cost-effective*.
  - *Average* is *proportional*.
  - *Contribution ratio* is *equitable*.

- Sanction abuse: This specifies if the CPRI will sanction when the difference between electricity provisioned and the electricity allocated reaches a threshold. The value of the threshold is also customised at each CPRI.
- Expel free-riders: This establishes if the CPRI will remove players that do not provision to the common pool. The number of tolerated rounds is also defined by a customisable threshold.
- Nested: This enables the CPRI to participate in upper-level CPRIs provisioning their surplus and demanding their needs.
- Democratised: This allows its members to vote on certain managing aspects. The voting situations implemented are: allow a new player to join, expel a free-rider and decide the CPRI leader. Each vote needs a simple majority to succeed.

6.2.3 *Reputation sources*

The reputation was modelled using "the information provided by other members about the interactions they had in the past." (see subsection 3.2.3). The reputation source used described as follows.

institutional reputation Each CPRI keeps a local reputation for its members, which can be checked by any of them. A triplet value is generated for any given member; this includes the provision, demand and allocation reputation. Each value is given by the deviation of each player's arithmetic mean from the CPRI whole mean. With this information, players can know if other players are provisioning, demanding or receiving allocations above the average of the CPRI. This information is completely relative to each individual CPRI and it is dependent on the CPRI's features that will affect players' actions.

6.2.4 *Social players*

Social players implement the social capital framework with our new decision module. Algorithm 6.2 describes their behaviour in a round of the game. At

```
E ← sequence of e events
C ← set of c CPRI available
J ← set of c CPRI joined by the agent
dm ← decision-making
feedback(dm)
for each event e ∈ E do
    update social capital sc_e
    if e = vote then
        vote(dm(v))
    end if
end for
if size(J) ≠ max then
    rc ← random(C)
    if dm(rc) then
        join(rc)
    end if
end if
for each CPRI j ∈ J do
    if not dm(j) then
        leave j
    end if
end for
g_i ← resource available
r_i ← resource need
for each CPRI c ∈ J do
    provision(g_i/γ)
    demand(d_i/γ)
end for
e_i ← calculate total available energy
consume(r_i − e_i)
```

Repeat for each CPRI leading

Algorithm 6.2: Social player in EPGG the beginning of each round, *social* players update their social capital with the relevant events from the CPRIs. The possible events are: a member joined, left, was sanctioned, or was expelled, what allocations were performed, reputation updates or calls for voting. In the case of voting, the actions are also performed. Management of the CPRI membership is done after all the events have been processed. Players will join a new CPRI based on the *social capital decision* for that CPRI (only when the maximum numbers of CPRIs defined by *max* is not reached; this is limited to 12 in this implementation). The same occurs when players must decide if they stay or leave. The following action is to *provision* to and *demand* from to the common pools. Players will calculate the difference between g and r to decide if they will provision to or demand from to the pool (surplus electricity or need for electricity). A factor γ is defined to distribute the provision or demand proportional to all the CPRIs the player is participating in. Now all the actions previously explained are repeated for all the CPRIs the player is leading, because each CPRI becomes a player in the upper level and the leader acts on its behalf. Resource allocation follows, and players receive the electricity assigned to them. Finally, players calculate the difference between the total available electricity and the amount they need; the difference is consumed from the grid.

6.2.5 *Free-rider players*

Free-rider players always save the energy not consumed and participate in all the CPRIs without provisioning. Their objective is to reduce the higher 'cost' of storing resources by free-riding. They will vote 'yes' in any vote, as more players in the CPRI will facilitate free-riding and expelling other free-riders allows them to receive more resources.

6.2.6 *Random players*

Random players participate in the EPGG by deciding all their actions through random selection. This includes joining or leaving CPRIs, any vote decision and how much to demand or provision.

6.3 machine learning for contextualised decision-making

6.3.1 *Specification*

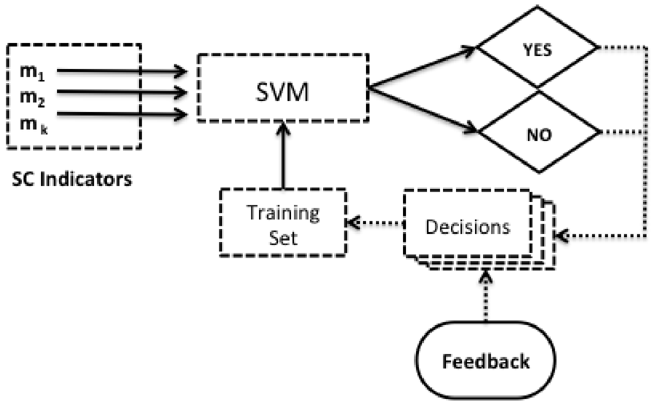

Figure 6.2: Decision-making using SVM

In previous experiments with the framework, social capital indicators were combined using the arithmetic mean to generate a unique output value $o \in [0, 1]$. It has been shown that, when multiple forms of social capital are present, distinct indicators might produce mixed results; using only a few forms of social capital could outperform using more forms in some contexts. Therefore, we propose an alternative approach for decision-making based on machine learning.

Most of the decisions using the framework are *binary*. For example, in the EPGG game, players decide if they will contribute to the common pool, when to join or leave a CPRI or to vote on different aspects of CPRI management. Essentially, these are *yes/no* decisions. The method chosen is a nonlinear classifier implemented using a Support Vector Machine (SVM) (Boser et al. 1992). We chose this solution due to the familiarity of the authors using SVM in previous work. We used the SVM provided by the machine learning library Encog (Heaton 2015) with the radial basis function as kernel. Figure 6.2 shows a schematic view of the designed module.

To be able to train the SVM without any prior data (agents start the simulation without any collected data), a dataset with all the discretised values for the metrics' inputs and the arithmetic mean was defined. The values were discretised to two decimals and a threshold of 0.5 was defined to split the output values into two categories. This initial configuration enables the decision-making module to initially work as in the previous versions. Every time a decision is taken, the module will save the output (indicators and output) for later learning; this can be adjusted by a learning rate discarding some of the decisions and was initially set to 50%. These decisions will be added to the training set based on the feedback received regarding the outcomes of the decision. If the decision was 'ok', it will be included in the training set as it is; in the contrary, it will be included switching the category.

The re-training of the SVM is triggered by a threshold defined by a percentage of values changed – in the size of the training set – mainly to avoid the computing costs associated with a considerably big dataset. By default the SVM is retrained when 1% of the dataset is changed.

6.3.2 *Implementation*

Figure 6.3 shows the Unified Modelling Language (UML) diagram of the main classes implemented for the extension of the decision-making module (it only shows the most relevant classes with only their public methods for simplicity).

The class *DecisionMaker* replaced the old *Decision* interface introduced in section subsection 3.4.3. When instantiating this class, the forms of social capital and the learning rate must be passed as an argument. The interface *LearnerInterface* defines the methods *learn* and *getValue*; that are used by the *DecisionMaker* class. Three classes implement this interface providing different learning capabilities to the decision module:

- *AverageNoLearner:* This class implements the same functionality as the first implementation of the decision module; it calculates the average of all the metrics produced by the forms of social capital. It does not have any learning capabilities and was used to verify that the new decision module worked as expected.

- *SVMLearner:* This class implements the SVM – explained in the previous section – for decision-making.

Other types of learners could be implemented using the *LearnerInterface*. For the purpose of our experiments the SVM provided enough flexibility to the agents to 'spot' which social capital indicators were relevant in each context.

6.4 EXPERIMENTS

In order to evaluate the nested structure of CPRIs, it was used in the EPGG game defined above with the three different types of players already mentioned (*free-rider*, *random* and *social*). *Social* players with different combinations of two or three forms of social capital were also used.

Four different experiments were undertaken with the following purposes:

- Experiment 1: Investigates the performance of the new decision module using SVM by comparing the outcome of the social players using different forms of social capital. It also serves as well as a baseline for more complex scenarios following.
- Experiment 2: Investigates the effects of adding the nested structure to the system, in terms of increased performance in contrast to increased free-riding.
- Experiment 3: Explores residual social capital after players are moved to another location in the structure and compares the cost of re-organising the CPRIs memberships.

In all the following experiments, the cost of saving energy was set at 50% ($\alpha = 0.5$) and the cost of providing to the common pool was set at 10% ($\beta = 0.9$). Those values were chosen based on the assumption that the cost of storage is higher than the cost of the network ($\alpha > \beta$) (see Section 6.2).The values were considered suitable to model an actual electricity market scenario and provide enough 'difference' to analyse the results.

Before analysing the results of the experiments, we would like to examine the electricity consumed from the grid in two simple ideal scenarios: one location with all players storing the energy surplus, and one location with all players contributing their surplus to the common pool.

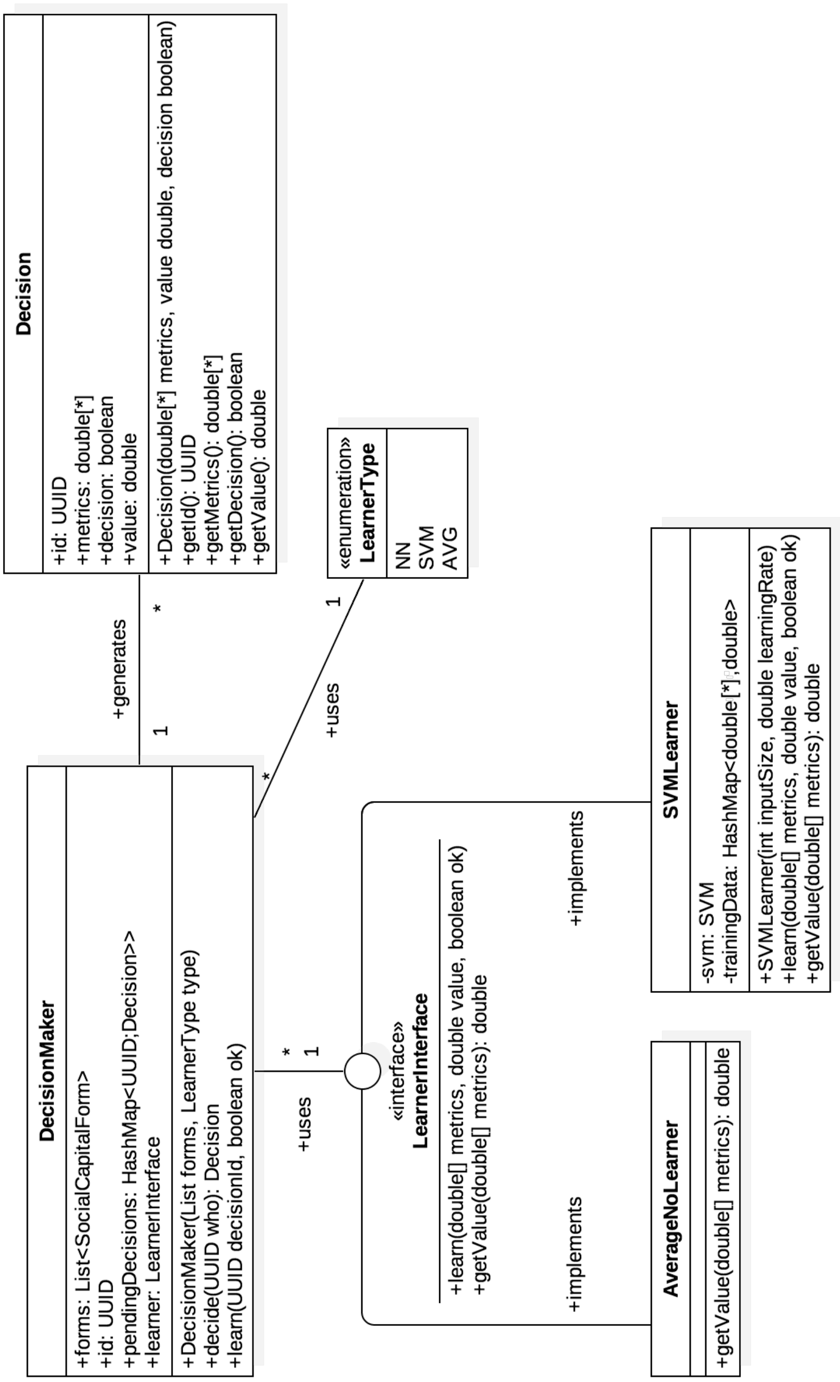

Figure 6.3: Extended decision-making module UML class diagram

The average electricity consumed from the grid for an agent i is given by:

$$\overline{c_i} = \frac{\frac{L}{3} \cdot \frac{1}{2} \cdot \phi}{\overline{g_i}}$$

where $\frac{L}{3}$ denotes the average difference between $g_i \in [0;1]$ and $r_i \in [0;1]$ (available and needed), $\frac{1}{2}$ discards half of them which are surplus of energy, ϕ denotes the loss by the costs of the energy storage or network and $\overline{g_i}$ is the average energy generated. In the first scenario where all players store their surplus, they will consume $16.\widehat{6}\%$ of the energy from the grid. Instead, if they all contribute to the common pool, they will consume $3.\widehat{3}\%$.

6.4.1 *Experiment 1: Multiple forms of social capital*

For the first experiment, the simulation was populated with *random*, *free-rider* and *social* players. *Social* players using the three forms of social capital and combinations of only two were included. The form *institutions* is present in all of the *social* players, as it is essential to manage and evaluate the CPRIs. Thirty players of each type participate in an EPGG game with one location and one level of CPRIs. The average values of 10 simulations have been used for the results of these experiments.

Figure 6.4 shows the average electricity consumed from the grid for each type of players. At the beginning, *free-rider* players outperform the rest by free-riding the system. *Social* players have no social capital information and start exploring the different institutional settings. As the game evolves, social capital starts 'learning' which institutions are beneficial and agents start providing to the right common pool. By their interactions in these institutions, they also vote to expel free-riders, maximising their own allocations of resources. *Random* players also benefit by some free-riding (behaving half of the time like *free-rider* players) which is also diminished after a few rounds. Having more forms of social capital is moderately beneficial in this scenario. Players using three forms outperform the others by reducing the institutions that allow free-riding at a faster speed. After 500 rounds, the results are close to the ideal scenarios presented in the previous section. The *free-rider* players slightly benefit by the initial situation, but it shows that is not sustainable as the simulation runs.

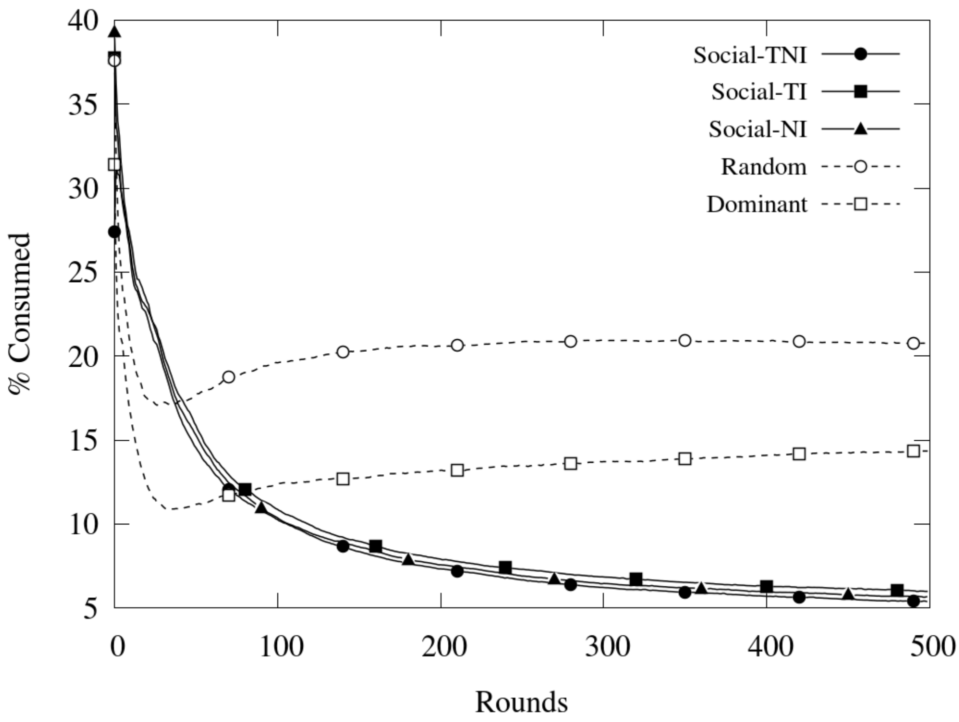

Figure 6.4: Consumption using multiple forms of social capital

6.4.2 *Experiment 2: Nested* CPRIs

For the next experiment, we ran the simulation with 3 locations using *social*, *free-rider* and *random* players and with I, II and III levels of nested CPRIs. With this setup, we investigated how beneficial it was to add the structure in terms of reducing the consumption from the main grid.

Figure 6.5 shows the players' consumption. *Random* players are not included as no changes were manifested. The creation of a nested structure of CPRIs allows the three locations to balance their energy surplus and needs, which benefits all the players. However, this gain is not equally distributed among all the players. *Free-rider* players decrease their consumption marginally, whereas *social* players considerably reduce their energy consumption. Table 6.1 compares the relative energy consumed at round 50 and 500 for the I, II and III nested levels. The results of *social* players and *free-rider* players at round 50 with one nested level are used as baseline for comparison purposes. As shown

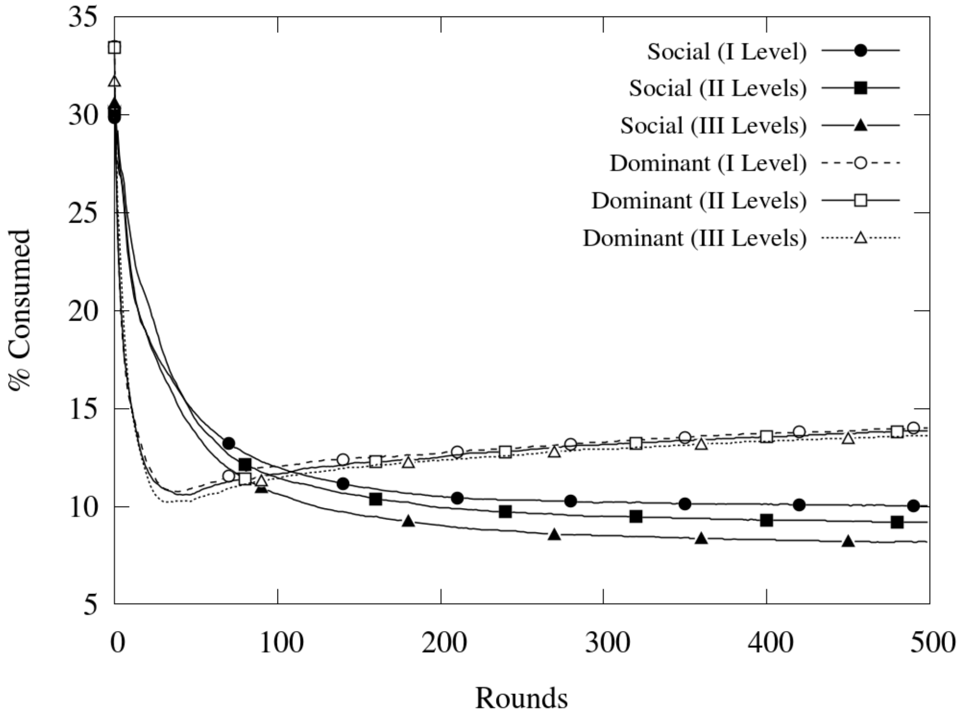

Figure 6.5: Consumption for I, II and III levels of nested CPRIs

in the table, at round 50, *social* players reduce their consumption by 7% when three layers are added. In contrast, the benefit for the *free-rider* players is only 3%. These results are more prominent at round 500, where the reduction is 20% (0.55/0.68) for the first and 2.5% (1.24/1.27) for the second.

Table 6.1: Electricity consumption evolution in EPGG

Player Type	*Round*	*I-Level*	*II-Levels*	*III-Levels*
Social	50	**1**	0.98	0.93
	500	0.68	0.62	0.55
Free-rider	50	**1**	0.99	0.97
	500	1.27	1.26	1.24

6.4.3 *Experiment 3: Locations reorganisation*

In the next scenario, we examine the effects of randomising the players' locations every 150 rounds. Our experimental hypothesis is: we expect that the self-organisation into the CPRIs should be done faster as players have learnt from the previous experience.

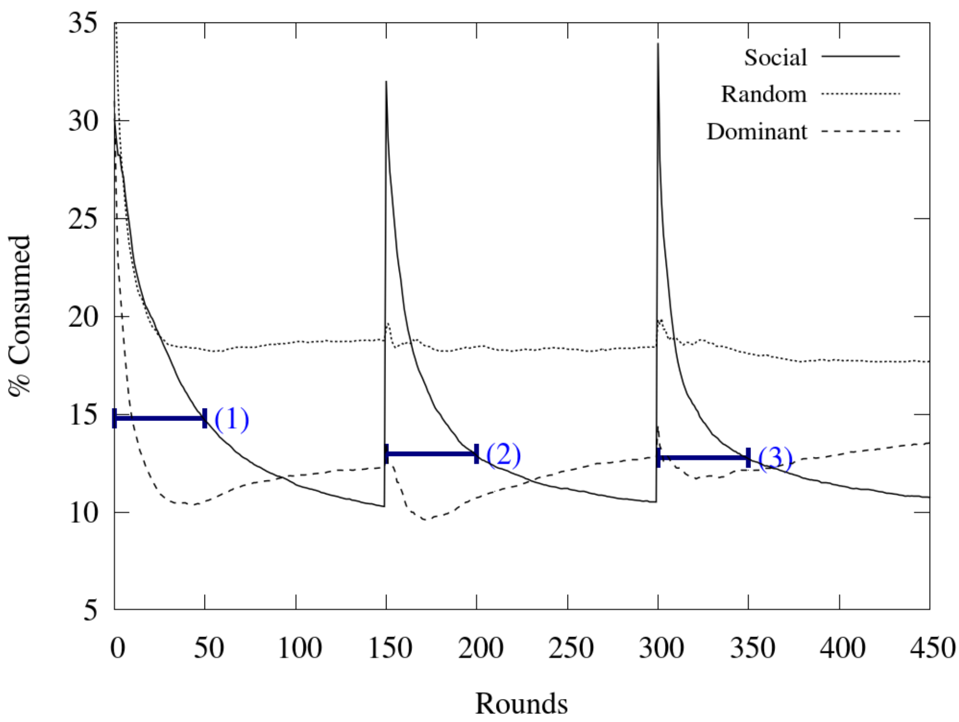

Figure 6.6: Consumption when locations are randomised

Figure 6.6 shows the average electricity consumed from the grid for *social*, *free-rider* and *random* players. We will only analyse the *socials* since are the ones we are interested in this experiment.

The first 50 rounds after the randomisation are particularly relevant, as it shows how fast players adapt to the change, and is highlighted in the figure. After 50 rounds from the beginning of the simulation, the average consumption is 14.8% (1). Comparing it to the first reorganisation, where the average consumption is 12.96% (2), there is a slight consumption decrease. If we compare the second and the third reorganisation, 12.7% (3), the decrease is marginal.

The results show that when changing locations players reorganise faster, achieving better outcomes (less consumption). The results of these experiments were averaged over 10 simulations of 20 agents of each type, for a total of 600 agents. The marginal increase on the third change of locations arise the question: is this trend of consumption reduction feasible in simulations with more agents or locations? This needs further experiments to explore it and opens up another line of further work.

6.4.4 *Applications to dCES*

The results of these experiments can be related to dCES in the following ways:

- Experiment 1: Tests how social capital can be used in an isolated energy community and analyses the levels of free-riding in the community.
- Experiment 2: Explores how a nested structure of CPRIs can be used to connect these communities between them, and compares the free-riding as the structure grows.
- Experiment 3: Investigates the outcomes of changing the location of the participants. This is not only relevant if the participants move home, but especially important if electric cars participate in the communities.

6.5 RELATED WORK

In the fields of self-organising, multi-agent and legal systems, there are many notations (and often associated tools) for describing organisations and institutions. This includes MOISE (an organisational model for MAS) (Hübner et al. 2007), OMACS (organisational model for adaptive computational systems) (DeLoach 2009), LAO (logic and organisations) (Dignum and Dignum 2012), LGI (law governed interaction) (Minsky 2005) and electronic institutions (García-Camino et al. 2005), amongst others. However, to the best of our knowledge, none of these works address the issue of multiple institutions whose interactions create social capital and reasoned with as part of the decision-making processes of the component entities of the institutions (or the entities acting on behalf of (i.e. empowered by) the institutions that they represent).

The issue of multiple interacting institutions has been addressed (Patel et al. 2005; Cliffe et al. 2006), but these works do not consider the concepts addressed here: norm-governed institutions, Ostrom's institutional design principles, and most significantly, inter-institutional social capital. However, recent work on teams, team structures and team coordination (Franco et al. 2016) could offer some useful insights into *structuration* (the duality of agency and structure, in that structures are made up of agents, and agents have memory of structures (Giddens 1984)). This could prove highly relevant to the formal conception of nested enterprises and entities, because this memory should perhaps includes elements of the reputational and relational economies which are, in fact, the essence of social capital.

Social capital has also been studied in the context of MAS. For example, it was stated that "there is a big interest in literature about 'social capital' and its powerful effects on the wellbeing of both societies and individuals, often it is not clear enough what is it the object under analysis" (Falcone and Castelfranchi 2011). In that paper, they proposed 'trust' as the capital of agents. In this work, we have followed the Ostrom and Ahn approach and formalised social capital in its different forms and trust as the decision-making process that leads to successful collective action outcomes.

6.6 SUMMARY

In the previous chapters, a generic framework for electronic social capital was specified and evaluated in 2-player and simultaneous n-player games. It reported that social capital reduces the complexity of decision-making and optimises the outcomes of collective decision-making.

The current chapter builds on this previous work but makes the following valuable stand-alone contributions.

Firstly, it presents an enhanced version of the social capital framework in which the decision-making module implements contextualised machine learning, providing for effective action selection across multiple individual and institutional collective action situations.

Secondly, it proposes a new experimental testbed that features

- Social agents that reason with this enhanced social capital framework.

- A new game EPGG, a model of the collective action situation exhibited by aggregated dCES.
- The creation and use of inter-institutional social capital between nested enterprises.

Lastly, new experimental results were reported that show the significance of inter-institutional social capital in the self-organisation of sustainable structures of nested enterprises, which: facilitates the transmission of prosocial behaviour from 'lower' levels to 'higher' levels; beneficially affect the development and maintenance of the nested structure; and possibly assists the members' transition within the structure.

In summary, it has been shown how the aggregation of agents into CPRI (that are embedded in a system of nested enterprises) promotes prosocial behaviour which can be propagated throughout the system by inter-institutional social capital.

In conclusion, we would argue that the (formal) representation of social capital is a critical aspect to effective coordination in open MAS, especially where issues of scalability demand that the system organises itself into a layered structure of nested enterprises. In particular, social capital acts as a 'mediating enabler', facilitating inter-institutional collective action and complementing the institutional design principle P7 – minimal recognition of rights to organise – which acts as a 'mediating constraint' balancing the rights and powers of such nested enterprises.

7 SUMMARY, CONCLUSIONS AND FURTHER WORK

7.1 OVERALL SUMMARY

In this thesis, we explored if the use of social capital with Self-Organising Electronic Institutions (SOEIs) can enhance the ability to solve collective action problems in Multi-Agent Systems (MASs); with the aim to develop an Electronic Social Capital (ESC) framework and test it in different scenarios.

In chapter 2 we discussed the different types of open MAS and *collective action*, especially those subject to the problem of resource allocation. Ostrom's analysis on self-governing institutions was introduced and the chapter examined how these institutions have been formalised electronically. The notion of social capital was introduced in this same chapter and the concept was explored from sociological analyses of human societies to Ostrom's ideas emphasising that social capital could be leveraged to solve collective action problems. Some examples of basic forms of social capital used in computer systems were described to conclude the chapter.

Based on sociologists' work analysing social capital and especially Ostrom and Ahn's findings of three forms of social capital, an ESC was developed and specified in chapter 3. The specification also included a MAS to describe the environment in which the experiments were conducted. An instantiation of the framework was described based on the formal specification, which included all its internal components and how they were developed in Java to use them in the PRESAGE2 platform.

In section 1.1 we introduced an example of an irrigation system described by Ostrom where the use of social capital facilitates the co-existence of two institutions that manage Common Pool Resources (CPRs) and are codependent. To test the framework in a similar scenario, we planned a set of experiments using three different scenarios that were incremental and built on the complexity as we progressed. The first testbed used 2-player games; later, we added more complexity into n-player games and the last testbed used was a CPR scenario with nested institutions.

The first experiments that used the ESC framework in 2-player games were presented in chapter 4. This testbed was a strategic game where a population of agents was repeatedly randomly paired to play a game against each other. We called it the Cooperation Game (CG). In these experiments, we could analyse the ESC framework in situations where agents only had to choose an action against only one opponent, simplifying the analysis of the results. These experiments also provided results on the outcome of agents using social capital to select institutions with different characteristics. A visualisation tool was developed during those experiments to make easier the task of analysing and understanding the results.

In chapter 5 experiments using the ESC framework in n-player games were presented. This new testbed was called Unscrupulous Diner's Dilemma Game (UDDG) and, in this setting, agents had to choose an action where the results were dependent on the choices of more than one other player. In these experiments, the reputation of the MAS was extended to include the information provided by other members regarding the interactions they had in the past (see subsection 3.2.3 (2)).

The last testbed was presented in chapter 6. The role and nature of social capital using agents with learning capabilities in the context of Ostrom's institutional principles was investigated, particularly, the eighth Ostrom's principle – i.e. multiple layers of nested enterprises. Both inter-agent and inter-institutional interactions were then situated in an extension of a Public Goods Game (PGG) for an electricity scenario. We called it the Electricity Public Goods Game (EPGG). This scenario was chosen due to its shared characteristics with the irrigation system described by Ostrom. In the experiments performed, the results of agents' action choices were dependent on more than one other agent as in n-player games, and were also part of a layered structure of nested institutions. For this testbed, the ESC framework was extended to include machine learning techniques in the *decision module* allowing agents to adapt to different levels of the structure and different decision situations.

7.2 SUMMARY OF CONTRIBUTIONS

The main contributions of this work are given below in more detail.

- Presented a generic framework to represent and reason with social capital. The framework decouples event processing and the updating of the social capital information from decision-making, thus providing a modular architecture.
- Provided one implementation of the Electronic Social Capital framework using *Trustworthiness*, *Networks* and *Institutions* forms of social capital. The implementation can be extended to include other forms of social capital.
- Presented the design and implementation of the Cooperation Game and the experiments performed. The results showed that social capital:
  - Facilitates to achieve win-win situations (two agents involved in a pairwise interaction benefit from behaving cooperatively).
  - Acts as a catalyst for self-organisation, where agents decided with whom to interact according to their social capital.
  - Can be used to decide which institutions to join or leave.
- Presented the design and implementation of the Unscrupulous Diner's Dilemma Game and the experiments performed. The results showed that social capital:
  - Optimises the outcomes (in terms of long-term satisfaction and utility) of collective decision-making in competitive environments, compared to alternative simplistic strategies.
  - Reduces the complexity of that decision-making.
  - Scales with the size of the population, as its complexity is independent of population size.
- Presented the design and implementation of the Electricity Public Goods Game and the experiments performed. An enhanced version of the social capital framework in which the decision-making module uses con-

textualised machine learning was developed for these experiments. The results showed that social capital:

- Facilitates the transmission of prosocial behaviour from 'lower' levels to 'higher' levels of a nested structure of Common Pool Resource Institution.
- Beneficially affects the development and maintenance of the nested structure.
- Possibly assists the members' transition within the structure.

7.3 LIMITATIONS

This section discusses some limitations in the presented work. These can be divided into two categories; theoretical limitations, which are inherited from the social capital analysis in human societies and implementation limitations, which are related to our particular implementation of the framework.

7.3.1 *Theoretical limitations*

FORMS Social capital is a broad field and there is still no agreement regarding a single concrete and formal definition of the concept. Researchers seem to agree on the qualitative value of connections and/or relations among individuals. The ESC framework is based on only one dimension: Ostrom and Ahn's *trustworthiness*, *networks* and *institutions* forms of social capital. Even if the formal definition of the ESC allows the implementation of different forms, our premise is that more forms of social capital enrich the framework and can enhance resilience in different scenarios. Using only other forms in the same experiments we have performed could produce different results.

COMMODIFICATION OF SOCIAL CAPITAL There is a risk of the commodification of social capital into currency value. When individuals achieve a significant amount of social capital, they could be tempted into performing actions (only possible by their accumulated social capital in the system) on behalf of other individuals and in exchange for any economic reward. By performing actions that are not in the *common* interest, the individual will be diminishing its social capital when exchanging it for 'money', as seen by the other participants. Although this situation is not 'sustainable', as the leveraged position will not endure, the goal of using social capital was to represent and reason using qualitative, instead of traditional quantitative, values. This situation must be taken into consideration when designing a system that uses social capital.

7.3.2 *Implementation limitations*

EVENTS The implementation of the ESC framework uses a set of predefined events with their corresponding weights. All the weights are assigned by perceived significance. This means that the 'importance' of each event must be known in advance by the system designer. The numeric specific value of each weight is not as decisively determinant as the relations among them – i.e. being sanctioned 10 times in an institution has the same weight as being expelled once.

SCALABILITY AND CAUSE AND EFFECT The results in subsection 5.3.2 showed that social capital scales with the size of the population. But, does it scale with the complexity of the environment? During the experiments, the difficulty of relating one particular event with its provoking cause increased with the number of actions and events happening in the testbed. Our solution was to individually test each part of the framework and the testbeds to ensure that the results were consistent. This solution might not be applicable to more complex scenarios that can not be tested modularly.

7.4 FURTHER WORK

In the foundational work on the eight institutional design principles (Ostrom 1990), Ostrom did not highlight the role of social capital much. However, a multi-institutional case study, involving water basins in California, was examined in some detail (Ostrom 1990, pp.133–136). It is implicit in her analysis that, in addition to all eight features of successful CPR management systems, elements of 'social capital' were also in play between the various public enterprises that sustained the water basins.

The last scenario discussed in this thesis is intended as an abstraction of the decentralised Community Energy System idea (Pitt et al. 2014a), which has also been analysed as a "polycentric set" of different actors (Diaconescu and Pitt 2015). However, even in abstract terms, it demonstrates two elements of social capital evident in the water basin scenario: an interplay of relations at multiple inter-institutional layers; and the learning of effective relationships by individual entities acting on behalf of, or as the representative of, the institutions. What is missing, perhaps, is the interaction between multiple heterogeneous institutions in a "polycentric set". This missing element, however, points the way towards further research in three directions: the relationship between social capital and the institutional design principles (particularly principle P7); the relationship between social capital and polycentric governance; and the role of social capital in what have been called *holonic institutions* (Diaconescu and Pitt 2015).

In the first direction, Ostrom's institutional design principle P7 states that there should be a "minimal recognition of the right to self-organise". Recent work (Pitt et al. 2016) has tried to formalise this principle in terms of empowerment and entitlement relations between *nested* enterprises as a trade-off between the rights and powers of 'inner' and 'outer' institutions. Thus, the principle acts as a kind of 'mediating constraint' between institutions that serves to limit excess of autonomy on the one hand and excessive interference (from the outside) on the other. By contrast, we believe that social capital acts as a kind of 'mediating enabler' that reinforces the trust between institutions and enables successful collective action at each institutional layer. Therefore, an intriguing next step would be to converge these experiments in inter-institutional social capital with investigations into the formalisation of Principle P7.

A successful convergence of these ideas would enable further research into the second direction, and an analysis of a polycentric set of institutions. In the experiments described in chapter 6, the institutions at each of the nested layers are essentially homogeneous, especially in terms of their intended goals. We also need to set up and analyse a set of nested institutions where there are different (and even conflicting) aims, ownership models, and participation strategies. This will allow the modelling of scenarios where polycentricity is clearly an important feature, such as the irrigation systems studied by Ostrom, but also smart grids or community energy systems, and other large-scale infrastructures.

Finally, this study of social capital and polycentric governance would contribute to the third direction of future research, namely that of holonic institutions. The idea of holonic institutions is to converge the benefits of voluntary regulation according to mutually agreed rules (as found in institutions) with the multi-scale, multi-criteria optimisation offered by holonics. The critical question here is the extent to which social capital can contribute to innovation in, of and between institutions. This is, arguably, essential to any development in, for example, smart cities, where the building of such a city from scratch is effectively impossible and needs to proceed from a potential melange of pre-existing organisations, institutions and other vested interests. We believe that social capital is not just the precondition to successful collective action, but also the precondition for successful collective (institutional) innovation.

7.5 CONCLUDING REMARKS

In this final chapter an overall summary of the thesis was presented. Furthermore, the main contributions and the limitations of the work were detailed, including some interesting ideas for further work to be carried out. This thesis presented a new electronic social capital framework, as well as three different scenarios where the experiments were performed and can be used as testbeds for further research.

The work carried on in this thesis has opened other lines of work, particularly:

- The electronic social capital framework developed will be distributed as a module of PRESAGE2 simulation platform. This will allow further research using the framework to be conducted, and to improve the current implementation by other people's collaboration – the simulation platform is available as open-source software.

- The EPGG is currently being implemented as a physical demonstrator by students at Imperial College London. The ideas developed in the testbed, will be empowered by a micro-model of a smart-house – including different appliances – and a mobile app to manage the different features while away from home.

The research presented in this thesis is part of the EPSRC programme "The Autonomic Power System", which focused on the electricity network of 2050. Especially, my research was pointed on a radically different perspective in which participants could conduct the exchange of the generation and the consumption of electricity using a distributed approach – rather than the current centralised energy generation and distribution. As the research project progressed and the results were shown at the project meetings, the interest in decentralisation and user-participation for self-organisation of electricity grids was sparked. The potential of these approaches was acknowledged by the project partners, who expressed their willingness to continue this line of work in further research.

BIBLIOGRAPHY

Gerrit Anders, Jan-Philipp Steghöfer, Florian Siefert, and Wolfgang Reif. A trust- and cooperation-based solution of a dynamic resource allocation problem. In *7th IEEE International Conference on Self-Adaptive and Self-Organizing Systems (SASO)*, pages 1–10, 2013. (Cited on page 23.)

Paul S. Andrews, Fiona A. C. Polack, Adam T. Sampson, Susan Stepney, and Jon Timmis. The CoSMoS process, version 0.1: A process for the modelling and simulation of complex systems. Technical Report YCS-2010-453, Department of Computer Science, University of York, 2010. (Cited on page 17.)

Aristotle. *Nicomachean Ethics*. Ross, W. D., 350BC. (Cited on page 27.)

Alexander Artikis. Dynamic specification of open agent systems. *Journal of Logic and Computation*, 22(6):1301–1334, 2012. (Cited on pages 18, 37, and 44.)

Robert Axelrod and D Hamilton. The evolution of cooperation. *SCIENCE*, 21 (1):1390, 1981. (Cited on pages 29 and 77.)

Mathieu Bastian, Sebastien Heymann, and Mathieu Jacomy. Gephi: An Open Source Software for Exploring and Manipulating Networks. In *International AAAI Conference on Weblogs and Social Media*, 2009. (Cited on page 77.)

Yvonne Bernard, Lukas Klejnowski, David Bluhm, Jörg Hähner, and Christian Müller-Schloer. *Self-Organisation and Evolution for Trust-Adaptive Grid Computing Agents*, pages 209–224. Springer Berlin Heidelberg, Berlin, Heidelberg, 2014. (Cited on page 23.)

Ken Binmore. *Natural Justice*. Oxford University Press, 2005. (Cited on pages 15 and 30.)

Vincent D. Blondel, Jean-Loup Guillaume, Renaud Lambiotte, and Etienne Lefebvre. Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2008(10):P10008, 2008. (Cited on page 77.)

Bernhard E. Boser, Isabelle M. Guyon, and Vladimir N. Vapnik. A training algorithm for optimal margin classifiers. In *Proceedings of the Fifth Annual Workshop on Computational Learning Theory*, COLT '92, pages 144–152, New York, NY, USA, 1992. ACM. (Cited on page 104.)

Pierre Bourdieu and Loïc J.D Wacquant. An invitation to reflexive sociology. *Chicago University Press*, 1992. (Cited on page 34.)

Samuel Bowles and Herbert Gintis. Social capital and community governance. *The Economic Journal*, 112(483):F419–F436, 2002. (Cited on page 93.)

Michael Bowling and Manuela Veloso. Multiagent learning using a variable learning rate. *Artificial Intelligence*, 136(2):215 – 250, 2002. (Cited on page 77.)

Felix Brandt, Vincent Conitzer, and Ulle Endriss. Computational social choice. *Multiagent systems*, pages 213–283, 2012. (Cited on page 30.)

Hendrik Van Brussel, Jo Wyns, Paul Valckenaers, Luc Bongaerts, and Patrick Peeters. Reference architecture for holonic manufacturing systems: {PROSA}. *Computers in Industry*, 37(3):255 – 274, 1998. (Cited on page 25.)

Estelle Cantillon and Martin Pesendorfer. Chapter 22 auctioning bus routes: The london experience, 2004. (Cited on page 25.)

Cristiano Castelfranchi, Rino Falcone, and Francesca Marzo. Being trusted in a social network: Trust as relational capital. In *International Conference on Trust Management*, pages 19–32. Springer, 2006. (Cited on page 30.)

Yann Chevaleyre, Paul E Dunne, Ulle Endriss, Jérôme Lang, Michel Lemaitre, Nicolas Maudet, Julian Padget, Steve Phelps, Juan A Rodriguez-Aguilar, and Paulo Sousa. Issues in multiagent resource allocation. *Informatica*, 30(1), 2006. (Cited on page 25.)

Yann Chevaleyre, Ulle Endriss, Jérôme Lang, and Nicolas Maudet. *A Short Introduction to Computational Social Choice*, pages 51–69. Springer Berlin Heidelberg, Berlin, Heidelberg, 2007. (Cited on pages 27 and 30.)

Owen Cliffe, Marina De Vos, and Julian Padget. Specifying and reasoning about multiple institutions. In *Coordination, Organizations, Institutions, and*

Norms in Agent Systems II, number 4386 in LNCS, pages 67–85, 2006. (Cited on page 113.)

James S. Coleman. Social capital in the creation of human capital. *American Journal of Sociology*, 94:S95–S120, 1988. (Cited on page 34.)

National Research Council. *Visionary Manufacturing Challenges for 2020*. The National Academies Press, Washington, DC, 1998. (Cited on page 25.)

Ben Kei Daniel, Juan-Diego Zapata-Rivera, and Gordon McCalla. A Bayesian computational model of social capital in virtual communities. In M. Huysman, E. Wenger, and V. Wulf, editors, *Proc. of Int. Conf. on Communities and Technologies*, pages 287–306. Kluwer Academic Publishers, 2003. (Cited on page 92.)

Constantinos Daskalakis, Paul Goldberg, and Christos Papadimitriou. The complexity of computing a Nash equilibrium. In *38th Annual ACM Symposium on Theory of Computing*, pages 71–78, 2006. (Cited on page 58.)

Scott A. DeLoach. OMACS: A framework for adaptive, complex systems. In V. Dignum, editor, *Multi-Agent Systems: Semantics and Dynamics of Organisational Models*. IGI Global, 2009. (Cited on page 112.)

Yiannis Demiris. Prediction of intent in robotics and multi-agent systems. *Cognitive Processing*, 8(3):151–158, 2007. (Cited on page 18.)

Daniel C. Dennett. *The Intentional Stance*. MIT Press, 1987. (Cited on page 18.)

Ada Diaconescu and Jeremy Pitt. Holonic institutions for multi-ccale polycentric self-governance. In *COIN'14 Proceedings*, number 9372 in LNAI. Springer, 2015. (Cited on page 120.)

Virginia Dignum and Frank Dignum. A logic of agent organizations. *Logic Journal of the IGPL*, 20(1):283–216, 2012. (Cited on page 112.)

Rino Falcone and Cristiano Castelfranchi. Trust and relational capital. *Comput. Math. Organ. Theory*, 17(2):179–195, 2011. (Cited on page 113.)

Rainer Feldmann, Martin Gairing, Thomas Lücking, Burkhard Monien, and Manuel Rode. *Selfish Routing in Non-cooperative Networks: A Survey*, pages 21–45. Springer Berlin Heidelberg, Berlin, Heidelberg, 2003. (Cited on page 25.)

Mariana Ramos Franco, Gustavo A. L. Campos, and Jaime Sim ao Sichman. An empirical approach for relating environmental patterns with agent team compositions. In *COIN@AAMAS2016 Pre-proceedings*, 2016. (Cited on page 113.)

Simon Gächter. Conditional cooperation: Behavioral regularities from the lab and the field and their policy implications. In *Psychology and Economics: A Promising New Cross-Disciplinary Field*, pages 19–50. MIT Press, Cambridge, MA, 2007. (Cited on page 97.)

Aram Galstyan, Karl Czajkowski, and Kristina Lerman. Resource allocation in the grid with learning agents. *Journal of Grid Computing*, 3(1):91–100, 2005. (Cited on page 25.)

Andrés García-Camino, Pablo Noriega, and Juan Antonio Rodríguez-Aguilar. Implementing norms in electronic institutions. In *Proceedings of the Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 667–673. ACM Press, 2005. (Cited on page 112.)

Anthony Giddens. *The Constitution of Society*. University of California Press, 1984. (Cited on page 113.)

Natalie S. Glance and Bernardo A. Huberman. The dynamics of social dilemmas. *Scientific American*, 270(3):76–81, 1994. (Cited on page 79.)

Uri Gneezy, Ernan Haruvy, and Hadas Yafe. The inefficiency of splitting the bill*. *The Economic Journal*, 114(495):265–280, 2004. (Cited on page 79.)

Garrett Hardin. The tragedy of the commons. *Science*, 162:1243–1248, 1968. (Cited on pages 15 and 30.)

Jeff Heaton. Encog: Library of interchangeable machine learning models for java and c. *Journal of Machine Learning Research*, 16:1243–1247, 2015. (Cited on page 104.)

Jomi Hübner, Jaime Sichman, and Olivier Boissier. Developing organised multiagent systems using the MOISE$^+$ model. *IJAOSE*, 1(3/4):370–395, 2007. (Cited on page 112.)

Leonid Hurwicz. The design of mechanisms for resource allocation. *The American Economic Review*, 63(2):1–30, 1973. (Cited on page 15.)

Trung Dong Huynh, Nicholas R. Jennings, and Nigel R. Shadbolt. An integrated trust and reputation model for open multi-agent systems. *Autonomous Agents and Multi-Agent Systems*, 13(2):119–154, 2006. (Cited on page 23.)

Andrew J. I. Jones. On the concept of trust. *Decision Support Systems*, 33(3): 225–232, 2002. (Cited on pages 35 and 91.)

Andrew J. I. Jones and Marek Sergot. A formal characterisation of institutionalised power. *Journal of the IGPL*, 4(3):427–443, 1996. (Cited on page 31.)

Andrew J. I. Jones, Alexander Artikis, and Jeremy Pitt. The design of intelligent socio-technical systems. *Artificial Intelligence Review*, 39(1):5–20, 2013. (Cited on pages 17 and 18.)

James Konow. Which is the fairest one of all? a positive analysis of justice theories. *Journal of Economic Literature*, 41(4):1188–1239, 2003. (Cited on pages 27 and 28.)

Ludwig Lachmann. Capital and its structure. *Kansas City: Sheed, Andrews and McMeel*, 1978. (Cited on page 33.)

Kate Larson and Tuomas Sandholm. Strategic deliberation and truthful revelation: An impossibility result. In *Proceedings of the 5th ACM conference on Electronic commerce*, pages 264–265. ACM, 2004. (Cited on page 30.)

Samuel Macbeth, Dídac Busquets, and Jeremy Pitt. Principled operationalization of social systems using presage2. In D. Gianni, A. D'Ambrogio, and A. Tolk, editors, *Modeling and Simulation-Based Systems Engineering Handbook*. CRC Press, 2014. (Cited on page 46.)

George J. Mailath and Larry Samuelson. *Repeated Games and Reputations: Long-Run Relationships*. Oxford University Press, 2006. (Cited on page 77.)

Roger C. Mayer, James H. Davis, and F. David Schoorman. An Integrative Model of Organizational Trust. *The Academy of Management Review*, 20(3): 709–734, 1995. ISSN 03637425. (Cited on page 77.)

Naftaly Minsky. *Law-Governed Interaction (LGI): A Distributed Coordination and Control Mechanism*. Rutgers University, 2005. (Cited on page 112.)

Jeonghoon Mo and Jean Walrand. Fair end-to-end window-based congestion control. *IEEE/ACM Trans. Netw.*, 8(5):556–567, 2000. (Cited on page 27.)

Roger B. Myerson. *Game theory - Analysis of Conflict*. Harvard University Press, 1997. (Cited on page 77.)

John B. Nagle. On packet switches with infinite storage. *IEEE Transactions on Communications*, 35(4):435–438, 1987. (Cited on page 27.)

John F. Nash. Equilibrium points in n-person games. *Proc. of the National Academy of Sciences*, 36:48–49, 1950. (Cited on page 29.)

Brendan Neville and Jeremy Pitt. A computational framework for social agents in agent mediated e-commerce. In R. Falcone, editor, *AAMAS Trust Workshop*, pages 83–91, 2004. (Cited on page 18.)

Brendan Neville and Jeremy Pitt. *PRESAGE: A Programming Environment for the Simulation of Agent Societies*, pages 88–103. Springer Berlin Heidelberg, Berlin, Heidelberg, 2009. (Cited on page 46.)

Mancur Olson. *The Logic of Collective Action: Public Goods and the Theory of Groups*. Number v.124 in Harvard economic studies. Harvard University Press, 1965. (Cited on pages 29 and 93.)

Elinor Ostrom. *Governing the Commons*. Cambridge University Press, 1990. (Cited on pages 15, 16, 30, 31, 32, 45, 64, 82, 95, and 119.)

Elinor Ostrom and TK. Ahn. *Foundations of Social Capital*. Edward Elgar Publishing, 2003. (Cited on pages 16, 31, 34, and 95.)

Christos H. Papadimitriou and Tim Roughgarden. Computing equilibria in multi-player games. In *Proc. 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 82–91, 2005. (Cited on pages 29 and 93.)

Jigar Patel, W. T. Luke Teacy, Nicholas R. Jennings, Michael Luck, Stuart Chalmers, Nir Oren, Timothy J. Norman, Alun Preece, Peter M. D. Gray, Gareth Shercliff, Patrick J. Stockreisser, Jianhua Shao, W. Alex Gray, Nick J.

Fiddian, and Simon Thompson. Agent-based virtual organisations for the grid. *Multiagent and Grid Systems*, 1(4):237–249, 2005. (Cited on page 113.)

Patricio E. Petruzzi, Dídac Busquets, and Jeremy Pitt. Self organising flexible demand for smart grid. In *7th IEEE International Conference on Self-Adaptive and Self-Organizing Systems (SASO)*, 2013. (Cited on pages 4 and 36.)

Patricio E. Petruzzi, Dídac Busquets, and Jeremy Pitt. Experiments with social capital in multi-agent systems. In *17th International Conference on Principles and Practice of Multi-Agent Systems*, pages 18–33, 2014a. (Cited on page 4.)

Patricio E. Petruzzi, Dídac Busquets, and Jeremy Pitt. Visualisation of social capital. In *8th IEEE International Conference on Self-Adaptive and Self-Organizing Systems (SASO)*, pages 189–190, 2014b. (Cited on page 4.)

Patricio E. Petruzzi, Dídac Busquets, and Jeremy Pitt. Social capital as a complexity reduction mechanism for decision making in large scale open systems. In *8th IEEE International Conference on Self-Adaptive and Self-Organizing Systems (SASO)*, pages 145–150, 2014c. (Cited on page 4.)

Patricio E. Petruzzi, Dídac Busquets, and Jeremy Pitt. A generic social capital framework for optimising self-organised collective action. In *9th IEEE International Conference on Self-Adaptive and Self-Organizing Systems (SASO)*, pages 21–30, 2015. (Cited on page 4.)

Patricio E. Petruzzi, Jeremy Pitt, and Dídac Busquets. Inter-institutional social capital for self-organising 'nested enterprises'. In *10th IEEE International Conference on Self-Adaptive and Self-Organizing Systems (SASO)*, 2016. (Cited on page 5.)

Rosalind W. Picard. *Affective Computing*. Cambridge, MA: MIT Press, 1997. (Cited on page 18.)

Isaac Pinyol and Jordi Sabater-Mir. Computational trust and reputation models for open multi-agent systems: a review. *Artificial Intelligence Review*, 40(1):1–25, 2013. (Cited on pages 77 and 91.)

Jeremy Pitt and Ada Diaconescu. Structure and governance of communities for the digital society. In *Workshop on Self-Improving System Integration*, 2015. (Cited on page 93.)

Jeremy Pitt and Andrzej Nowak. The reinvention of social capital for socio-technical systems. *IEEE Technol. Soc. Mag.*, 33(1):27–33, 2014. (Cited on page 36.)

Jeremy Pitt, Julia Schaumeier, and Alexander Artikis. Axiomatisation of socio-economic principles for self-organising institutions: Concepts, experiments and challenges. *ACM Trans. Auton. Adapt. Syst.*, 7(4):39:1–39:39, 2012. (Cited on pages 16, 24, and 37.)

Jeremy Pitt, Dídac Busquets, Aikaterini Bourazeri, and Patricio E. Petruzzi. Collective intelligence and algorithmic governance of socio-technical systems. In Daniele Miorandi, Vincenzo Maltese, Michael Rovatsos, Anton Nijholt, and James Stewart, editors, *Social Collective Intelligence: Combining the Powers of Humans and Machines to Build a Smarter Society*, pages 31–50. Springer Verlag, 2014a. (Cited on pages 5 and 120.)

Jeremy Pitt, Dídac Busquets, and Samuel Macbeth. Distributive justice for self-organised common-pool resource management. *ACM Trans. Auton. Adapt. Syst.*, 9(3):14:1–14:39, 2014b. (Cited on pages 28, 89, and 100.)

Jeremy Pitt, Jie Jiang, and Aada Diaconescu. On the minimal recognition of rights in holonic institutions. In *COIN@AAMAS2016 Pre-proceedings*, 2016. (Cited on page 120.)

Robert D. Putnam. The prosperous community: Social capital and public life. *The American Prospect*, pages 35–42, 1993. (Cited on page 34.)

Anand S. Rao and Michael P. Georgeff. BDI agents: From theory to practice. In *Proceedings First International Conference on Multi-Agents Systems (ICMAS)*, 1995. (Cited on page 18.)

Anatol Rapoport and Albert M Chammah. *Prisoner's dilemma: A study in conflict and cooperation*, volume 165. University of Michigan press, 1965. (Cited on page 29.)

Nicholas Rescher. *Distributive Justice*. Bobbs-Merrill, 1966. (Cited on page 28.)

Jordi Sabater and Carles Sierra. Reputation and social network analysis in multi-agent systems. In *Proceedings of the First International Joint Conference* *on Autonomous Agents and Multiagent Systems: Part 1*, AAMAS '02, pages 475–482, New York, NY, USA, 2002. ACM. (Cited on page 43.)

Jordi Sabater and Carles Sierra. Review on computational trust and reputation models. *Artificial intelligence review*, 24(1):33–60, 2005. (Cited on page 30.)

Theodore W. Schultz. Investment in human capital. *The American economic review*, 51(1):1–17, 1961. (Cited on page 33.)

Matthew Smith. *Social Capital in Online Communities*. PhD thesis, Brigham Young University, 2011. (Cited on page 92.)

Toshiharu Sugawara. Emergence and stability of social conventions in conflict situations. In *22nd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 371–378, 2011. (Cited on page 78.)

Asimina Vasalou and Jeremy Pitt. Reinventing forgiveness: A formal investigation of moral facilitation. In *iTrust*, volume 3477 of *LNCS*, pages 146–160, 2005. (Cited on page 36.)

Asimina Vasalou, Astrid Hopfensitz, and Jeremy Pitt. In praise of forgiveness: Ways for repairing trust breakdowns in one-off online interactions. *Int. J. Hum.-Comput. Stud.*, 66(6):466–480, 2008. (Cited on pages 18 and 91.)

Jun Wu, Xin Xu, Pengcheng Zhang, and Chunming Liu. A novel multi-agent reinforcement learning approach for job scheduling in grid computing. *Future Generation Computer Systems*, 27(5):430 – 439, 2011. (Cited on page 23.)

Chao Yu, Minjie Zhang, Fenghui Ren, and Xudong Luo. Emergence of social norms through collective learning in networked agent societies. In *12th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 475–482, 2013. (Cited on page 78.)